INTERTEMPORAL CONSIDERATIONS IN SUPPLY OFFER DEVELOPMENT IN THE WHOLESALE ELECTRICITY MARKET

A thesis
submitted in partial fulfilment
of the requirements for the Degree
of
Doctor of Philosophy in Management Science
at the
University of Canterbury
by

PAUL ANDREW STEWART

University of Canterbury
New Zealand
2006

To Dad

Michael Armstrong Stewart

27/4/1946 - 28/1/2006

ACKNOWLEDGEMENTS

This thesis would not have been possible without the help and support of many people.

Firstly, thanks must go to my supervisors, Dr Grant Read and Dr Ross James, for their excellent supervision, guidance and support over the past few years. Both have been never-ending sources of new ideas and helpful suggestions whenever needed. In addition, I would like to thank other staff and students of the Department of Management with whom I have developed great friendships. Without the stress relief of the MANGlers indoor cricket games each Thursday and visits to the Foundry on Friday evenings over the past few years, I don't know how I could have gotten through the weeks.

I would also like to acknowledge the financial support of the Foundation for Research, Science and Technology, the Department of Management and Charles River Associates (especially Deb Chattopadhyay), without which this research could not have been completed.

To my family, Mum, Dad, Mark and Natalie, thank you for always supporting me and my education and thanks especially for your help and support over the last 12 months.

Finally, thank you Christina for always being there for me. You have been a great source of encouragement, help and enjoyment.

ABSTRACT

Traditionally, electricity markets around the world have been operated centrally by an agent of the local government, in order to minimise the cost of supplying electricity to all users. However, over the last 20 years, these markets have gradually been deregulated, producing wholesale markets in which generating companies compete for the right to supply electricity, through an offering system. These major structural changes to the operation of the electricity systems have significantly affected the behavioural incentives of the players in the market, and have thus presented a whole new realm of problems for researchers to address. Specifically, generating companies now operate in order to maximise their own profits, subject to various behavioural constraints and other considerations. In particular, these considerations include intertemporal constraints that the generator may face over time, such as fuel limitations and thus conservation, market uncertainty and correlation, uncertain fuel inflows, and unit operational rules.

It is the optimisation of the offering process subject to these intertemporal constraints that we address primarily in this thesis. Some of these constraints and considerations have been considered previously in the literature, but they have never been dealt with simultaneously, or in a computationally efficient manner. We present a multi-dimensioned dynamic program for constructing optimal offers over a planning horizon,

which uses a two-stage approach whose computational time increases at a substantially slower rate than previously published algorithms as the problem size increases, and thus is able to solve more practical problems.

Another significant contribution of this thesis is the combination of dynamic programming with a decision analysis or branching structure, where the branches represent overall *macro-states* that either the market or the generator can be in at any given time. This technique has been developed as a general solution approach, and as such its application is not limited to the offer optimisation problem explored in this thesis.

TABLE OF CONTENTS

List of Figures

LIST OF TABLES

LIST OF ABBREVIATIONS

AL	Attractor Line
ANOVA	Analysis of Variance
BOP	Beginning-of-Period
CfD	Contract for Differences
CIP	Critical Intersection Point
CL	Contributory Line
DADP	Decision Analysis Dynamic Programming
DMVC	Direct Marginal Value Curve
DP	Dynamic Program/Programming
DS	Number of "Down" Unit States/Minimum Down Time
EMV	Expected Marginal Value
EMVS	Expected Marginal Value of Storage
EOP	End-of-Period
EP	Expected Payoff
EPOC	Electric Power Optimisation Centre
EPRC	Electric Power Research Centre
GLM	General Linear Model
HERO	Hydro-Electric Reservoir Optimization
HVDC	High Voltage Direct Current

infl	Inflow
ISO	Independent System Operator
MC	Marginal Cost
MDF	Market Distribution Function
MDP	Markovian Decision Process
MFV	Marginal Fuel Value
MFV	Marginal Future Value
MIP	Mixed Integer Program
MOC	Marginal Opportunity Cost
MPEC	Mathematical Program with Equilibrium Conditions
MPV	Marginal Present Value
MR	Marginal Revenue
MV	Marginal Value
MVR	Marginal Value of Release
MVS	Marginal Value of Storage
MW	Megawatt
MWh	Megawatt Hour
NEM	National Electricity Market (Australia)
NZEM	New Zealand Electricity Market
OCF	Offer Curve Family
P	Price
PO	Pseudo-Offer
PP	Pre-Processing
PV	Present Value
Q	Quantity
R&A	Rajarman and Alvarado
RD	Residual Demand
RDVP	Residual Demand Vertex Path
RT	Real-Time
SD	Shut-Down
SFE	Supply Function Equilibrium

SOE	State-Owned Enterprise
SS	Number of Start-Up States
SU	Start-Up
SUG_{ss}	Required Generation Level in Phase *ss* of the Start-Up Phase
UMS	Uncertain Market Scenario
US	Number of “Up” Unit States/Minimum Up Time
VC	Value Curve
VTG	Value-To-Go

Chapter 1

INTRODUCTION

1.1 Introduction

Since the mid 1980s, centrally managed electricity systems around the world have gradually been deregulated, forming markets in which multiple generation companies compete for the right to supply electricity to meet demand. The idea behind such deregulation is that the introduction of competition will increase incentives to improve operational costs of these facilities and is also likely to produce competitive results, whereby electricity is provided at marginal cost (Gross & Finlay (2000)). However, under a deregulated framework, there is no requirement for the generators to reveal their true cost structures through offers at this level, and hence there is potential to "game" the market by inflating their supply offers above costs, thus forcing the market price above the desired competitive levels (Hao (1999)).

In the New Zealand context, deregulation began in 1996 with ownership of generation capabilities being split between multiple state-owned enterprises (SOEs) and private companies, and contestability introduced to the distribution sector by allowing open access to the network. This produced three groups of market participants; generators, distributors, and retailers. Under this new market structure, the role of the government is to ensure that electricity is delivered in an efficient, reliable, and environmentally sustainable manner to all classes of consumer (Barton (2000a)).

There were many reasons behind the introduction of restructuring and deregulation in New Zealand, in addition to the desire to produce competitive market outcomes. Barton (1998) proposes that one of the foremost motivations was to avoid repeating the mistakes of the past. Debatably, one of these key mistakes was the implementation of the six "Think Big" energy development projects proposed in the late 1970's and early 1980's, by the then Minister of Energy for the National Party, Bill Birch. In 1988, it was calculated by Treasury that these projects had made the nation $1.3 billion poorer. The problem was that uneconomical investments had been made for political reasons (attempting to recover economically from the oil shocks of the 1970's) that would not have been made under commercial scrutiny of a privatised generation company (the clearest example being the Clyde Dam project).

Naturally, market restructuring alters the problems and issues that are faced by the participants, and hence new research is required to explore these issues. Specifically, electricity market optimisation modelling work has traditionally focused on the cost minimisation of the entire market as a whole (for example, Read (1979)). Clearly this goal is no longer appropriate, and as such, the focus of much research work has changed to the optimisation of behaviour by various players in the market.

1.2 Background

As explained in Mco (2003), the New Zealand market is predominantly hydro-based, with hydro generation accounting for 60% of the total generation capacity and the remainder made up through gas, coal, and geothermal resources. Although there are many market nodes, there are two distinct market regions, corresponding to the two main islands of the country (North and South). These regions are joined by a single inter-connector, known as the high voltage direct current, or HVDC, link. Resulting from the deregulation of the New Zealand market, there are eight generation companies (supply-side) along with seven electricity retailers and several large industrial companies (demand-side) competing in the market, with many companies falling into both the demand and supply categories. In every half-hour period of the day, each of the generators provides a set of offer stacks[1] and each of the demand-side participants provides a set of bid stacks to the central market coordinator, for each of many coming periods[2]. From a generator's point of view, for each of these periods they are allowed to provide a five-stepped offer stack for every generation unit that they operate, or alternatively, they can aggregate their offers to a station–level, or even a block-level[3]. These offer stacks are effectively a supply of different volumes of electricity at various prices, which may or may not be accepted by the market coordinator. This acceptance depends on many factors, including electricity demand from the retailers and some large consumers, the offers from competing generators, and certain physical constraints on the electricity network. Once all offer and bid stacks are submitted, the market coordinator, or New Zealand Electricity Market (NZEM[4]), takes all this information for each single

[1] An offer stack is a set of prices at which the generator is prepared to supply successive blocks of electricity to the market.
[2] Offers and bids are made for all periods from 2 hours away from the current time (offers can only be changed within two hours of real-time if there is a good reason behind the change), to the end of the current day if it is before 1pm, or to the end of the following day if it is after 1pm. The only binding offers or bids are those that are made on the nearest period, as there will be opportunities for all the others to be changed.
[3] A block is a set of units or stations that may be dispatched together. In other words, dispatch that is allocated by the central coordinator to a unit within the block may in fact be generated by any of the stations or units within that block. This concept is known as "block dispatch".
[4] The NZEM is a commodity exchange for wholesale electricity trading in the New Zealand market, in which the bulk of electricity generated in New Zealand is sold to large commercial end users and to

period individually and produces a dispatch (or pre-dispatch) schedule to minimise the overall cost of meeting demand for the country as a whole, while accounting for transmission constraints and losses. A simplified version of this model is presented in Alvey, Goodwin, Ma, Streiffert, & Sun (1998). If we ignore these spatial, transmission aspects of the market, and thus consider the entire market to be located at a single market node, then this can easily be explained as the intersection of the aggregated market demand and supply stacks, as demonstrated in Figure 1.

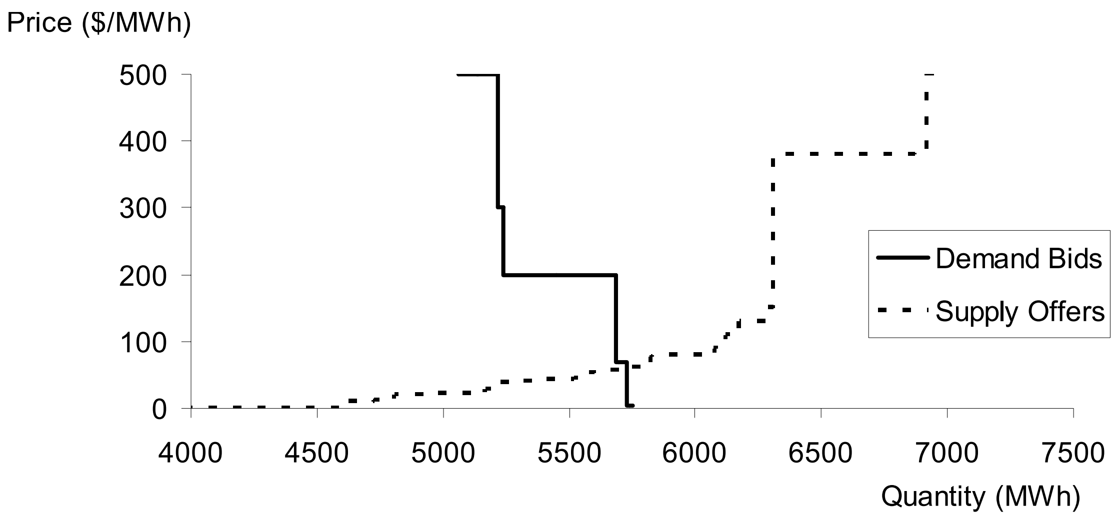

Figure 1.1: Single-Node Electricity Market Clearing Process

Because the market coordinator optimises the market for a single period at a time, any internal constraints faced by the generators must be accounted for within their own offers, contrary to what is the case in some other markets (such as the pre-2001 England and Wales market, as discussed in Gross, Finlay, & Deltas (1998) and Finlay (1995)).

1.3 Scope

When modelling the electricity market, there are several unique features that distinguish it from other standard commodity or financial markets. For example:

---

energy retailers who in turn sell their electricity to their relatively low-volume customers in the retail market.

- Electricity cannot be easily stored for any length of time, and so flows and consumption/supply must be continuously balanced in real-time in order to maintain the correct system frequency.
- Exact directional flows of electricity cannot be decided by a human controller, as they are governed by the laws of physics.
- Ownership of power on communal transmission lines cannot be distinguished.
- In the short-run at least, there is little or no substitution for electricity, so demand is highly inelastic within this horizon.
- Generation investment is very capital intensive, lumpy, and requires a long lead time, which makes entry into the market very risky.
- Reservoir levels in a river system are highly inter-related, with release from one generator resulting in an increased level for the next reservoir downstream.
- There are complicated joint constraints in the network. For example, the use of 1 MW of incremental generation at one location could increase flow on a particular, constrained, transmission line elsewhere in the network. This could potentially block production of more than 1 MW at another location, depending on their relative contributions to electricity flow on the constrained line. Hence, increase in generation at one node in the network can decrease capacity in another.
- Both supply and demand are very weather dependent, often in an inversely correlated manner. For example, colder weather in the winter might increase demand, while inhibiting snow melt and thus reducing supply.

These unique features are discussed in many papers, including Rudkevich, Duckworth, & Rosen (1998), Hogan (1997), Counsell & Evans (2003) and MED (2000).

In fact, the New Zealand market is even more challenging in its uniqueness. For one thing, a significant portion of generation is located far from the areas of major demand, making transmission especially important. Our high hydro dependency (about 65% of capacity) combined with low maximum storage levels of about 15%[5] of national demand or 8 weeks use (as opposed to some countries overseas with years worth of storage capacity) make New Zealand particularly vulnerable to sustained variations in conditions. In addition, the market is small and isolated, as opposed to areas in Europe and the Americas (such as discussed in Hammons, Rudnick, & Barroso (2002)), which have vast interconnected networks between countries.

As previously mentioned, the new goal of electricity market research is commonly the optimisation of behaviour of various market participants, and in particular, of the generators providing electricity into this market. These generators face strategic and operational issues on varying time scales:

Long-term planning models consider time horizons in the scale of many years, and involve decisions about the commission of new plant, in addition to other strategic issues.

Medium-term planning models consider time horizons in the scale of many weeks or months. Decisions may be either strategic or operational, and may involve the overall plan for the use of limited fuel resources or commitment decisions for inflexible plant.

Short-term decision making models consider time horizons in the scale of hours or days, and decisions involve the specific commitment and offering decisions of all generation units under their control.

This thesis addresses the optimisation of short-term offers provided to the market by generators with fuel restrictions, both hydro and fuel-limited thermal. In particular, we deal with many of the intertemporal issues and constraints that they face, including start-up and shut-down restrictions, minimum up and down times, stochastic rest-of-market behaviour, and the allocation of their limited fuel over the planning horizon. Our approach is to develop a stochastic dynamic program to optimise the offering strategy

[5] Source: Tipping, Read, & McNickle (2004)

for the planning horizon, subject to market uncertainty. In the modern market context, it is of great importance to the generation companies to develop such tools to help select the offer stacks that will lead to the greatest profits in any given circumstance.

In order to derive advanced analytical results in the optimisation of offering behaviour with respect to intertemporal issues, various simplifications are assumed throughout this thesis. The main simplification is in our exclusion of spatial effects. These spatial effects can be split into two distinct groupings; locational-type spatial effects in the transmission system and hydro-type spatial effects, which can possibly exist at single nodes of the network. Locational effects relate primarily to transmission constraints and losses in the network, Kirchoff's Voltage and Current Laws (relating to the physics of electricity flow, and explained in Hobbs, Metzler, & Pang (2000)) and nodal pricing. This thesis makes the assumption that transmission is lossless and unconstrained, thus ignoring these locational considerations. As such, potential for congestion gaming (as reported by Seeley, Lawarree, & Liu (2000)) is also not considered. The hydro-type spatial effects relate to flow-of-river considerations. Physical distances, as such, between consecutive hydro-reservoirs or generators need not be considered, however the time (or flow) delays between consecutive reservoirs are important. Li, Hsu, Svoboda, Tseng, & Johnson (1997), for example, reports that the release from one reservoir affects the capacity of downstream reservoirs and therefore that next unit's generating options, depending on flow delays. This relationship is not perfectly correlated; that is, some of this downstream generation can be deferred or brought forward as preferred by the owner of the system. This is a well-defined area of literature known as reservoir management, and as such is addressed only briefly in this thesis. Additionally, we consider energy-only offers, ignoring the complicating factors of the markets for reserve and other ancillary services, which are discussed in David & Wen (2000).

1.4 Research Contribution

Despite the importance of intertemporal issues in the electricity offering problem, which will be demonstrated in this thesis, relatively little work in the literature has been devoted to dealing with these issues in conjunction with market uncertainty. This thesis presents two new stochastic dynamic programming algorithms that optimise the supply behaviour over a short-term planning horizon of a generator with limited fuel, where the stochastic rest-of-market behaviour and pricing outcomes is correlated over time. The difference between the two algorithms is in the way that they consider correlation between market states in consecutive periods.

The first algorithm considers a single Markov chain for each period, which provides the probability of observing a particular residual demand curve in this period, given that we know the residual demand curve that occurred in the previous period. The main contribution of this algorithm is in the application of a technique known as marginal cost patching. This technique enables us to separate the relatively straightforward, but highly inefficient two-level dynamic programming algorithm into a significantly more efficient two-phase dynamic program which has computational times that are practical with respect to real-life application.

The second algorithm that we develop combines the methods of dynamic programming and decision analysis in a new approach with a wide range of potential applications, beyond the offering problem faced by generators that is primarily addressed in this thesis. The idea of the technique is to partition the state of the system of interest at each stage into discrete groups of possible system states, each associated with a particular *macro-state.* These macro-states represent overall, high-level states of the system at the given stage, where transitions between these macro-states are either probabilistic or based on an internal decision by the decision maker. With respect to the application of this technique to the electricity offering problem, sets of residual demand curves that could occur under a particular macro-state are grouped together. These macro-states could be based on external market scenarios (such as a significant interconnector being binding or non-binding) or on internal operating decisions that the generator has made (offering aggressively or defensively, for example). Under the external scenarios, water value curves for the alternative future macro-states are combined probabilistically at the point in time where the uncertainty is to be resolved, while under internal decisions, water value curves are combined depending on the optimal decision from each reservoir level at the point in time where the decision is to be made. Dividing the overall set of possible residual demand curves in the method proposed by this second model has many benefits, most notably that of an improved representation of the uncertainty and decision structures in the real-world.

In summary, the key contributions of this thesis are:

- The development of a marginal cost based offer patching theory, which can significantly reduce the computational complexity of the offer strategy development problem. Based on this theory, we will develop and implement a two-phase dynamic programming algorithm, which separates the construction of offers into separate Pre-Processing and Real-Time phases, in a manner not previously suggested in the literature.
- We will determine theoretical extensions to this algorithm to deal with a combination of intertemporal constraints and considerations, including market correlation and uncertainty, fuel limitations, and unit rules (including ramp rates) that have not previously been considered together, while still maintaining feasible computational times, with respect to realistic marketplace requirements. These extensions will also be implemented.
- The development of a new technique that combines the concepts of dynamic programming and decision analysis by building a branching structure into the overall state of the generator and the market at any point in time. As well as being able to better represent a scenario that has this type of structure, this branching will significantly reduce the computational complexity of the algorithm.

1.5 Outline of Thesis

As detailed in Section 1.3, the focus of this thesis is on the development of offering strategies by electricity generators, particularly dealing with intertemporal issues and constraints that they face in the real-world. Chapters 2 and 3 provide details on the context of the problem and its associated literature, reviewing electricity market structures, along with factors that affect supply behaviour, including intertemporal issues. Chapters 4 and 5 address the problem of constructing optimal single-period offers under market uncertainty, from an analytic perspective. In particular, Chapter 5 looks at how optimal offer segments corresponding to fixed marginal cost levels can be patched together to produce the optimal offer for a generator with a stepped marginal cost curve.

Chapter 6 shows how the marginal cost patching concepts can be used to separate the obvious dynamic programming algorithm for offer construction into a pre-processing and a real-time phase, with improved computational efficiency over those currently presented in the literature. The algorithm is designed to consider such complexities as correlated stochastic market outcomes and limited fuel availability. Chapter 7 develops a Decision Analysis Dynamic Programming (DADP) framework and applies this new technique to the problem of electricity market offer optimisation. The technique enables us to better represent the decision and uncertainty processes of the real-world, when the electricity system can at any given point be considered to be in one of a set of high-level "macro-states". Chapter 8 presents a detailed experimental design and tests the two algorithms that are developed in Chapters 6 and 7, first against a model from the literature that considers similar complexities, and then against simplified algorithms, to determine the value of considering each of these complexities.

Chapters 9 and 10 then address various possible extensions to the algorithms that are presented in earlier chapters. In particular, Chapter 9 presents modifications to the algorithms to deal with a further set of intertemporal constraints, known as generation unit operational rules. These include:

- Minimum feasible generation level
- Start-up and shut-down costs
- Minimum up and down times

- Fixed start-up process
- Ramp rate restrictions

Experimental results under these models are also presented, confirming hypotheses regarding offering behaviour presented in Chapters 3 and 4. Chapter 10 considers how we might go about dealing with more difficult extensions which are considered beyond the scope of this thesis, including multiple reservoir problems, and gas market interaction.

In summary, the remainder of this thesis is organised as follows:

- Chapter 2 provides context to the electricity market offer optimisation problem, including a discussion of intertemporal issues and their importance.

- Chapter 3 outlines the literature regarding optimisation of offering behaviour, with an emphasis on approaches to offer construction over short-term horizons that deal with various intertemporal issues.

- Chapter 4 looks at how optimal monotone offers for a single period can be constructed analytically under various forms of market uncertainty.

- Chapter 5 develops a marginal cost patching technique, which enables segments of existing optimal offers to be patched together to create the new optimal offer under changed circumstances.

- Chapter 6 shows how the marginal cost patching concepts can be used to separate the obvious dynamic programming algorithm for offer construction into a pre-processing and a real-time phase, with improved computational efficiency over those currently presented in the literature.

- Chapter 7 develops a Decision Analysis Dynamic Programming (DADP) framework and applies this new technique to the problem of electricity market offer optimisation.

- Chapter 8 presents computational results from the two algorithms developed in Chapters 6 and 7.

- Chapter 9 presents extensions to the algorithm presented in Chapters 6 and 7 to deal with the intertemporal constrains implied by generation unit operational rules.

- Chapter 10 outlines possible extensions to the algorithms presented in the thesis, to consider complexities such as multiple reservoirs, and gas market interaction.

- Chapter 11 draws conclusions and proposes areas for future research.

Chapter 2

ELECTRICITY MARKETS, INFLUENCING FACTORS ON SUPPLY AND INTERTEMPORAL ISSUES

2.1 Introduction

Many electricity markets around the world are currently heading in the direction of deregulation, and are thus facing major decisions on the structure of the new markets that they will create. These decisions will have a direct influence on the offering strategy that a participant will choose to employ, and thus their potential to gain market power. Therefore, we begin this chapter with a discussion of the various types of market structures and their likely influence on market power. This discussion will also provide a context for the offer optimisation problem faced by generators in the New Zealand market. We then go on to explore the various factors that influence supply and demand decisions, before identifying and explaining the different intertemporal constraints and issues that are faced by generators providing electricity into a market of the described structure.

2.2 Market Structure and Market Power

2.2.1 Market Structures

Throughout the world there are, and have been, many different deregulated market designs implemented and/or proposed, with ongoing debate about which particular design will obtain the closest real-world results to the ultimate goal of a completely efficient generation process (for example, Bower & Bunn (1999), Drayton-Bright (1997) and Denton, Rassenti, & Smith (2001)). This complete efficiency refers to the situation whereby the aggregate supply curve accurately reflects system marginal costs and thus the load is being met with the lowest total generation cost. The market governance body should be designed to minimise the loss of efficiency that occurs when we stray from the merit-order dispatch[6], as reported by Baron (1972). For example, the market coordinator gets the impression that the generator offering the lowest price is the most efficient generator, but if all generators are offering at above marginal cost, then it may be that they are just the most risk averse firm, trying to ensure their own dispatch.

Markets should also be designed to give pricing signals to encourage consumers to conserve electricity at the most important times to do so, and to encourage generators to build new capacity where and when it is forecast to be needed and not before[7]. The concept of nodal pricing (having a different price corresponding to the dual of the energy conservation constraint at each node in the transmission network) is one method towards giving correct operational pricing signals and signals for siting new plants (accounting for all costs and transmission constraints).

[6] Merit-order dispatch is where electricity with the lowest generation cost, rather than electricity with the lowest offer price is dispatched first
[7] Barton (1998)

Some of the key differences between various market designs are as follows:

- Pool vs. Bilateral Exchange Markets. In a pool market, the Independent System Operator (ISO) determines all dispatch levels at a fixed market clearing price, whereas in a market based on bilateral exchanges, each participant must arrange its own trades at a negotiated price. Singh, Hao, & Papalexopoulos (1998) suggests that the bilateral market is motivated by free market competition as opposed to the pool, which is motivated by the need to include special characteristics of electricity networks such as transmission constraints, in the trading process. Ignoring these characteristics is a problem with the bilateral system as, for example, individual firms have no reason or incentive to consider the transmission constraints. A type of financial instrument known as a contract for differences (or CfD) gets around this problem by enabling bilateral-type trades in the pool system
- Double Auctions vs. One-Sided Auctions. A one-sided auction is one where the demand is considered to be completely inelastic (that is, elasticity = 0), and thus the total quantity traded will be independent of the shape of the aggregate market supply curve (only market price will be affected). On the other hand, a double auction allows demand-side bidding in addition to supply-side and so both quantity traded and market price are dependent on the shape of the supply curve. Demand-side bidding is investigated by Lamont & Rajan (1997) and Barton (1998), with respect to the limitations on efficient bidding due to lack of information. Barton suggests that the ultimate goal of disaggregation of electricity industries is to allow homes and businesses to shop around for electricity and benefit by changing consumption patterns. To do this would require smart metering and an ability to convey greater levels of information to these customers in real-time.
- Static vs. Dynamic Auctions. This differentiation refers to whether the offer stack information is sealed permanently or if the market participants can observe their rival's offers and update their own offers accordingly for the next (repeated) round of offering.
- Required structure of the offer stacks (that is, complex offers vs. single-part offers). In some markets such as the old, pre-2001 England and Wales market, offers were

requested in a complex 27-part form which included offers for start-up costs and other complex features. However, in markets such as that of New Zealand, Australia and Norway, offers are submitted in the form of a single-part offer stack with a single price offered for each block of energy. Generators in these markets must therefore internalise into their offers all the complex information that is defined explicitly in the complex offers[8].

- Offer Submission and Adjustment Timelines. Some markets (again, such as the old, pre-2001 England and Wales market) require that a single offer stack is submitted for each generating unit the day before its associated period, and that this offer will remain in place for 24 hours. The offer can only be adjusted in the case of a breakdown (such changes are scrutinised by the market operator) or under special circumstances such as unit start-up. Alternatively, other markets allow specification of separate ("sculpted") offers for each period of the corresponding day, which can be changed at any time up to not long before the real-time dispatch model is run. This changing of offers is known as re-offering, and may be restricted in the parameters of the offer function that you can change (that is, by allowing changes to either the price level, or quantity offered). This re-offering may also be restricted in terms of the variation allowed between consecutive offers, in a manner such as that proposed for market activity rules by Wilson (1997). These rules help to ensure that correct pricing signals are given in the early offering iterations (that is, without them, there is little incentive for a generator to give a meaningful offer until the latest possible time).
- Discriminatory vs. Uniform Pricing. Discriminatory pricing means that generators are paid the amount at which they offered their dispatched quantity (leading to much flatter offers, which attempt to estimate what the market clearing point will be), whereas uniform pricing pays all generators the market clearing price for their dispatched generation quantity (adjusted for marginal loss factors).

[8] There is actually a whole spectrum of possible market requirements. In the NZEM for example, the market clearance model does actually take into account forward-looking ramp-rates of generation units.

To put these market types into context, both the New Zealand and Australian electricity markets (NZEM and the NEM, respectively) operate a pool system, with a double auction (that is, demand-side bidding is possible), allowing re-offering of the offer stacks up towards real-time. Both the markets are repeated dynamic games, whereby dispatch and pre-dispatch information is released in real-time back to the market. Although both use single-part offers, the NEM requires specification of ten price steps, whereas the NZEM allows only five.

Many authors have published work on this topic based on the markets of various countries, such as Barton (2000a) (New Zealand – structuring of market governance bodies), Bower & Bunn (1999) (England/Wales market – bilateral trading vs. pooled system), Sakk, Thomas, & Zimmerman (1997) (Generic market – iterative offering vs. single offers), etc. There have also been many studies into the way that markets should be or should have been restructured (for example, Sheble (1996a), Drayton-Bright (1997), Bower & Bunn (1999) and Drayton-Bright (1997)).

2.2.2 Determining Market Dispatch and Pricing

Market dispatch, transmission and pricing decisions in most pool markets are made by the central coordinator or ISO, by solving a large linear program in order to minimise the total cost of meeting demand at all the system nodes. This optimisation is performed given the offer stacks (supply) from the generators and the demand bids from electricity retailers and some large firms or institutions (aluminium smelters, for example). If transmission issues are set aside, this is simply explained as the coordination of aggregate demand and aggregate supply (both step functions). This concept is demonstrated in Figure 2.1, whereby the heavy blue line represents the combined offer stacks of the generators in the market, and the light blue line represents the combination of all demand. The large black dot denotes the dispatch point at a market clearing price of $15, and a total quantity traded of around 39 MWh's.

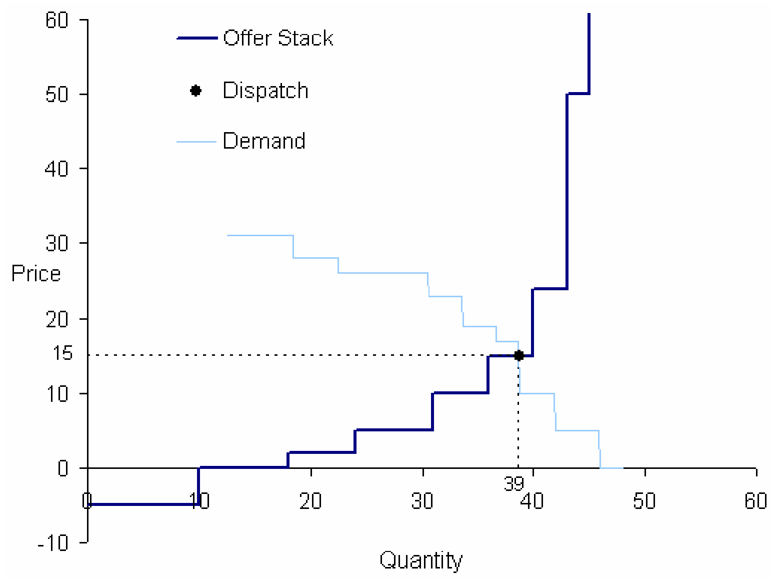

Figure 2.1 Example of dispatch by equating aggregate demand and supply stacks

Note that in a centralised market, the total cost minimised by the dispatch model is based on actual costs of generation, whereas in a deregulated market, the total cost is based on prices submitted by generating units, which may well be above their generation costs for the corresponding blocks of power. It is also important to note that although only around 10% of all power is traded at the market clearing price as determined above (the rest being accounted for with pre-priced contracts for differences), the spot price has a great impact on the level of new contract prices.

2.2.3 Market Power in an Oligopoly Market

The literature generally considers the electricity market to take the form of an oligopoly (or duopoly in the case of the England and Wales market) due to the limited number of producers, the transmission constraints, the large investment costs, and transmission losses (which discourage some seller/buyer trades). Examples of papers that explicitly make this assumption are Klemperer & Meyer (1989), Garcia, Campos-Nanez, & Reitzes (2005), Newberry (1991) and Green (1996). Due to this structure, operators can be expected to wield some degree of market power and can therefore expect to profit from offering in a strategic manner, either by pushing up the market price or by withholding capacity. The extent to which this market power exists depends greatly on the structure of the given electricity market.

Hogan (1997) defines this market power as "the ability to profitably maintain prices above the competitive level by restricting output to lower than competitive levels". According to Culy & Read (1994), there are numerous economic consequences from gaming (or the use of market power), including the following:

- The structure of spot prices may not reflect the structure of marginal costs (bad signalling),
- The benefits of hydro coordination and merit order operation are compromised,
- Long-term contract prices may not reflect the long-run marginal costs of investment in new supply, and
- Incentives for entry may be compromised.

In addition, Anderson & Philpott (2001) note that most classical financial asset-pricing models are unable to deal with electricity spot prices, as they assume price is exogenous to the offering decisions that are made. Therefore, they cannot be used to optimise offers of a generator with sufficient market power to affect prices (such as an oligopolistic firm).

Due to the negative consequences of market power described in Culy & Read (1994), markets should be designed in such a way to mitigate the potential use of market power as much as possible. Throughout the literature, many factors are identified as having potential to affect the degree to which market power can be exercised. These include the following:

- Market concentration (as defined by the Hirschman-Herfindahl Index) and nature of the generation break-up (Rudkevich, Duckworth, & Rosen (1998)).
- Market design, structure and rules (for example, the required offer form), Rudkevich, Duckworth, & Rosen (1998).

- Elasticity of demand/Presence of demand-side bidding (Rudkevich, Duckworth, & Rosen (1998)).
- Nature of the product (electricity has unique features that can be exploited), Rudkevich, Duckworth, & Rosen (1998).
- Level of contracts held by the generation company. Both Kelman, Barroso, & Pereira (2001) and Villar & Rudnick (2003) demonstrate that market power is largely mitigated by adding physical/financial contracts (changing the objective function).
- Capacity constraints of generators (Kelman, Barroso, & Pereira (2001)).
- Threat of not being scheduled (Rudkevich, Duckworth, & Rosen (1998)).
- Presence of a regulator (Green (1996)).
- Transmission bottlenecks/constraints and their locations, as well as transmission pricing (Hobbs, Metzler, & Pang (2000)).
- "Caution factors". Culy & Read (1994) explains that it is not worth it to game for a small possible return, as there is risk involved and the firm's reputation may be harmed.

Many authors have researched this area, looking either at a) factors that affect the level of market power wielded by a firm, or b) choosing between a set of market structure choices in order to minimise the ability of firms to exercise market power. Rudkevich, Duckworth, & Rosen (1998) took the first approach, finding that more generating firms, higher capacity availability, and increased accuracy of demand forecasts lead to reduced potential for market power. In addition, to other papers take the first approach. Barroso, Fampa, Kelman, Pereira, & Lino (2002) consider a dynamic market model, where hydro generators attempt to maximise their revenues over multiple periods, finding that prices and revenues were increased by reducing output and decreasing water transfers from wet to dry seasons (compared with a centralised operation policy). Green & Newberry (1992) compares market power over a range of demand elasticities, and also under the assumptions of symmetric and asymmetric duopoly firms, finding that asymmetric firms lead to higher market prices than do symmetric. His analysis suggests that the UK government had seriously underestimated the scope for exercising market power, prior to their restructuring. Alternatively, Green (1996) took the second approach, comparing three policies for decreasing market power in the duopolistic UK market (partial divestiture, breaking up of large firms, and encouraging entry). He found that the first option is the best, as the second is politically infeasible and would lead to a loss of economies of scale, and the third option is inefficient if done in advance of need. Culy & Read (1994) also took this approach, comparing outcomes for different rules regarding the day-ahead form of offers (sculpted over each period versus. fixed for the whole day) over various contracting levels. They find that the sculpted offers actually appear to result in reduced gaming due to the limited ability to game up prices without explicit collusion.

At this point, gaming and other effects on offering strategies should be distinguished. A generator is said to be “gaming the market” when it is attempting to influence the market clearing price so as to increase its own profits. This is essentially what the generators are attempting to achieve through the optimisation of their offers and re-offering process. Such gaming, which could be described as exploitation of market power, is a normal part of business behaviour. In an electricity marketplace, a regulator should only be concerned with gaming when it is suspected that a generator is abusing their market power. At an extreme level, market power abuse (or anti-competitive behaviour) may involve such strategies as capitalising on transmission constraints to raise regional prices to levels at which super-normal profits are achieved, or temporarily lowering prices to such levels that force competitors out of the market and discourage entry of new competitors. Although market power exploitation is legal, and abuse is not, a line between the two situations is very difficult and arbitrary to draw. Analytically thinking, it would be convenient to distinguish between any kind of gaming intended to increase profits and other drivers of offering patterns, such as a generator’s attitude towards risk, or internal linkages, which may imply similar effects for offers, but should not be considered gaming. For a discussion of the potential for market power abuse within the New Zealand market, refer to Murray & Stevenson (2005).

2.3 Determinants of Supply and Demand

When a generator provides an offer into a pool electricity market, their resulting generation level, price received, and thus payoff, will clearly depend on the level of demand that occurs, in addition to the level and shape of the supply offers from the rival generators. Hence, in constructing optimal offers, these external decisions and the factors that influence them must be considered.

2.3.1 Determinants of Demand

Electricity demand is highly variable over different times of the year, different periods of a single day, and different days of the week. For example, demand in New Zealand peaks in the winter months as heating equipment is used, as opposed to Queensland, where demand peaks in the summer months as air conditioning use increases. Also, peak daily demand is usually around 6 – 6:30pm, as people arrive home from work, start to cook their dinners, heat/air condition their houses, and so on. In the short-run, demand is highly inelastic as there is very little time to move to other fuels such as gas, or for the general user to react at all. However, some large companies do bid demand stacks to the market and are thus sensitive to spot market prices (David & Wen (2000)). The general consumer is not affected by the spot price, as they are on fixed retail contracts, and so have no incentive to substitute their load away from periods of high spot prices (Barton (1999)). Therefore, determinants of demand levels include the time of year, time of day, weather, events, economic growth (long-term demand), and the availability of alternate fuels. There already exist many established demand forecasting methods, and as such, this thesis does not aim to contribute in this area.

2.3.2 Determinants of Supply

When constructing offers, there are a vast number of factors that need to be considered. Of course, the objective of the generating firm will determine the purpose for which offers are constructed[9]. Throughout this thesis it will be assumed that the generators are rational, risk neutral and incumbent (established) firms who seek to maximise their expected long-run profit. The various determinants of supply can be separated into those that affect the long-term supply strategy, and those which affect either the medium-term day-to-day planning problem, or the detailed hourly offering problem that the generators face. Some factors may fall into multiple categories.

Long-Term Strategic Supply Determinants

- Risk attitudes. A more risk averse generator would make lower-priced offers (they may even self-commit[10] their contracted generation level and remove the possibility of their offers being on the margin) - Mielczarski, Michalik, & Widjaja (1999a).
- Gaming attitudes. For example, a government-owned generation company may be less inclined to attempt to game the market than a privately-owned company[11].
- Threat of regulatory interventions. This relates to whether there is a credible market power monitoring process in place. In other words, does the regulator "punish" generators for behaving in an undesirable manner (from a market efficiency point of view)? For example, this is discussed in Barton (2000b).
- Threat of new entry. There are many relevant factors that will determine the ease of entry for a new generator into the market. These relate to whether the political environment facilitates the new entry and whether existing companies can make life hard for the new entrants (for example, if a monopoly fuel supplier also owns a major existing generator). This is discussed in Newberry (1991)
- Vertical market power issues. If a generation company also owns retail or fuel sources[12], it may have implications on its wholesale offering strategy (Sheble (1996a)).

[9] Many possible objectives can be hypothesised, such as maximising short-term profit, maximising expected value of profit in the long-run, recovering fixed costs, satisfying contract load, establishing a market share or presence, and signalling to competitors.
[10] Self commitment involves a generator offering part of its capacity at a very low negative price (e.g. -$5000/MWh) to ensure that at least this level of their capacity is dispatched.
[11] As a government-owned company are expected to be acting in the interests of the nation as a whole, rather than purely in the interests of their own profits.

- Supply mix of peakers[13], intermediate and base load plants[14] and the associated fuel prices. For example, South Australia has a lot of high cost peakers while Victoria has predominantly base load generation.
- Location of the generators. It may be the case that a generator can take advantage of transmission bottlenecks, as discussed in Singh, Hao, & Papalexopoulos (1998) and Borenstein, Bushnell, & Stoft (2000). For example, if generation from the unit is essential to curing low voltage problems, the generation company may have a stranglehold on the market. This depends on the market design (e.g. there may be separate must-run reliability/voltage type arrangements in place for the system operator to be able to control the dispatch of the unit).
- Ancillary service (reserve) issues. There may be implications on a generator's energy offering strategy if they are a critical player in the reserve market. These implications are considered in both Singh (1999b) and Wen & David (2002)

Medium-Term Planning Supply Determinants

- Maintenance schedules. If a unit is broken down or having maintenance performed, then generation is unlikely from this unit, and will thus limit the capacity that can be offered (Mielczarski, Michalik, & Widjaja (1999a)).
- Forced outages of generation capacity or breakdowns. This can truncate the aggregate offer stack by removing one or more of its (possibly infra-marginal) steps, potentially increasing the market clearing price.
- Resource consents. Some hydro generators may be required to generate between certain levels for environmental reasons, such as to keep within reservoir level bounds.
- Technical limitations (maximum generation levels, ramp rates, etc - Mielczarski, Michalik, & Widjaja (1999a))

---

[12] Joint retail-generator ownership is a major issue in New Zealand and Australia. For example, Tarong Energy in Australia owns both a critical gas pipeline and generators at the two ends of the pipeline.

[13] Peakers are high-cost generation units that operate only in times of peak demand and thus high prices.

[14] Base load plants are low-cost generation units that should operate almost constantly, no matter what the market price.

- Market experience. Experience in the market may mean that a generator is able to recognise certain games being played by rival participants or possibly have greater knowledge of how demand patterns change.
- Cost/availability of fuel. Availability of fuel is a big problem in New Zealand at present. Water reservoir levels are low due to lack of rainfall, affecting hydro capacity, coal stocks are too low to keep up with the desired levels of coal-fired generation, and it has recently been discovered that the Maui gas field is likely to run out of gas much sooner than expected. Because of this factor, prices in a hydro-dominated market such as that of New Zealand or Norway tend to be highly correlated with weather patterns and thus hydro inflows. This issue is addressed for a generator in a deregulated market in Rajaraman & Alvarado (2003).
- Intertemporal considerations of expected price variation. For example, if a hydro generator expected very high demand and thus high prices later in the day, it may want to hold back its limited generation capacity until that time, if it has insufficient capacity to supply in all periods. Although, if all generators were to take this approach, of course the supply in that forecasted high price period would increase, resulting in lower prices. This is an issue that should be accounted for.
- Offering strategy of rivals. This is of course a major determinant of a generator's strategy, as it is these rival strategies, in combination with the market demand and their own strategies that determine the market dispatch quantities and prices, which determine profit levels. If a given generator were to persist with a given strategy over an extended period, it must realise that the rival's strategies may respond accordingly. This is widely accepted in the literature, in papers such as Hobbs, Metzler, & Pang (2000), and Ferrero, Shahidehpour, & Ramesh (1997b).

Short-Term Operational Supply Determinants

- Recent market clearing prices. High-priced offers tend to follow periods with high prices[15] (Mielczarski, Michalik, & Widjaja (1999a)).

[15] Note that this point brings up the issue of causality. In other words, while high spot prices do drive up bids for future periods, in the reverse, it is also true that high bids cause high spot prices.

- Predicted peak (MW) and energy (MWh) patterns. A generator is likely to offer at higher prices if they expect higher demand, or that rivals will offer less capacity (Mielczarski, Michalik, & Widjaja (1999a)).
- Government policies such as regulatory involvement in the form of price-controls/caps, or actual market restrictions in terms of the shape of supply offers (Barton (2000b)).
- Generation costs.
- Contract levels. Many studies have looked at the effect of differing contract levels on the optimal choice of offer stack. Mielczarski, Michalik, & Widjaja (1999a) believe that contract levels are in fact one of the main determinants of an offer. It is clear that as contract levels increase, the generator actually receives the market clearing price for less of its dispatched quantity and therefore the incentive to raise prices decreases. In particular, Wolak (2000) demonstrates that in the extreme case where a generator is over contracted, then they may actually gain by encouraging the market clearing price below their own marginal costs. In this situation, they gain because they are required to buy power from the spot market and re-sell it at their contracted price (thus, the lower the market clearing price, the better off that they are).

In addition, if a generator has a positive probability of being on the margin[16], they have an incentive to increase the price on the corresponding offer step, in order to increase the price received for all infra-marginal units[17] supplied. Thus, the generator is trying to select prices for each of its generation blocks in order to balance the gain from getting a higher price for all units dispatched against the increased probability of not being dispatched at all on the given block. Similar incentives apply to non-marginal offers, whereby reducing the quantity offered at lower (non-marginal) prices may force the dispatch point to a higher step on the aggregate supply curve and thus result in a higher price level (see Figure 2.1). Reflecting this, Read (2001) notes that the role of offer curves is to, in a given period, automatically deal with variations around the expected

[16] On the margin means to be providing the final and most expensive unit of electricity and thus setting the price.
[17] Infra-marginal units are those that do not set the price

conditions. Therefore, beliefs about likely variations in the residual demand[18] will have an impact on the shape of the curves. Additionally, David (1993) states that "prices offered by suppliers for blocks of generation reflect the portion of load that the supplier hopes to win".

Importantly, we note that many of the supply factors considered for the medium and short-term decisions are effectively moving the short-run marginal cost curve for generators with fuel limitations. For example, if particularly high prices have been observed in recent periods (or expected hydro inflow levels fall), then the expected level of available fuel will be lower than anticipated, thus increasing the value that should be placed on the fuel that remains.

Supply Determinants to Consider

While it is important to take into account as many of the above factors as possible, in any model created of a deregulated electricity market the factors taken into account will depend heavily on the perspective from which the research is approached. This thesis focuses on the medium and short-term issues in order to create models that can develop good offers in a real-time scenario from a generator's perspective (thus potentially also enabling regulators to understand how and why the generators behave in the way that they do).

2.4 Intertemporal Issues

The focus of this thesis is on the impact of intertemporal linkages[19] and considerations on participant supply offers, particularly with respect to those linkages that arise in hydroelectric systems such as that of New Zealand or Norway, or with inflexible fuel-limited thermal units used elsewhere.

[18] Residual demand is the net demand function obtained by subtracting the supply offers of rival participants from the market demand function.
[19] The term "intertemporal linkages" is taken to mean the connections between multiple periods of the electricity market, particularly with respect to supply offers made by the generators.

The nature of these intertemporal linkages is diverse, and can be broken into three categories:

1) Linkages within a generating company. This relates to technical costs or constraints of the generating units, such as start-up and shut-down costs, ramp rate restrictions and minimum up- and down-times (these issues are more of a factor for thermal generators). In addition, for a hydro generator, linkages also relate to water considerations such as limited total water availability over time, flow delays between consecutive hydro reservoirs, and the need to account for reservoir balance constraints, given stochastic inflows over time.
2) Effects due to forecasting of market clearing prices and system load. This relates to efficiently planning and preparing for upcoming peak price/load periods. For example, if a hydro generator forecasts that the price will be high later in the day, then it may wish to hold back its (limited) generation capacity until that time.
3) Reactions of other participants in the offering process. Following from point 2, a particular generator may be able to estimate how its rivals will be forecasting prices and load, and behaving as a result. There is also a much more subtle linkage that involves participants hypothesising about the rival's response to the offers that it might make during the re-offering[20] process as real-time is approached.

Song (2000a) partially sums up the motivation for this thesis when he says "The general literature proposes that the optimal strategy is one that maximises expected return for one bidding period. But, in the daily electricity market, the decision makers bid may influence the future market or their own market position in the future". In other words, the current period's offer will affect the state of the market in the future and the generator's own future offering capabilities. Given that market participants are rationally concerned with expected payoff in the long run rather than pay-off in a particular period, such intertemporal linkages are very important.

[20] Re-offering is the process of re-submitting theoretical offers to the market coordinator, up until the time that the market is going to be cleared (at which time, offers becoming binding).

While previous work has considered some of the above-mentioned features (for example, Rajaraman & Alvarado (2003), Philpott (2004) or Scott (1998)), no published research has successfully dealt with any of these features comprehensively, while also considering market uncertainty. The particular issue that this thesis addresses is that of the practical construction of optimal offers of electricity in the short-term from hydro or fuel-limited thermal stations, while considering correlated and uncertain market scenarios, and the various physical intertemporal costs and constraints that these units face. The remainder of this chapter presents some examples of the importance of dealing with intertemporal issues in the offering context.

2.4.1 The Importance of Intertemporal Issues

For both hydro and thermal generation units, the importance and significance of intertemporal characteristics and considerations can easily be demonstrated. Below, we demonstrate through a couple of simple examples how ignoring such considerations can be detrimental to the performance of a generating company. These examples consider deterministic scenarios for ease of comprehension, but the results apply equally (or possibly even more so) to stochastic scenarios. This is demonstrated later, in Chapter 10 of this thesis. Also, they make the assumption that beyond the monotonically non-decreasing requirement, there are no restrictions on the form of the offers.

Non-Monotonicity of the Optimal Offer Curve

It has been shown in Drayton-Bright (1997) that for a hydro operator it is not possible to construct a monotonically non-decreasing supply function that will achieve an efficient outcome over multiple periods. This is due to the linked nature of the periods, and therefore means that intertemporal linkages make the electricity offering problem non-trivial even before gaming or risk are considered. The following example demonstrates why this is the case.

Figure 2.2 illustrates two potential price paths over a 48 period horizon, while Figure 2.3 provides the associated optimal generation paths[21] for a single hydro generation unit, taking into account many types of intertemporal constraints. Note that the price paths account for the fact that the unit's generating level will affect the price, using fairly simple linear residual demand curves for each period. Specifically, in Figure 2.2, the blue line shows a long-run expected daily price path. Consider that we observe a price in the current period (period 1) near the start of the day that is above the long-run average. If we assume that this means we expect it to be a high price day (a reasonable assumption), then our estimates of price will increase for all periods in the horizon, up to the pink line. This particular example assumes that if the price is *x*% above the long-run average in this period, then it will be *x*% above the long-run average for all periods remaining in the horizon, and thus the magnitude of the increases will be more extreme for peak periods.

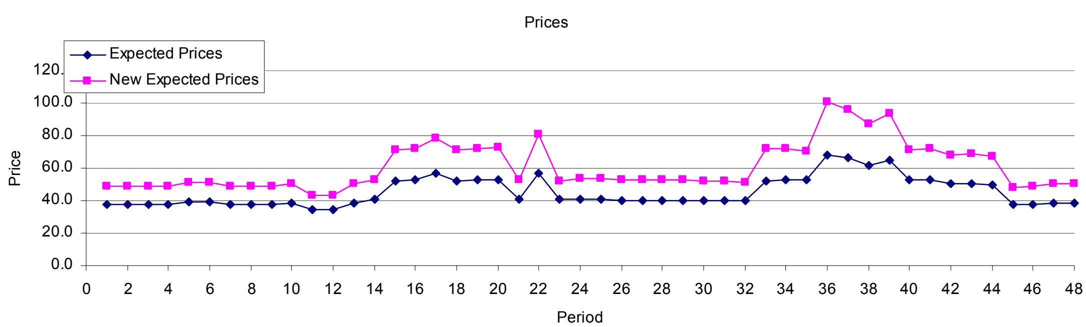

Figure 2.2 Two Potential Price Paths over a 48 Period Horizon

[21] These paths have been determined through the use of a deterministic dynamic program that takes into account limited fuel, in addition to many intertemporal unit operational rules.

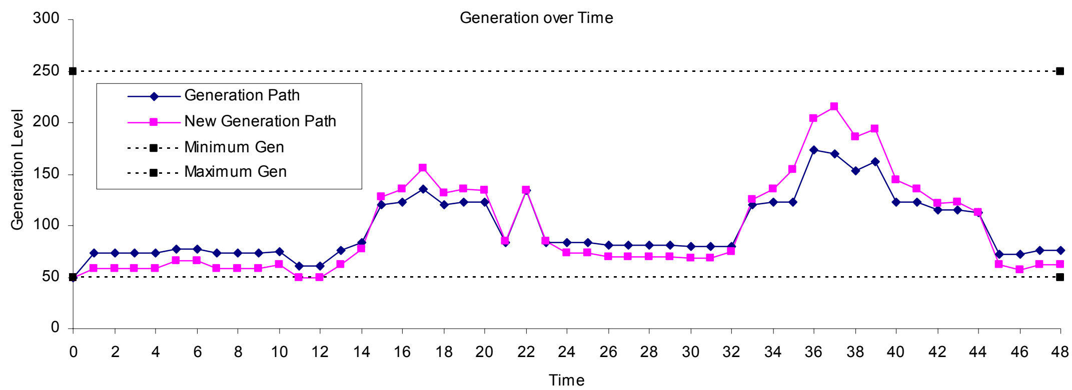

Figure 2.3 Corresponding Optimal Generation Schedules

If the unit being considered is hydro, then there may be a limited total generation capacity over the day associated with water restrictions. As a result of the increases in expected prices over the horizon, this unit's optimal generation schedule will change. It will want to generate more at the peak periods (which now have even more extreme prices), and so must trade off on generation at the off-peak periods in order to still meet the limited total generation constraints. In particular, Figure 2.3 shows that in period 1, the optimal generation level has fallen from its previous level. In other words, despite the fact that price has increased for this period, the optimal generation level has fallen. Or alternatively, the higher the price is in this current off-peak period, the less that the unit will desire to generate. This implies a backward sloping optimal supply function rather than a traditional forward sloping one, as required by market conditions. If we take this example further and consider multiple potential price paths ranging from all prices being 50% above the long-run average down to 40% below the long-run average, we can form an "offer curve" for each period that will produce an optimal dispatch schedule if any one of these potential scenarios were to actually occur. These offers can hence be considered to account for the stochasticity in price under very simple circumstances. Take the "offer curve" for period 30 for example, as shown in Figure 2.4. We observe that as the price for this period increases from 30% below, up to 50% above the long-run average, the optimal quantity to generate in that period reduces. Thus, a backward-sloping offer curve is produced. Note that moving from 40% to 30% below the long-run average, the total water restriction is not binding, and so the offer appears as forwards sloping. Examining all of these “offers” together enables an interesting, and logical observation of a “fan effect”. The lower the overall vertical position of the offer curve (and thus the more off-peak the period is), the more backward sloping the optimal offer curve is for that period. As we move through to periods with mid-range prices, the offers become close to vertical, and then for the peak periods, the offers are very forward-sloping. Hence producing a set of offers that resemble a fan.

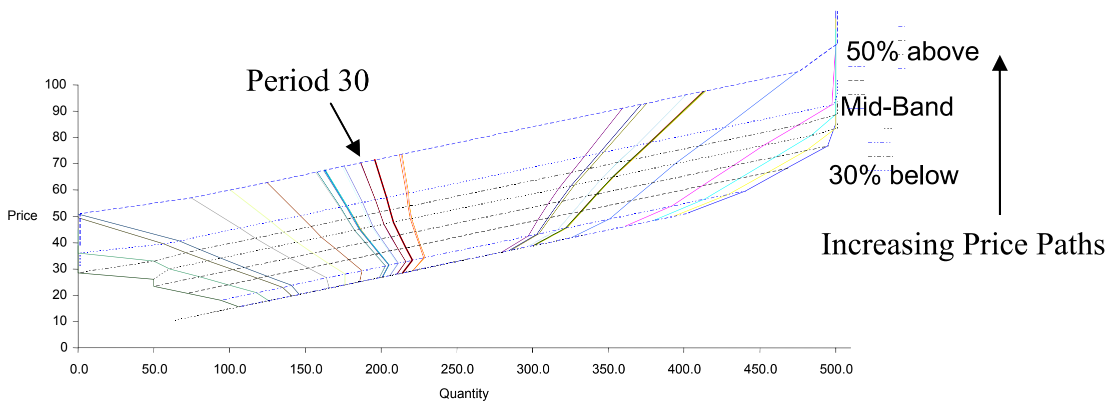

Figure 2.4 Optimal Offers for all Periods Considering Eight Possible Price Paths

The problem with this though, is that backward sloping offers are not generally permissible under market restrictions, and so the closest that can be achieved in these off-peak or shoulder off-peak periods is to make the offers vertical at some generation quantity (a Cournot-style offer), and move these quantities as time progresses, depending on what has been learnt about the state of the market in that day. This confirms the findings of Drayton-Bright (1997), that one cannot construct a monotonically increasing supply function that will be efficient over multiple periods.

Joint Effect of Many Intertemporal Constraints on a Hydro Unit

This second demonstration (Figure 2.5) of the importance of considering intertemporal characteristics of the generating units and of the marketplace compares three generation paths over a given price series:

1. The optimal generation path, accounting for intertemporal linkages (blue)
2. The naïve desired generation path produced by optimising each period individually (brown)
3. The actual generation path that would result if one tried to follow the desired generation path but ended up hitting many intertemporal constraints (turquoise).

The intertemporal constraints considered in this example include ramp-rates, minimum up and down times of the unit, minimum and maximum generation levels, limited total water (and hence generation) capability over the horizon, and intra-horizon reservoir bounds.

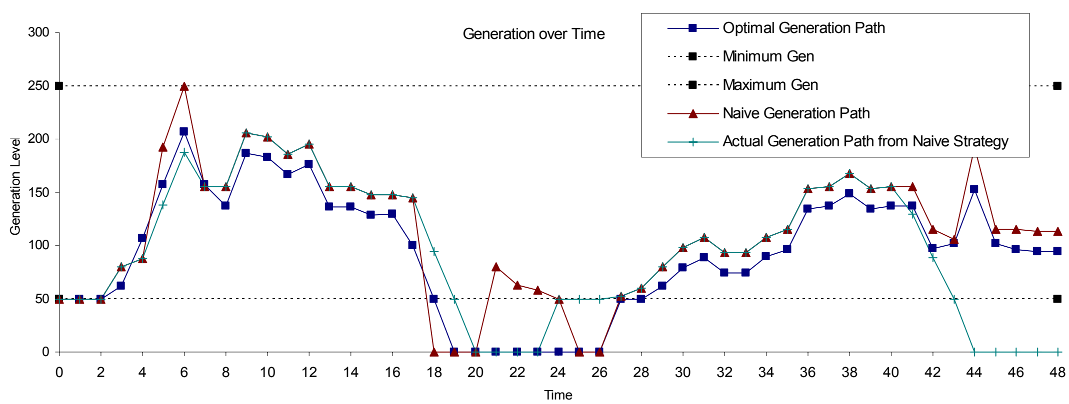

Figure 2.5 Optimal, Desired, and Actual Generation Paths

There are many points of interest in this figure:

1. The naïve generation path is above the optimal path for most of the horizon because it does not consider the overall water restrictions.
2. Observing periods 4 and 5, the naïve schedule implicitly assumes that the unit can ramp up very quickly, and as a result, the ramp up is left too late and the generators are unable to take full advantage of the high prices in these periods. This can be seen in that the line marked with crosses ("actual") is below the line marked with squares ("optimal") in these periods.

3. Observe the low prices periods 18 through to 26. The naïve schedule again ignores the restrictions of ramp rates and as such leaves it too late to shut down and thus gets caught generating at substantial losses in period 19.
4. From period 40 onwards, the actual schedule that results from the naïve desired schedule starts to run out of water, in order to meet the end of horizon reservoir targets. Therefore it must ramp down and shut off for the last five periods, missing out on some potentially profitable generation.

To illustrate the sub-optimality of ignoring these intertemporal linkages by attempting to follow the naïve desired path, we can compare the profits that would result over the horizon under this deterministic scenario. The profit from the optimal schedule, for this given problem instance, was around $151,000. The profit that would have resulted from the naïve schedule, had it been feasible was almost $158,000. However, the actual profit that resulted from trying to follow the naïve schedule was just over $132,000. In other words, 12.5% of the possible revenue was foregone by not explicitly considering intertemporal constraints when planning the generation path over the horizon, in advance. Therefore, this example reiterates the importance of considering the whole planning horizon when planning generation schedules and therefore offers to the market, even under deterministic assumptions.

2.5 Summary

In this chapter, we have discussed the various structures that deregulated electricity markets may take, and identified the forms of market that exist in the New Zealand and Australian context. It is clearly very important to understand the market form when developing offering strategies for the generators that exist within these markets. We then explored the various factors that determine how a generator should form an offering strategy over the long-term and specific offers (or offer sets) in the medium and short-runs. Finally, the concept of intertemporal issues, the main focus of this thesis, was introduced, and the importance of dealing with these costs and constraints presented.

The following chapter explores in much more detail the literature on short-term generator behaviour optimisation under a deregulated electricity structure, which has been presented briefly in this chapter.

Chapter 3

LITERATURE ON GENERATOR OFFER OPTIMISATION

3.1 Introduction

Many different methods can be used to develop generation offers for the wholesale electricity market. The most simplistic of these is to provide an offer curve corresponding to the unit's own marginal cost curve. This approach is not particularly satisfactory, if only because intertemporal links such as start-ups and shut-downs mean that a simple marginal cost may not be valid. Most offer formation methods seek to maximise profits by accounting for some form of estimation of offering behaviours of rival competitors, either explicitly, or implicitly through expectations on prices. Although the specifics of successful real-life applications of Operations

Research/Management Science (OR/MS) in development of offering strategies often cannot be made public for reasons of commercial sensitivity, the literature provides us with a plethora of analytic and empirical studies of the electricity market performed by (primarily) academics around the world. In David & Wen (2000), a comprehensive survey is provided of the various offer strategies that had been tried at that time, and ways of mitigating the abuse of market power. In addition, Hobbs (2001) provides a survey of approaches to optimal unit commitment for generators, with a focus on future research directions that may be desirable to pursue.

Due to the recency of deregulation in electricity markets around the world, there has been a distinct lack of data to develop models that reflect actual market behaviours. Accordingly, much research has been purely theoretical, and based on common economic conjectures of rival behaviour, such as Cournot, Bertrand, or Stackleberg behaviour. It is important to realise that almost all models developed in the literature, even those not based on standard game theoretic conjectures, do assume that participants are gaming the market in some way. If they were to assume otherwise, then the implicit assumption would be that the electricity market was perfectly competitive (or at least that the player considered themselves to be a perfect competitor), and they would thus optimally choose to supply an offer corresponding to their marginal cost curve[22].

Until recently, there has been surprisingly little research that focuses on how best to account for intertemporal linkages and other realistic market characteristics that are met by firms facing the task of providing offers into a market of the type found in New Zealand. Hobbs (2001) is a text based on a September 1999 workshop that involved many leaders in the field of electricity market modelling, and presents a prioritised list of research needs with respect to unit commitment modelling. This list includes many of the features and considerations of offering that are addressed within this thesis, including:

---

[22] This idea is well defined in the literature, and is proven in both Ferrero, Rivera, & Shahidehpour (1998) and Gross, Finlay, & Deltas (1998)

- Uncertainty, variability (multiple scenarios for, e.g., demand)
- Single versus multiple bids/rebidding (ability to revise and resubmit bids in a multiple round auction process)
- How to choose your bids (bids as decision variables)
- Demand responsiveness to price (price elasticity)
- Dynamic feasibility (limited energy plants, fuel constraints)

In the last few years, research in this area has increased, and addressed some of these issues, but generally independently from one another, rather than in a combined fashion as we will present in this thesis. Considering the issues independently ignores the possible interaction effects that may occur.

This chapter begins with a brief discussion of assumptions that have traditionally been quite prevalent in the offering strategy literature. We then go on to explore the game theoretic approaches that authors have taken to equilibrium modelling of the electricity market, and the appropriateness of these methods in the electricity market context. The bulk of this chapter is then devoted to modelling approaches found in the literature that specifically deal with the various intertemporal issues the generators face, and which is the focus of this thesis. We begin at that point with a short discussion of pre-deregulation models that deal with intertemporal considerations, specifically those dealing with the reservoir management problem, followed by a discussion of models designed to deal with intertemporal issues in modern, deregulated markets. In particular, we discuss the work performed by the Electric Power Optimisation Centre at the University of Auckland, New Zealand, and at the Power Systems Engineering Research Centre, made up of multiple collaborating universities in the USA. The work in this thesis extends the significant research contributions made by these two groups.

3.2 Prevalent Assumptions in the Literature

There are many assumptions or simplifications made by authors in the literature in order to enable detailed studies of the particular facet of electricity markets that interest them.

However, such assumptions often mean that the methods developed could not be effectively applied in the real-world due to inefficiency or possibly even infeasibility. These assumptions are the reason that further research needs to be performed in the area. Some of these assumptions are as follows:

- Form/shape of the supply offer. Many papers assume a linear supply function or simple offers (one step only), rather than a fully stepped offer stack. Natarajan (2003) notes that empirically, offer stacks are usually flatter at lower levels of load and get steeper as load increases towards the maximum capacity of the generating unit. Thus, some authors such as Zhang, Wang, & Luh (2000) choose to approximate the offer curve as a quadratic function instead. These approximations decrease the number of decision variables and thus make analytic solutions more tractable.
- Ignoring network considerations such as transmission constraints and losses, thus assuming that all trading occurs at a single node, is a common simplification. For example, Rajaraman & Alvarado (2003) and Anderson & Philpott (2002a), amongst many other papers.
- Sometimes, such as in Sugianto & Widjaja (2001) and Gross, Finlay, & Deltas (1998), the assumption is made that the player believes they are unable to affect the market clearing price (that is, they consider themselves to be a perfect competitor or a price-taker). Alternatively, the assumption may be made that rivals do not offer strategically (as explained previously, they therefore would offer corresponding to marginal cost).
- When optimising over a linear supply curve, authors often decide to fix one parameter, either the intercept (for example Ferrero, Shahidehpour, & Ramesh (1997b)) or the slope (for example Hobbs, Metzler, & Pang (2000)), and optimise over the other. However, this ignores the gaming implications of the shape of the supply curve[23].

[23] Read (2001) notes that strategic reactions from rivals may be affected by the shape of our curve. For example, it is well recognised in the literature that Cournot and Bertrand result in very different outcomes, even if the curve offered passes through the same point for any particular period.

- Ignoring technical constraints such as Kirchoff's voltage and current laws is yet another common simplification found in many papers, including Rajaraman & Alvarado (2003) and Anderson & Philpott (2002a). By ignoring these, a unique opportunity that electricity markets provide for market manipulation is disregarded.
- The assumption of independent or one-off auctions, rather than considering a repeated game with response by rivals to our offers is also very common, such as in Philpott, Pritchard, Neame, & Zakeri (2002).
- Finally, but very importantly, many intertemporal constraints and considerations, such as those discussed in Section 2.4 of Chapter 2, are bypassed by authors. Overcoming this key group of assumptions is the main focus of this thesis.

3.3 Game Theoretic Equilibrium Analysis and Simulation Modelling Approaches

Game theoretic models are those that describe how rational economic agents would act given their beliefs on possible rival player strategies, and the resulting equilibrium outcome given the assumptions associated with the particular model. This category includes models based on a wide variety of economic conjectures of rival's responses to our own actions. These include the Cournot conjecture (rivals hold output fixed), Bertrand (rivals hold prices fixed), Stackleberg ("leader" affects "follower's" decisions but not vice versa) and Supply Function Equilibrium (SFE, rivals hold their upward-sloping supply function fixed).

There is much debate as to whether such game theoretic models are sophisticated enough to successfully model the electricity market. Taking into account the fact that generators are likely to be aware of their mutual dependencies, the oligopoly assumption would lead us to believe that game theoretic models are likely to be suitable for modelling such markets. However, we note that while game theory can be good for modelling uncertainty, long-term dynamic issues are not able to be easily incorporated, as game theory ignores intertemporal effects. Also, Sheble (1996a) claims that with the size and complexity of the markets to be analysed in the electricity industry, analysis of micro interactions between players, with game theory, is not feasible. We do not seek here to settle the debate, only to present a review of work that has been performed, primarily with regard to the SFE model. We consider only the SFE model, as the offer-shape assumptions of the Cournot, Bertrand and Stackleberg models are substantially more different to the real-world markets that we wish to consider in this thesis, than they are for the SFE model.

3.3.1 Appropriateness of the Supply Function Equilibrium Model

The Supply Function Equilibrium model looks to find the equilibrium market clearing price and quantities, given a set of generator offer curves and a demand curve. It does this by setting the derivative of each of the *n* individual firm's profit equations to zero, and solving the set of these *n* equations simultaneously to give the slope and intercept parameters of each firms' supply function (which can be used to find the dispatch point).

The motivation for the Supply Function Equilibrium models, as developed by Klemperer & Meyer (1989) was the consideration that it is not optimal to commit to a given price or quantity (Bertrand or Cournot) when demand is stochastic. Like most areas of the literature, there is debate as to whether this modelling approach is appropriate for representing an electricity market. Although Green & Newberry (1992) claims that the electricity market is probably the best example of a market characterised by a Supply Function Equilibrium model, however all other reviews are unfavourable due to the major assumptions implicit in the model (such as linear supply curves, one-off trades, and single-location markets). Read (2001) indicates that the models are less tractable analytically than Cournot, and are of dubious relevance when independent offers are made for each period and updated dynamically. However, he notes that they were originally used in the old England and Wales market, where fixed offers were supplied for the whole day (which puts Green's claim in context). A lot of the advantage of the supply function model comes from its ability to provide good solutions over a variety of uncertain demand outcomes. However, this benefit is largely lost in systems where half-hourly or five-minute offering reduces these demand forecast errors (as forecasts for a given period are made closer to real-time). Day, Hobbs, & Pang (2002) believes that Supply Function Equilibrium is not a practical modelling method if the modeller wants to incorporate realistic details on demand, generation, and transmission characteristics. It is generally accepted that these models are very difficult to compute, except for very simple systems, for larger systems with only a few discrete strategies allowed, or for offers restricted to linear with fixed slope or intercept. This is as opposed to complementarity models based on the Cournot conjecture, which have been solved for very large systems. In addition, von der Fehr & Harbord (1993) demonstrate that Supply Function Equilibrium does not generalise from linear to step functions. In other words, users don't get good/optimal results by approximating the linear results with step functions, due to the large size of generating units.

3.3.2 Applications of Supply Function Equilibrium Models

Many applications of the supply function concept can be found in the literature, developing or applying the original Klemperer and Meyer model. For example, both Hobbs, Metzler, & Pang (2000) and Hogan (1997) develop mathematical programs with equilibrium conditions (MPEC's) for calculating oligopoly equilibrium under supply function theory, with each firm optimising over the intercept parameter of a linear offer curve. A development in these papers is that both of the formulations incorporate Kirchoff's power laws. Rudkevich, Duckworth, & Rosen (1998) developed the original Klemperer and Meyer model to relax the convexity/differentiability requirements, and thus allow the incorporation of step functions, in order to identify factors that affect the ability to use market power. However, in order to relax these requirements, many additional assumptions are made, such as identical firms/costs/supply curves, demand elasticity = 0, perfect information on rival's costs, and so on. Conditions for a smooth offer that is an optimal response to a rival's offers over a range of demands are developed by Anderson & Philpott (2001), who then attempt to approximate this offer with an offer that meets market rules (stepped, for example).

3.3.3 Identifying Potential Coalitions or Collusion

Yet another area of the game theoretic literature relates to finding potentially stable coalitions or collusive partners, somewhat under the assumption that such behaviours will not be regulated. This falls into the category of cooperative game theory, as opposed to non-cooperative assumed by most of the literature. Bower & Bunn (1999) observe that "opportunity for generators to learn about each others bidding behaviour, and adapt their bidding strategies accordingly is provided by daily repetition of the day-ahead market and by the provision of extensive bid information". Thus, (tacit) collusion is a much greater problem when participants meet repeatedly. Note that collusion or coalitions are both likely to hurt those who demand the electricity. Therefore, market regulations should be structured in such a way to prevent any coalitions other than the grand coalition from forming[24].

Hobbs & Kelly (1992) presents the idea of a "core" or "policy-restricted core", which is a set of feasible profit distributions among a coalition, for which the coalition will be stable. This is related to the idea of the stand-alone test, as described by Singh (1999a). The idea of this test is simply that profits should be allocated within a stable coalition in a manner such that all coalition members are at least as well off as they would be if they weren't part of the coalition. Krishna & Ramesh (1997a) and Krishna & Ramesh (1997b) develop intelligent agent approaches for performing negotiations to identify potentially beneficial coalitions, and for then suggesting strategies. Finally, Garcia, Campos-Nanez, & Reitzes (2005) address the issue of learning about a rival's behaviour and adjusting a generating company's own offers accordingly, under the assumption that hydro generators are engaging in dynamic price-based competition.

[24] The grand coalition refers to the situation whereby all market participants are cooperating with one another, and the central authority, in order to produce the maximum system-wide benefits.

3.4 Optimal Analytic Construction of Single Period Offers

As stated earlier, there are vast quantities of literature available on the topic of electricity market modelling and offer optimisation methods. This section provides a brief review of some of the authors and their work.

Members of the Electric Power Research Centre (EPRC) at Iowa State University have looked at the application of artificial life techniques such as neural networks, genetic algorithms and genetic programming to electric power optimisation. Both Sheble (1996b) and Richter & Sheble (1998) look at evolving offering strategies using genetic algorithms, under a double auction with multiple periods and discriminatory pricing. Each population member corresponds to a generator in the market. While being more computationally efficient than having a separate population corresponding to each generator, it seems that this would require a very small population size and would thus restrict the solution diversity able to be achieved and therefore not be a particularly sophisticated approach. Reflecting this issue, and the fact that different participants have different cost curves and capacity limits, both Doty (2002) and Petrov (2002) apply a separate population of offers for each market participant represented, in their respective Masters theses. Within each of the population members in the work of Sheble and Richter, there are three evolving (mutation and crossover) parts. These are, 1) the number of megawatts to offer as contracts in each market round, 2) "Bid Multipliers" in the form of 5-bit binary coding, which are used in combination with costs and expected market clearing price to produce a simple offer, and 3) Which of a set of pre-defined prediction techniques to use for forecasting the price. Doty and Petrov differ in their population representation, incorporating only a 10-bit bid multiplier for each member. It appears that the crossover process and the use of the bid multiplier are actually very simplistic ideas, which do not incorporate a large amount of information. Richter, Sheble, & Ashlock (1999) develops these ideas further by incorporating a single-step offering strategy that is adaptive, and able to respond to changes in the behaviour of rivals. He uses a decision-tree ("parse-tree") style approach to determine the price that will be offered, while the genetic algorithm crossover works on the shape and contents of multiple trees.

Wolak (2000) proposes two tools that he calls the best-response bidding model and the best-response pricing model. He looks empirically at how much profits for a given firm could potentially have increased using the best-response pricing model, over a set of historical data. Noting that “without the ability to replicate actual market prices using actual generator bid functions, it is impossible to compare with any degree of confidence market outcomes under any alternative bidding strategies”, the generated historical results are based on a detailed model of the Australian electricity market. The analysis makes the major assumption that rival’s offering strategies will not change in response to our strategy changes (although more recent, unpublished work may expand on this as he claims this factor could easily be incorporated). The paper demonstrates graphically the theoretical best response prices with different levels of contracts and different levels of demand elasticity, in (partly) an attempt to identify whether there are changes in the contract position of the firm under investigation that could lead to increased expected profits. The conclusion from this work is that fewer contracts would lead to such an increase, but at the expense of much increased volatility on return. In detail, the fundamental determinant of the optimal contract level is the elasticity of residual demand (more incentive to sell contracts if residual demand is flatter).

The following section deals with an approach to single period offer optimisation using a function known as the Market Distribution Function (MDF), developed by the Electric Power Optimisation Centre (EPOC) at the University of Auckland, New Zealand.

3.4.1 Market Distribution Function

The MDF, denoted by ψ, is defined over the valid quantity and price offer combinations and takes a value of between 0 and 1, corresponding to the probability that an offer at each (q, p) point will not be dispatched[25]. In other words, $\psi(q, p)$ is the probability that if the supply curve passes through (q, p), the point of dispatch will lie below (q, p). Because the probability of being dispatched a given quantity, q, at an offer price, p, cannot be greater than the probability of being dispatched the same quantity at a higher price, $p + \delta$ (and similarly for increasing dispatch quantities at a given price), this function clearly must be monotonically non-decreasing in both its dimensions. The market distribution function is essentially a stochastic representation of a residual demand curve for a single period of time. For example, if we knew the residual demand function for certain, then the market distribution function would have values of 0 below the curve, and 1 above the curve. Figure 3.1 and Figure 3.2 demonstrate sample non-linear MDF's, in 3-D and 2-D form respectively. In Figure 3.2, the MDF is represented by contours corresponding to values of 0.1, 0.2, … , 0.8, 0.9, rather than explicitly showing the entire function.

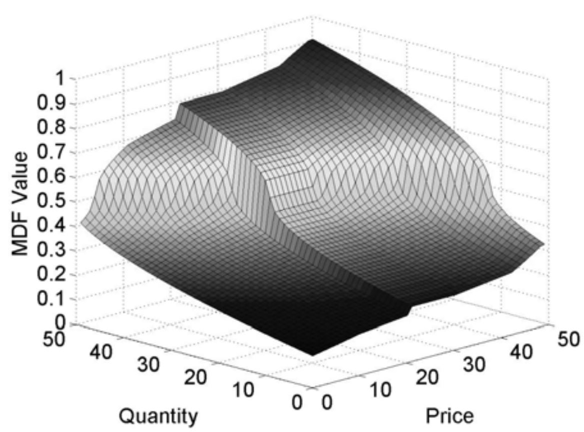

Figure 3.1 Example 3D MDF

[25] The foundation of this approach is in the work of Friedman (1956), which estimates the probability of winning a bid against competitors based on market data, then uses this probability to determine the optimal bid.

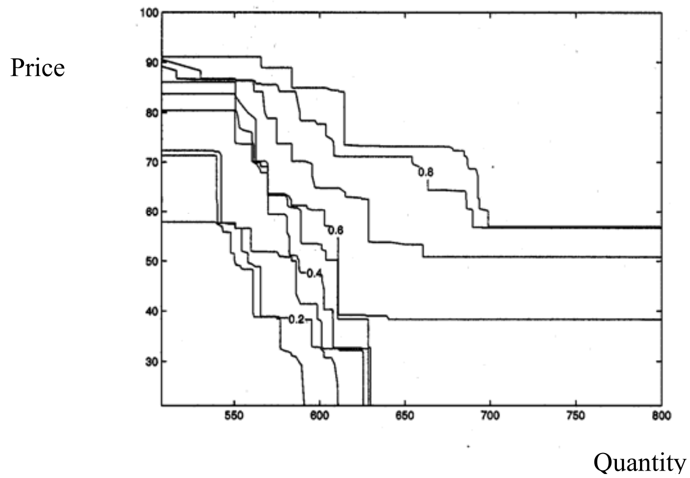

Figure 3.2 Example 2D MDF, Philpott (2002)

In taking this market distribution function approach, there are two issues that need to be dealt with. The first is how to create the market distribution function from historical data and how to update it using new data as it becomes available. The second is, given the market distribution function, how to find the optimal offer of the required form for the market. The following two sections describe the approaches that EPOC have proposed to each of these issues.

Creating the Market Distribution Function

In recent papers, the EPOC group propose multiple ways for estimating and updating the market distribution function.

The first of these papers is based on Bayesian estimation and updating of the function (Anderson & Philpott (2001)). In this paper, they present a general Bayesian approach to this issue, recognising that the exact form of Bayesian update will depend on the market structure, and in particular, the amount of information that is made publicly available. In developing this approach, they make the assumption that very little information is actually fed back to the generator. Specifically, only their own dispatch price and quantity is made available. This is contrary to the market conditions in New Zealand, in which the entire supply offer and demand bid stacks for all pre-dispatch and final dispatch periods are available to the market participants. As a result of this assumption, sufficient information is gained about the regions of (q, p) space where dispatch often occurs, but very little is known about other areas, and so the estimation of the MDF in these regions is likely to be unreliable. The paper aims to construct separate, smooth market distribution functions for each period of the day (or group of periods), reflecting the fact that conditions and hence dispatch probabilities vary significantly at different times. For each period or group of periods, they consider there to be a whole family of possible market distribution functions, characterised by their varying parameter values over a common function. When a dispatch occurs, the probability that each member of this family is the true underlying market distribution function is updated, given the likelihood of the observation for each MDF.

The second of these papers Pritchard, Zakeri, & Philpott (2002) uses a maximum likelihood estimation approach to create a discretised market distribution function, based on historical data of dispatches under conditions similar to that faced in the current period. They state that for a given period, you only need a single market distribution function, because ψ takes the same value for all offer stacks passing through a particular (q, p) point, regardless of stack shape. A grid is created in (q, p) space based on the locations of the observed dispatch points of an individual generator. The approach then attempts to determine the ψ values for each of the cells that are created in order to maximise the likelihood of observing the given sample.

The most recent of these papers, Pritchard, Zakeri, & Philpott (2005), recognises that smooth parametric models are not easy to adapt to the discontinuous market distribution functions that they wish to consider, and as such they present a non-parametric estimation method. In addition, the paper reports that in a real electricity market, a generator must be thought of as facing not a single ψ function, but many ψ functions corresponding to different times of day, hydrological conditions, etc.

Applying the Market Distribution Function

In recent papers, the EPOC group have proposed methods for constructing "optimal" offers for single periods, with a given market distribution function.

Neame, Philpott, & Pritchard (2003) consider the case of a perfectly competitive market. It is well known in the literature that if the form of the offer is unrestricted in this market situation, then the optimal choice of offer will be to exactly reflect your marginal cost of generation. However, if stepped offers are required, then this may not be feasible, and so an optimal approximation must be determined. Under the assumption of piecewise linear marginal costs, this paper demonstrates that the breakpoints of the offer steps will occur only at breakpoints of the piecewise marginal cost function, and in particular, develops optimality conditions on the location of the steps. To then limit the number of steps to match market requirements, they use dynamic programming to determine where the steps should best be placed. Their results show, as would be expected, that the stack will follow the marginal cost curve more closely in regions where the price probability density function has high values. Unfortunately, using this approach, these results cannot easily be extended to the case of a price-maker.

Taking another approach, Anderson & Philpott (2002a) create a set of optimality conditions for offer stacks, and then seek to maximise the expected return over the line integral of the offer curve. The focus of the paper is the derivation of the necessary conditions for optimality and, as such, it does not fully address the important practical question of efficient computation of the optimal offer. Anderson & Philpott (2002b) expand on these optimality conditions and demonstrate that, in the case of a convex cost function and inverse log concave rest-of-market supply curve, the optimal unrestricted offer will be monotonically non-decreasing in both the price and quantity dimensions, where the uncertainty in the market is in relation to the level of inelastic demand. This implies that, given these particular conditions, the optimal offer under uncertainty is simply the path, or locus, traced out by identifying the optimal dispatch point under each possible resulting position of the residual demand curve. However, it is likely that in a real market situation, these conditions will not be met. Thus, in Philpott, Pritchard, Neame, & Zakeri (2002), a practical approach to more realistic scenarios of rest-of-market supply and uncertainty is suggested. As with Anderson & Philpott (2002a), this paper seeks to maximise expected return over the line integral of the supply function, given the payoff at each point and the probability of that payoff actually occurring. In this paper, it is assumed that for realistically sized problems, an explicit ψ may be impractical. As such, they break the (q, p) plane into a grid, with known payoffs on the edges of the grid. The offer curve is then constructed under the restriction that it must follow the discrete edges of this grid. For each vertex in the grid, the maximum expected additional payoff of any curve passing through the vertex is calculated dynamically by selecting the maximum of:

a) The expected payoff on the 'up' edge + maximum expected payoff from the 'up' vertex
b) The expected payoff on the 'right' edge + maximum expected payoff from the 'right' vertex

Once these payoffs are all known, then it is a simple process to determine the offer curve, by using dynamic programming to determine the optimal path through the MDF surface.

Limitations of the Market Distribution Function Approach

This market distribution function approach, as described above, successfully considers the ability of a unit to affect the price, and the fact that price levels will differ at different times[26]. However, it has not yet been fully incorporated in a multiple period model that deals optimally with the various intertemporal linkages (including market correlation over time) that occur in the marketplace. More recent papers, which broach this topic, are discussed in Section 3.5.2.1 of this chapter.

[26] Although it does not consider the impact which one party's offers may have on offers from other parties, as in a "gaming" model.

3.4.2 Rival Generator Modelling Alternatives

The market distribution approach brings to light a distinction that can be made with regard to literature types. Some approaches look to individually model the behaviours of each of the rival players, while some look at the net effect of demand combined with all rival behaviours, by considering a residual demand function or something similar. The EPOC work falls into this second category, as does work by Read (2001), David (1993), Green & Newberry (1992), Yang, Wen, Wu, & Ni (2002), Wolak (2000), Baillo, Ventosa, Rivier, & Ramos (2004), Rajaraman & Alvarado (2003) and so on. Read (2001) looks at creating optimal offer curves over a residual piece-wise linear demand curve, assuming that rivals behaviour is fully captured in the residual demand curve, and its variations occur independently of, rather than in reaction to, our offers. David (1993) combines all rivals into a single competing rival when looking at stepped offer stack construction and considers all the opposing offers that may be in competition with any of his own firm’s individual steps. Green & Newberry (1992) consider the effect of contracts on the shape of the residual demand curve in the duopolistic UK market, and its meaning in terms of optimal offering response. Baillo, Ventosa, Rivier, & Ramos (2004) considers both the ability of a generation company to affect the price and uncertainty about rival’s behaviour for a generating company operating in a day-ahead, 24-period market. We will discuss the focus of the Rajaraman & Alvarado (2003) paper in the following section, as it considers RD curves in a multi-period, rather than just a single-period model.

Note that if rivals respond to a generator’s offers, then this corresponds to a shift of the residual demand curve rather than a movement along it.

3.5 Literature Addressing Intertemporal Issues

With respect to approaches to hydro generator operation over multiple periods, we can partition the literature into three broad groups; dynamic programming approaches, mathematical programming approaches (including stochastic linear and non-linear programming) and heuristic approaches.

In addition, the vast majority of models found in the literature that deal with intertemporal issues (such as unit commitment and reservoir management models) are based on the pre-deregulation style structure and aim to minimise costs of meeting demand rather than maximise profits of offering into an electricity market. We begin this section with a short discussion on such models, but as the specific application to a centralised situation is not appropriate to this thesis, the bulk of this section focuses on models appropriate for a deregulated market structure.

3.5.1 Pre-Deregulation Reservoir Management Models

There exists a vast field of literature pertaining to pre-deregulation reservoir management and unit commitment models. A comprehensive survey of literature at the time is found in Yeh (1985), covering a broad range of techniques that had been applied to the reservoir management problem, including simulation models, heuristic models, mathematical programming models (both discrete and stochastic), and stochastic dynamic programming models. Yakowitz (1982), also provides a good summary of literature on the problem, with a focus on dynamic programming applications. A more recent summary of a variety of approaches to the problem of unit commitment, accounting for start-up and shut-down issues and other intertemporal constraints is presented in Read, George, Lamar, & Rosenthal (1994). In addition to dynamic programming, they discuss integer programming, Lagrangian relaxation, network optimisation, and heuristic approaches. Although much of this literature applies specifically to hydro systems, the problems faced by energy-limited coal or gas are basically similar in structure, so that adaptations of hydro optimisation techniques can often be applied to deal with these too.

Since this thesis is not proposing any new developments with respect to the inter-temporal optimisation of power systems, per se, it does not seem appropriate to survey this literature in any depth. The basic technique applied here is simply Stochastic Dynamic Programming, as first proposed by Bellman & Dreyfus (1962), assuming an embedded Markov chain. However, one section of the thesis does employ a technique known as "Dual Dynamic Programming", which may require some explanation.

Internationally, one of the primary models applied to the problem of reservoir management in more recent times has been the Stochastic Dual DP (SDDP) approach presented in Pereira & Pinto (1991) and Pereira, Campodonico, & Kelman (1999). This technique is similar to, but different from an SDDP approach presented in Read (1985) and Read & George (1986). The SDDP concept of Read and George is dual to a standard DP in that it finds optimal primal variables corresponding to a set of critical shadow prices, rather than vice versa. The technique is able to characterise the solution strategy without discretising the primal state-space, and produces a piece-wise linear approximation to the marginal value function, and it is this technique from Read & George (1986) which has been applied here.

On the other hand, the SDDP technique developed in Pereira & Pinto (1991) comes from a mathematical programming base. Their technique also avoids discretising the primal state space, but constructs a value function which is piecewise linear, so that the marginal value function is made up of a set of level planes. This technique is able to be applied to systems with a large number of reservoirs, for which it would be impossible to define a complete operating strategy under the technique of Read & George (1986). As both of these techniques avoid discretising the primal state space, they reduce the computational issues caused by the "curse of dimensionality" as the problem size increases. For a more detailed comparison of the two techniques, refer to Yang (1995).

3.5.2 Models Addressing Intertemporal Issues under a Market Structure

We consider three different groups of electricity offering models under a market structure:

1. Those that recognise the behavioural and gaming effects of the repeated nature of the electricity market,
2. Those that deal with unit commitment decisions (primarily concerned with thermal generation units), and
3. Those that consider the fuel conservation issues of hydro generation units.

The focus is on the third group, which is the main area of research in this thesis.

3.5.2.1 Behavioural Feedback Models

In this section, we consider a number of models that recognise the feedback and learning mechanisms that are present in a realistic market.

Gajjar, Khaparde, Nagaraju, & Soman (2003) presents a system using Markov Decision Processes to select from a range of actions from a generator's policy (set of possible actions). Reinforcement learning is used to simulate intelligent agents that can learn how to make good decisions by observing their own behaviour (allowing Markov transition probabilities to evolve as a function of temporal difference error). Petrov (2002) develops the Roth-Erev reinforcement learning algorithm (Roth & Erev (1995) and Erev & Roth (1995)) such that agents learn how to offer to the auction and respectively develop a winning strategy. However, this is restricted by its assumption that at any given time, a trader knows only their last offer and the amount of profit resulting, thus cannot take advantage of information from prior rounds of offering.

Sakk, Thomas, & Zimmerman (1997) develops a sequential offering scheme with feedback of information, so that generators can modify their offers after pre-dispatch notification, but without knowledge of the profit functions of their competitors. A simple market clearance model is then used to determine the combined outcome and market clearing price from the set of linear bids. The approach attempts to predict other generators offers based on past dispatch/pre-dispatch and offering histories, but does not account for stepped offers, start-up costs, and many intertemporal constraints. Stothert &

MacLeod (2000) presents offering as a control problem, with feedback of system marginal price and supply offer accept/reject information over multiple offering rounds. The model finds a Pareto optimum over the conflicting objectives of generator profit-maximisation and consumer price minimisation. This paper makes many assumptions though, such as simple price based offers, a single offering period, no spatial aspects incorporated, and all or nothing of each offer must be accepted by the market. Likewise, Visudhipan & Ilic (1999) models a feedback system with generators learning market opportunities and evaluating their offers using available information about past market clearing prices. The model used is a closed-loop dynamic system, in which previous and current information are used as feedback signals, and outcomes are evaluated by a simple market clearing model. They consider single-step and linear supply functions, elastic and inelastic demand, and two different offering strategies, in order to establish determinants of the ability to exercise market power. Li, Svoboda, Guan, & Singh (1999) uses parametric dynamic programming to produce hourly offer stacks in order to meet certain revenue-adequacy conditions in an iterative offering market (but gaming considerations are ignored due to a requirement that bids cannot increase in price in an iteration, unless the revenue adequacy conditions are not met).

Finally, Oren & Rothkopf (1975) shows that the concept of bidding into repeated auctions is neither a new concept, nor limited to the electricity market. This paper addresses the issue of current bidding strategies affecting rival behaviour in future auctions, and the gaming implications of such an effect, and develops the optimal infinite-horizon equilibrium strategy as a result. In this thesis, we do not look to find equilibrium outcomes of auctions, and so we will not consider this paper or the corresponding wealth of literature on this topic in a general context, any further. The reason for this assumption is that, as discussed in Section 3.3, game theoretic models that look for equilibrium outcomes are not able to produce sufficiently detailed offer strategy results in the context discussed in this thesis.

3.5.2.2 Unit Commitment Models

The Unit Commitment problem faced by a generator is to determine when each of their generation units should be turned on and off over a given planning horizon.

In the context of the deregulated electricity market structure, Hobbs (2001) identifies that in addition to dynamic programming, many techniques have been applied to the unit commitment problem, including branch-and-bound mixed integer programming, linear and network programming, Benders decomposition, and various metaheuristic approaches such as genetic programming, simulated annealing, and neural networks. In this section, we will briefly present the literature related to these and other approaches.

Pokharel, Shrestha, Lie, & Fleten (2004) identifies that *"… optimal scheduling is a non-convex highly constrained nonlinear mixed integer optimization problem, which can only be fully solved by complete enumeration. Moreover, the solution must be obtained within a time that makes it useful for the intended purpose"*. In an attempt to address these issues, they go on to present a new heuristic technique in coordination with dynamic programming and nonlinear programming that determines the optimal unit state decisions and dispatch quantities (rather than offers) over a planning horizon, under expected price levels.

Pokharel, Shrestha, Lie, & Fleten (2005) uses dynamic programming to optimise behaviour of a set of thermal units with significant intertemporal constraints related to the physical operation of the units, under a stochastic price path for the planning horizon (but ignoring any market power). In this model, the decisions in each period are whether each unit should be switched on or off, and at what level the units that are 'on' should be generating.

Baillo, Ventosa, Ramos, & Rivier (2001) describes an MIP for the unit commitment problem that includes a set of constraints ensuring a minimum market share in each period, while Baillo, Ventosa, Rivier, & Ramos (2004) creates an MIP solved with

Bender's decomposition, that considers residual demand uncertainty within each period (much like a discrete market distribution function). O'Neill, Sotkiewicz, Hobbs, Rothkopf, & Stewart (2005) discusses the issue of non-convexities in the unit commitment problem in depth and presents a simple MIP that deals with this issue.

In Section 3.4.1 of this chapter, we discussed the market distribution function approach to offer optimisation within a single period, developed by the EPOC research group. In addition to developing this method, they have also carried out parallel work on unit commitment problems and hydro scheduling over time, not using the MDF approach. Philpott, Craddock, & Waterer (2000) formulates a large mixed-integer program to optimise unit commitment and release behaviour over a short-term horizon where each unit faces simple operational rules and costs in addition to generation with non-linear efficiency.

In more recent work, the EPOC group have begun to bring together the two branches of work described here, to consider market uncertainty within each period under a multiple-period horizon. Philpott (2004) considers a non fuel-limited single unit offering into a pool market under a known set of market distribution functions for the horizon. The basic DP model that the paper presents considers start-up and shut-down costs as well as feasible operating ranges. As such, there are two decisions needing to be made at each period over the horizon; whether the unit will remain in its current on/off state, and what the offer will be if the unit is switched on. They then extend this model to consider ramp-rate restrictions, increasing the dimension of the state space and thus requiring a different offer to be chosen, depending on the dispatch level in the previous period. Philpott & Schultz (2006) extends the approaches further, to consider multiple units that are offered collectively into the pool. The two key extensions that this thesis will provide on the single unit model, with respect to market complexities considered, are to take into account fuel limitations and market correlation between consecutive periods, the latter of which is suggested as a valuable extension in Philpott & Schultz (2006). Fuel limitations substantially increase the complexity of the problem, as the offers are no longer independent between periods, but rather, are highly dependent on one another through opportunity cost of fuel use considerations. This increases the dimension of the DP, as a different offer will need to be provided for each reservoir level in each period. Market correlation again considerably increases the complexity of the problem, as there is now a probabilistic set of MDFs for each period connected by way of a Markov chain, rather than a single known MDF for each period. Again the dimension of the DP increases, as a different offer will need to be provided depending on the market outcome of the previous period (which thus presents the generator with a different MDF for the coming period).

3.5.2.3 Reservoir Management Models

Set just before the deregulation of the New Zealand market, Drayton-Bright (1997) develops a model for coordination of energy and reserve offers, incorporating a market clearing model that includes many important intertemporal linkages, such as water flow balance, storage bounds, ramp rates, and others. His work is based not on gaming, but on understanding the implications of various intertemporal constraints. The thesis is focused on comparisons between various market designs and rules, such as the form that the offers will take (fixed versus sculpted over each period, for day-ahead offering versus re-offering). He notes that intertemporal constraints may make Linear Programming subproblems (decomposed by market player) respond to energy prices in such a way that a monotonically non-decreasing step supply function cannot be drawn through the observed price and response pairs. Thus, there is no stable offer curve for which repeatedly solving static models, (thus ignoring the presence of intertemporal linkages) can achieve an efficient outcome. In other words, it is not possible to create a single stack that will be optimal for many periods. Simulating a periodic update market with a rollover of planning horizon every half hour and incorporating an end-of-day target storage water level (to avoid end effects), he finds that re-offering along with sculpted offer stacks produced the most efficient results (that is closest to those that are expected from a perfectly competitive market). The thesis does not address the use of the market power associated with being in a monopoly situation, gaming, contracts, and some spatial aspects (transmission constraints).

In Scott (1998), a dual dynamic program is developed in order to deal with the management of a mixed hydro-thermal system in the context of the newly created competitive electricity market in New Zealand. The approach incorporates a Cournot oligopoly representation of the energy spot market. In the model, all players are considered to offer in a Cournot manner in order to optimise their own position, and thus equilibrium within each period is identified. This is similar in concept to the game theoretic approaches discussed in Section 3.3, and is distinct from most of the approaches described in this section, which consider the optimisation of a single player's behaviour, subject to expectations about the behaviour of rival market participants (and which do not necessarily assume all players to have the same bidding constraints and structures).

A dynamic programming simulation model is developed by Villar & Rudnick (2003), simulating a mixed hydro/thermal market based on simple (Cournot, quantity based) hourly offers. They recognise the importance of current decisions on future possibilities, particularly with respect to hydro. They consider all (large) firms in the market to offer strategically, but that only hydro stations use a dynamic model (thermals use a static model). The intertemporal consideration incorporated is the water usage/availability restriction over the short-term (24-hour) planning horizon, and thus the states of the dynamic model are reservoir levels. A longer term model could be developed, but the model would need to be adjusted to take account of stochastic inflows (the authors report that they are currently in the process of developing longer-term simulators). The paper, however, does not take into account many of the other important intertemporal constraints.

Ladurantaye, Gendreau, & Potvin (2006) presents an MIP model for a price-taking generator, where price path scenarios are organised into a tree structure to reflect different overall states that the market could be in. The paper demonstrates the use of a branching method, in a different context to that which we present in Chapter 7 of this thesis. However, the paper allows only for probability nodes in the branching structure, and not decision nodes (as we will present).

Pereira, Campodonico, & Kelman (1998) provides an overview of the possible methodologies and tools that can be used to address the new set of operational decisions that need to be made under deregulated market structures. The paper considers both thermal systems, which are decoupled in time, and hydrothermal systems, which are coupled in time (face intertemporal issues). The paper discusses the internal cost minimising dispatch process faced by a generator with multiple units in order to meet load, as well as discussing the process of offering generation prices and capacities to a market, in order to maximise revenue. The efficiency of bidding schemes is then considered (i.e. whether it is possible to achieve the same operating efficiency as an ideal centralised dispatch), and they conclude that in a hydrothermal system, such a scheme is inherently inefficient, but that the specific results will be system-dependent.

A range of reports have also been produced by consulting firm PSR, which introduce and progressively develop their techniques for optimising the hourly bidding process of a generator with multiple generating units, within a hydrothermal system. In Kelman & Pereira (1998) and Pereira, Granville, Dix, & Barroso (2004a), the bids for each of the units are a simple (price, quantity) pair, and a series of MIPs and non-linear formulations are constructed in order to address this problem. Pereira, Granville, Dix, & Barroso (2004c) compares these two methods, recognising that even in its most basic form (ignoring transmission, unit commitment, and so on) strategic bidding under uncertainty is a nonlinear programming problem, and hence that the two options are to solve it directly using a general nonlinear optimiser, or to transform it into a MILP (or MIP). They conclude that while the nonlinear approach is simpler to apply and more efficient with respect to CPU time, the non-convexities can cause the non-linear optimiser to reach locally optimal solutions that are not globally optimal. On the other hand under the MILP (which is obtained by a binary expansion scheme), accuracy can be controlled through the level of discretisation of the binary representation (where a more refined representation will of course increase CPU time).

Extensions to the simple model are presented in Pereira, Granville, Dix, & Barroso (2004a) and include dealing with uncertainty in load and in rival generator bids (through a set of possible scenarios), dealing with transmission networks (which is considered only indirectly in this thesis), dealing with unit commitment, and the assumption of Cournot or Bertrand bids. Note also that Pereira, Granville, Dix, & Barroso (2004b) applies the binary expansion scheme mentioned above, to the problem of determining Nash equilibrium solutions.

Both Neame, Philpott, & Pritchard (2001) and Neame, Philpott, & Pritchard (2005) model price, rather than market, uncertainty under the assumption that the generator is unable to influence the market clearing price, and thus present DP algorithms that find the optimal offering stacks within each half-hour period over a short-term planning horizon. These papers then consider a longer-term horizon, with weekly periods, and deal with the combined problem of firstly selecting a mean and variance release for each week, and then constructing stacks for use throughout each week that have this mean and variance. Pritchard, Zakeri, & Philpott (2004) then presents the HERO (Hydro-Electric reservoir optimization) model that implements this approach.

The remainder of this section presents a brief summary of the method presented in Rajaraman & Alvarado (2003). There is further discussion of this algorithm in later chapters, as it is this benchmark algorithm to which we will compare the algorithm developed in Chapter 7 of this thesis.

In this paper, Rajaraman & Alvarado (2003) present a method for producing a complete offering strategy over a short-term planning horizon, accounting for complex intertemporal considerations that were not previously considered together in the literature. The intertemporal considerations in the main algorithm that they present are limited fuel and uncertain market outcomes (residual demand curves), which are correlated between periods by way of a lag-1 Markov chain. A complete offering strategy is therefore constituted by a set of offers that define the optimal decisions under all states that could occur throughout the planning horizon, where the two-dimensional state-space is defined by the fuel level and the previous period's residual demand curve outcome.

In particular, the paper addresses three key concerns of generator offering decisions:

1. The impact of offer monotonicity requirement
2. The effect of intertemporal constraints
3. The effect of price/market uncertainty

The approach used is a multiple-level nested DP optimisation, where the upper DP level is the two-dimensional state transition for the planning horizon, and the lower DP level is the feasible offer construction for a given state in a single period. This lower level DP is similar to that discussed for finding the optimal offer over a MDF in Section 3.4.1, however, rather than having incremental payoffs located on the edge of a discrete grid, the incremental payoffs are located on the vertices of the discrete grid. This turns out to cause a computational error, and is discussed further in Chapter 7.

The approach used is a purely primal method, where the incremental payoffs on each vertex include both the within-period payoff and the expected future value of dispatch at the given vertex. Such a primal method turns out to be quite computationally inefficient, and this is explored further in Chapter 7.

In the formulation presented, all fuel is required to be exhausted within the planning horizon. In reality, this is clearly not often going to be the optimal behaviour. Fortunately, it would be very easy to generalise their approach to consider an end-of-horizon fuel value curve.

3.6 Summary

In this chapter we have discussed literature on generator behaviour and optimisation that covers a very wide range of assumptions regarding the structure of the market, the types of decisions that a generator must make, and the behaviours of other generators in the market. In this thesis we focus on the optimal half-hour energy-only offering decisions by fuel-limited generators over a short-term horizon in a deregulated market such as that found in New Zealand. We look specifically to extend on the work presented in two papers, Philpott & Schultz (2006) and Rajaraman & Alvarado (2003). The former paper presents a DP for unit commitment and offer optimisation over a short-term horizon, but ignores fuel limitations and market outcome correlation between periods, while the latter considers these extra complications, while ignoring unit switching considerations, in a computationally inefficient manner. In Chapters 7 to 9 we will present various algorithms that combine and extend on these complexities using a very efficient algorithm.

Chapter 4

OPTIMAL SINGLE-PERIOD ANALYTIC OFFERS UNDER MARKET UNCERTAINTY

4.1 Introduction

Given an uncertain market scenario for a single period, there are two possible ways of determining the optimal offer to provide under a given marginal cost curve. The first of those is to perform a dynamic program over an arbitrary grid on the offering space, where each edge (or vertex) of this grid contains the expected marginal payoff associated with an offer passing through this segment (or point). This approach is suggested in many papers discussed in Chapter 3, including Rajaraman & Alvarado (2003) and Philpott, Pritchard, Neame, & Zakeri (2002). The alternative is to apply a more accurate analytical approach to the construction of optimal offers, possibly involving optimality conditions on the location of the offer. Such an approach is the focus of much of the market distribution function work by the EPOC group, as demonstrated particularly in Anderson & Philpott (2002a). As we will show in Section 4.2, these analytic results and optimality conditions are highly complex to prove and apply, and as such, much of this chapter is devoted to producing much more intuitive analytical forms of optimal offers, under relatively simple residual demand (RD) curve forms and types of market uncertainty.

4.2 MDF Optimality Conditions

In this section, we will provide a very brief overview of the mathematics behind the optimality conditions used for constructing optimal offers in the market distribution function theory of the EPOC research group. For an in depth presentation and proof of the theorems discussed here, please refer to any of the papers referred to below. We will use the following terms in this discussion:

$V(s)$ = Expected return of any offer s

$R(q, p)$ = The return from dispatch at any (q, p) point

$\psi(q, p)$ = The market distribution function defined over (q, p) space

R_q = The partial derivative of $R(q, p)$ with respect to q

R_p = The partial derivative of $R(q, p)$ with respect to p

ψ_q = The partial derivative of $\psi(q, p)$ with respect to q

ψ_p = The partial derivative of $\psi(q, p)$ with respect to p

Ψ = Region of offering space in which $0 < \psi(q, p) < 1$

t = The index over which the offer is defined parametrically

T = The maximum value of the index t, associated with the end point of the offer

t_0 = The intersection between the current offer and the border of Ψ and the region where $\psi(q, p) = 0$

t_M = The intersection between the current offer and the border of Ψ and the region where $\psi(q,p)=1$

$w(t)$ = The integral along an offer path to a certain point t

q_M = The maximum quantity that can be offered

p_M = The maximum price that can be offered

$x(t)$ = The x coordinates of an offer parametrically defined over t

$y(t)$ = The y coordinates of an offer parametrically defined over t

In Anderson & Philpott (2002a), the expected return from any offer, *s*, through the feasible offering space is derived and defined as the line integral along the offer curve, as shown in Equation 4.1.

$$V(s)=\int_s R(q,p).d\psi(q,p)$$

Equation 4.1 Expected Return of any Offer, *s*

In order to construct an optimality condition in the absence of monotonicity conditions, the paper describes a function, $Z(q,p)$, as shown in Equation 4.2.

$$Z(q,p)=\begin{cases} R_q\psi_p - R_p\psi_q, & (q,p)\in\Psi \\ 0, & otherwise \end{cases}$$

Equation 4.2 *Z(q,p)* Function

Since both R and ψ are continuously differentiable in Ψ, $Z(q,p)$ is a continuous function on Ψ. As such, when the optimisation of the offer curve is viewed as a problem in the calculus of variations, $Z(q,p)$ = 0 is the Euler (optimality) condition for $(q,p)\in\Psi$. The paper explains that this condition can be derived using Green's theorem in the plane, showing that any deviation (or 'feasible variation') from a curve that meets this optimality condition must be suboptimal. In other words, if the curve passes through a region in which $Z(q,p) > 0$, then the line integral of $Rd\psi$ along the curve may be improved by adding an anti-clockwise traversal of a closed contour lying in this region (and vice versa).

Therefore, in the absence of monotonicity conditions, the maximal curve lies on the common boundary of two regions having $Z(q,p) > 0$ and $Z(q,p) < 0$. This curve need not be either monotonic or even connected. Therefore in a real market situation, adjustments must be made to these to ensure monotonicity in the offers chosen.

In order to derive new local optimality conditions in the face of these monotonicity constraints, Anderson & Philpott (2002a) redefines the offer curves parametrically over the range Ψ shown in Figure 4.1, with respect to a variable *t*, as shown in Equation 4.3

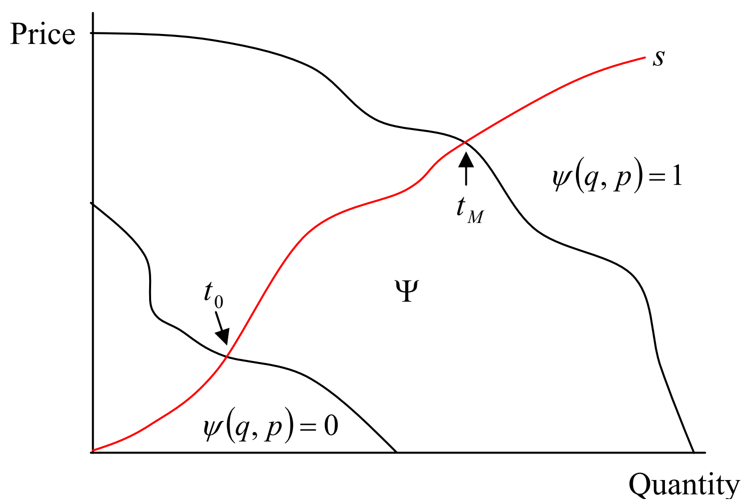

Figure 4.1 Offering Region of Interest

$$s = \{(x(t), y(t)), 0 \le t \le T\}$$

Equation 4.3 Parametric Representation of an Offer, *s*

The optimality of the offer is then described by their following theorem:

THEOREM 4.1. Suppose $s = \{(x(t), y(t)), 0 \le t \le T\}$ *is an offer stack, and let w(t),* t_0 *and* t_M *be defined as above. The following conditions are necessary for s to be a local optimum for P* (the offering problem).

1. *If* $x(t_M) = q_M$*, then for every* $t \in [t_M, T], w(t) \ge w(t_M)$
2. *If* $y(t_M) = p_M$*, then for every* $t \in [t_M, T], w(t) \le w(t_M)$
3. *If* $y(t_0) = 0$*, then for every* $t \in [0, t_0], w(t) \le w(t_0)$
4. *If* $x(t_0) = 0$*, then for every* $t \in [0, t_0], w(t) \ge w(t_0)$
5. $w(t_0) = w(t_M)$
6. *For every* $t \in [t_0, t_M]$*,* $x'(t)(w(t) - w(t_0)) \le 0 \le y'(t)(w(t) - w(t_0))$

For further explanation of these optimality conditions, the reader is referred to Anderson & Philpott (2002a). Note that in the conclusion of Anderson & Philpott (2002a), it is stated that they do not consider "the important practical question of the efficient computation of an optimal solution".

In another paper, Philpott, Pritchard, Neame, & Zakeri (2002), derive and apply the simple form of the optimality condition in a case where offers are naturally monotone. Anderson & Philpott (2002b) state formally in their Theorem 4 that *"if C(.) is a non-decreasing convex function and S(.) is a differentiable inverse log concave function, then there is a supply-function response that is optimal for any* $h \in H$*."*, where *C*(.) is the generator's cost function, *S*(.) is the rest-of-market supply curve, and *h* is a possible demand level from the set of all possible demand levels, *H*. In other words, there will be a naturally monotone offer when there is a known aggregate rest-of-market supply function that is inverse log concave and the uncertainty comes from a stochastic level of inelastic demand. Philpott, Pritchard, Neame, & Zakeri (2002) state that the optimal offer must at all points be vertical, horizontal, or satisfy the first-order condition shown in Equation 4.4 (equivalent to the Euler condition above).

$$\frac{\partial R}{\partial q}\frac{\partial \psi}{\partial p}-\frac{\partial R}{\partial p}\frac{\partial \psi}{\partial q}=0$$

Equation 4.4 Optimality Condition

For a general market distribution function shape, the derivatives involving ψ would make this very difficult to solve efficiently, or at all. For example, the approach is not able to be applied when there are discontinuities in ψ, as would be the case if another player had an offer curve with a horizontal section.

However, in this simpler case, they are able to define the market distribution function analytically, as shown in Equation 4.5 (where $S(p)$ is the aggregate supply curve), and so the partial derivatives needed are thus easy to find, and shown in Equation 4.6.

$$\psi(q,p)=P(D<q+S(p))=\int_0^{q+S(p)} f(\eta)d\eta$$

Equation 4.5 Market Distribution Function

$$\frac{\partial \psi}{\partial q}=f(q+S(p)),\frac{\partial \psi}{\partial p}=f(q+S(p))S'(p)$$

Equation 4.6 Partial Derivatives

For a simple payoff function, $R(q,p)=qp$, the optimality condition produce the optimal offer form shown in Equations 4.7.

$$\frac{\partial R}{\partial q}\frac{\partial \psi}{\partial p} - \frac{\partial R}{\partial p}\frac{\partial \psi}{\partial q} = f\left(q + S(p)\right)\left(pS'(p) - q\right) \overset{set}{=} 0$$

$$q = pS'(p)$$

Equations 4.7 Optimal Offer Form

In other words, the quantity level to offer is a function of the price and the slope of the aggregate rest-of-market supply curve at each price.

This section briefly discussed the simple optimality condition produced in the market distribution function research performed by the EPOC research group, and demonstrated that for a very restrictive set of conditions, under which the optimal offer is naturally monotone, this is easy to apply. However, under other, possibly more realistic forms of uncertainty and RD curve forms, the conditions are difficult or not possible to apply. In the remainder of this chapter we first provide a more intuitive derivation of a similar offer form to that shown in Equations 4.7, and then use this intuition to determine the optimal analytic form of offers under more varied situations, overcoming non-monotonicity with simple optimality conditions where necessary.

4.3 Optimal Offer Forms under Single-Kinked RD Curves

In this section, we establish the desired form of offers that generators should provide to a wholesale electricity market for a single period. We consider various scenarios including different RD forms and uncertainty as to the position of the RD curve. We make the assumption that the entire RD curve will move simultaneously in a single direction. It will do so in such a manner that different possible positions of the RD curve do not cross over one another. We also make the assumption in this section that the RD curves are piecewise linear with two segments, forming a RD curve that is either convex or concave to the origin.

We will observe that in certain circumstances, these desired offer forms do not meet conditions required by the market for valid offers. In particular, the condition that offers must be monotonically non-decreasing in both the price and quantity dimension is commonly not met, especially when the RD curve is convex. We determine the optimal way of dealing with this problem, thus restoring monotonicity to the offer while maximising the expected level of profit.

4.3.1 Terminology

This section is devoted to defining the key terms and concepts that will be used throughout the remainder of the chapter.

Residual Demand Vertex Path (RDVP): In a market situation, a generator will face uncertainty as to the position of the RD curve for all periods in which they are required to offer. In particular, we will define the direction of RD curve uncertainty with a line referred to as the *Residual Demand Vertex Path (RDVP)*. For example, if the entire RD curve will move in the quantity dimension[27], as demonstrated in Figure 4.2a, then the RDVP will be horizontal[28]. If the entire curve could move up in both price and quantity dimensions simultaneously, then the RDVP would be positively sloped as demonstrated in Figure 4.2b.

[27] This would be the case if we knew rest of market supply, but were unsure about the level of inelastic demand

[28] Note that the path does not need to be based on the path of the vertex, but rather could be based on *any* specified point on the RD curve. The vertex has been chosen for simplicity of representation.

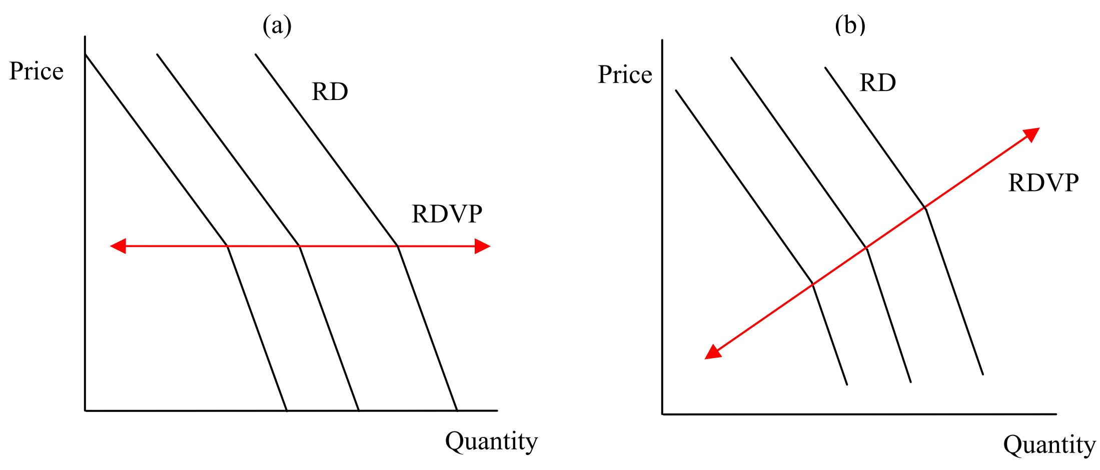

Figure 4.2 Horizontal and Positively Sloped RDVPs

In this section we restrict the RDVP so that alternative RD curve positions may not overlap, as shown in Figure 4.3. This will be generalised in Section 4.5.

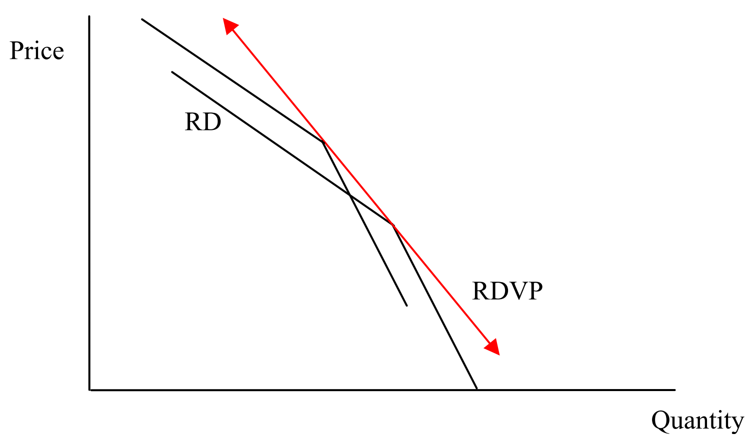

Figure 4.3 Overlapping Residual Demand Curve Instances

Contributory Line (CL): Paths formed by the generator's optimal dispatch points (where marginal revenue = marginal cost) under all possible positions of RD, assuming a constant slope of the RD curve. Assume that marginal cost is a linear function of the form $MC = MC_{\text{int}} + MC_{sl}q$ (i.e. a quadratic cost function), where MC_{int} and MC_{sl} are the intercept and slope terms of the marginal cost curve respectively. For each slope, $-b_i$, that exists in the RD curve, there is a corresponding contributory line of form $p = MC_{\text{int}} + (b_i + MC_{sl})q$ (see Figure 4.4). For proof, refer to Appendix A.

Attractor Lines (AL): The paths formed by any dispatch points of local maxima under all possible positions of RD. These are subsets of the contributory lines, possibly combined with a connecting segment along the RDVP (if this segment would be monotone). In particular, the sections of the contributory lines that occur on the same side of the RDVP as their corresponding section of the RD curve will be part of the attractor line (see Figure 4.5).

Pseudo-Offer (PO): The path formed by the set of globally optimal dispatch points under all possible realisations of RD (see Figure 4.5). This may or may not be monotone. If it is monotone in both the price and quantity dimensions, then this is simply the optimal offer for the generator to provide to the market. Sections 4.3.4 and 4.3.5 deal with establishing optimal monotonic transformations for pseudo-offers that are not naturally monotone.

At this point, we note that the optimal condition in Equation 4.4 is equivalent to our contributory line equation, while the plot of all points meeting this condition is equivalent to our attractor line. However, we define these terms and the simple process required to create them, as they can be used much more easily to establish the form of the desired offer under more general scenarios of uncertainty direction and RD form.

4.3.1.1 Interpretation of Key Terminology

To understand the CL, AL, and PO terms fully, consider the example shown in Figure 4.4. Assuming a linear marginal cost curve, take a RD curve with a constant slope $-b_1$ and plot the locus of optimal dispatch points (marginal revenue = marginal cost) under all possible realisations of RD (regardless of direction of movement). The result is a CL that has the form $p = MC_{\text{int}} + (b_1 + MC_{sl})q$. Under this very simple form of RD, this CL is equivalent to the PO and hence also the AL.

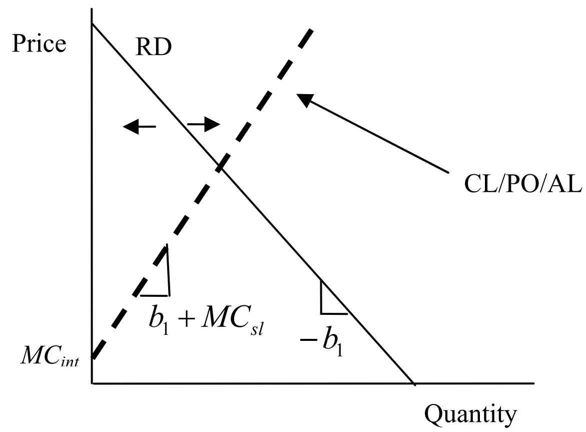

Figure 4.4 CL/PO/AL for a Simple Residual Demand Curve

Now consider a concave to the origin RD curve as shown in Figure 4.5, with slope $-b_1$ above the RDVP and slope $-b_2$ below. When we plot the locus of optimal dispatch points, we get a PO (black) that is monotonically non-decreasing in both the quantity and price dimensions. As defined above, the attractor line (shown in black) is composed of CL2 below the RDVP and CL1 above, along with a segment along RDVP connecting the two. In this particular case we can therefore observe that the PO is identical to the AL.

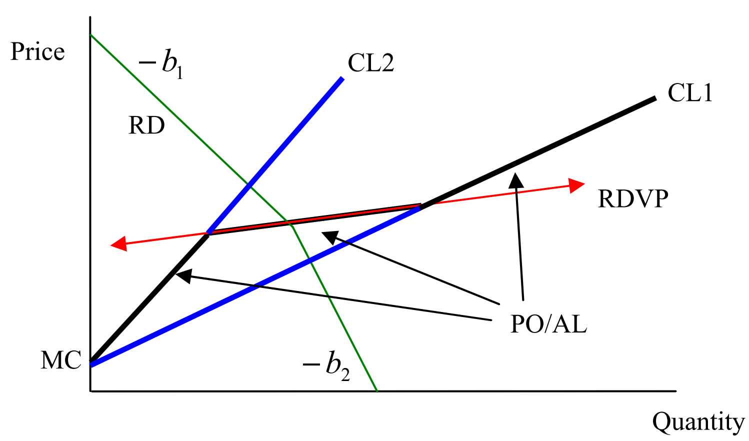

Figure 4.5 CL/PO/AL for a Concave Residual Demand

The interpretation of this AL, as demonstrated in Figure 4.6 is that any current offer (such as those shown in red) can be improved by moving it towards this line. In other words, it is a point wise attractor for market outcomes.

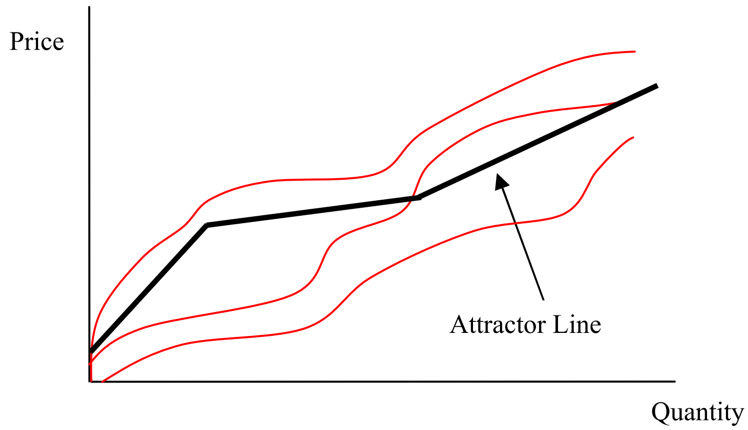

Figure 4.6 Attractor Line Concept

To understand why a shift of the current offer in the direction of this attractor line will give a guaranteed improvement, we must consider the form of the marginal revenue (MR) and marginal cost (MC) curves for each RD scenario. When the RD curve kinks in this concave manner, the MR curve is monotonically non-increasing (see Figure 4.7), and thus for a constant or non-decreasing MC, there is only a single intersection possible between MR and MC, and thus a single point of local maximum in the profit, this being the global optimal. Therefore, we know that whatever the intersection between a particular RD curve and the current offer, the profit achieved under that outcome of RD would have to be improved by moving that point of the offer curve towards the attractor line, which is the location of the optimal dispatch point under that scenario.

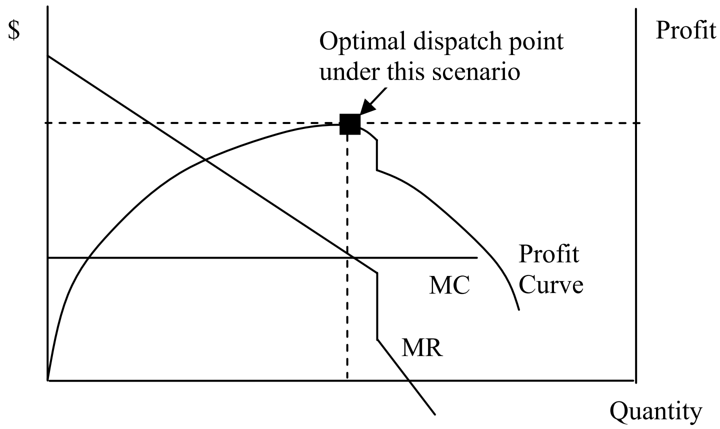

Figure 4.7 Optimal Dispatch Point under a single Residual Demand Scenario

Finally, consider a convex to the origin RD curve as shown in Figure 4.8. For lower levels of RD, the optimal dispatch position is on CL2, until a certain level (RDc), where the optimal dispatch positions jump to CL1. Therefore, the PO is as shown by the solid black lines. As before, the attractor line is composed of CL2 below the RDVP and CL1 above, as shown by the combination of the black solid and dotted lines.

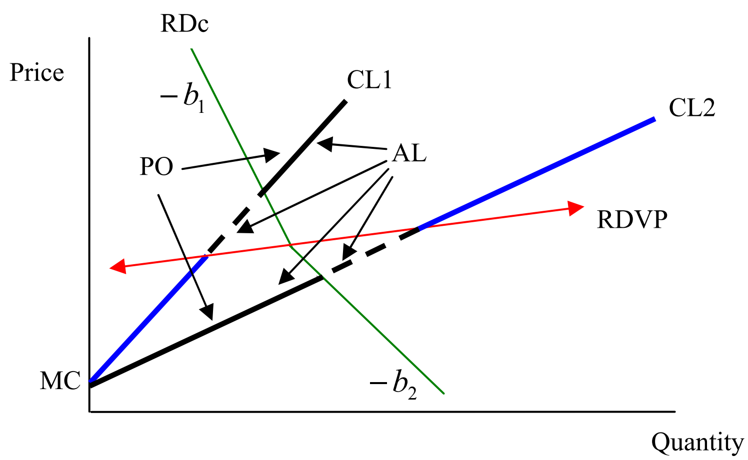

Figure 4.8 CL/PO/AL for a Convex Residual Demand

We can observe that the PO is clearly not monotone in this case, and as such, does not meet requirements for an offer to be made to the market. We define this particular type of non-monotonicity as Type I non-monotonicity.

Type I Non-Monotonicity: If the pseudo-offer has a step up and to the left, then define that pseudo-offer as having Type I non-monotonicity. This could be considered to be non-monotonicity in the quantity dimension.

At this point, we also define another type of non-monotonicity.

Type II Non-Monotonicity: If the pseudo-offer has a step down and to the right, then define that pseudo-offer as having Type II non-monotonicity. This could be considered to be non-monotonicity in the price dimension.

To illustrate Type II Non-Monotonicity, consider a concave RD curve. The black line in Figure 4.9 shows the form of the PO in this case, which contains a non-monotone step down and to the right.

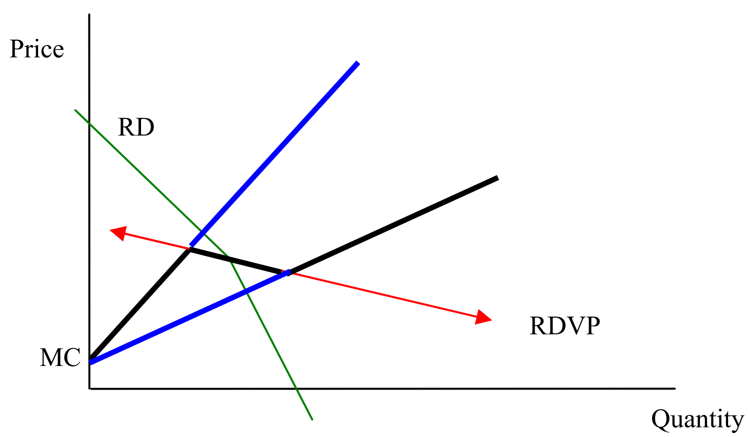

Figure 4.9 Concave RD under RDVP Slope Range 4

4.3.2 Form of Pseudo-Offers with respect to RDVP Slope

Figure 4.5 and Figure 4.8 demonstrated possible RDVP scenarios under concave and convex RD curves where the RDVP slope was positive, but less than the slopes of either CL. This led to monotone and non-monotone POs, respectively. There are many other

RDVP slope ranges that could occur, as demonstrated in Figure 4.10. Note that for two of these ranges, labelled 2 and 5, meaningful POs are not able to be constructed, and hence we consider these to be invalid RDVP slope ranges.

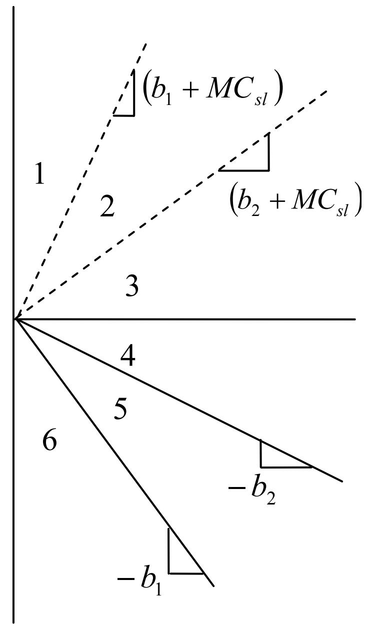

Figure 4.10 Possible RDVP Slope Ranges

The form of the PO will depend on the slope of the RDVP relative to the slopes of the two segments of the RD curve and the slopes of the two CLs, and on whether the RD is concave or convex. Consider the two slopes of the RD to be $-b_1$ and $-b_2$, while $(b_1 + MC_{sl})$ and $(b_2 + MC_{sl})$ are the slopes of the associated CLs, as indicated in Figure 4.10. Therefore, we can consider six different possible RDVP slope ranges (labelled 1-6), as shown in Table 4.1.

RDVP Slope Range	Convex Residual Demand	Concave Residual Demand
1	Non-Monotonicity Type II	Monotone
2	*Invalid RDVP Slope Range*	*Invalid RDVP Slope Range*
3	Non-Monotonicity Type I	Monotone
4	Non-Monotonicity Type I	Non-Monotonicity Type II
5	*Invalid RDVP Slope Range*	*Invalid RDVP Slope Range*
6	Non-Monotonicity Type II	Non-Monotonicity Type I

Table 4.1 General Forms of Pseudo-Offers

Clearly there are certain (limited) RD shape and uncertainty direction (RDVP) combinations under which the optimal set of dispatch points for all possible RD curve outcomes would form an offer that is naturally monotonic in both the price and quantity dimensions. The remainder of Section 4.3 deals with the remaining RD shape and uncertainty direction scenarios, where the set of optimal dispatch points are non-monotonic, and establishes the best way to resolve the non-monotonicity in these cases. We begin by determining the possible forms of the profit curve, defined over the dispatch level.

4.3.3 Form of the Single-Scenario Profit Curve

In order to determine how best to overcome the issues of non-monotonicity in POs, we need to consider the possible forms of the profit curve for a single RD scenario, under both convex and concave RD curves. In particular, in this section we will show that the profit curve for a concave RD curve must be *unimodal*, and that the profit curve for a convex RD curve (with two segments) must be either unimodal or *bimodal*. By understanding the forms of the profit functions, we can then understand how the non-monotonicity in the POs comes about, and hence how best to deal with it. Sections 4.3.4 and 4.3.5 show how knowledge of this unimodularity or bimodularity can be used to create an optimal monotonic transformation of Type I and Type II non-monotonic POs, respectively.

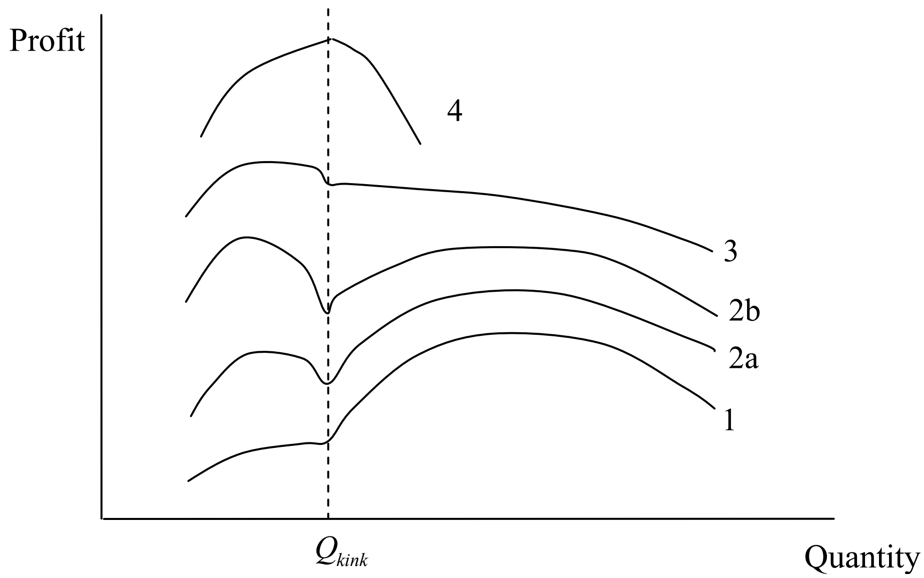

Figure 4.11 Possible Forms of the Profit Curve

Figure 4.11 demonstrates possible forms of the profit curve around the quantity level at which the RD curve kinks (Q_{kink}). These will be referred to throughout this section.

4.3.3.1 Bimodularity of Profit Function for Convex Residual Demand Curves

We will begin by proving the unimodularity/bimodularity of the profit curve for any convex to the origin RD scenario with a single kink. We will also demonstrate that if we are faced with a RD curve with multiple local maxima, then the local minimum, between the two points of local maxima, must be located at the quantity where the kink occurs in the RD curve.

We can consider the RD curve to have the form:

$$P = a(q) - b(q)q$$

where *a(q)* and *b(q)* are the intercept and slope terms respectively, defined in the quantity dimension and with respect to the particular RD scenario that is being considered (a_i and b_i are the intercept and slope associated with a particular RD segment *i*). We note that these are step functions as demonstrated in Figure 4.17, where the values change only at Q_{kink}.

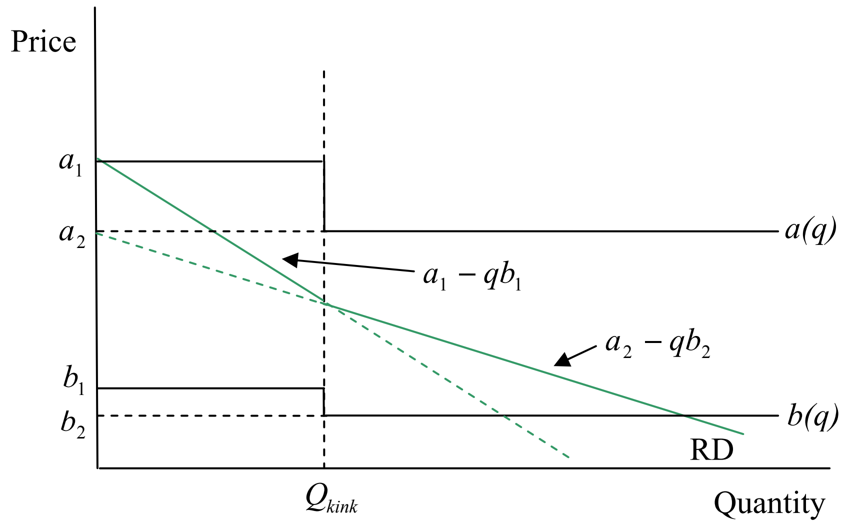

Figure 4.12 Convex Residual Demand Curve and the *a* and *b* functions

For such a RD curve, we also note that marginal revenue has the form:

$$MR = a(q) - 2b(q)q$$

For a linear marginal cost, we can calculate the profit at any quantity, using the formula:

$$\begin{aligned}\text{Profit} &= q(a(q) - b(q)q) - \left(qMC_{\text{int}} + \frac{1}{2}MC_{sl}q^2\right)\\ &= qa(q) - q^2b(q) - \left(qMC_{\text{int}} + \frac{1}{2}MC_{sl}q^2\right)\end{aligned}$$

The derivative with respect to *q* gives the rate of change of profit (or the slope of the profit curve) as we move in the quantity dimension:

$$\frac{\partial \text{Profit}}{\partial q} = a(q) - 2qb(q) - (MC_{\text{int}} + MC_{sl}q) = MR - MC$$

As we move from $q = 0$ up to $q = Q_{kink}$, we can observe that $a(q)$, $b(q)$, MC_{int}, and MC_{sl} are all constant, hence the only terms in the profit derivative that change are the second and fourth, and thus the derivative (the slope) must be constantly decreasing (assuming that $b(q) \neq 0$). This means that this segment of the profit function must contain at most a single local maximum.

Likewise, beyond $q = Q_{kink}$, $a(q)$, $b(q)$, MC_{int}, and MC_{sl} again remain constant, and thus the derivative (and slope) must again constantly decrease. This means that this segment of the profit function also can contain only a single local maximum.

At $q = Q_{kink}$, both $a(q)$ and $b(q)$ change. This can result in four different possible scenarios. The first is the case where both the right and the left derivatives of profit at this quantity level are positive values (i.e. positive slope either side of the breakpoint). Mathematically, this is described:

$$a_1 - 2b_1Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink}, a_2 - 2b_2Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink} > 0$$

This corresponds to profit form 1 in Figure 4.11, and indicates that the curve must be unimodal, with the global maximum lying to the right of Q_{kink}.

The second scenario is the case where both the right and the left derivatives of profit at this quantity level are negative values:

$$a_1 - 2b_1Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink}, a_2 - 2b_2Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink} < 0$$

This corresponds to profit form 3 in Figure 4.11, and also indicates that the curve must be unimodal, this time with the global maximum lying to the left of Q_{kink}.

Finally, the left derivative may be negative, while the right derivative is positive, at this point (as in profit forms 2a and 2b in Figure 4.11). This indicates that at Q_{kink}, the slope of the profit curve is changing from negative to positive, hence this point is a local minimum. This also implies that the profit curve will be bimodal, but we are unsure as to which side of Q_{kink} the global maximum will lie. Mathematically, these cases are described:

$$a_1 - 2b_1Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink} < 0 < a_2 - 2b_2Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink}$$

Note that under this convex RD scenario, case 4, where the global optimum occurs on the RDVP, can not exist.

Hence, we have shown that for a convex RD curve, the profit curve must be either unimodal or bimodal, and if it is bimodal, then the local minimum of the curve must be located at Q_{kink}. It is interesting to understand both how this bimodularity can occur in the profit function for a convex RD curve, and the meaning of the non-globally optimal local points of optima.

In order to comprehend these issues, consider the case of a convex RD curve under a RDVP in slope range 3, as shown in Figure 4.13. As the RD curve moves out to the right, it is initially optimal to be dispatched on the lower segment of the RD curve. In particular, for these RD curves, the profit curve is unimodal. As the RD curve moves to the right, there comes a point (RDmin) where the profit curve becomes bimodal (develops multiple local optima), while retaining a single global optimum on the lower segment of the RD. Once the RD passes RDc, the global optimum jumps up to the upper segment of the RD curve, thus causing the Type I non-monotonicity observed in the PO. Once the RD gets beyond RDmax, the profit curve returns to being unimodal, but with the global optima now on the upper segment of the RD curve. It is important to note that for the range of RD curves that produce bimodal profit curves (RDmin to RDmax), the non-globally optimal points of local optima in the profit curve (dotted segments) correspond to points on the ALs that are not part of the PO.

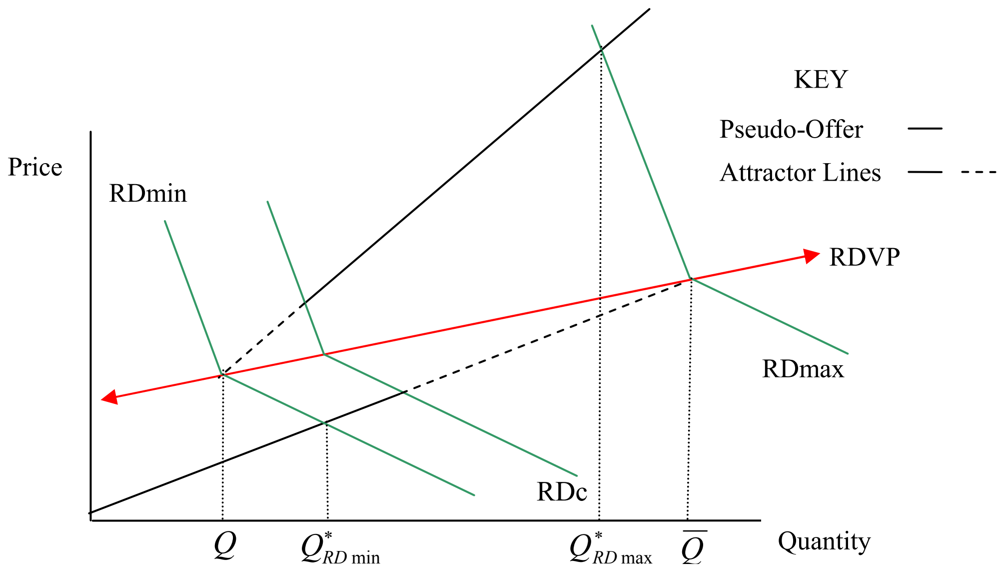

Figure 4.13 Source of Local Points of Optima under Convex RD

With respect to the profit curves shown in Figure 4.11, observe that in Figure 4.13:

- RDmin is the position of the RD curve where the profit curve shifts from form 1 to form 2a
- RDc is the position of the RD curve where the profit curve shifts from form 2a to form 2b
- RDmax is the position of the RD curve where the profit curve shifts from form 2b to form 3.

The multiple local optima described above for RD curves between RDmin and RDmax are caused by the shape of the marginal revenue (MR) curve for this form of RD. The particular issue is that the MR curve is not monotonically non-increasing, and thus, can potentially have multiple intersections with the monotonically non-decreasing marginal cost curve. To demonstrate, the leftmost pair of diagrams in Figure 4.14 show the situation where we have the RD curve RDmin, and the rightmost pair show the situation where we have RDmax. This case assumes that the marginal cost is constant and equal to zero, but the concepts are equally valid for any non-decreasing marginal cost function. For any RD curve in between these two, there will be multiple local optima in the revenue curve. Note that for a concave RD curve, the MR curve will be monotonically non-increasing (as demonstrated in Figure 4.7, and proven in Section 4.3.3.2), and as such, there will be no points of local optima that are not globally optimal. Hence, under that RD shape, the PO will simply be the same as the ALs.

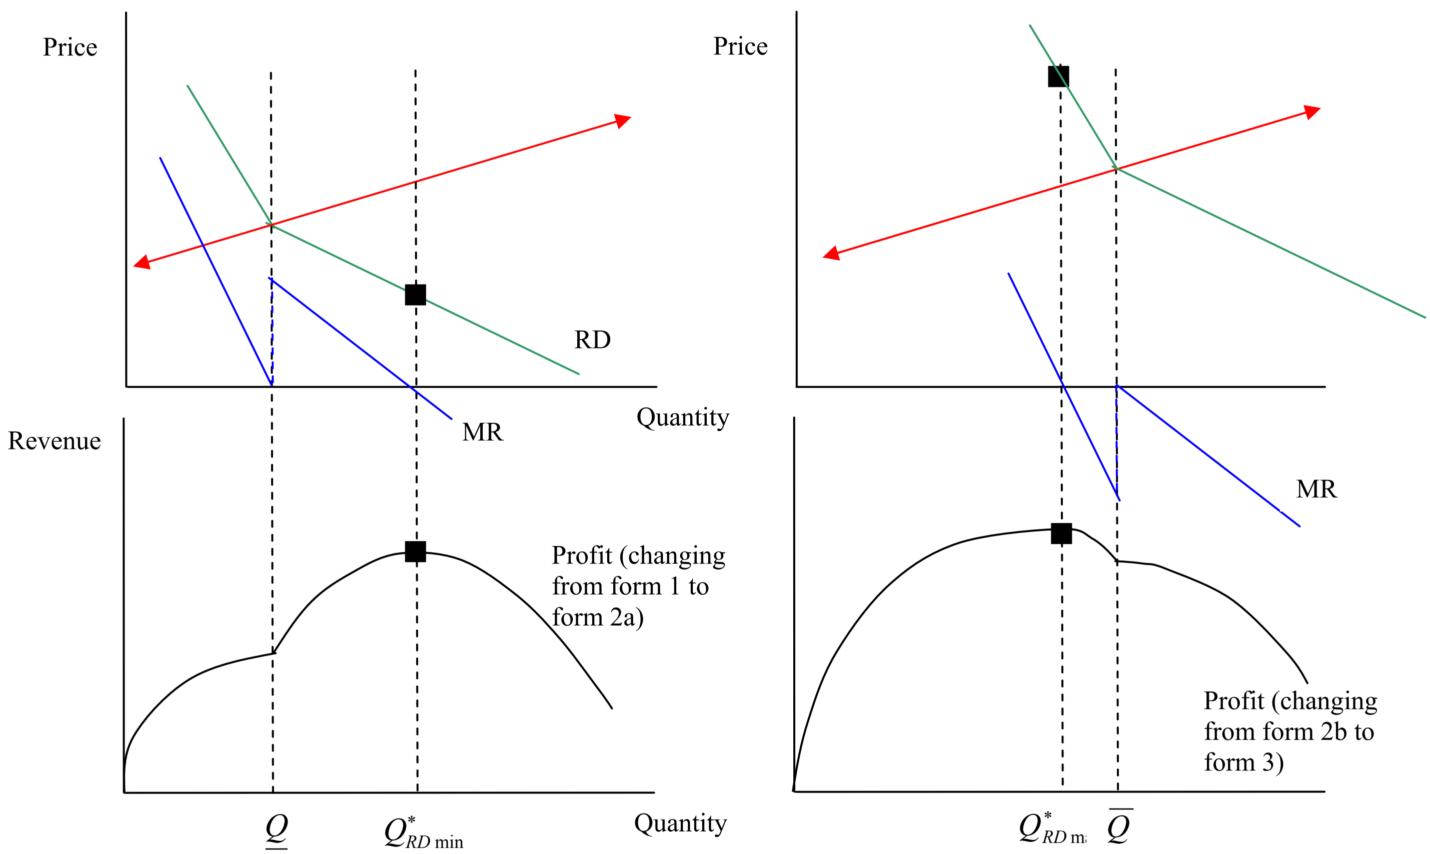

Figure 4.14 Marginal Revenue and Profit Curves

Equivalent figures could similarly be drawn for other RDVP slope ranges.

4.3.3.2 Unimodularity of Profit Function for Concave Residual Demand Curves

We will now prove the unimodularity of the profit curve for any concave to the origin RD scenario with a single kink.

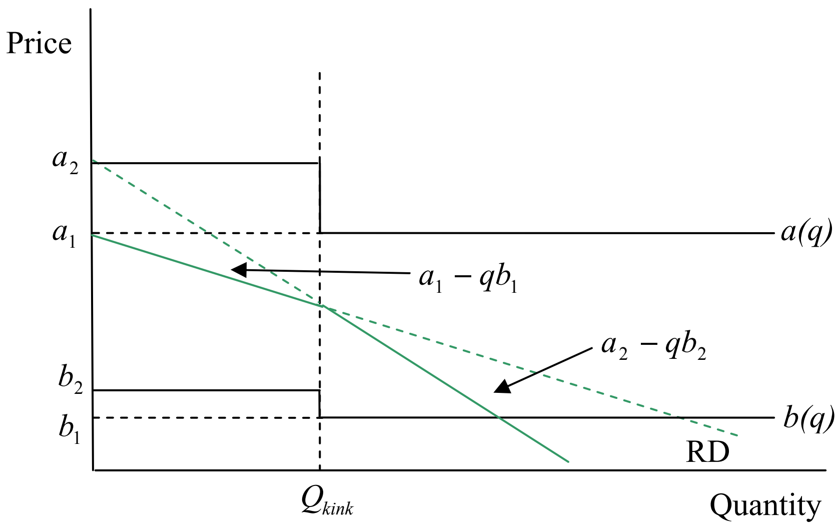

Figure 4.15 Concave Residual Demand Curve and the *a* and *b* functions

As in the convex case, we can easily show that the slope of the profit curve up to and beyond Q_{kink} must be constantly decreasing.

For a concave RD curve, both *a(q)* and *b(q)* increase at Q_{kink}, as shown in Figure 4.15. In this RD case, the slope of the profit curve immediately before Q_{kink} is equal to $a_1 - 2b_1Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink}$, while the slope immediately after Q_{kink} is given by $a_2 - 2b_2Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink}$. Therefore, to know what happens to the slope of the profit curve at Q_{kink}, we need to establish a relationship between these two slope equations. Start by re-writing the equations as follows:

$$\begin{aligned}
a_1 - 2b_1Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink} &= a_1 - b_1Q_{kink} - MC_{\text{int}} - MC_{sl}Q_{kink} - b_1Q_{kink} \\
&= RD_1(Q_{kink}) - MC_{\text{int}} - MC_{sl}Q_{kink} - b_1Q_{kink}
\end{aligned}$$

and

$$a_2 - 2b_2 Q_{kink} - MC_{\text{int}} - MC_{sl} Q_{kink} = a_2 - b_2 Q_{kink} - MC_{\text{int}} - MC_{sl} Q_{kink} - b_2 Q_{kink}$$
$$= RD_2(Q_{kink}) - MC_{\text{int}} - MC_{sl} Q_{kink} - b_2 Q_{kink}$$

By definition, $RD_1(Q_{kink}) = RD_2(Q_{kink})$, as this is the location where the kink occurs in the RD curve. Therefore, given that the marginal cost terms in each equation are equivalent and that $b_2 > b_1$, we can see that the slope of the profit function must be lower to the right of the breakpoint than it is to the left.

In other words, we can conclude that the slope of the profit curve is decreasing over the entire valid quantity range for a concave RD curve, hence proving that the profit curve is unimodal. If both slopes are negative at Q_{kink}, then the profit maximum will lie to the left of this point, and vice versa if both slopes are positive at Q_{kink}. If the slope to the left of Q_{kink} is positive, while the slope to the right is negative, then we have case 4 shown in Figure 4.11, where the maximum is at Q_{kink}, or in other words, on the RDVP.

4.3.4 Optimal Monotonic Transformations – Type I Non-Monotonicity

So far in this section, we have established two forms of non-monotonicity that need to be adjusted to produce monotone offers that could be provided by a generator to the market. Such an adjustment is required as most markets require that submitted offers must be monotonically non-decreasing in both the price and quantity dimensions.

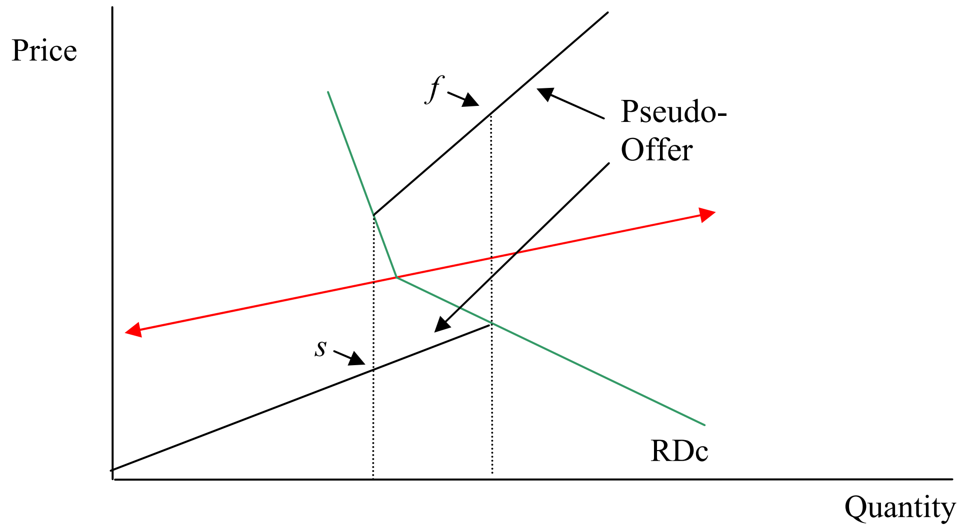

Figure 4.16 Type I Non-Monotonicity – Range to be Transformed

Type I non-monotonicity involves a PO that jumps up and to the left, as shown in Figure 4.16. The question is therefore, how can this pseudo-offer best be transformed to a feasible offer that meets market monotonicity requirements? To be precise, we need to know how optimally to connect the points *s* and *f*. Before addressing this issue, we note that the optimal offer up to point *s* and beyond point *f* must follow along the PO. This is because there is just a single, global optimum corresponding to each of the RD scenarios that could lead to dispatch points in these regions, and these optimal points exist along the attractor lines stated.

4.3.4.1 Form of the Optimal Monotonic Transformation

In this section, we will hypothesise and prove the optimal form of the monotonic transformation of a PO that is subject to Type I non-monotonicity.

Theorem 4.1:

Consider a pseudo-offer with Type I non-monotonicity, such as that in Figure 4.16. The optimal offer will trace along the *lower* (*right*) segment of the pseudo-offer curve from point *s*, and then at some *quantity* point, will jump up to the *upper* (*left*) segment of the pseudo-offer curve, following this along to point *f*. If we take any existing offer segment from point *s* to point *f* that does not consist of part of each of the lower and upper attractor lines connected by a vertical segment in this manner, then the expected payoff must be improved by transforming the offer to this form, such that the vertical segment passes through the intersection of the original offer with RDVP. This will *not* necessarily be the optimal position for the vertical jump, but the new offer must be better than the original offer.

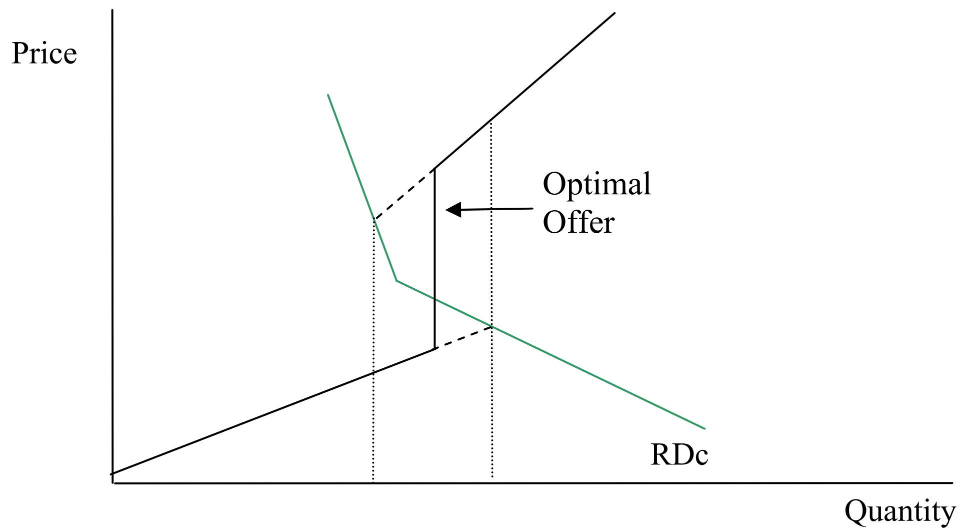

Figure 4.17 Demonstration of Vertical Jump

Proof Using Direction of Improvement Theory:

To show that the optimal connection must indeed contain a vertical jump from the lower to the upper AL, we need to consider the direction of improvement from any given point of dispatch. Take any monotone offer path from *s* to *f*, such as that indicated in blue in Figure 4.18, and consider the resulting dispatch points along this blue segment for each of the RD scenarios and how they could be improved.

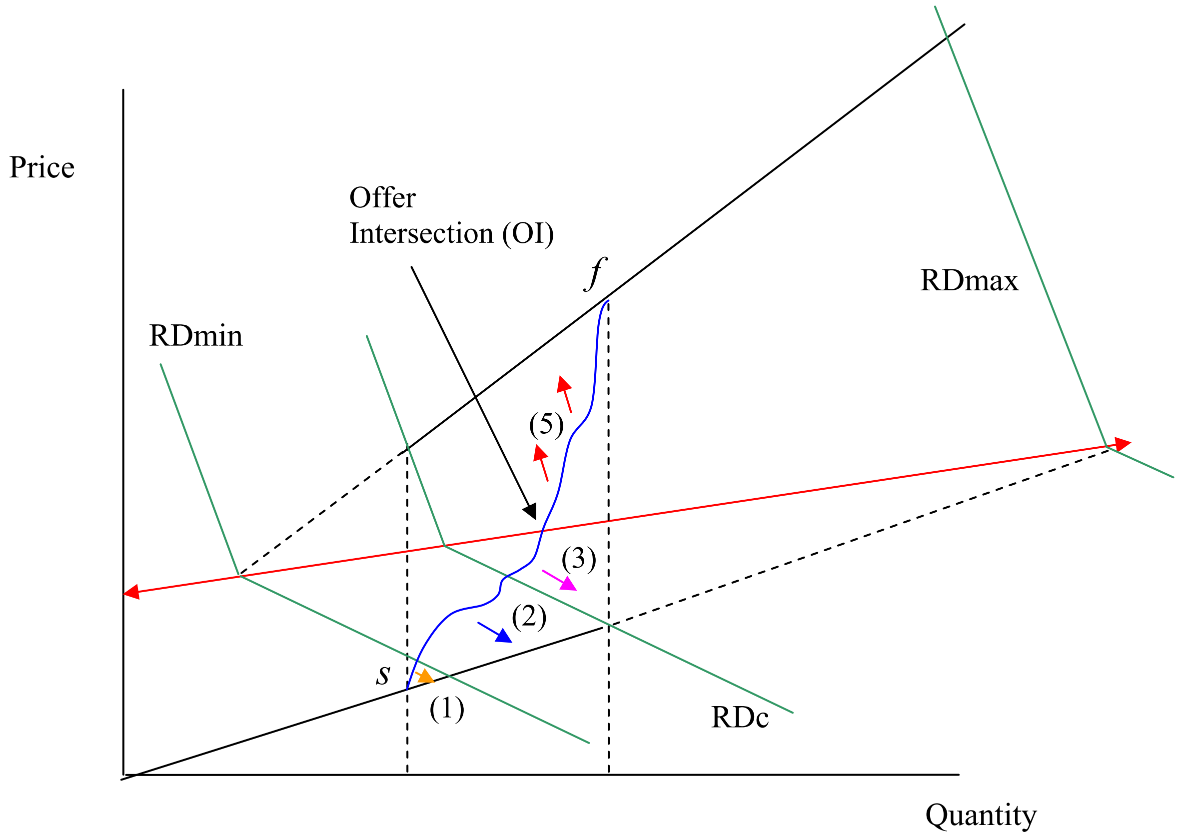

Figure 4.18 A Non-Optimal Offer and Associated Directions of Improvement

1) For the points of dispatch on the RD curves up to RDmin, we are currently up and to the left of the global optimum on this unimodal profit curve, and so know that the direction of improvement along the RD curve (as indicated by the yellow arrow) must be down and to the right. This is demonstrated by the revenue diagrams in Figure 4.14.

2) The profit curves corresponding to RD curves between RDmin and RDc are bimodal and hence there are multiple locally optimal points of dispatch. We are presently up and to the left of the global optimum, and to the right of the RDVP (local minimum in profit curve). Hence the local direction of improvement (as indicated by the blue arrow) is down and to the right.

3) For the points of dispatch on the RD curves between RDc and the point where the current offer (blue line) intersects RDVP (i.e. the point labelled Offer Intersection, OI),

there are again multiple local optima. We are now down and to the right of the global optimum, *but* also to the right of the RDVP. Hence, although the global direction of improvement is up and to the left, the *local* direction of improvement (as indicated by the pink arrow) is down and to the right.

4) If the current offer were to intersect RDVP to the left of RDc, then the points of dispatch on the RD curves between these points would also have multiple local optima. We would be up and to the left of the global optimum, and also to the left of the breakpoint. Hence, the direction of improvement (not shown on this figure) would be up and to the left.

5) For the points of dispatch on the RD curves between the point where the current offer (blue line) intersects RDVP and RDmax, there are multiple local optima. We are down and to the right of the global optimum and to the left of the RDVP. Hence, the direction of improvement (as indicated by the red arrows) is up and to the left.

6) If the current offer were to intersect RDmax before reaching point b, then for the points of dispatch on the RD curves beyond RDmax, we would be down and to the right of the global optimum, and so we know that the direction of improvement (not shown on this figure) would be up and to the left.

To summarise, we know that for all points of dispatch that are currently below the intersection with the RDVP, the direction of improvement is down and to the right, while the direction of improvement for all points above this intersection is up and to the left. Put another way, a point of dispatch is attracted to the intersection between its RD curve and the AL on the same side of the RDVP as that point is currently located.

As demonstrated in Figure 4.19, this implies that we can construct an offer that follows along the lower AL up until the level Q_{jump} (which is the quantity level at which the original offer intersected RDVP), then jumps up to the upper AL, and carries on along this line. The new offer is then guaranteed to have a better payoff than the original offer, as all the dispatch points have moved in their local direction of improvement, towards their attractor lines.

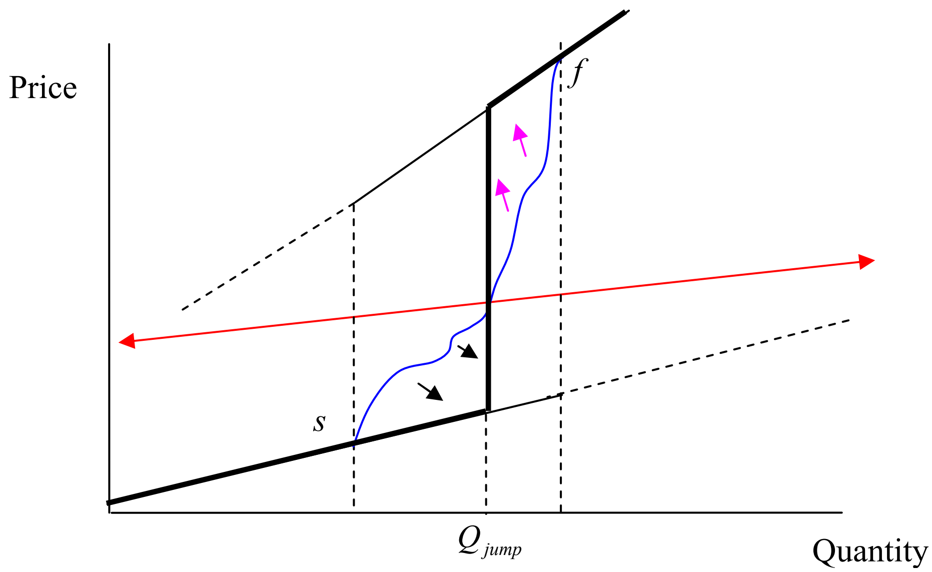

Figure 4.19 A Transformation of the Offer, Guaranteed to be an Improvement

This same logic can be applied to any curve through this range that is not of this form. Therefore, for any such curve, there must be a feasible improvement to be made by adjusting the offer to make it vertical between the lower and upper attractor lines, passing through the offer intersection point with RDVP.

Q.E.D.

4.3.4.2 Vertical Jump Location Optimality Condition

We have shown above that the optimal monotone offer must involve a segment of the lower AL, a segment of the upper AL, and a vertical jump between the two. The question remains as to where exactly this jump should occur. Clearly, the optimal location of the jump will depend on the payoffs and probability distribution of the particular case that is being considered. In general though, we can form an optimality condition on the derivative of the profit with respect to the quantity level of the jump, which will determine the optimal position of the jump in all cases.

The profit equation that we need in order to create an optimality condition on the position of the vertical jump is the integral over all possible RD curve scenarios.

Term Definitions

sc is a RD position scenario index used to define the level of the RD curve.

pr(sc) is the probability density function of RD scenario level.

a(sc,q) is a surface defined over the dimensions of RD scenario level and quantity, which provides the intercept term for the RD curve.

b(sc,q) is a surface defined over the dimensions of RD scenario level and quantity, which provides the slope term for the RD curve.

u(q) is the highest RD scenario level that leads to a dispatch on the vertical segment of the offer curve for the current quantity position of that segment.

l(q) is the lowest RD scenario level that leads to a dispatch on the vertical segment of the offer curve for the current quantity position of that segment.

q is the quantity level where the vertical connection segment will be placed in the offer curve.

For a given scenario, we know the functions *a(sc,q)*, *b(sc,q)* and *pr(sc)* in advance, *q* is the decision that we are making, while *l(q)* and *u(q)* will directly depend on this decision.

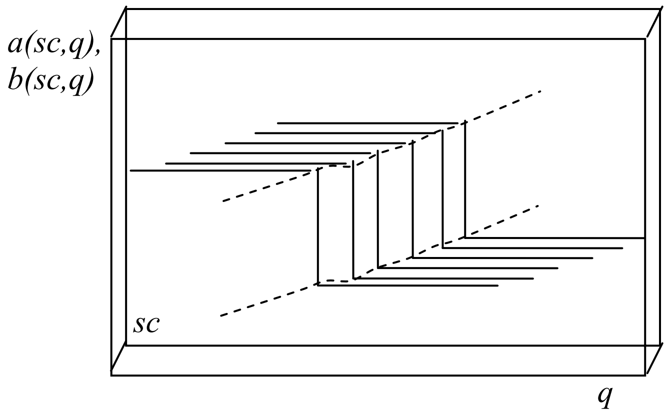

Figure 4.20 The *a(sc,q)* and *b(sc,q)* Functions

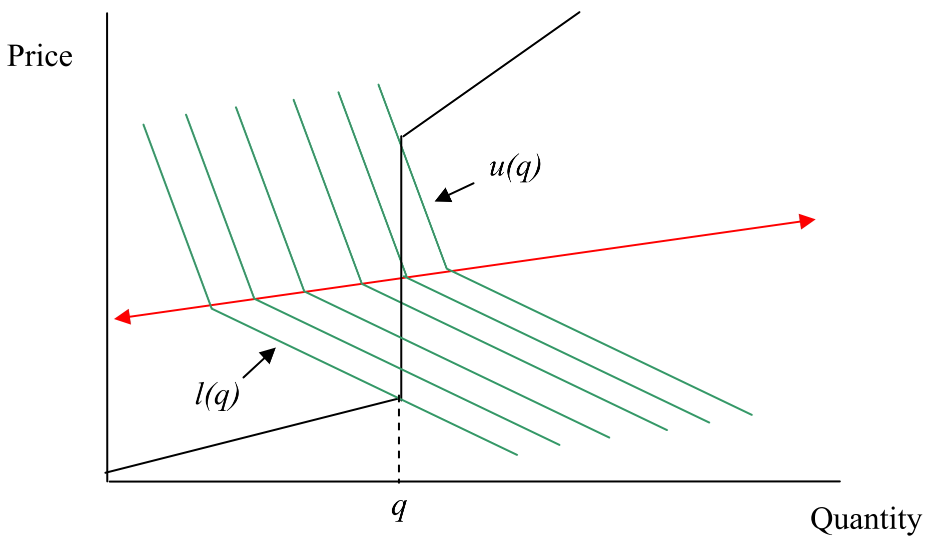

Figure 4.21 The *l(q)* and *u(q)* Terms

Note: *l(q)* and *u(q)* are the RD scenario levels that lead to these RD curves

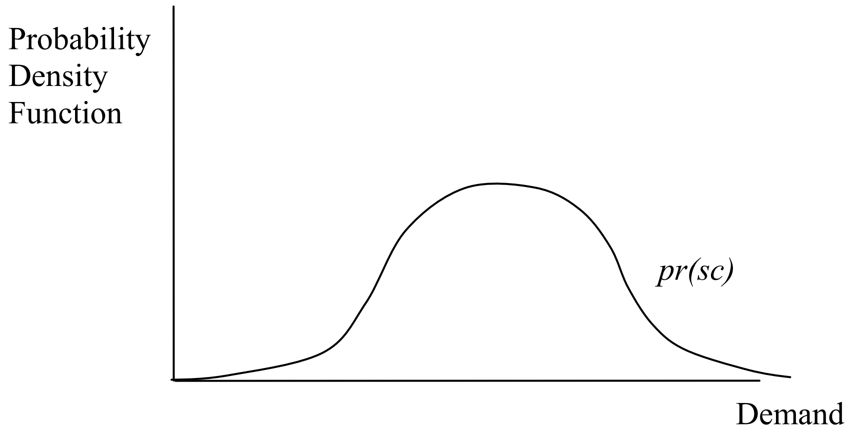

Figure 4.22 The *pr(sc)* Probability Density Function

Optimal Dispatch Points before Vertical Jump

From the concept of contributory and attractor lines, we know that the AL below the RDVP has slope equal to the negative of the slope segment of RD that is below RDVP plus the marginal cost curve slope. Likewise, we know that the AL above RDVP has slope equal to negative of the slope segment of RD that is above RDVP plus the marginal cost curve slope. We can therefore denote the slope of the AL below RDVP as $b(sc,\inf)+MC_{sl}$, and the slope on the AL above it as $b(sc,-\inf)+MC_{sl}$. This is because -*b(sc,inf)* is the slope that occurs at all points beyond the kink in the RD curve (and hence must always occur at q = infinity), and -*b(sc,-inf)* is the slope that occurs at all points before the kink (and hence must occur at q = -infinity). This information is shown in Figure 4.23.

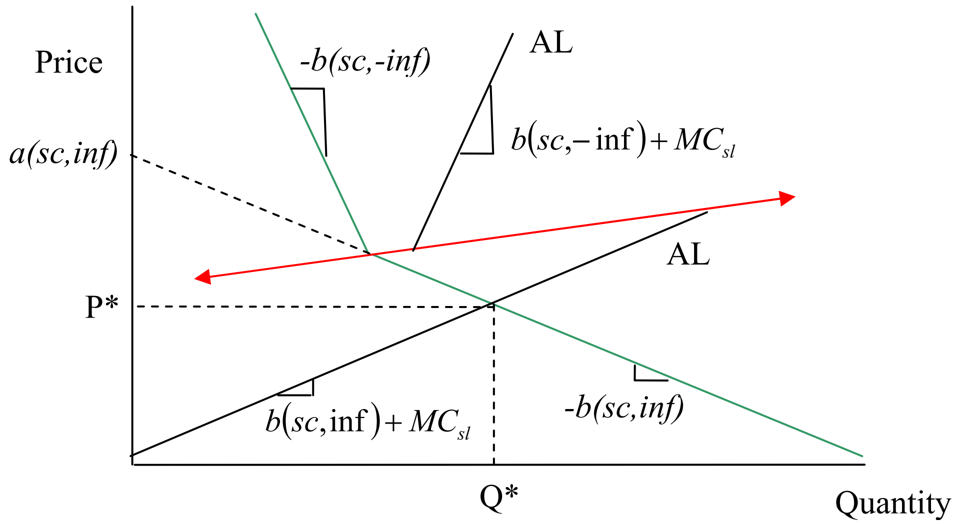

Figure 4.23 Attractor Line Slopes

If, under the current position of the RD curve, we are to be dispatched on the lower AL, we can analytically determine our dispatch quantity and price levels. We know from Section 4.3.1.1 that the optimal point is at the intersection of the AL (form $p = MC_{\text{int}} + q(b(sc,\inf) + MC_{sl})$), and the associated RD segment (form $p = a(sc,\inf) - qb(sc,\inf)$). Therefore, we find that the optimal dispatch quantity and price are:

$$Q^* = \frac{a(sc,\inf) - MC_{\text{int}}}{2b(sc,\inf) + MC_{sl}}$$

$$P^* = a(sc,\inf) - b(sc,\inf)\left[\frac{a(sc,\inf) - MC_{\text{int}}}{2b(sc,\inf) + MC_{sl}}\right]$$

Therefore, given a cost function $c(q)$ and a simple revenue term pq , we can show that the profit contribution from the RD scenario levels up to *l(q)* is as follows (with all dispatch points on the lower AL):

$$profit = \int_{0}^{l(q)} pr(sc)\left[\begin{array}{l}\left[a(sc,\inf) - b(sc,\inf)\left[\dfrac{a(sc,\inf) - MC_{\text{int}}}{2b(sc,\inf) + MC_{sl}}\right]\right]\left[\dfrac{a(sc,\inf) - MC_{\text{int}}}{2b(sc,\inf) + MC_{sl}}\right] \\ -c\left(\dfrac{a(sc,\inf) - MC_{\text{int}}}{2b(sc,\inf) + MC_{sl}}\right)\end{array}\right].dsc$$

Optimal Dispatch Points after Vertical Jump

Similarly, we can show that the profit contribution from the RD scenario levels beyond *u(q)* is as follows (with all dispatch points on the upper AL):

$$profit = \int_{u(q)}^{\inf} pr(sc)\left[\begin{array}{l}\left[a(sc,-\inf) - b(sc,-\inf)\left[\dfrac{a(sc,-\inf) - MC_{\text{int}}}{2b(sc,-\inf) + MC_{sl}}\right]\right]\left[\dfrac{a(sc,-\inf) - MC_{\text{int}}}{2b(sc,-\inf) + MC_{sl}}\right] \\ -c\left(\dfrac{a(sc,-\inf) - MC_{\text{int}}}{2b(sc,-\inf) + MC_{sl}}\right)\end{array}\right].dsc$$

Optimal Dispatch Points on the Vertical Jump

For all RD scenario levels between *l(q)* and *u(q)*, the form of the profit equation is slightly different, due to the fact that we know the quantity level in this range (as it is the vertical section of the offer):

$$profit = \int_{l(q)}^{u(q)} pr(sc)[q[a(sc,q) - b(sc,q)q] - c(q)].dsc$$

Profit Formula and Derivative

Hence, the total profit equation is:

$$profit = \int_{l(q)}^{u(q)} pr(sc)\big[q\big[a(sc,q)-b(sc,q)q\big]-c(q)\big].dsc$$

$$+\int_{0}^{l(q)} pr(sc)\left[\begin{array}{l}\left[a(sc,\inf)-b(sc,\inf)\left[\dfrac{a(sc,\inf)-MC_{\text{int}}}{2b(sc,\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(sc,\inf)-MC_{\text{int}}}{2b(sc,\inf)+MC_{sl}}\right]\\ -c\left(\dfrac{a(sc,\inf)-MC_{\text{int}}}{2b(sc,\inf)+MC_{sl}}\right)\end{array}\right].dsc$$

$$+\int_{u(q)}^{\inf} pr(sc)\left[\begin{array}{l}\left[a(sc,-\inf)-b(sc,\inf)\left[\dfrac{a(sc,-\inf)-MC_{\text{int}}}{2b(sc,-\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(sc,-\inf)-MC_{\text{int}}}{2b(sc,-\inf)+MC_{sl}}\right]\\ -c\left(\dfrac{a(sc,-\inf)-MC_{\text{int}}}{2b(sc,-\inf)+MC_{sl}}\right)\end{array}\right].dsc$$

The derivative of profit is therefore:

$$\frac{\partial profit}{\partial q} = \int_{l(q)}^{u(q)}\left(pr(sc)\left[a(sc,q)-b(sc,q)q+q\left[\frac{\partial a(sc,q)}{\partial q}-\frac{\partial b(sc,q)}{\partial q}q-b(sc,q)\right]-\frac{\partial c(sc)}{\partial q}\right]\right).dsc +$$

$$\frac{\partial u(q)}{\partial q} pr(u(q))\big[q\big[a(u(q),q)-b(u(q),q)q\big]-c(q)\big]$$

$$-\frac{\partial l(q)}{\partial q} pr(l(q))\big[q\big[a(l(q),q)-b(l(q),q)q\big]-c(q)\big]$$

$$+\frac{\partial l(q)}{\partial q} pr(l(q))\left[\begin{array}{l}\left[a(l(q),\inf)-b(l(q),\inf)\left[\dfrac{a(l(q),\inf)-MC_{\text{int}}}{2b(l(q),\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(l(q),\inf)-MC_{\text{int}}}{2b(l(q),\inf)+MC_{sl}}\right]\\ -c\left(\dfrac{a(l(q),\inf)-MC_{\text{int}}}{2b(l(q),\inf)+MC_{sl}}\right)\end{array}\right]$$

$$-\frac{\partial u(q)}{\partial q} pr(u(q))\left[\begin{array}{l}\left[a(u(q),-\inf)-b(u(q),-\inf)\left[\dfrac{a(u(q),-\inf)-MC_{\text{int}}}{2b(u(q),-\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(u(q),-\inf)-MC_{\text{int}}}{2b(u(q),-\inf)+MC_{sl}}\right]\\ -c\left(\dfrac{a(u(q),-\inf)-MC_{\text{int}}}{2b(u(q),-\inf)+MC_{sl}}\right)\end{array}\right]$$

Of the five terms in the derivative above, the first three are the derivative of the profit term for demand levels from *l(q)* up to *u(q)*, while the fourth and fifth terms are from the derivatives of the profit terms for *0* to *l(q)* and *u(q)* to *infinity* respectively. Hence, the only terms in the profit derivative for demand levels up to *l(q)* and beyond *u(q)* relate to the breakpoints themselves, as the rest of the terms are unaffected by a marginal change in the level of the quantity jump-point. These terms then cancel out with the final two terms of the derivative of profit between *l(q)* and *u(q)*. In other words,

$$-\frac{\partial l(q)}{\partial q} pr(l(q))[q[a(l(q),q)-b(l(q),q)q]-c(q)]$$

cancels out with

$$+\frac{\partial l(q)}{\partial q} pr(l(q))\left[\begin{array}{l}\left[a(l(q),\inf)-b(l(q),\inf)\left[\frac{a(l(q),\inf)-MC_{\mathrm{int}}}{2b(l(q),\inf)+MC_{sl}}\right]\right]\left[\frac{a(l(q),\inf)-MC_{\mathrm{int}}}{2b(l(q),\inf)+MC_{sl}}\right] \\ -c\left(\frac{a(l(q),\inf)-MC_{\mathrm{int}}}{2b(l(q),\inf)+MC_{sl}}\right)\end{array}\right]$$

and

$$\frac{\partial u(q)}{\partial q} pr(u(q))[q[a(u(q),q)-b(u(q),q)q]-c(q)]$$

cancels out with

$$-\frac{\partial u(q)}{\partial q} pr(u(q))\left[\begin{array}{l}\left[a(u(q),-\inf)-b(u(q),-\inf)\left[\frac{a(u(q),-\inf)-MC_{\mathrm{int}}}{2b(u(q),-\inf)+MC_{sl}}\right]\right]\left[\frac{a(u(q),-\inf)-MC_{\mathrm{int}}}{2b(u(q),-\inf)+MC_{sl}}\right] \\ -c\left(\frac{a(u(q),-\inf)-MC_{\mathrm{int}}}{2b(u(q),-\inf)+MC_{sl}}\right)\end{array}\right]$$

Although the terms appear to be slightly different, by substituting $\frac{a(sc,q)-MC_{\text{int}}}{2b(sc,q)+MC_{sl}}=q$ (as previously defined), it can easily be shown that the two pairs of terms are equivalent. Therefore, the profit derivative reduces down to the equation:

$$\int_{l(q)}^{u(q)}\left(pr(sc)\left[a(sc,q)-b(sc,q)q+q\left[\frac{\partial a(sc,q)}{\partial q}-\frac{\partial b(sc,q)}{\partial q}q-b(sc,q)\right]-\frac{\partial c(q)}{\partial q}\right]\right).dsc$$

Optimality Condition

In the specific case that we are considering here, with a single kink in the piecewise linear RD curve, this equation can be simplified significantly. The derivative of both the *a* and *b* functions with respect to *q* is equal to zero at all points, giving the optimality condition:

$$\int_{l(q)}^{u(q)} pr(sc)\left[a(sc,q)-2b(sc,q)q-\left[MC_{\text{int}}+MC_{sl}q\right]\right].dsc=0$$

Given that the RD curve is equal to $a(sc,q)-b(sc,q)q$, we know that the marginal revenue curve must be given by the equation $a(sc,q)-2b(sc,q)q$, as observed within the optimality condition. Therefore, the optimality condition says that we take the integral of the probability distribution of demand multiplied by *MR* - *MC* over the range of RD curve scenarios that would lead to dispatch on the vertical segment of the offer, and find the quantity for which this equation is equal to zero.

$$\int_{l(q)}^{u(q)} pr(sc)\left[MR(q)-MC(q)\right].dsc=0$$

Or

$$\int_{l(q)}^{u(q)} pr(sc)[\text{Marginal Profit}(q)].dsc = 0$$

If, rather than having a continuous probability density function for the RD level, there was actually a set of discrete demand levels that could occur (with a corresponding probability of each occurrence being *pr(sc)*), we could rearrange the original profit equation to be a summation of the individual expected profit terms. It can be shown that the optimality condition would be:

$$\sum_{l(q)}^{u(q)} pr(sc)[\text{Marginal Profit}(q)] = 0$$

The optimal position of the step could thus be solved by a simple algorithm that iterates through the possible range of positions.

4.3.5 Optimal Monotonic Transformations – Type II Non-Monotonicity

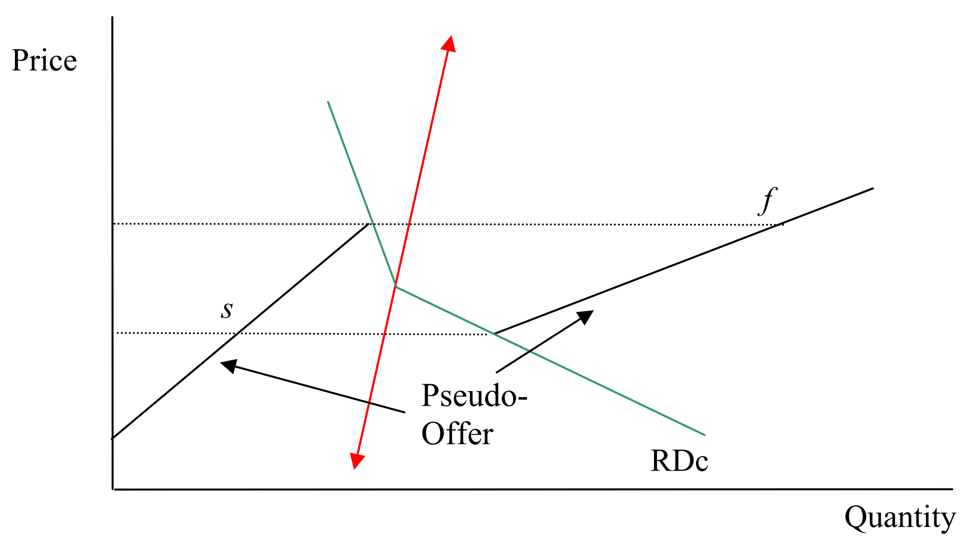

Figure 4.24 Type II Non -Monotonicity – Range to be Transformed

Type II non-monotonicity involves a PO that jumps down and to the right, as shown in Figure 4.24. The problem is again how best to transform this pseudo-offer into a feasible offer that meets market monotonicity requirements (i.e. how do we optimally connect points *s* and *f*). Just as with the Type 1 non-monotonicity case, the optimal offer up to point *s* and beyond point *f* will follow along the PO.

We propose that the optimal offer will trace along the *upper* (*left*) segment of the pseudo-offer curve from point *s*, and then at some *price* point, will step across to the *lower (right)* segment of the pseudo-offer curve, following this along to point *f*. This is demonstrated in Figure 4.25, and can easily be shown to be true by applying direction of improvement theory on any alternative offer form, similar to that demonstrated in Section 4.3.4.1. Stated more formally,

Theorem 4.2:

Consider a pseudo-offer with Type II non-monotonicity, such as that in Figure 4.24. If we take any existing offer segment from point *s* to point *f* that does not consist of part of each of the left and right attractor lines connected by a horizontal segment, then we know that we can make an improvement to the expected payoff by transforming the offer to this form, such that the horizontal segment passes through the intersection of the original offer with RDVP. This is *not* necessarily the optimal position for the vertical jump, but we know that it is a better offer than the original.

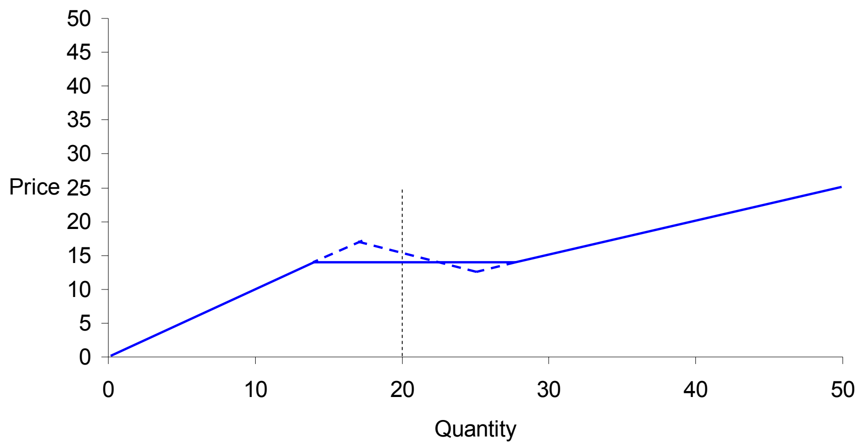

Figure 4.25 Demonstration of Horizontal Jump

Given that the optimal monotone offer in this case must involve a segment of the left AL, a segment of the right AL, and a horizontal jump between the two, we must form an optimality condition on the position of the jump in a similar way to that presented in Section 4.3.4.

4.3.5.1 Profit Formula and Derivative

The following additional terms are needed to define the optimality condition on the position of the horizontal jump.

a(sc,p) is a surface defined over the dimensions of RD scenario level and price, which provides the intercept term for the RD curve.

b(sc,p) is a surface defined over the dimensions of RD scenario level and price, which provides the slope term for the RD curve.

u(p) is the highest RD scenario level that leads to a dispatch on the horizontal segment of the offer curve for the current *price* position of that curve.

l(p) is the lowest RD scenario level that leads to a dispatch on the horizontal segment of the offer curve for the current *price* position of that curve.

p is the price level where the horizontal connection segment will be placed in the offer curve.

For a given scenario, we know the functions a*(sc,p)*, *b(sc,p)* and *pr(sc)* in advance, *p* is the decision that we are making, while *l(p)* and *u(p)* will directly depend on this decision.

The total profit equation is the sum of optimal dispatch points on the horizontal jump (with a known price level), before the horizontal jump (on the left AL) and after the horizontal jump (on the right AL):

$$
\begin{aligned}
profit =& \int_{l(p)}^{u(p)} pr(sc)\left[p\left[\frac{a(sc,p)-p}{b(sc,p)}\right]-\left[MC_{\text{int}}\left[\frac{a(sc,p)-p}{b(sc,p)}\right]+\frac{1}{2}MC_{sl}\left[\frac{a(sc,p)-p}{b(sc,p)}\right]^2\right]\right].dsc \\
&+\int_{0}^{l(p)} pr(sc)\left[\begin{array}{l}\left[a(sc,-\inf)-b(sc,-\inf)\left[\dfrac{a(sc,-\inf)-MC_{\text{int}}}{2b(sc,-\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(sc,-\inf)-MC_{\text{int}}}{2b(sc,-\inf)+MC_{sl}}\right] \\ -c\left(\dfrac{a(sc,-\inf)-MC_{\text{int}}}{2b(sc,-\inf)+MC_{sl}}\right)\end{array}\right].dsc \\
&+\int_{u(p)}^{\inf} pr(sc)\left[\begin{array}{l}\left[a(sc,\inf)-b(sc,\inf)\left[\dfrac{a(sc,\inf)-MC_{\text{int}}}{2b(sc,\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(sc,\inf)-MC_{\text{int}}}{2b(sc,\inf)+MC_{sl}}\right] \\ -c\left(\dfrac{a(sc,\inf)-MC_{\text{int}}}{2b(sc,\inf)+MC_{sl}}\right)\end{array}\right].dsc
\end{aligned}
$$

The derivative of profit is therefore:

$$\frac{\partial profit}{\partial p}=\int_{l(p)}^{u(p)} pr(sc)\left[\begin{array}{l}\dfrac{a(sc,p)-p+(p-MC_{\text{int}})\left(\dfrac{\partial a(sc,p)}{\partial p}-1\right)}{b(sc,p)}-\\ \dfrac{[p-MC_{\text{int}}][a(sc,p)-p]\left[\dfrac{\partial b(sc,p)}{\partial p}\right]-MC_{sl}[a(sc,p)-p]\left[\dfrac{\partial a(sc,p))}{\partial p}-1\right]}{b(sc,p)^2}\\ +\dfrac{MC_{sl}[a(sc,p)-p]^2\left[\dfrac{\partial b(sc,p)}{\partial p}\right]}{b(sc,p)^3}\end{array}\right].dsc$$

$$+\frac{\partial u(p)}{\partial p}\left[pr(u(p))\left[p\left[\frac{a(u(p),p)-p}{b(u(p),p)}\right]-\left[MC_{\text{int}}\left[\frac{a(u(p),p)-p}{b(u(p),p)}\right]-\frac{1}{2}MC_{sl}\left[\frac{a(u(p),p)-p}{b(u(p),p)}\right]^2\right]\right]\right]$$

$$-\frac{\partial l(p)}{\partial p}\left[pr(l(p))\left[p\left[\frac{a(l(p),p)-p}{b(l(p),p)}\right]-\left[MC_{\text{int}}\left[\frac{a(l(p),p)-p}{b(l(p),p)}\right]-\frac{1}{2}MC_{sl}\left[\frac{a(l(p),p)-p}{b(l(p),p)}\right]^2\right]\right]\right]$$

$$+\frac{\partial l(p)}{\partial q}pr(l(p))\left(\begin{array}{l}\left[a(l(p),-\inf)-b(l(p),-\inf)\left[\dfrac{a(l(p),-\inf)-MC_{\text{int}}}{2b(l(p),-\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(l(p),-\inf)-MC_{\text{int}}}{2b(l(p),-\inf)+MC_{sl}}\right]\\ -c\left(\dfrac{a(l(p),-\inf)-MC_{\text{int}}}{2b(l(p),-\inf)+MC_{sl}}\right)\end{array}\right)$$

$$-\frac{\partial u(p)}{\partial q}pr(u(p))\left(\begin{array}{l}\left[a(u(p),\inf)-b(u(p),-\inf)\left[\dfrac{a(u(p),\inf)-MC_{\text{int}}}{2b(u(p),\inf)+MC_{sl}}\right]\right]\left[\dfrac{a(u(p),\inf)-MC_{\text{int}}}{2b(u(p),\inf)+MC_{sl}}\right]\\ -c\left(\dfrac{a(u(p),\inf)-MC_{\text{int}}}{2b(u(p),\inf)+MC_{sl}}\right)\end{array}\right)$$

As in the vertical jump case, using the substitution $\frac{a(sc,p)-MC_{\text{int}}}{2b(sc,p)+MC_{sl}}=q$ enables the final four terms of the above condition to cancel out. Therefore, the profit derivative reduces down to the equation:

$$\int_{l(p)}^{u(p)} pr(sc)\left[\begin{array}{l} \dfrac{a(sc,p)-p+[p-MC_{\text{int}}]\left[\dfrac{\partial a(sc,p)}{\partial p}-1\right]}{b(sc,p)}- \\ \dfrac{[p-MC_{\text{int}}][a(sc,p)-p]\left[\dfrac{\partial b(sc,p)}{\partial p}\right]-MC_{sl}[a(sc,p)-p]\left[\dfrac{\partial a(sc,p))}{\partial p}-1\right]}{b(sc,p)^2} \\ +\dfrac{MC_{sl}[a(sc,p)-p]^2\left[\dfrac{\partial b(sc,p)}{\partial p}\right]}{b(sc,p)^3} \end{array}\right].dsc$$

4.3.5.2 Optimality Condition

In the specific case that we are considering here, with a single kink in the piecewise linear RD curve, this equation can be simplified significantly. The derivative of both the *a* and *b* functions with respect to *p* is equal to zero at all points, giving the optimality condition:

$$\int_{l(p)}^{u(p)} pr(sc)\left[\frac{a(sc,p)-2p}{b(sc,p)}+\left[MC_{\text{int}}\frac{1}{b(sc,p)}-MC_{sl}\frac{(a(sc,p)-p)}{b(sc,p)^2}\right]\right].dsc=0$$

Let us consider the meaning of this condition. If there is a linear marginal cost, then the profit for a particular RD curve that intersects the horizontal segment is:

$$\begin{aligned} \text{Profit} &= p\left[\frac{a(sc,p)-p}{b(sc,p)}\right]-c\left(\frac{a(sc,p)-p}{b(sc,p)}\right) \\ &= p\left[\frac{a(sc,p)-p}{b(sc,p)}\right]-\left(MC_{\text{int}}\left[\frac{a(sc,p)-p}{b(sc,p)}\right]+\frac{1}{2}MC_{sl}\left[\frac{a(sc,p)-p}{b(sc,p)}\right]^2\right) \end{aligned}$$

Therefore, marginal profit with respect to price is:

$$\text{Marginal Profit} = \frac{a(sc,p) - p + [p - MC_{\text{int}}]\left[\frac{\partial a(sc,p)}{\partial p} - 1\right]}{b(sc,p)}$$

$$+ \frac{[MC_{\text{int}} - p][a(sc,p) - p]\left[\frac{\partial b(sc,p)}{\partial p}\right] - MC_{sl}[a(sc,p) - p]\left[\frac{\partial a(sc,p)}{\partial p} - 1\right]}{b(sc,p)^2}$$

$$+ \frac{MC_{sl}[a(sc,p) - p]^2\left[\frac{\partial b(sc,p)}{\partial p}\right]}{b(sc,p)^3}$$

Which simplifies to:

$$\text{Marginal Profit} = \frac{a(sc,p) - 2p}{b(sc,p)} + \frac{1}{b(sc,p)}\left[MC_{\text{int}} + MC_{sl}\frac{[a(sc,p) - p]}{b(sc,p)}\right]$$

In other words, as we increase the price, the revenue increases at a rate of $\frac{a(sc,p)-2p}{b(sc,p)}$, while the quantity decreases at a rate of $\frac{1}{b(sc,p)}$. Therefore, as we increase the price level, the total cost decreases at a rate of $\frac{1}{b(sc,p)}\left[MC_{\text{int}} + MC_{sl}\frac{[a(sc,p)-p]}{b(sc,p)}\right]$ (where $\left[MC_{\text{int}} + MC_{sl}\frac{(a(sc,p)-p)}{b(sc,p)}\right]$ is the marginal cost at the quantity where the RD curve crosses the horizontal), thus increasing profit at the same rate. So, effectively, our optimality condition is:

$$\int_{l(p)}^{u(p)} pr(sc)[\text{Marginal Profit}(p)].dsc = 0$$

Therefore, the optimality condition says that we take the integral of the probability distribution of demand multiplied by the marginal profit with respect to price and find the price for which this equation is equal to zero.

Recall that the optimality condition for the location of the vertical step under horizontal uncertainty was:

$$\int_{l(q)}^{u(q)} pr(sc)[\text{Marginal Profit}(q)].dsc = 0$$

This is clearly a very similar optimality condition, the only difference being the direction of the derivatives (in the *p* direction for the horizontal jump and in the *q* direction for the vertical jump).

Again, if rather than having a continuous probability density function for the RD level, there was actually a set of discrete scenario levels that could occur (with a corresponding probability of each occurrence being *pr(sc)*), we could rearrange the original profit equation to be a summation of the individual expected profit terms. It can be shown that the optimality condition would be:

$$\sum_{l(p)}^{u(p)} \mathrm{p}r(sc)[\text{Marginal Profit}(p)] = 0$$

The optimal position of the step could thus be solved by a simple algorithm that iterates through the possible range of positions.

4.4 Optimal Offer Forms under Multiple-Kinked RD Curves

In this section we develop the concepts of Section 4.3 further, to consider how the optimal monotone offers will look under RD curves that have multiple kinks, potentially of both convex and concave nature.

When we consider a RD curve with multiple kinks, there will be multiple RDVPs. In particular, if there are *n* segments to the RD curve, there will be (*n*-1) kinks in the curve and hence (*n*-1) RDVPs. Just as in the single kink case, we can establish the appropriate attractor lines in each range between the consecutive RDVPs, corresponding to the slope of the RD curve in that region. The combination of these ALs will then determine the pseudo-offer (or desired offer) for that RD scenario.

As a result of each of these (*n*-1) kinks that occur in the RD curve, there is potentially a region of uncertainty as to how the optimal offer should be constructed (from s_i to f_i). These regions of uncertainty can come about from either Type I or Type II non-monotonicity.

For example, Figure 4.26 demonstrates a RD curve with three kinks, and hence four different RD curve slopes, denoted b_1 to b_4, where b_1 is the flattest segment and b_4 the steepest. Therefore, each of these slope segments has a corresponding CL, as with the single kink case.

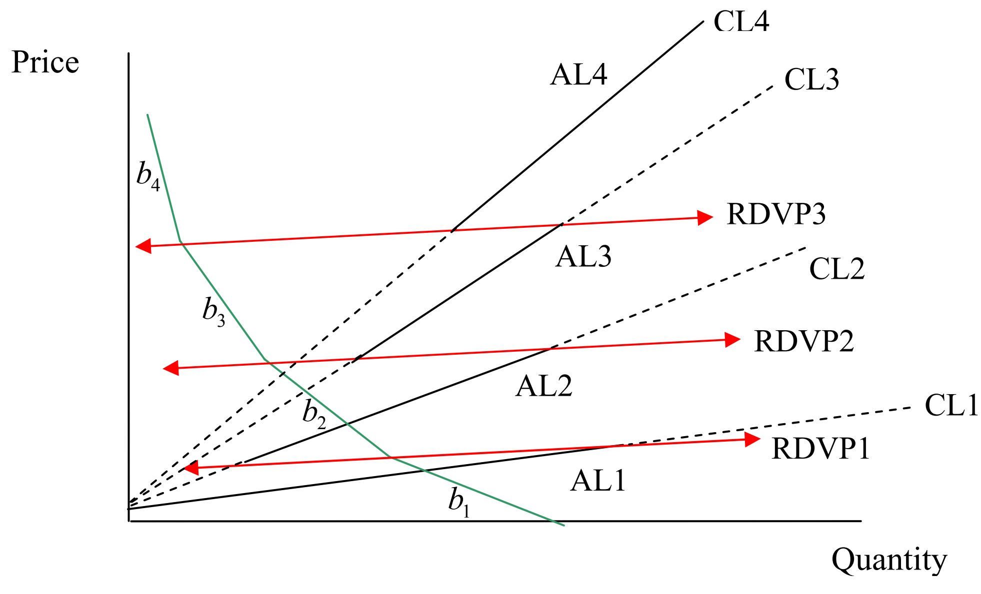

Figure 4.26 Construction of the ALs for a Multi-Kink Residual Demand Curve

The ALs are then determined by the regions in which each of these CLs are valid. In other words:

- AL1 is the segment of CL1 below RDVP1 (i.e. where b_1 is valid)
- AL2 is the segment of CL2 between RDVP1 and RDVP2 (where b_2 is valid)
- AL3 is the segment of CL3 between RDVP2 and RDVP3
- AL4 is the segment of CL4 above RDVP3

Therefore, in this case we encounter three kinks that lead to Type I non-monotonicity in the PO. This produces three regions of uncertainty, with start and finish points denoted s_1, f_1, s_2, f_2, s_3 and f_3, as shown in Figure 4.27.

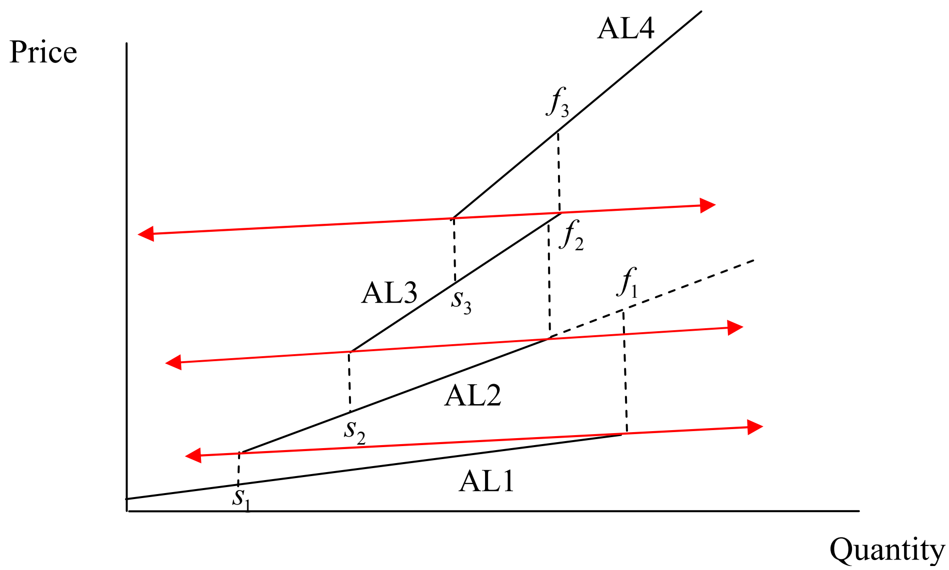

Figure 4.27 Regions of Uncertainty

4.4.1 Direction of Improvement Theory under Multiple Kinks

The direction of improvement theory in Section 4.3 was largely based on the fact that for a single kink in a given RD curve, the payoff function with respect to quantity supplied was either unimodal or bimodal. When it was bimodal, the local minimum between the two points of local maxima was located at the kink in the RD curve (i.e. at the RDVP).

This implied that any current dispatch point would be attracted to the AL on the same side of the RDVP as it was currently located. Likewise, it can easily be shown that if there are (n-1) kinks in the RD curve, then the payoff function will be at most n-modal, with points of local minima again located on the RDVPs. This implies that any current dispatch point will be attracted to the AL in the same region (between consecutive RDVPs) as it is currently located.

Returning to our example, consider any current offer, such as that shown in blue in Figure 4.28. We produce a new offer (bold black) that has a vertical segment passing through each intersection between the current offer and one of the RDVPs. We know that this new offer must be an improvement on the original, as all dispatch points have moved in their local directions of improvement (i.e. towards the AL in the same region).

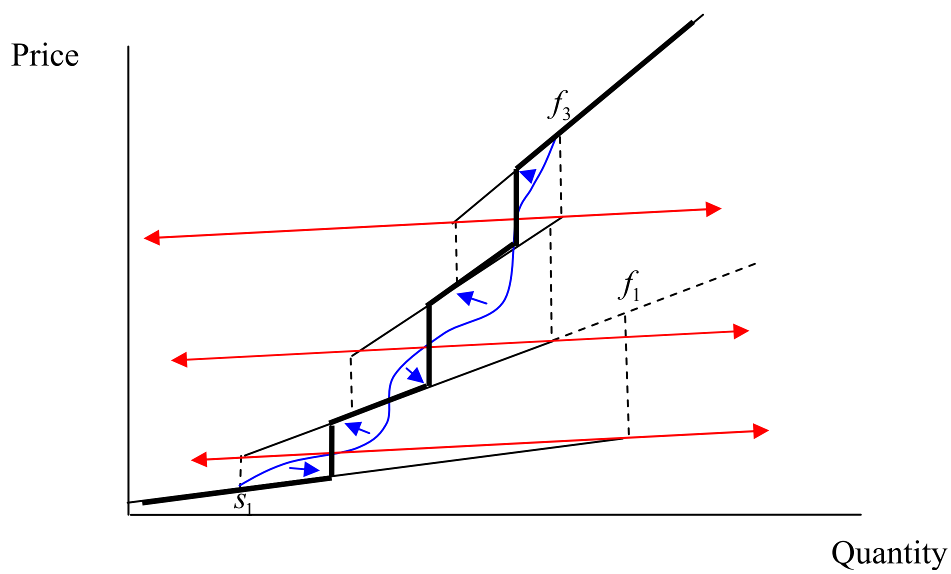

Figure 4.28 New Offer with Guaranteed Improvement

This example shows that the connections between the ALs to overcome these Type I non-monotonicities must be vertical segments, but does not define where the optimal position of these individual segments will be, and whether the combination of them will be monotone.

Consider all the regions of non-monotonicity that result from kinks in the RD curve. There are four possible outcomes:

1) All kinks lead to Type I non-monotonicity in the PO
2) All kinks lead to Type II non-monotonicity in the PO
3) Some kinks lead to Type I non-monotonicity and some lead to Type II non-monotonicity, but *there are no areas* where the regions of uncertainty corresponding to the different types of non-monotonicity overlap.
4) Some kinks lead to Type I non-monotonicity and some lead to Type II non-monotonicity, and *there are areas* where the regions of uncertainty corresponding to the different types of non-monotonicity overlap.

If any of outcomes 1-3 occur, then the optimal offer can be formed relatively simply, using the same logic and optimality conditions of Section 4.3. In these cases, the regions of Type I and Type II non-monoticities are dealt with separately. This is explained in Section 4.4.2. If outcome 4 occurs, then, although the same logic of Section 4.3 still applies, the optimality condition on the placement of the corresponding vertical and horizontal segments must be adjusted (as they may overlap). This is explained in Section 4.4.3.

4.4.2 Dealing with Multiple Type I or All Type II Non-Monotonicity

When faced with consecutive regions of Type I or Type II non-monotonicity, the first step is to find the optimal resolution of each region of uncertainty independently. To demonstrate, we will return the example started above, which deals with consecutive Type I non-monotonicity. In general, there are three possible scenarios that could result from this first step.

Scenario 1: The combination of the independent steps is naturally monotone. The optimal offer is therefore determined by the combination of these steps and the AL segments in between.

Scenario 2: The combination of a subset of the first or last *i* independent steps is guaranteed to be monotone (all to the left or all to the right of all other possible steps). All that remains is to determine the optimal monotonic resolution of the remaining *non-monotone* steps (this may effectively lead you to Scenario 3).

Scenario 3: Multiple independent steps are non-monotone. We must determine how to optimally resolve this non-monotonicity.

Let us consider the possible scenarios that could be encountered when trying to resolve the non-monotonicity in our example. Figure 4.29 – Figure 4.32 show in blue, a selection of possible combinations of optimal vertical step positions between each pair of consecutive ALs, one associated with each of the three possible scenarios. The possible resolutions of non-monotonicity discussed are shown in red. We shall denote the quantity level of the optimal vertical step between AL_i and AL_j as qv_{ij}.

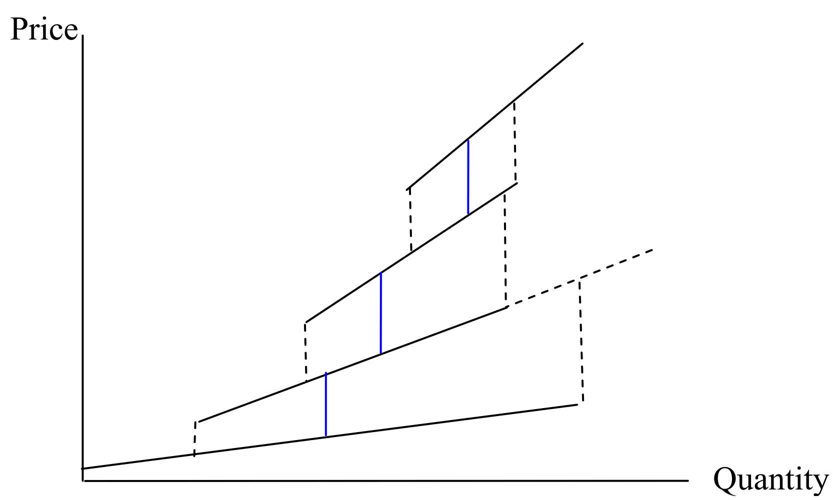

Figure 4.29 Scenario 1

Figure 4.29 demonstrates Scenario 1, where $qv_{12} \le qv_{23} \le qv_{34}$. Therefore, we immediately have a monotone offer.

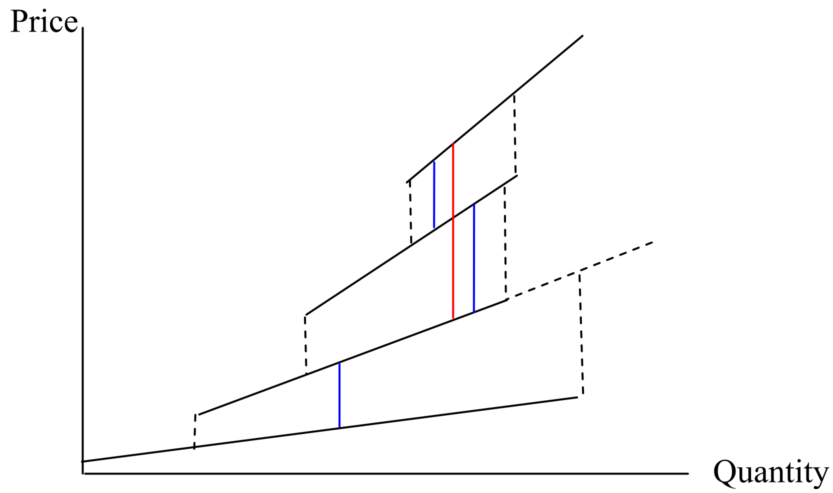

Figure 4.30 Scenario 2 – Instance 1

Figure 4.30 shows one possible occurrence of Scenario 2. Here, the first vertical step, qv_{12}, is to the left of both the other vertical steps. Therefore, we immediately know that it must be part of the optimal offer. However, qv_{23} is to the right of qv_{34}, and so this non-monotonicity must be resolved. The optimal solution to this problem must be a single vertical step, qv_{24}, from the second AL up to the fourth. Clearly, the quantity position of this vertical step must be between the quantities of the two vertical steps whose non-monotonicity it is overcoming. i.e. $qv_{34} \le qv_{24} \le qv_{23}$.

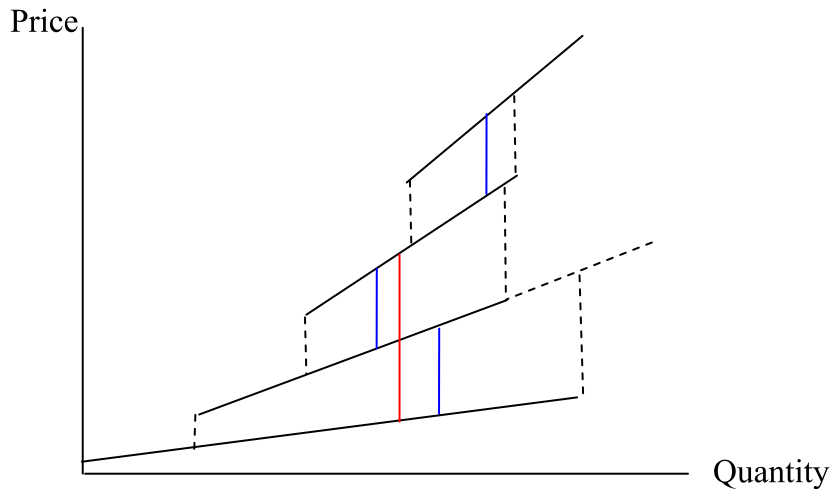

Figure 4.31 Scenario 2 – Instance 2

Figure 4.31 shows another possible occurrence of Scenario 2, but this time the last vertical step, qv_{34}, is to the right of both other vertical steps. Similarly to the previous case, we therefore immediately know that this will be part of the optimal offer. The non-monotonicity of qv_{12} and qv_{23} must therefore be resolved, with a step qv_{13} that lies between the two.

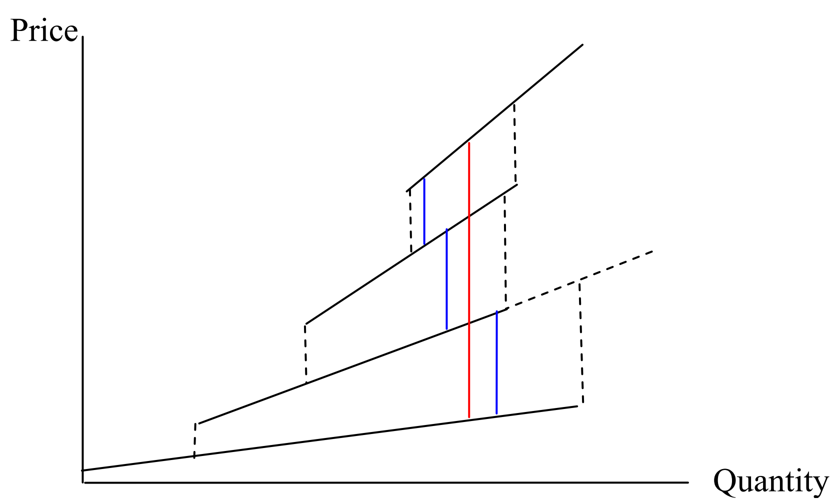

Figure 4.32 Scenario 3

The case demonstrated in Figure 4.32 demonstrates Scenario 3, where there is non-monotonicity between all three steps. We know the optimal monotonic resolution of any two consecutive non-monotonic steps will be a new vertical that falls between the original two verticals. Expanding on this, we therefore know that the resolution of any *n* consecutive non-monotonic verticals is a single new vertical lying somewhere between the first and last original verticals.

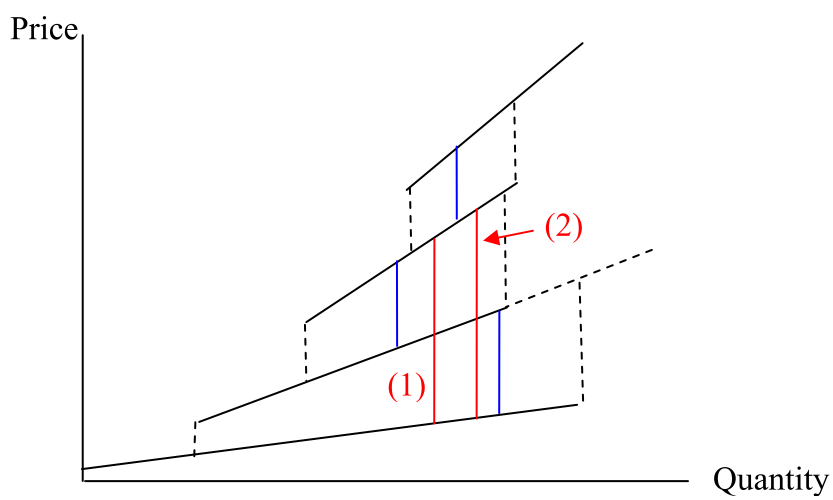

Figure 4.33 Scenario 3 – Instance 2

This final case, as shown in Figure 4.33, is slightly different as there are two possible results. As with the previous case, we start by resolving the non-monotonicity of the *n* consecutive non-monotone vertical steps ($n = 2$ in this case). If the optimal resolution is a qv_{13} that lies to the left of qv_{34} (as indicated by line 1), then we have a monotone offer at this point. If, however, the optimal resolution lies to the right of qv_{34}, (as indicated by line 2), then we have a new non-monotonicity between qv_{13} and qv_{34}. The process of resolving non-monotonicities therefore repeats until the resulting offer is completely monotone. If there are many consecutive sections of Type I or Type II non-monotonicity, this may take a few iterations.

4.4.3 Dealing with Overlapping Regions of Type I and Type II Non-Monotonicity

This section determines the process that should be followed if there are overlapping regions of Type I and Type II non-monotonicity.

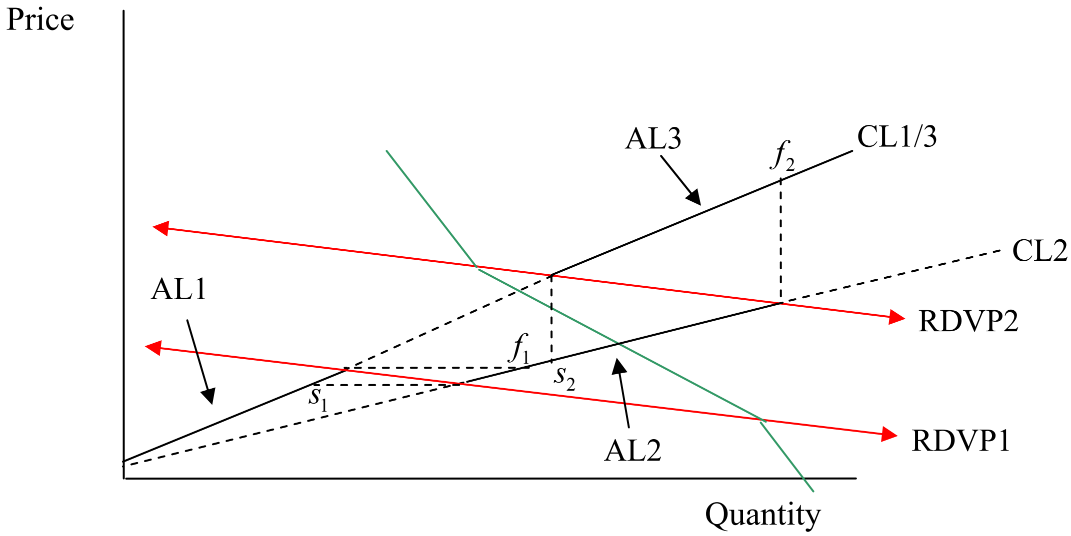

Figure 4.34 Non-Overlapping Regions of Type I and Type II Non-Monotonicity

Consider Figure 4.34 In this case, we have both a Type I and a Type II non-monotonicity in the PO. However, because their regions of uncertainty do not overlap, the optimal resolutions of their non-monotonicities can be performed independently. This is an example of PO outcome 3, as described in Section 4.4.2.

Now consider a case such as that shown in Figure 4.35, where the regions of uncertainty do overlap. Assume that we have a current offer from s_1 to f_2, such as that shown in blue. We use the term Critical Intersection Points (CIPs) to refer to the intersections between the current offer and the RDVPs.

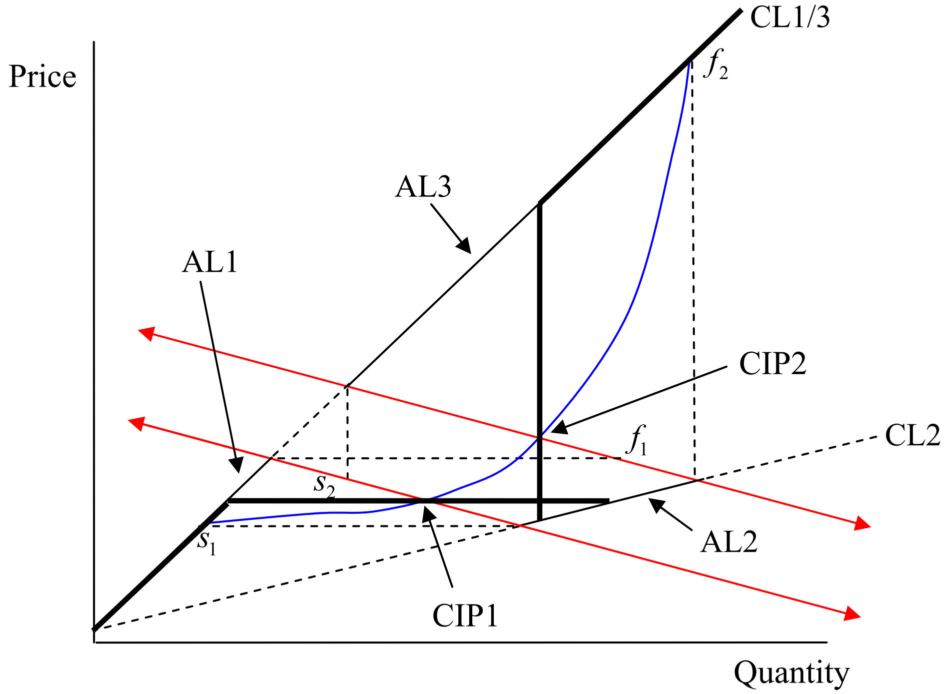

Figure 4.35 Overlapping Regions of Type I and Type II Non-Monotonicity

Recall that, as explained in Section 4.4.2, a point of dispatch on any current offer is attracted to the AL in the same region as it (i.e. between the same pair of RDVPs). For this reason, in this case, we can observe that the bold offer in Figure 4.36 is a guaranteed improvement over the original offer. Again though, we need an optimality condition to determine the exact position of the optimal vertical and horizontal steps.

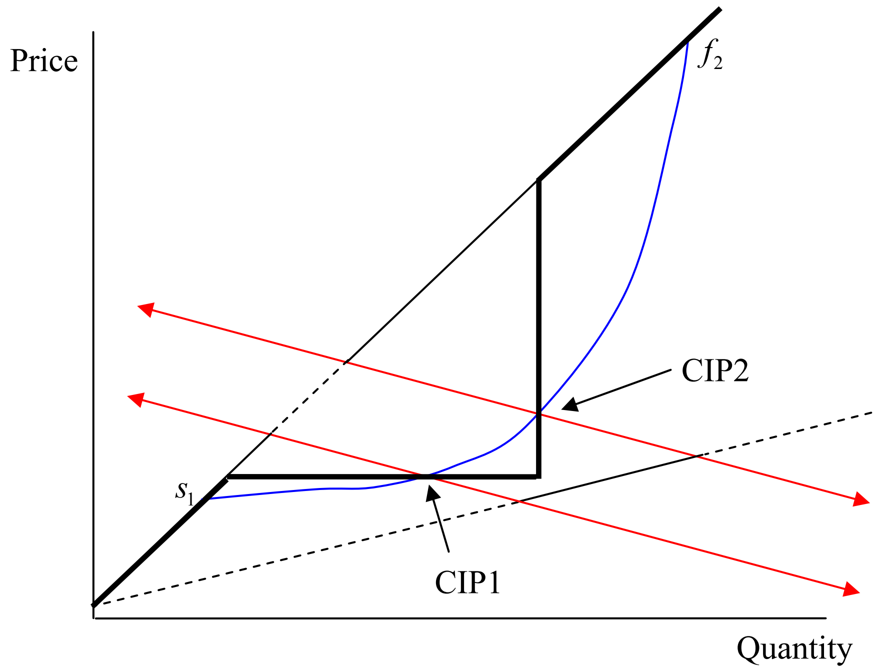

Figure 4.36 Overlapping Regions of Type I and Type II Non-Monotonicity – Improved Offer

In general, for multiple overlapping regions of Types I and II non-monotonicity, we can define an algorithm for the generation of an improved offer over any current offer that does not have the optimal form. This process will tell us the general form that the optimal resolution of the non-monotonicity should take. The algorithm is as follows:

1) Create the vertical and horizontal segments for each Type I and II non-monotonicity.
    a. These are lines through the CIPs – see Figure 4.36.
    b. The vertical and horizontal segments exist only between the RDVPs in which the associated kink is relevant, and within the two associated CLs.
2) The improved offer must pass through all the CIPs, so remove all parts of the vertical and horizontal segments that could not be included in a monotonic offer that meets these requirements. For example, we are left with the bold lines in Figure 4.36.
3) The remaining vertical and horizontal segments, along with the monotonic segments of the ALs are the new, improved offer.

Figure 4.37 shows the range of possible optimal vertical and horizontal segment positions for this example. Note that some positions have them overlapping, and others do not. For the cases where they are optimally not overlapping, this can be found by solving each separately first, and then leaving offer as is, if they don't overlap. If they do overlap, then a new optimality condition is needed, and can be constructed under the assumption that the optimal horizontal and vertical steps will definitely intersect. The goal of this optimality condition will be to find the Turning Point (TP) of the optimal horizontal and vertical segments, as shown in Figure 4.37.

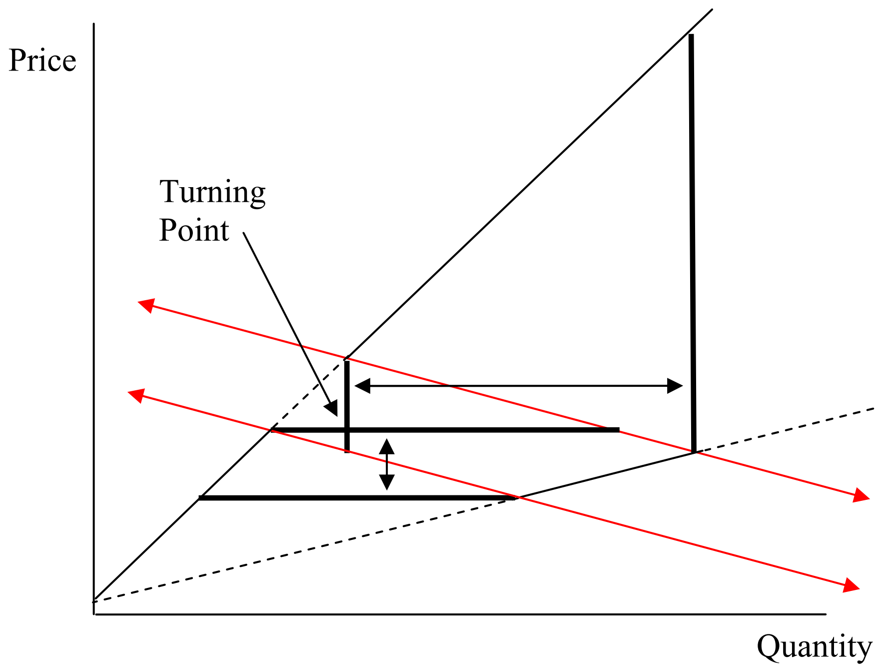

Figure 4.37 Overlapping Regions of Type I and Type II Non-Monotonicity – Possible Optimal Offer Positions

4.4.4 Optimality Condition for Overlapping Regions of Types I and II Non-Monotonicity

Because we have already established that the vertical and horizontal segments must intersect, the optimality conditions for the horizontal and vertical segments are linked via the requirement that $u(p) = l(q)$. In other words the highest RD curve that intersects the horizontal segment is the lowest RD curve that intersects the vertical segment. This could be generalised so that they were consecutive RD curves, rather than the same one, but if we assume a continuous (or close to) distribution of RD curves, there will effectively be no difference.

Recall that our optimality conditions for independent horizontal and vertical steps were:

$$\int_{l(p)}^{u(p)} pr(sc)[\text{Marginal Profit}(p)].dsc = 0 \text{ and } \int_{l(q)}^{u(q)} pr(sc)[\text{Marginal Profit}(q)].dsc = 0$$

So, our new optimality condition is:

$$\int_{l(p)}^{u(p)} pr(sc)[\text{Marginal Profit}(p)].dsc + \int_{l(q)}^{u(q)} pr(sc)[\text{Marginal Profit}(q)].dsc = 0$$

where $u(p) = l(q)$

Therefore, the optimality condition says that (quantity, price) space must be searched for a TP that satisfies this equation.

If rather than having a continuous probability density function for the RD level, there was actually a set of discrete scenario levels that could occur (with a corresponding probability of each occurrence being *pr*(*sc*)), we could rearrange the original profit equation to be a summation of the individual expected profit terms. It can be shown that the optimality condition would be:

$$\sum_{l(p)}^{u(p)} \mathrm{pr}(sc)[\text{Marginal Profit}(p)] + \sum_{l(q)}^{u(q)} \mathrm{pr}(sc)[\text{Marginal Profit}(q)] = 0$$

4.4.5 Offer Construction Algorithm under Multiple Kinked RD Curves

To summarise, in order to produce the optimal offer under a multiple-kinked RD curve, this process should be followed:

1) Break the two-dimensional (quantity, price) space into regions by the multiple RDVPs. In each of these regions, establish the appropriate ALs and hence produce the overall PO (as in Figure 4.26).
2) Independently determine the optimal vertical or horizontal step to use for each Type I and Type II non-monotonicity respectively, that occurs in the PO.
3) If any of these steps are clearly valid, then include that step automatically and consider only the ranges that have non-monotone or crossing steps. A step is clearly valid if it is further to the left or the right than any other step, and retains monotonicity, no matter what the optimal resolution of the remaining region of uncertainty (for example, as explained in Scenario 2 of Section 4.4.2).
4) If there are optimal steps from Types I and II non-monotonicity crossing one another, then find the optimal Turning Point to determine the position of the horizontal and vertical steps, as described in Section 4.4.4.
5) If there are optimal steps from non-monotonicity of the same Type (I or II) forming a non-monotone offer, then find the optimal position of the monotone step(s) to replace them, as described in Section 4.4.3.

4.5 Optimal Offer Forms under MDF Contours rather than RD Curves

In Sections 4.3 and 4.4, we have considered the optimal analytic forms of offers based on alternative possible forms of the RD curves. We first considered the optimal dispatch points under each possible curve, and then resolved the non-monotonicities that occurred. We did this by forming CLs corresponding to each possible slope that occurred in the RD curve, and then defining ranges for which these CLs were valid, based on the direction of movement of the RD curve. The combination of these valid CL segments produced the AL or PO, and from here, monotonicity was restored if need be.

The offer was constructed such that it passed through each RD curve once, as we know must occur.

However, this approach was quite restrictive. The alternative RD curves could not cross one another which meant that they all had to have an identical shape, and that the uncertainty had to be in a single direction.

Now, rather than considering alternative RD curves and the optimal dispatch point on each, we consider each of the contours of an MDF, and the optimal dispatch point on each of these (i.e. the point with the highest payoff). As with the RD curve case, we know that the optimal offer must pass through each of these contours once, and only once. Therefore, if we know the optimal point at which each contour is intersected, then the combination of these is the PO. If the PO is monotone, then this must be the optimal offer. If the PO is not monotone, as it will likely not be, then we must resolve this non-monotonicity as we did under the RD problem.

4.5.1 Constructing Pseudo-Offer under MDF Contours

In Section 4.3 we showed that if we have a linear RD curve segment, then the payoff along that segment as it sloped down through (q, p) space is unimodal, with its maximum point at the intersection with its CL. This must be the same for a straight segment of a MDF contour. This helps us determine the direction of improvement along a given contour from any existing offer, and will help us determine the shape of the guaranteed improved offer. If the contours were allowed to take very general, non-linear shapes, then the payoff function would not have such nice properties. However, this is not the case if the implied MDF is being constructed from a set of piecewise linear RD curves (the contours will also be piecewise linear).

Consider an example, as shown in Figure 4.38, where there are only two slopes to the contours, and therefore only two possible CLs. The optimal dispatch points, marked on each contour, do not form a monotonic set. Specifically, we can observe “Type I” non-monotonicity (resolved with a vertical step in the RD case). As with the cases in Section 4.3, we can show the direction of improvement along each contour (blue arrows) from any given offer (such as the blue one shown). From these, we can see that a vertical step between the two CLs must clearly be the optimal resolution of the non-monotonicity, where the vertical step is placed somewhere between *s* and *f*.

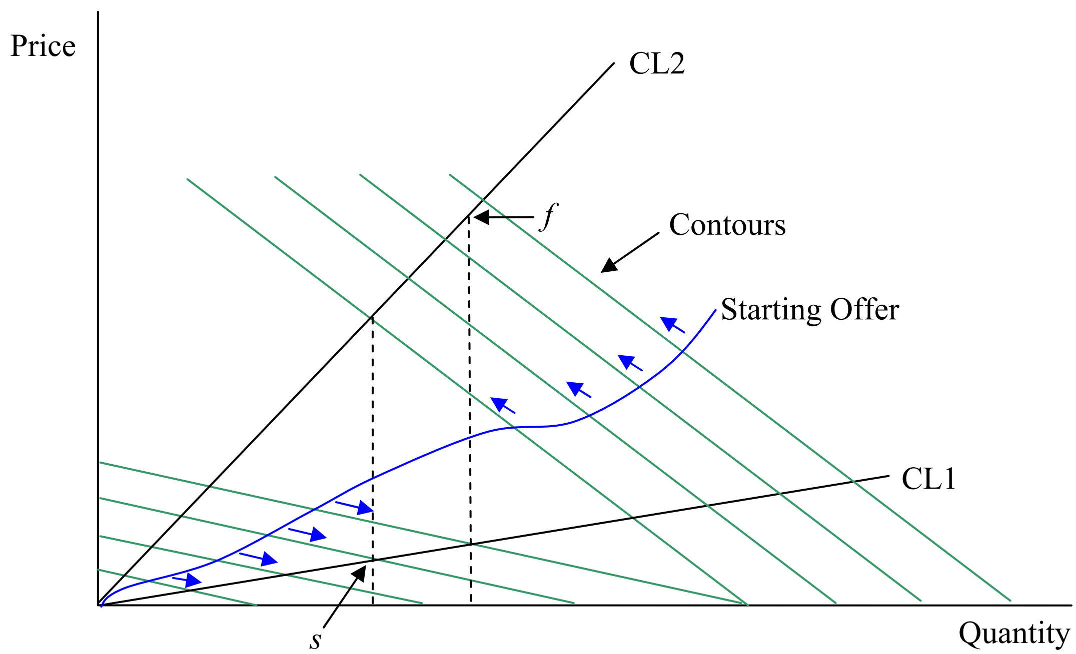

Figure 4.38 Simple Set of MDF Contours - Two Slopes Only

4.5.2 Optimality Condition

As with the analysis in Section 4.3, we can construct an optimality condition on the position of the vertical step within this range of uncertainty. We know from this previous analysis that the terms for expected payoff outside and on the start and end of the vertical step will cancel one another out, so we just need to consider the expected payoff on the contours that intersect the vertical segment of the offer. For these, q is given, and price is given by $p = \mathrm{int} - slope * q$, where *slope* is the local slope of the contour at the intercept point, and *int* is the price axis intercept of a line with that slope, extended out.

Therefore, expected payoff from dispatches on these contours is given by:

$$payoff = \sum_{l(q)}^{u(q)} pr(ct) * \left(q\left(\text{int} - slope * q\right) - c\left(q\right)\right)$$

Where *pr(ct)* is the difference in probability between the current contour and the previous contour (i.e. marginal probability that dispatch would occur on that contour, in a discrete sense), and *u(q)* and *l(q)* are the highest and lowest contours that would intersect a vertical at the given quantity level.

Take the derivative with respect to *q*, under the assumption that MC is fixed:

$$\frac{\partial profit}{\partial q} = \sum_{l(q)}^{u(q)} pr(ct)\left[\text{int} - slope * q + q\left[\frac{\partial\,\text{int}}{\partial q} - \frac{\partial slope}{\partial q} q - slope\right] - \frac{\partial c(q)}{\partial q}\right]$$

$$\frac{\partial profit}{\partial q} = \sum_{l(q)}^{u(q)} pr(ct)\left[\text{int} - 2slope * q - MC\right]$$

$$\frac{\partial profit}{\partial q} = \sum_{l(q)}^{u(q)} pr(ct)\left[MR - MC\right]$$

The optimality condition is therefore:

$$\sum_{l(q)}^{u(q)} \text{pr}\left(ct\right)\left[\text{Marginal Profit}(q)\right] = 0$$

This aligns with that found in Section 4.3.4.2, the only difference being in the definition of the probability terms used.

4.5.3 More Complex MDF Contour Examples

The examples shown in Figure 4.39 and Figure 4.40 give directions of improvement and guaranteed improved offers under two differing original offers for the same MDF. These guaranteed improved offers have been created by modifying the original offer by moving in the direction of improvement along all contours of the offering space. Note that the offer in each case is only one specific choice of many possible improved offers.

In these examples, the MDF is a much more general shape, allowing for RD curves that cross one another, and oddly shaped (but still linear) contours. For simplicity, the example MDF contains only two different slopes in the contours, flat and steep, each with their own CL (CL1 and CL2 respectively).

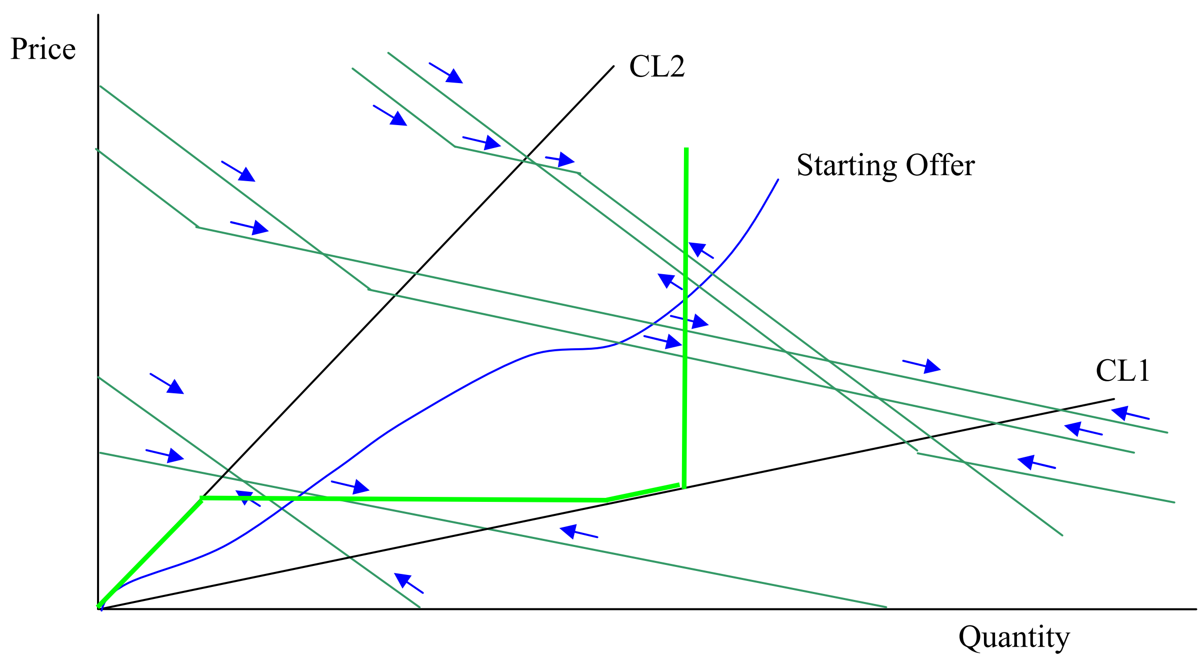

Figure 4.39 Complex Set of MDF Contours – Starting Offer 1

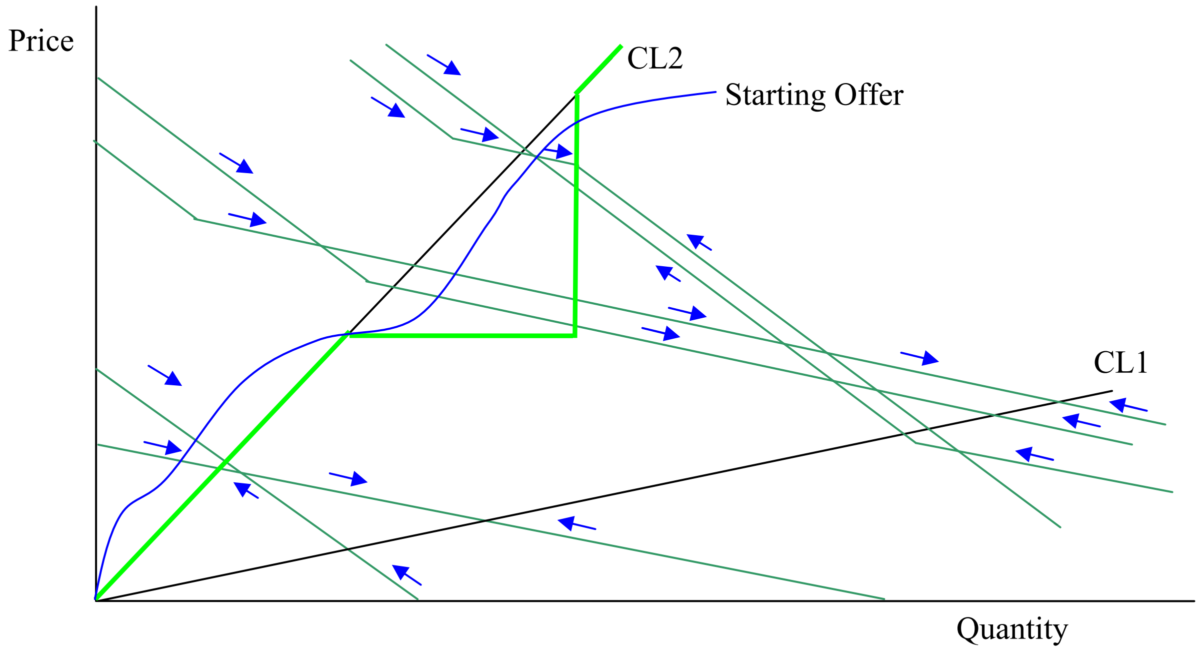

Figure 4.40 Complex Set of MDF Contours – Starting Offer 2

We can see that in both of these cases, the optimal offer forms, with examples shown in bright green, must indeed consist of parts of the positively sloped CL's, connected by horizontal and/or vertical segments.

4.5.4 Summary of Offer Forms under MDF Contours

In this section, we have briefly shown that the techniques for constructing optimal analytic offer forms apply to MDF contours, as well as to sets of possible RD curves. This extension enables the technique to deal with any set of possible RD curves, including alternatives that cross one another. However, we also note that as the number of different contour slopes and overlapping regions of Types I and II non-monotonicity increase, the optimality conditions on the position of the horizontal and vertical segments would quickly become very complex to define and solve. As such, in these more complex cases, it is likely that a DP on a pre-defined grid, while losing some accuracy, would be significantly more computationally efficient to apply.

4.6 Optimal Offer Shifts under Progressive MC Curve Shifts

In this section, we will identify rules for translation of (shifting) an existing offer, if the marginal costs were to move horizontally, under the simple forms of RD curve uncertainty presented in this chapter. This is highly relevant for operators of generation units with limited fuel resources. In such situations, the marginal cost curve for the current period can be determined directly from the marginal fuel value (MFV) curve for the following period, for the range of fuel levels that could potentially be reached. This is demonstrated in Figure 4.41, where the black lines indicate the initial situation, and the blue lines indicate the possible situation if the storage level at the current time was one unit higher. We can observe that the range of possible storage levels at the end of the period has moved to the right, and as such our marginal cost curve has also shifted horizontally to the right. Note though that a new segment of marginal cost has appeared (as shown in red), corresponding to the marginal fuel value at the new end-of-period storage level that could potentially now be reached.

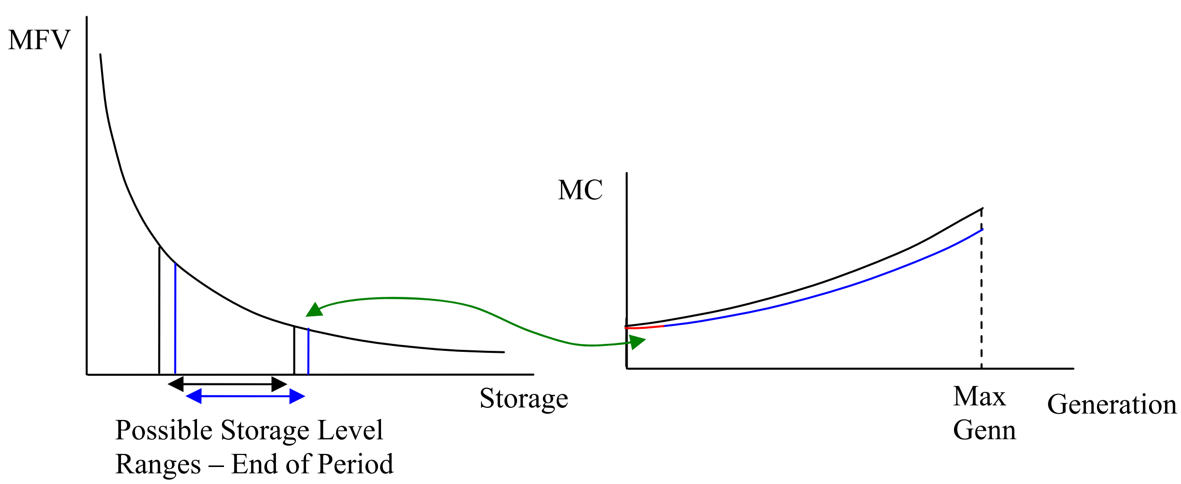

Figure 4.41 Shifting Marginal Cost Curve

Ideally, if we have constructed an optimal offer for a given storage level, *s*, in the current period, but it turns out that the storage level is actually *(s+1)*, we would prefer not to have to recalculate the optimal offer from scratch. By developing rules for translation of the previously optimal offer under such changes, there is potential for large computational benefits in the updating of such offers.

4.6.1 Explanation of Translation Rules

In Section 4.3, we developed simple analytic expressions for the form of the optimal offer under piecewise linear RD curves and linear marginal costs. In particular, we have shown that the optimal offer is made up of segments of line that are defined by:

$$p = MC_{\text{int}} + \left(b_i + MC_{sl}\right)q$$

in the appropriate ranges for the RD slope $-b_i$ and marginal cost curve:

$$\begin{aligned} MC &= MC_{\text{int}} + MC_{sl}q \qquad \text{(if linear MC)} \\ &= MC_{\text{int}} \qquad\qquad\quad\; \text{(if stepped MC)} \end{aligned}$$

4.6.2 Case 1 – Linear RD Curve and Stepped Marginal Cost

Consider a simple example, where there is a single step in the marginal cost curve at Q = 5, from MC = 5 to MC = 10, and the RD curve is linear (slope of -2), and exhibits uncertainty in a horizontal direction. Now, let us move the marginal cost curve to the right by one unit. This situation is depicted in Figure 4.42, where the original marginal cost and offer are shown in black, and the adjustments shown in blue.

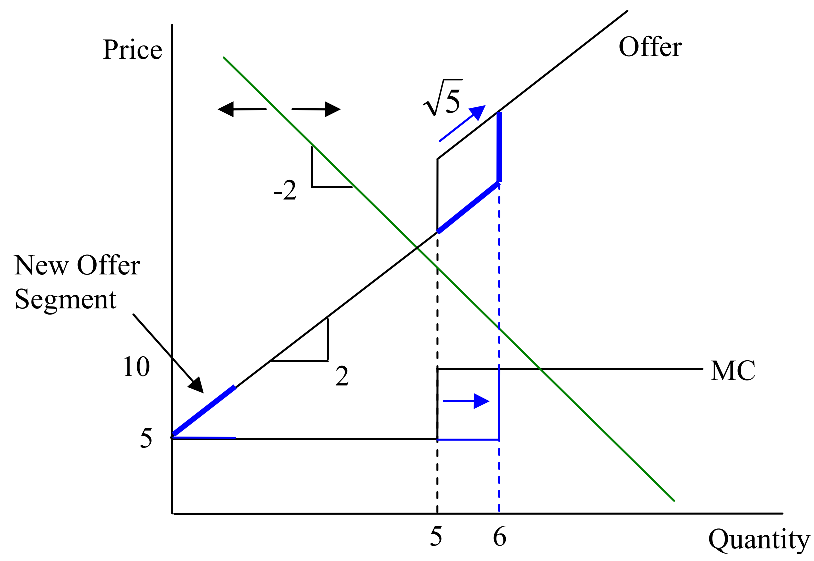

Figure 4.42 Change in Offer from Shifting Marginal Cost Curve – Case 1

Here, we can see that as a result of the one unit shift to the right of the marginal cost curve, the offer has translated up on an angle equivalent to the (negative of the) slope of the RD curve. The offer has shifted a total horizontal distance of one unit, meaning that the distance of the angled shift is $\sqrt{b_i^2 + 1^2}$, or in this case, $\sqrt{5}$. We also observe, as indicated, that there is a new segment of offer created from scratch. This is the offer segment corresponding to the part of the marginal cost curve that has just appeared.

4.6.3 Case 2 – Linear RD Curve and Linear Marginal Cost

Now consider a slightly more complicated case, where the marginal cost curve is linear rather than stepped. In particular, we will assume the same RD curve as above, but a marginal cost curve with the form $MC = 2 + 0.5q$. If we were to move the marginal cost curve to the right by one unit, the new marginal cost curve would have the form $MC = 1.5 + 0.5q$, and the offer would change as shown in Figure 4.43.

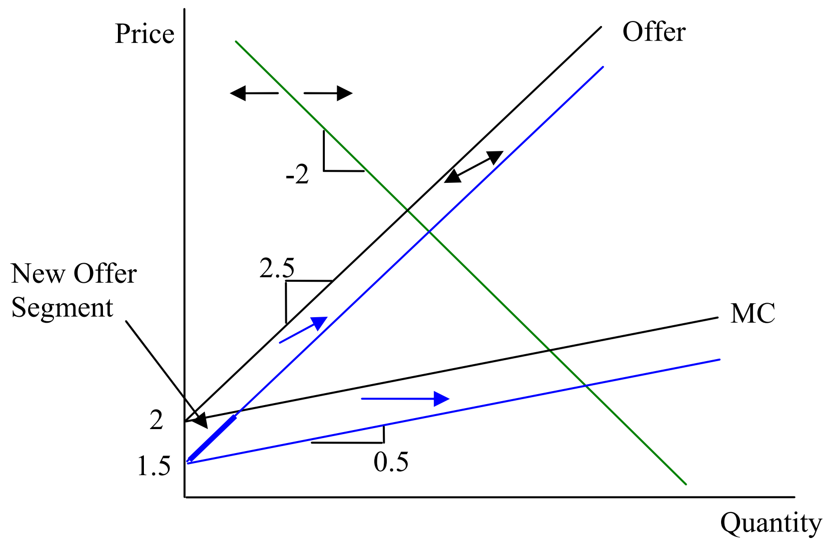

Figure 4.43 Change in Offer from Shifting Marginal Cost Curve – Case 2

Although not quite as obvious as in the previous case, we can see by inspection here that as a result of the one unit horizontal shift to the right of the marginal cost curve, the offer has again translated up on an angle equivalent to the (negative of the) slope of the RD curve. The offer has also again shifted a total horizontal distance of one unit, meaning that the distance of the angled shift is $\sqrt{b_i^2 + 1^2}$, or $\sqrt{5}$. To confirm this, consider the original offer point (0, 2).

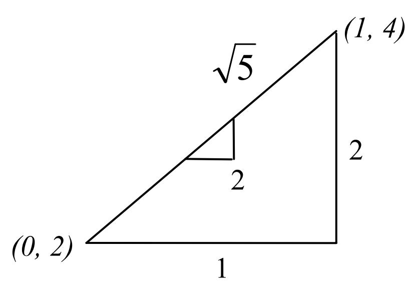

Figure 4.44 Translation Confirmation

Figure 4.44 shows us that this offer point, translated a distance of $\sqrt{5}$ at a slope of 2 to 1 (angle of $66.\dot{6}^o$), should now be located at *(1, 4)*. We can easily confirm that this point is indeed on the new offer line, $p = 1.5 + 2.5q$.

Note that as again indicated, a new segment of offer has had to be created from scratch.

4.6.4 Case 3 – Concave RD Curve and Stepped Marginal Cost

To generalise these results beyond a single-slope RD curve, consider the translation required under a concave RD curve. Assume that the RD curve has slope -1 above a price level of 30, slope of -2 below this price level, and is subject to uncertainty in the horizontal direction. In addition, assume that the marginal cost curve steps from 5 to 10 at a quantity of 5, and from10 to 15 at a quantity of 22. This case is shown in Figure 4.45.

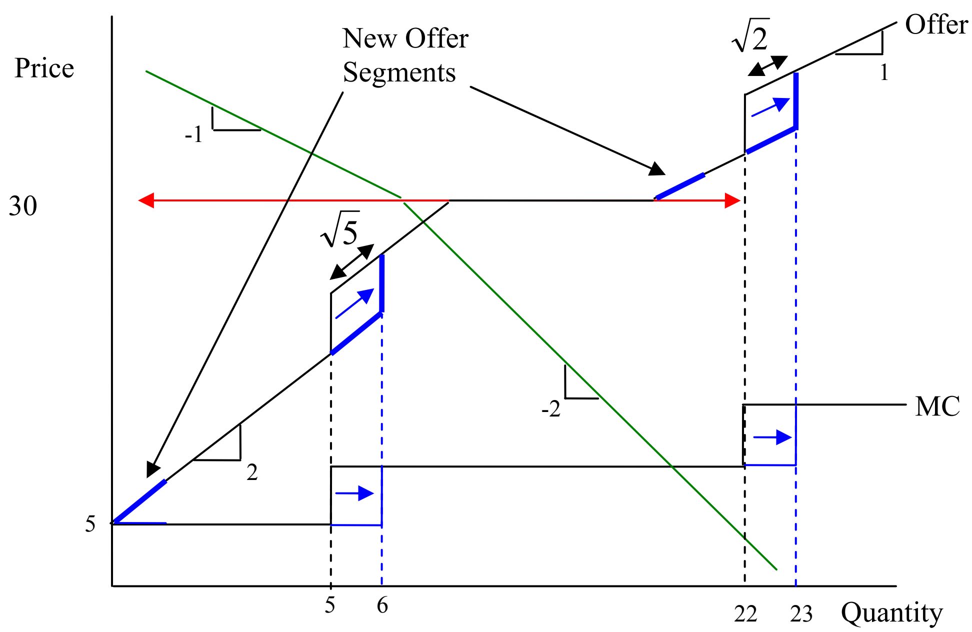

Figure 4.45 Change in Offer from Shifting Marginal Cost Curve – Case 3

We observe that the new optimal offer has been translated in sections bounded by the RDVP(s). In other words, the segment of the original offer that was below the RDVP has been translated at a slope of 2, by $\sqrt{5}$, while the segment of the original offer above the RDVP has been translated at a slope of 1, by $\sqrt{2}$. Additionally, we note that there are two new offer segments that have been created from scratch in this case, and they are indicated in Figure 4.45. The same results as shown here apply when we consider a *linear* marginal cost over a concave RD curve.

4.6.5 Case 4 – Convex RD Curve and Stepped Marginal Cost

Finally, consider the translation that is appropriate under a convex RD curve. Assume the same marginal cost curve as in case 3, but now the RD slopes are reversed, so that the slope is -2 above a price level of 30, and -1 below. This case is shown in Figure 4.46.

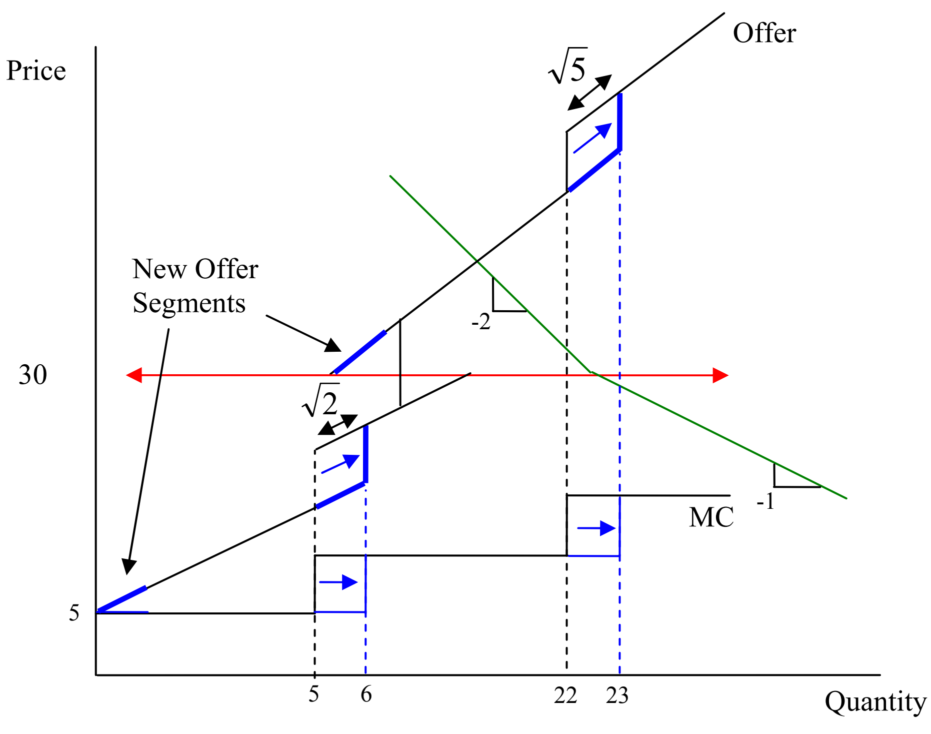

Figure 4.46 Change in Offer from Shifting Marginal Cost Curve – Case 4

We observe that the new optimal pseudo-offer has again been translated in sections bounded by the RDVP(s). The segment of the original offer that was below the RDVP has been translated at a slope of 1, by $\sqrt{2}$, while the segment of the original offer above the RDVP has been translated at a slope of 2, by $\sqrt{5}$. Again, there are two new offer segments that have been created from scratch in this case. The same results apply when we consider a linear marginal cost over a convex RD curve

4.6.6 Translation Summary

In summary, the above cases have demonstrated that if the marginal cost curve shifts horizontally, then we can translate the previously optimal offer to create most of the new optimal offer. In particular, between $RDVP_i$ and $RDVP_j$, the slope of the RD curve will be $-b_i$. When the marginal cost curve shifts horizontally by *n* units, the original offer segment in this range will be translated at a slope of b_i, a distance of $\sqrt{b_i^2 + n^2}$. The new offer segments that are adjacent to the price axis or each of the RDVPs must then be created separately to complete the offer.

We have seen that this process works successfully to translate the raw pseudo-offers, regardless of any non-monotonicity that they may contain. However, the optimal offer that we initially have for the original marginal cost curve position would likely be the monotone offer, possibly formed using the optimality conditions presented in Section 4.3. If the marginal cost has not changed at the position of the previously optimal vertical step or over any part of a horizontal step, the optimal position of that vertical or horizontal step will be unchanged. If however, the marginal cost has changed over any part of these ranges, the translation would need to be performed on the pseudo-offer, and the position of the vertical or horizontal steps recalculated with the new marginal cost information.

4.7 Summary and Conclusions

In this chapter, we have discussed the optimal analytic form of single-period offers under restricted forms of market uncertainty. We have shown that when there is a piecewise linear residual demand curve subject to uncertainty in a single direction only, the optimal desired offer can be constructed by combining segments derived from easily constructed analytic expressions, over appropriate regions of the offering space. Due to offer monotonicity conditions in the market, when the combined offer segments do not form a monotonic set, we can apply optimality conditions to determine the optimal position of vertical or horizontal steps through the non-monotone regions. For small regions of non-monotonicity, these optimality conditions can be quite simple. However, as the regions of non-monotonicity increase in size, the optimality conditions can quickly become significantly more complicated.

We have also shown that these analytic offer construction concepts apply to general forms of residual demand curve uncertainty, but that again, in these cases, the optimality conditions quickly become more complex as the structure of the uncertainty becomes less straightforward.

In the following chapter, we will use the optimal analytic offer concepts developed here to prove a marginal cost-based offer patching approach, which will later be embedded within a dynamic program to significantly improve optimal offer construction computational efficiency.

Chapter 5

MARGINAL COST PATCHING

5.1 Introduction

In real-life, offers made to the pool market by generators need to be constructed on the go, in real-time, so ideally the operators need an approach to offering that is very fast to apply. Hence, having a pre-computed table of optimal responses to various circumstances that can be searched quickly would be very helpful. In this chapter we present a marginal cost based “offer curve patching” approach. For a given (expected) form of uncertain RD curve, a family of offer curves corresponding to a range of constant marginal costs (or marginal fuel values) can be constructed. In a continuous sense, this family of curves would effectively form an offer surface.

Our basic proposition is that given the marginal cost curve for the current period and fuel storage level, segments of members of these offer curve families can be joined together, to produce the optimal offer for the coming period. Specifically, the motivation for the development of this patching approach lies in the production of a dynamic programming model for the generation of an offering strategy in real-time, for a short-term planning horizon. This model is described in Chapter 6.

In this chapter, we begin by demonstrating that marginal cost based offer patching is appropriate when the underlying pseudo-offers for the constant marginal costs are naturally monotonic. We then show that this approach is also appropriate under the more complex scenario of underlying non-monotonic pseudo-offers. In this case, the proposed patching is no longer just the combination of previously optimal dispatch points, but also the combination of the vertical or horizontal segments that have been applied to the underlying pseudo-offers to resolve their non-monotonicity (using the methods described in Section 4.3). In this chapter, we begin by retaining the assumption that the RD curve moves as a whole in any defined direction, but then extend the analysis to consider any form of RD curve uncertainty.

5.2 Patching Naturally Monotonic Offers

In Section 4.3.2, we showed that under certain RD shape and uncertainty direction (RDVP) combinations, the optimal set of dispatch points for all possible RD curve outcomes could form an offer that is naturally monotonic in both the price and quantity dimensions. These results concurred with Theorem 4 in Anderson & Philpott (2002b). Further to that theorem, we now define the following corollary.

Corollary 5.1:

When the set of optimal dispatch points under all possible residual demand curve outcomes forms a naturally monotonic offer, the optimal offer price at any particular quantity level will depend only on the local marginal cost level. It is therefore *independent of the marginal cost levels that have come before that quantity and those that will occur beyond it. In other words, the optimal offer for a firm with a marginal cost curve that steps from* MC_1 *to* MC_2 *at the quantity BP will be the combination of:*

- *The section of the offer from 0 to BP under the assumption that marginal cost is equal to* MC_1 *over the entire range, and*

- *The section of the offer from BP to Qmax under the assumption that marginal cost is equal to* MC_2 *over the entire range.*

Proof. The key to this proof is simply that the optimal price at which to offer electricity under each quantity level is dependent on the marginal cost at that quantity level, and no other. In Section 4.3.2, we explained that in order for the set of optimal dispatch points to be naturally monotonic, the RD curve must be concave. We know that under such a concave form, as the RD curve moves to the right, the optimal dispatch quantity will always increase. Therefore, there can only be a single RD curve that will lead to dispatch at any given quantity, and it is this curve that is used to determine the optimal price level at which to offer that quantity. Figure 5.1 shows such a single RD curve scenario, where the optimal dispatch point is at the quantity where marginal revenue (MR) is equal to marginal cost (MC). With a convex cost function (non-decreasing marginal cost function) and concave RD curve (non-increasing marginal revenue curve), there must be a unique point of intersection between these curves. Further, observe that this point of intersection is dependent only on the local marginal cost (i.e. the optimal point would not change if marginal costs at other generation levels were different, as indicated by the arrows). This is confirmed by the explicit restriction on the marginal cost curve to be non-decreasing.

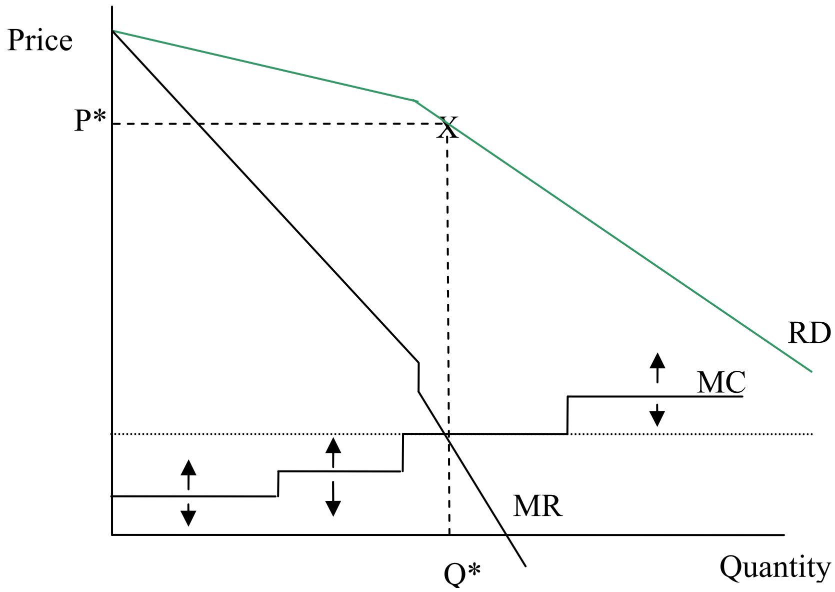

Figure 5.1 Optimal Dispatch Point for a Single Residual Demand Scenario

5.2.1 Examples of Patching Naturally Monotonic Offers

Assume that a generator faces an uncertain linear RD curve with slope $-b_i$ and a constant marginal cost of either MC_1 or $MC_2 (> MC_1)$. It was shown in Section 4.3.1 that the optimal form of the offer for each of these cases would be $p = MC_i + b_i q$ (lines with slope b_i and intercept MC_i). Figure 5.2 shows the offer for the lower marginal cost in blue, and for the higher marginal cost in black. Now consider the form of the new optimal offer if the marginal cost stepped from MC_1 to MC_2 at the quantity level *BP*. Corollary 5.1 tells us that when each of the original offers is naturally monotone, the optimal price to offer for any quantity level is dependent on the marginal cost at that quantity, but independent of the marginal cost at any other quantity. Therefore, the new offer should be the same as the original offer for MC_1, up to $q = BP$, and the same as the original offer for MC_2 beyond $q = BP$. In other words, to get the new optimal offer, we patch together the segments of the original offers for MC_1 up to the breakpoint, and for MC_2 beyond the breakpoint, as shown with the dotted line.

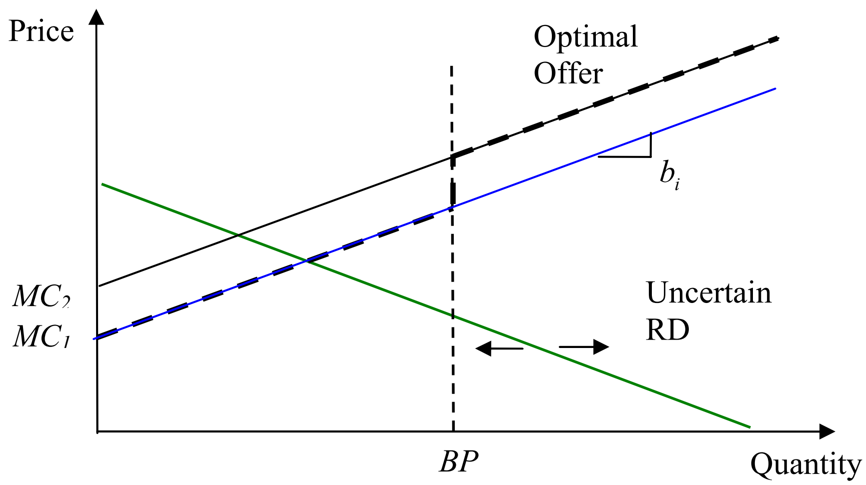

Figure 5.2 Patching Example – Linear Residual Demand

Now consider a slightly more complicated example, where the RD curve contains a concave kink, and is now subject to uncertainty in the direction indicated by the positively-sloped RDVP in Figure 5.3. The RD curve now has slope $-b_1$ below the RDVP, and $-b_2$ above (where $|-b_1| > |-b_2|$). Again, the results from Section 4.3.1 require that the pseudo-offer for each marginal cost will have the form $p = MC_i + b_1 q$ below the RDVP, and $p = MC_i + b_2 q$ above, with a connecting segment along the RDVP in-between. Again, if we plot the optimal offer for a marginal cost curve that steps from MC_1 to MC_2 at the *BP*, we observe that this is equivalent to patching together the segments of the original offers for MC_1 up to the breakpoint, and for MC_2 beyond the breakpoint. On examination, it can be shown that this marginal cost patching will work no matter where the breakpoint is located in the quantity dimension.

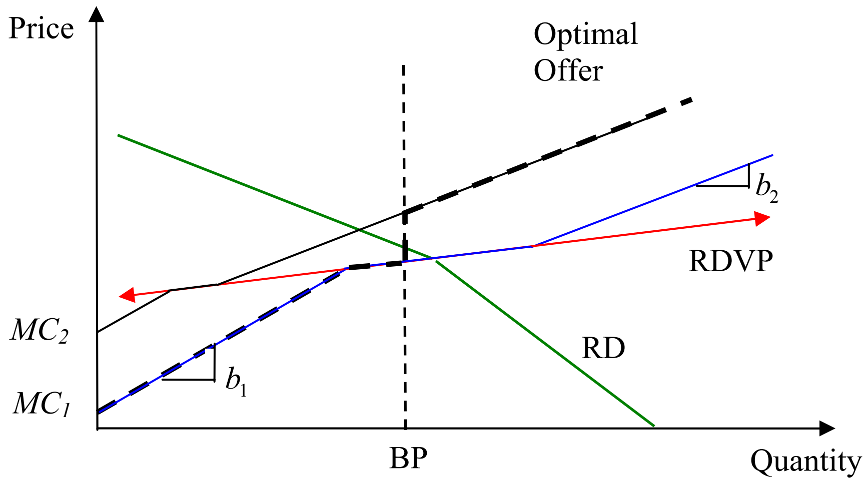

Figure 5.3 Patching Example – Concave Residual Demand

Therefore, the theorem from Philpott *et al* (2002), in conjunction with our corollary shows that under the assumption of an inverse log concave rest-of-market supply curve with demand uncertainty, offers corresponding to constant marginal costs can be patched together to produce the optimal offer for the case of stepped marginal costs. However, it is highly likely that the actual rest-of-market supply curve faced by a generator will not meet these conditions[29]. Therefore, in the following sections we investigate the applicability of marginal cost patching under offers corresponding to any set of uncertain RD curves.

5.3 Patching Naturally Non-Monotonic Offers

The proof provided for Corollary 5.1 is restricted to the case where the shape of the RD curve leads to naturally monotonic sets of optimal dispatch points. We now state and prove a new theorem, which says that the marginal cost patching approach will work regardless of the form of the set of optimal dispatch points. Therefore, this is applicable to both the non-monotonic cases presented in Section 4.3 or to any other set of uncertain RD curves that has a non-monotonic desired set.

[29] For example, any concave piecewise linear rest-of-market supply curve does not satisfy this condition.

Theorem 5.1:

Regardless of the form of the set of optimal dispatch points under all possible residual demand curve outcomes, segments of the optimal offers for constant marginal costs can be patched together to provide the optimal offer for a marginal cost curve that is stepped.

Proof. We have already shown that this theorem is true when the set of optimal dispatch points is monotonic. Therefore, all that remains is to prove it to be true when this set of points suffers from each of Type I or Type II non-monotonicity. In these cases, it is possible that the marginal revenue curve will not be monotonically non-increasing. As a result, we can no longer say that the dispatch quantity for a given RD curve will remain the same when the marginal cost in other quantity ranges change.

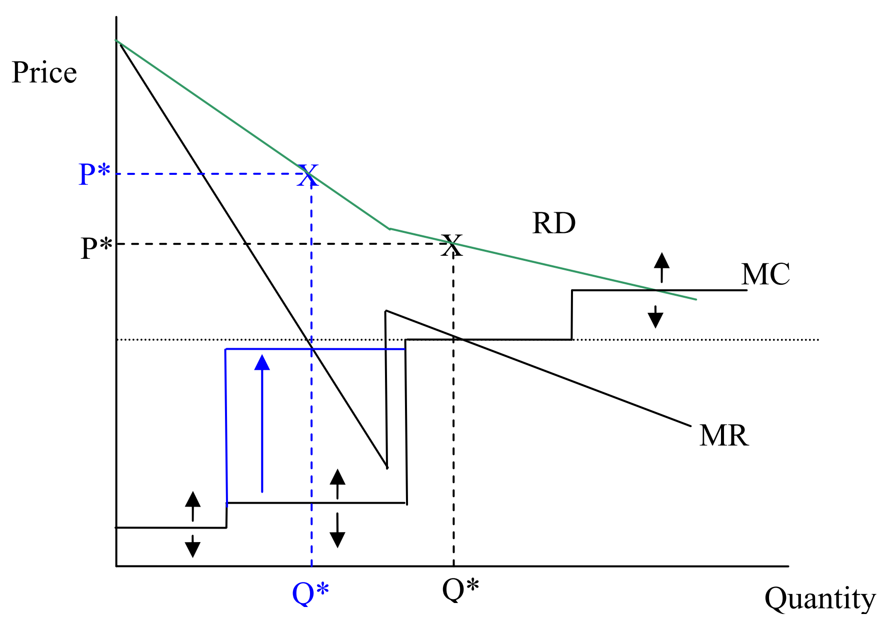

Figure 5.4 Dispatch Quantity Affected by Marginal Cost in other Ranges Changing

For example, Figure 5.4 shows the original optimal dispatch point for a given RD curve in black. The blue lines indicate a possible new optimal dispatch point if the marginal cost in a range other than that of the original dispatch quantity changes (there are now two possible local optima, the previously optimal point and the new intersection as market in blue). Clearly, under this scenario, the optimal dispatch point is not dependent only on the local marginal cost level, and hence a new proof approach is required for these cases.

Marginal Cost Patching under Type I Non-Monotonicity

Recall that the optimal resolution of Type I non-monotonicity is to insert a vertical offer segment between the lower and upper attractor lines, at a position such that the following optimality condition is met:

$$\sum_{l(q)}^{u(q)} pr(sc)[MR(q) - MC(q)] = 0$$

This optimality condition requires that we select a quantity level for the vertical segment such that the sum of the probability of a scenario occurring multiplied by its marginal profit at that quantity, for all RD scenarios that would pass through the vertical segment, must sum to zero. Importantly, we note that this formula relates to a given quantity, and is dependent only on the marginal cost at that quantity. We also observe that if the level of the constant marginal cost is increased over the entire range, the optimal position of the vertical step that is produced from this optimality condition must move to the left. This is clear when you consider that all terms will decrease under these circumstances.

Now let us consider the form of the offers that we are attempting to patch. Figure 5.5 shows two offers, each corresponding to a different marginal cost, where vertical steps ($V(MC_1)$ and $V(MC_2)$) have been used to overcome non-monotonicity in the pseudo-offers.

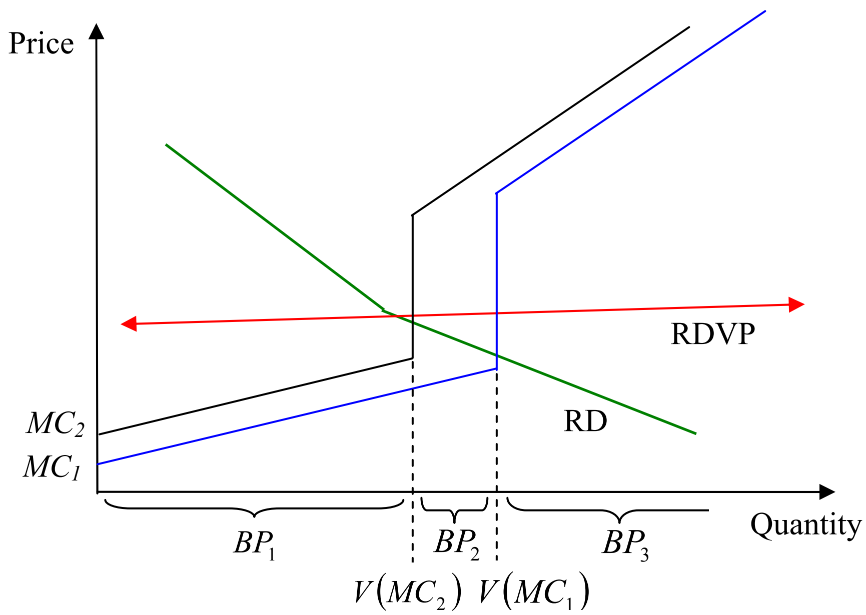

Figure 5.5 Offer Forms from Type I Non-Monotonicity

There are effectively just three different quantity regions for the marginal cost jump that we must consider. These regions are:

- $BP_1 < V(MC_2)$: before the optimal vertical step for the higher marginal cost,
- $V(MC_1) < BP_2 < V(MC_2)$: between the optimal vertical steps for the two marginal costs,
- $BP_3 > V(MC_1)$: beyond the optimal vertical step for the lower marginal cost.

Marginal Cost Breakpoint in Region BP_1:

Start by considering a marginal cost jump from MC_1 to MC_2 in region BP_1, as indicated by the dotted line BP in Figure 5.6. The heavy dashed lines indicate the new locally optimal points of dispatch associated with all possible RD curve scenarios, giving the new attractor line and hence defining the new region of offer uncertainty from *s* to *f*. This is clarified in Figure 5.7.

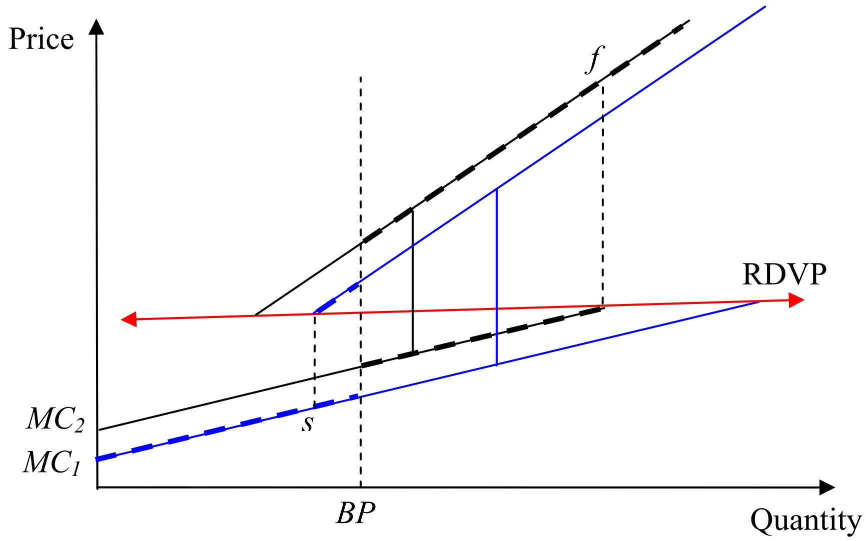

Figure 5.6 Marginal Cost Jump BP$_1$

The Direction of Improvement theory, discussed in Section 4.3, tells us that the optimal resolution of this region must involve a vertical step from the segment of the attractor line below the RDVP to the segment above. If we were to place the vertical step in the range labelled *a* in Figure 5.7, the lower marginal cost would be applied in the optimality condition. We know that the quantity at which the optimality condition is satisfied for this marginal cost, $V(MC_1)$, is to the right of BP, and hence forces would push this vertical to at least $q = BP$. For any current vertical step position in the ranges labelled *b* or *c*, the higher marginal cost would be used in the optimality condition. We know that when this marginal cost is used, the optimal position of the vertical step is at $V(MC_2)$. In other words, the optimal offer to provide under this stepped marginal cost curve (as shown in red), consists of the optimal offer for the lower marginal cost curve up to BP patched together with the optimal offer for the higher marginal cost curve beyond this point (including the vertical step).

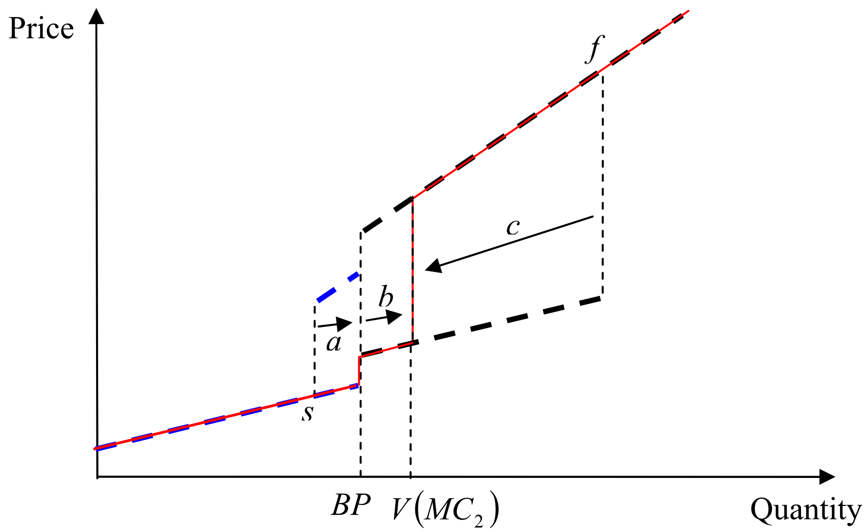

Figure 5.7 Marginal Cost Jump in Region BP$_1$ – Optimal Vertical Position

Marginal Cost Breakpoint in Region BP_3 **:**

Very similarly, we can show that the patching approach is appropriate for a marginal cost step in region BP_3, as indicated by the dotted line BP in Figure 5.8. This diagram shows the region of offer uncertainty under this new marginal cost curve. If the vertical step was placed to the right of BP, there would be forces pushing it to the left (as far as BP). If the vertical segment was placed anywhere below BP, then the forces would attract it to the same position as it would be under a constant lower marginal cost level. Hence, for a marginal cost step at BP, the new optimal offer consists of the original optimal offer for a constant marginal cost of MC_1 up to $q = BP$ (including the vertical step), patched together with the original optimal offer for a marginal cost of MC_2 beyond $q = BP$.

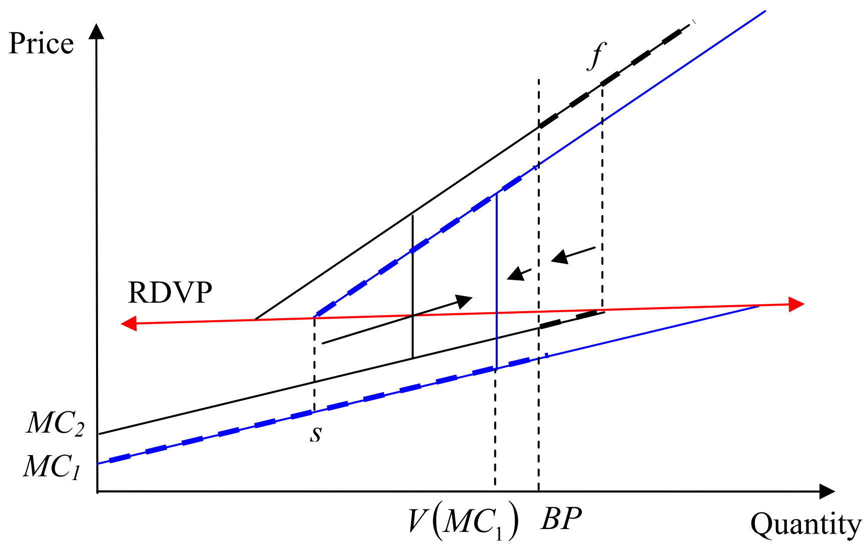

Figure 5.8 Marginal Cost Jump in Region BP$_3$

Marginal Cost Breakpoint in Region BP_2:

Again, for a marginal cost step in region BP_2, as indicated by the dotted line BP in Figure 5.9, we can prove that patching the optimal offer segments together is appropriate, by considering the forces acting on a vertical step placed at any position within the range of offer uncertainty. The heavy dotted lines on Figure 5.9 show the new locally optimal points of dispatch associated with all possible RD curve scenarios, under each of the two MC levels. A key observation at this point is that both:

- The optimal position of the vertical jump for the lower marginal cost, $V(MC_1)$, is not within the new range for which the lower marginal cost is valid, and

- The optimal position of the vertical jump for the higher marginal cost, $V(MC_2)$, is not within the new range for which the higher marginal cost is valid.

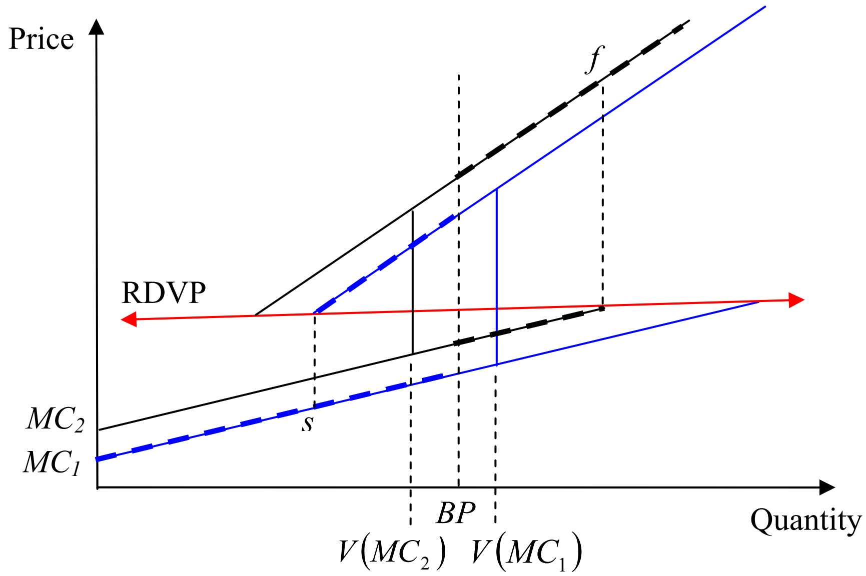

Figure 5.9 Marginal Cost Jump in Region BP_2

Therefore, if the vertical step was placed to the right of BP (in the range where MC_2 is valid), forces would push it left towards $V(MC_2)$, as far as BP (beyond which the marginal cost changes). Likewise, if the vertical step was placed to the left of BP (in the range where MC_1 is valid), forces would push it right towards $V(MC_1)$, as far as BP (beyond which the marginal cost changes. Therefore, the optimal position for the new vertical step must be at the quantity BP, and we can say that for a marginal cost step at BP, the new optimal offer consists of the original optimal offer for a constant marginal cost of MC_1 up to $q = BP$, patched together with the original optimal offer for a marginal cost of MC_2 beyond $q = BP$, along with a completely new vertical segment at BP.

Marginal Cost Patching under Type II Non-Monotonicity

Recall that the optimal resolution of Type II non-monotonicity is to insert a horizontal offer segment between the left and right attractor lines, at a position such that the following optimality condition is met:

$$\sum_{l(p)}^{u(p)} pr(sc)[MR(p) - MC(p)] = 0$$

This optimality condition requires that we select a price level for the horizontal segment such that the sum of the probability of a scenario occurring multiplied by its marginal profit at that quantity, for all RD scenarios that would pass through the horizontal segment, must sum to zero. We observe that if the level of the constant marginal cost is increased over the entire range, the optimal position of the horizontal step that is produced from this optimality condition must move up.

Now let us consider the form of the offers that we are attempting to patch. Figure 5.10 shows two offers, each corresponding to a different marginal cost, where horizontal steps ($H(MC_1)$ and $H(MC_2)$) have been used to overcome non-monotonicity in the pseudo-offers.

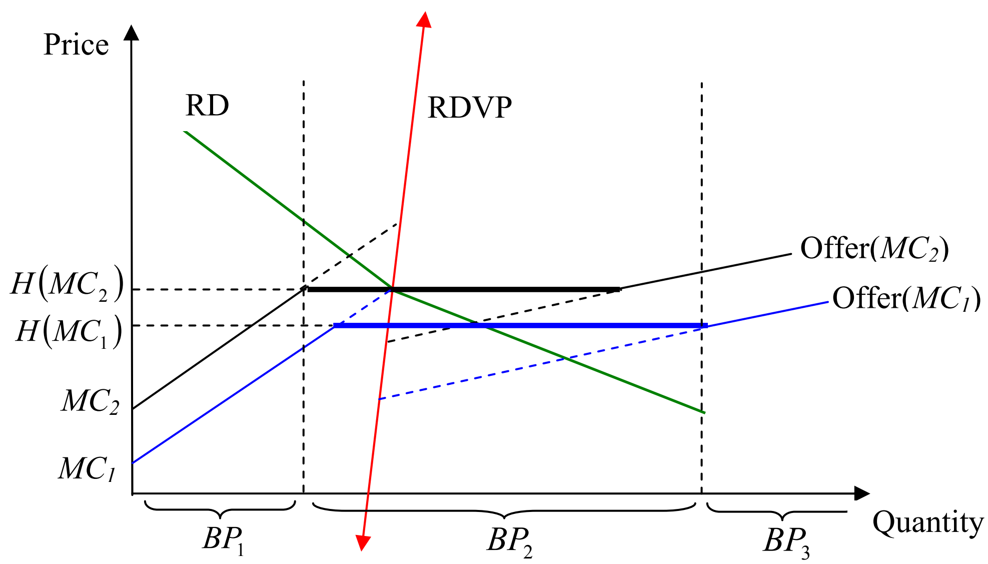

Figure 5.10 Offer Forms from Type II Non-Monotonicity

Again, there are effectively just three different quantity regions for the marginal cost jump that need to be considered. These regions are:

- BP_1: before the left end of both of the optimal horizontal segments, shown in bold blue and black at the levels of $H(MC_1)$ and $H(MC_2)$.
- BP_2: in the middle of either or both optimal horizontal segments,
- BP_3: beyond the right ends of both of the optimal horizontal segments.

Patching can be shown to give the new optimal offers under these three positions in the same way as for Type I Non-Monotonicity. For example, Figure 5.11 shows all new points of locally optimal dispatch associated with a marginal cost jump from MC_1 to MC_2 at *BP*, located in region BP_1. It is clear that the optimality condition on the position of the horizontal step will still be met at the same price, as it includes only terms on the step itself, and the MC and MR has not changed for any of these terms. Therefore, the optimal offer can be patched together in the same way as it was for a vertical step.

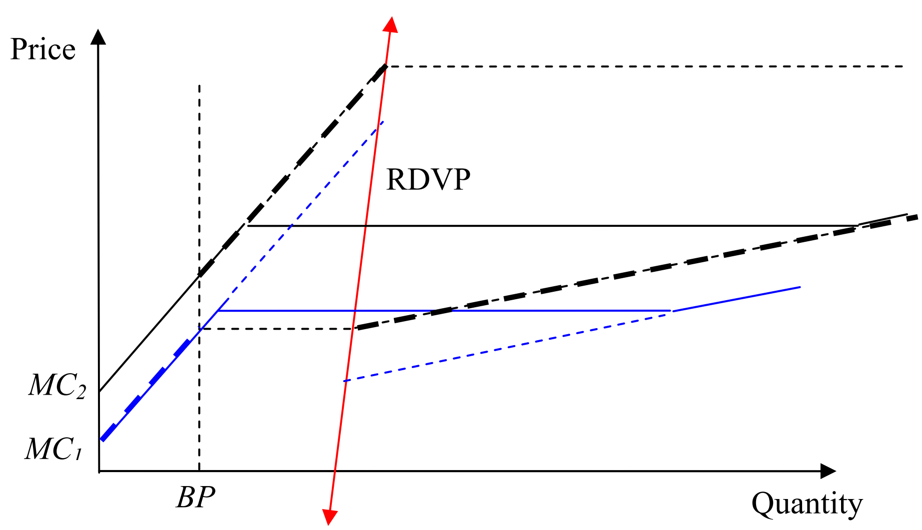

Figure 5.11 Marginal Cost Jump – BP1

When we consider general RD uncertainty and produce the optimal offer by performing a DP over a discrete grid on (price, quantity) space (as explained in Section 3.4), we are implicitly locating the optimal positions of any necessary horizontal or vertical steps. As discussed in Chapter 3, under anything other than very simple forms of market uncertainty, this is a simpler process than the analytical methods to achieving the same goal. As it is just a different approach to the same problem, then the MC patching presented in this chapter must hold true when the optimal offers are constructed using this DP approach. This information is applied in Chapter 6.

5.3.1 Examples of Patching Naturally Non-Monotonic Offers

In this section, we will demonstrate two examples of patching under naturally non-monotonic offers, where all offers (for both fixed and stepped marginal costs) have been constructed using a dynamic program on a discrete grid over the offering range, as discussed in Section 3.4. The first example considers Type I non-monotonicity, where the offers contain vertical steps (the results from the DP confirms our analytical expectations), while the second example considers both Types I and II non-monotonicity and the offers contain both vertical and horizontal steps.

Example 1: Type I Non-Monotonicity

Consider an uncertain RD curve that is subject to horizontal uncertainty and that corresponds to a rest-of-market supply curve that is not inverse log concave, and thus violates the assumptions of Theorem 4 of Anderson & Philpott (2002b). Specifically, the expected position of the RD curve is defined by:

$$p = \begin{cases} 30 - 0.5q & p < 16 \\ 44 - q & p \geq 16 \end{cases}$$

In Section 4.3.2, we showed that an uncertain RD curve of this (convex) shape will produce a non-monotonic attractor line, and the optimal resolution of this Type I non-monotonicity was to provide an offer containing a vertical segment (at a position determined by a simple optimality condition).

Consider Figure 5.12. The stepped red, black and blue lines represent offers corresponding to constant marginal costs, in this case, of 5, 10 and 25 ($/MWh) respectively. By considering a stepped marginal cost function which steps between these marginal cost levels at the breakpoints of 12 and 20 MWh, this numerical example, which has been used by solving a DP to maximise profit for any feasible offer over a discrete grid on the offering space, confirms that the offer obtained (green line) corresponds to patching together the appropriate segments of the offers for constant marginal costs.

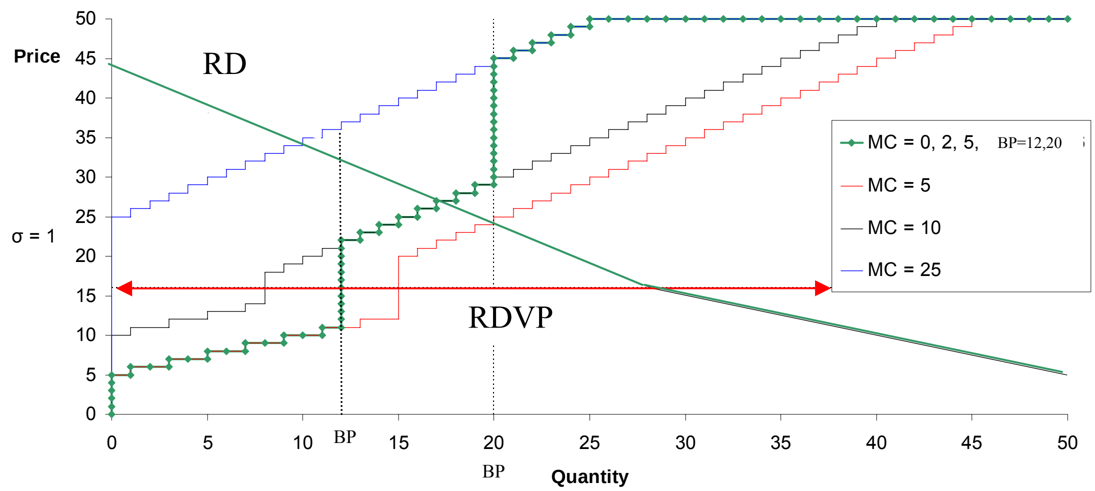

Figure 5.12 Optimal offers with fixed and stepped marginal costs

Example 2: Types I and II Non-Monotonicity

Consider a rest-of-market supply function that has a very general form, with concave and convex segments, as well as horizontal steps and vertical jumps, implying a very general expected RD curve (as shown in Figure 5.13). In this example, the red, black and blue lines represent the optimal offers under this RD curve form if the marginal cost was fixed at the levels of 0, 10 and 25 ($/MWh) respectively, starting from the lower-most curve. The green line shows the optimal offer if we knew in advance that the marginal cost was going to step between these three levels at the breakpoints of 12 and 31 MWh's. It is clear that this is the same as the offer that would be achieved under patching, and thus this confirms that marginal cost patching is appropriate for the offers produced under general RD curve forms.

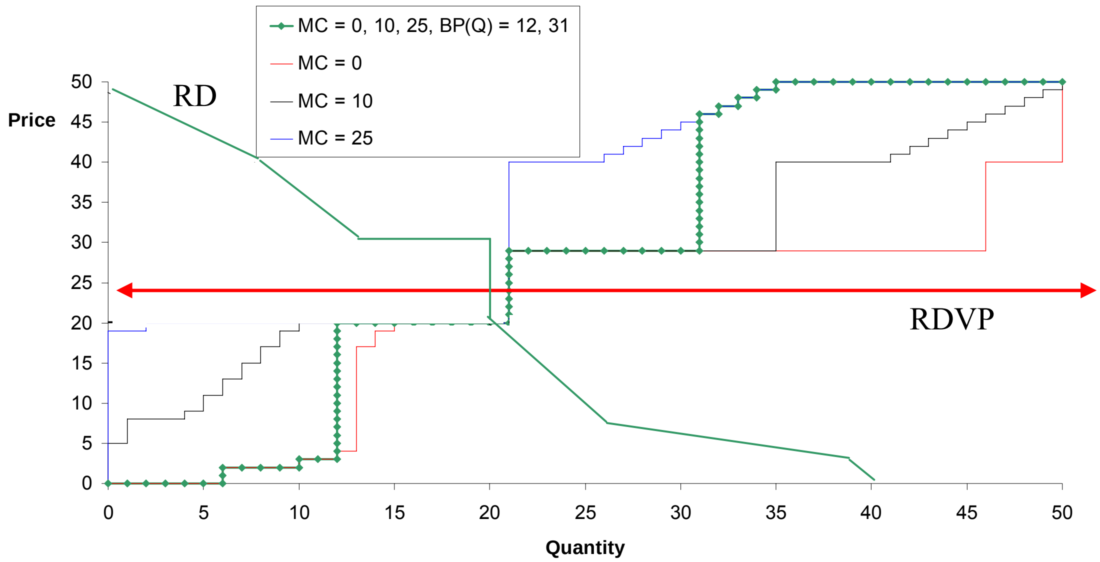

Figure 5.13 Optimal offers with fixed and stepped marginal costs

5.4 Summary and Conclusions

In this chapter we have introduced a concept that we have termed as marginal cost patching, which states that the optimal offer for a firm with a marginal cost curve that steps from MC_1 to MC_2 at the quantity BP can be obtained by joining together the section of the offer from 0 to BP under the assumption that marginal cost is equal to MC_1 over the entire range, and the section of the offer from BP to $Q_{\max}$ under the assumption that marginal cost is equal to MC_2 over the entire range. We began by presenting and proving a corollary on Theorem 4 from Anderson & Philpott (2002b) which shows that this marginal cost patching is appropriate under the restricted conditions of horizontal uncertainty in the RD curve, and log concave RD curve. We then went on to propose and prove a theorem which states that this marginal cost patching is appropriate under any set of possible RD curves. This means that the offer curve family concept can be used to produce the optimal offer under any marginal cost curve, quickly and efficiently.

In the next chapter, the marginal cost patching concept will be applied to the problem of constructing an offering strategy over a multiple-period horizon. We will show that it can be used to separate an existing *two-level nested* dynamic program into a significantly more computationally efficient *two-phase* dynamic program. In this two-phase DP, the optimal offers for fixed MC levels can be constructed in advance of real-time and then combined quickly when needed, to produce optimal offers for any marginal cost curve that is faced by a generator.

Chapter 6

OFFER CONSTRUCTION ALGORITHMS USING MARGINAL COST PATCHING

6.1 Introduction

In this chapter, we describe two algorithms for the construction of the optimal offer set for a generator. The algorithms we propose have significantly lower computational complexity than those found in the literature and hence increased potential for solving more complex scenarios within the limited available time-frame of online offer construction[30]. Computational complexity has been reduced by applying Theorem 5.1, regarding marginal cost-based offer patching, developed in Section 5.3.

We begin the chapter by describing the purpose of the algorithms and introduce a new approach to optimal offer construction over time, which we call the Two-Phase approach. We then describe the algorithms underlying each of these phases in detail, including proofs that the approaches will consistently produce concave value curves (i.e. decreasing MV curves).

6.2 Algorithm Purpose

The goal of the algorithms studied in this thesis, including those described in this chapter, is to produce an optimal company-wide offering strategy for an energy-limited generator that provides electricity into a wholesale market, in a solving time that would be practical for online implementation. Note that the term energy-limited does not restrict its application to just hydro generators, as it is relevant to any generator with a significant constraint on fuel with respect to the short-term planning horizon (a coal-powered generator with a contract to purchase a fixed amount of coal, for example). We define an offering strategy as a set of monotonic (and hence feasible) offers that would be provided under any state (with respect to reservoir level and market situation) that the operator can face over the planning horizon.

The models consider either a single generating unit or a set of units attached to the same reservoir. The extension to a hydro reservoir chain is discussed in Section 10.5. The reason that this problem is especially complex is due to the consideration of energy generation limitations over a planning horizon, reflecting either a resource-constrained thermal unit, or a normal hydro unit. This intertemporal constraint effectively links the decisions made in all periods, as trade-offs that have to be made to ensure the most profitable expected use of fuel. An additional major complexity considered is that the

[30] This specific time-frame is dependent on the market considered and the other processes that need to occur between market offering rounds. In the New Zealand market, offers must be provided once every 30 minutes, and as such, we define a reasonable computational time to be up towards 10 minutes.

uncertainty in the (discrete) market outcomes is correlated. This is modelled through a set of Time-Varying Markov Chains representing the probability that if residual demand curve "*A*" is observed in period *t*, then residual demand curve "*B*" will be observed in period *(t+1)*.

There are many offering strategy algorithms reported in the literature, but most of these either make significant unrealistic assumptions (such as a restriction to quantity-based offers, as in Stothert & MacLeod (2000) or Pereira, Granville, Dix, & Barroso (2004a)) or are computationally too complex to be used in any sort of practical setting (as in Rajaraman & Alvarado (2003)). In particular, the literature provides three particularly interesting DP models used in the supply of electricity over time. Philpott & Schultz (2004) considers market uncertainty within a single period, in addition to many realistic intertemporal constraints. Scott (1998) deals with inflow uncertainty, while trying to optimise the use of limited fuel over a horizon (an important constraint which the former ignore). Finally, one of the most comprehensive, although computationally demanding, algorithms in the literature to date is described in Rajaraman & Alvarado (2003). Not only does their algorithm consider offers restricted by monotonicity, it also considers correlated (discrete) market uncertainty. To the author's knowledge the latter extension has not been dealt with in any of the previous literature. However, the algorithm as reported is quite inefficient and, as a result, intractable within a feasible solving time[31] for all but very small models. Even improving the efficiency of the algorithm does not do enough to bring the computational time down to a practical level. Hence, a significantly new approach is required in order to achieve the goal of practicality. The algorithm described in Rajaraman & Alvarado (2003) is hereafter referred to as the R&A algorithm.

This chapter develops two alternate *marginalistic* DP approaches that combine the desirable features of each of these existing models, in order to achieve the goal of practicality. The new models have been designed as practical tools for developing

[31] We consider a maximum feasible period for solving an online model for offering would be around 10 minutes on a P4 3.2GHz, 1GB machine.

optimal (or near-optimal) offers, with minimal computational complexity. This low complexity will be made possible by applying the knowledge gained about analytic offer forms in earlier chapters, rather than using complex optimality conditions and maximising complicated path integrals as Philpott & Schultz (2004) requires.

To summarise, the algorithms discussed in this chapter consider the following key issues, or complexities of the modelling situation:

- Market uncertainty within each individual period
- Correlated market outcomes between periods, using a Markov Chain
- Fuel limitations over the horizon
- Fuel inflow uncertainty
- Construction of feasible (monotone) offers for every possible state throughout the horizon.
- Limited solving time for real-time calculations

6.3 R&A Algorithm

The R&A algorithm is structured as a two-level nested dynamic program. The upper level is two dimensional, and works backwards from the end of the horizon, establishing the value at each state, where a state is defined by the previously observed residual demand curve (hence implying a particular uncertain market scenario, or UMS[32] for the following period), and the current reservoir level. The lower level of the algorithm is implemented at each of these states in order to determine this value. The dynamic program at this level works backwards from (qmax, pmax) to (0, 0) through a grid of expected payoffs, constructing the feasible offer that gives the optimal expected value (current period payoff plus expected future value) for this (previous RD, reservoir level) state (demonstrated in Section 6.5.3). A feasible offer is one that is monotonically non-

[32] Note that a UMS is effectively a discrete market distribution function, as reported in Philpott, Pritchard, Neame, & Zakeri (2002).

decreasing in both its (p, q) dimensions. A simplified representation of this algorithm is presented in Figure 6.1.

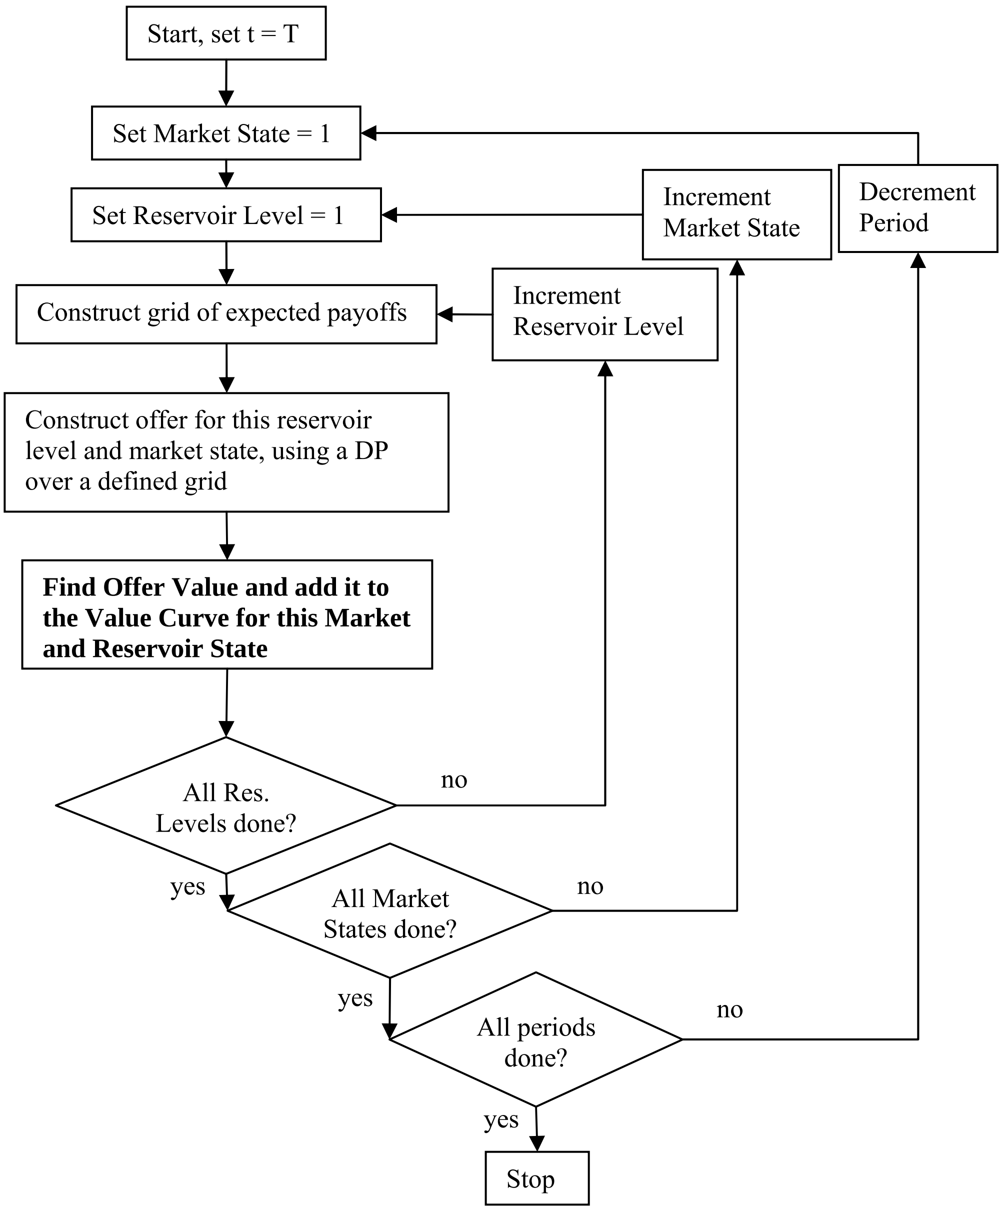

Figure 6.1 R&A Algorithm Flow Diagram

There are many ways in which we can compare algorithms discussed in this thesis with one another and previous models found in the literature. These include:

1. The complexities and issues considered by the algorithms

2. The theoretical computational complexity of the algorithms
3. The quality of the results that the models produce
4. The actual computational times required on test data sets

As discussed in the previous section, the three approaches compared in this chapter all consider the same issues. This chapter considers the theoretical computational complexity of the algorithms, while Chapter 8 is devoted to comparing the quality of results and actual computational times. Unfortunately, unlike a Travelling Salesperson and similar problems, where the size of a problem can be defined by a single parameter, here there are many parameters that affect the computational complexity of the algorithm.

We define the following five parameters for use in expressing the computational complexity of the R&A algorithm:

t = Number of periods in the horizon
r = Number of reservoir levels
d = Number of dispatch levels
p = Number of possible RD curves per period[33]
i = Number of possible inflow levels

From the source code for the R&A algorithm, we can show that the computational complexity can be approximately summarised by Equation 6.1.

$$\frac{t\left(10 + 2pi + 30pr + 4pri + 8r + 2rpd^2 + 32pdr + 18rp^2d + 16rp^2d^2 + 2rp^2d^3 + pdr^2\right)}{2}$$

Equation 6.1 Computational Complexity of the R&A Algorithm in Full

Simplifying this down to the terms of the largest orders of magnitude, we can see that the algorithm is of order of complexity $trp^2d^3 + tpdr^2$.

[33] for simplicity, we consider this to be constant over all periods

This shows us that the computational complexity of the R&A algorithm will increase polynomially as the values of the various parameters which define the problem size increase (excluding t, in which the computational complexity increases only linearly).

6.3.1 Problems with the R&A Algorithm

There are a couple of important issues with the R&A algorithm, with respect to computational burden, that have led us to develop the two-phase approach presented in the remainder of this chapter.

1. In the R&A approach to offer construction, a Markov chain is incorporated to describe the correlation between the occurrences of RD curves. As the horizon progresses, no changes are made to the Markov chain as market outcomes are observed, as this information is incorporated into the chain. Changes to the chain would occur only as a result of external events. Therefore, if the horizon was of a fixed length[34], and no external events were to occur over the horizon and thus no probabilities or expectations were to change over time, then there would be no need to re-solve the model throughout the horizon (as the initial solve would provide the offers for every reservoir and market state combination that we could observe through the horizon). However, it is highly likely that external events will occur and there *will* need to be changes made to some of these future expectations as time progresses. Such external events could be either:

    a. Pre-dispatch outcomes for future periods
    b. Information that comes from outside the model. For example, if we were to suddenly learn that maintenance was likely to shut down the main generation unit of a rival generator later in the day.

[34] In a practical sense, it is likely that there would be a discrete form of a rolling horizon, to reflect the real-life offering requirements. In the New Zealand case, this means that at 1pm every day, an additional 48 periods would be added to the horizon, and in all other periods, the horizon would just progressively shorten.

If such changes to expectations occur[35], then under the R&A algorithm, the entire model (including all of the lower-level, time-consuming offer construction DPs) would have to be completely re-solved before the new set of offers can be made to the market. This severely restricts the detail and complexity that can be considered by this approach, in order to get the new offers into the market in the required time available.

2. In the R&A algorithm, offers must be calculated for all possible system states in all periods, regardless of whether identical scenarios are considered to occur in multiple periods. For example, if a possible UMS from period 6 (off-peak) is the same as one of the possible UMSs in period 30 (also off-peak), the R&A algorithm is not be able to capitalise on this repetition due to its reliance on a purely primal DP approach. The algorithm that we will propose in this section will be able to avoid these wasteful lower-level DP calculations being performed. Specifically, if the same UMSs are likely to be repeated over longer periods of time, then the new approach that we will present would enable a *Solution/Offer Curve Family (OCF) Database* to be implemented. This would involve storing all offers produced in a long-term database, so that they could be recalled in future periods when needed.

6.4 Two-Phase Algorithm Concept

It is clear from investigating the computational complexity described in Section 6.3 that most of the computational complexity of the R&A algorithm is in the lower level DP (the two main terms, trp^2d^3 and $tpdr^2$, of Equation 6.1 come from the lower level DP). As such, one of the key features of the two-phase algorithms that we describe in this chapter is that this lower level DP is effectively brought out of real-time and done in advance for each UMS that could be faced throughout the horizon. However, it is not

[35] Note that with the R&A algorithms and the Two-Phase approaches described in this chapter, it would be relatively difficult to update these probabilities. For example, if you have 100 possible RD curves in each period, then the Markov Chain for each period has 10,000 elements, of which you would have to select the correct ones to update. This issue will be dealt with in Chapter 7.

until real-time that we will know the particular form of the marginal value of storage curve (or, effectively, the marginal opportunity cost, MOC, curve for this state) that we will face. This might make it seem like we would need to solve an infinite number of such DPs for each market scenario in order to cover all possible MOC outcomes in real-time. However, recall from Chapter 5 the concept of marginal cost patching:

Let us say that we have optimal offers under a given market scenario for fixed marginal costs of MC1 and MC2 (>MC1). If the true marginal cost jumps from MC1 to MC2 at a quantity of q1, then the optimal offer for this new MC curve will be that found by patching together the appropriate segments of the fixed marginal cost offers.

This concept implies that for each UMS, we can solve the lower level DP for a range of fixed MC levels, and then in real-time patch segments of these together as appropriate. This implies that the more fixed MC levels that are solved for in advance, the greater the accuracy that is able to be achieved with optimal offers in real-time. Of course though, there is a trade-off between the benefits of this accuracy and the solving time required for these advance calculations.

Therefore, we can reduce this algorithm to two separate phases; a Pre-Processing phase (PP phase) and a Real-Time phase (RT phase).

Phase 1: Pre-Processing Phase

The PP phase occurs before the planning horizon begins, as shown in Figure 6.2.

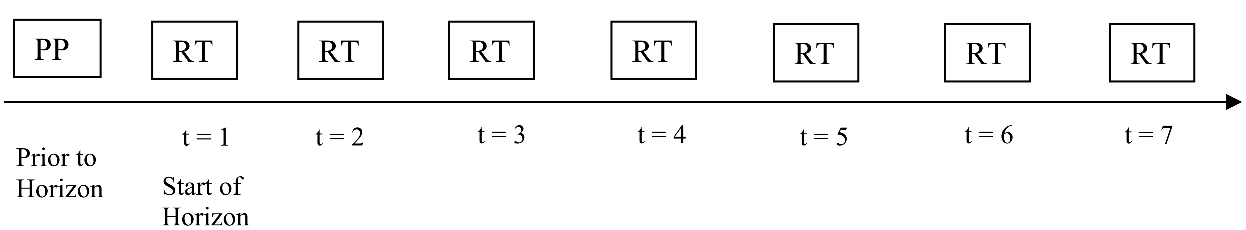

Figure 6.2 Overall Two-Phase Approach Solving Process

For each period of the horizon there is a set of p possible RD curves that may occur, whose probabilities are given by a Markov chain as described in Section 6.2. Hence, if there are p possible RDs in period t, then there will be p possible UMSs in period $t+1$. The PP phase produces an OCF for every one of these possible UMSs throughout the horizon (i.e. pt of them), where we define an OCF for a given UMS as the set (or family) of optimal offers for all MC levels that we consider. The construction of these optimal offers is described in Section 6.5, while an example OCF is shown in Figure 6.3. We can see that, as you would expect, when the fixed MC level increases, the pre-processed offer becomes more restrictive.

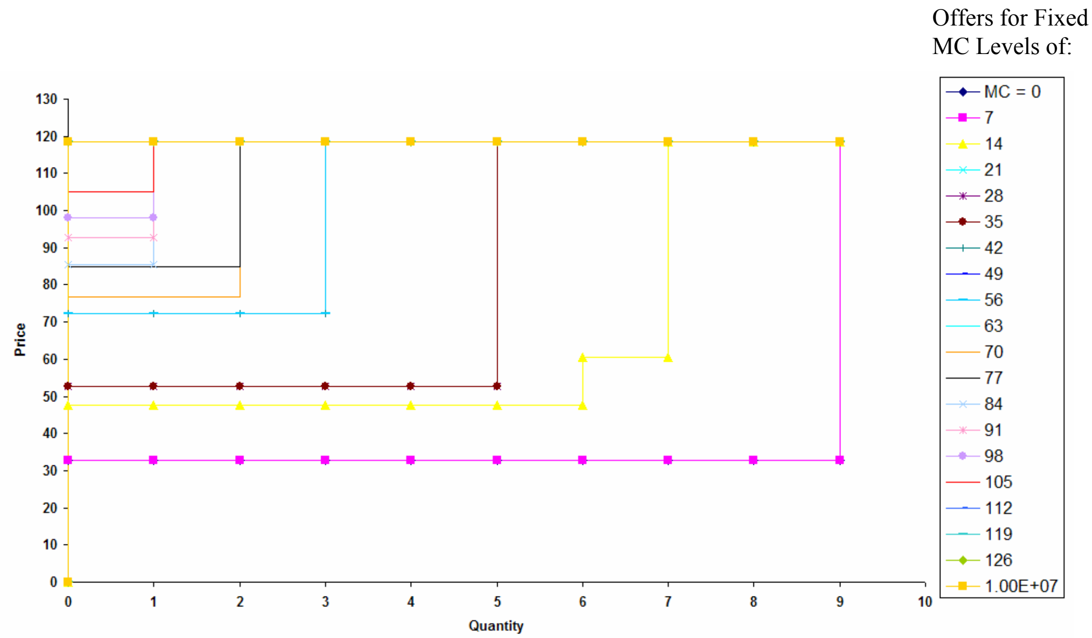

Figure 6.3 Example OCF

Note that although the upper level DP is two dimensional, we only need to produce an OCF for each UMS over the horizon, as opposed to producing one for each combination of UMS and reservoir level. The reason for this is that under a given UMS, as the reservoir level increases, all that changes is that the MOC curve slides horizontally to the right. Therefore, the same OCF is still valid (and further still, the same members of the OCF are still used, but just different portions of them). This idea is shown in Figure

6.4, where A and B are two possible beginning reservoir levels, $Q_{\max}$ is the maximum dispatch level of the generator, MV is the marginal value, and MOC(A) and MOC(B) are the MOC curves associated with the beginning reservoir levels A and B respectively.

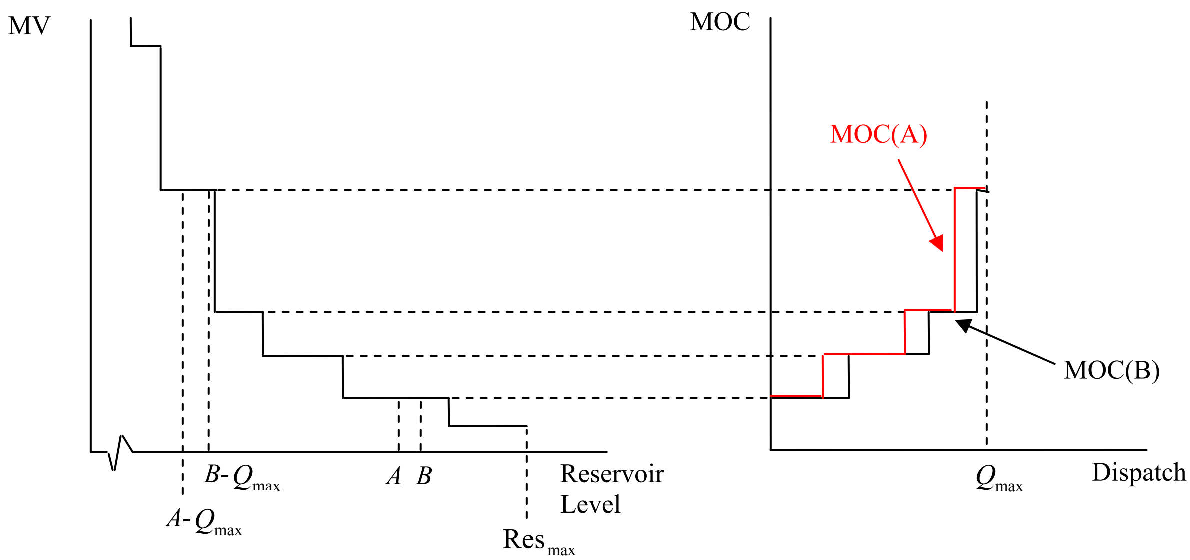

Figure 6.4 Shift in MOC Curve as Reservoir Level Increases

Phase 2: Real-Time Phase

The RT phase occurs at the beginning of each period in the horizon, as shown in Figure 6.2, and needs to be designed with near-linear computational complexity.

This phase takes any new information about our expectations for the remaining horizon, and works backwards through the remaining periods, constructing fuel value curves, expected marginal opportunity cost of generation curves (EMV curves), and then offers for each possible unit state at each period (where the two-dimensional unit state is defined by the reservoir level and the previously observed RD curve). These offers are constructed by patching together the appropriate sections of the pre-processed offers (as defined by the marginal opportunity cost of generating each additional unit), and thus can be designed with a much lower level of computational complexity than the R&A algorithm. Figure 6.5 shows an example set of offers over all reservoir level states, for a given previous RD curve, as produced by the RT phase, and is the equivalent of the output produced by the R&A algorithm. We can see that as the reservoir level increases, the offers become more generous. The reason for this is that the costs of generating fall as the reservoir level increase – this is explained in greater depth in Section 6.6.

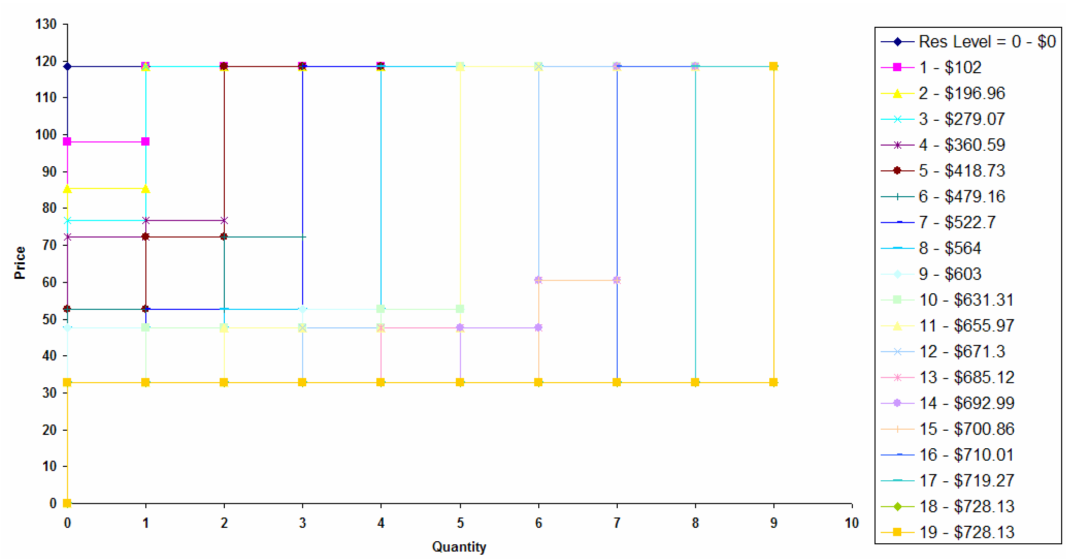

Figure 6.5 An Example Set of Offers

6.4.1 Dynamic Program Approaches

The R&A algorithm is a primal DP approach to offer construction. It directly considers the value of current and all possible future generation when making the offering decision at each state, and therefore has no need for marginal value or cost curves. Two algorithms are developed in this chapter. They are

1. Value Curve Approach (VC approach).
2. Direct Marginal Value Curve Approach (DMVC approach)

Both of these algorithms apply a DP approach that is different to that of the R&A algorithm. They are marginalistic DP methods, where at each stage and under each state we balance expected potential marginal returns within the period with the expected marginal opportunity costs of dispatch. In other words, this is equivalent to the condition used in many previous models (for example, Read, George, & Macgregor (1994)) for quantity-based dispatch, that:

(Expected) MV of generation = (Expected) MV of storage

The difference between these older models and our new model is that we are considering the effects of market uncertainty within a period, and hence how best to deal with this uncertainty (in addition to the correlated UMS structure). In these earlier models, a technique known as Dual DP (see Read, George, & Macgregor (1994) or Yang (1995)) was implemented which directly constructed MV curves working back through the planning horizon. This involved finding guidelines, or breakpoints in the storage space at which point it would prove optimal to shift from one particular strategy to another. As explained in both these references, dual DP was found to be a very efficient algorithm, where much of the benefit was gained through the fact that there were large sections of the reservoir level state space in which the optimal operating decision did not change. The efficiency of the approach fell back towards that of the primal DP as the decision changed more frequently. Traditionally, these decisions were either to turn an additional unit on or possibly related to a particular additional release quantity. However, in our model, the decisions to be made are the specific offer to supply to the market, and as this provides another whole dimension of alternatives, you would therefore expect the decision to change much more frequently as the storage level changes (in fact it is likely that it will change continuously). Thus, the potential benefits of a dual approach are negated[36]. Theoretically, we could create a set of possible offers to be used, and break up the storage space appropriately. However, this would either be too restrictive, or quite inefficient, depending on the number of alternatives made available. Our second algorithm, the DMVC approach, employs some of the facets of the Dual DP technique, in bypassing the construction of a value curve and thus directly constructing MV curves

[36] Note that the dual DP is considered to be more accurate too, as it does not require discretisation of the reservoir level state space. However, as we will show, the approach proposed is efficient enough that a very fine discretisation can be applied to this dimension, and this dual DP benefit is also negated.

for a given period from those for the following period, based on probabilistic expectations of the state in which the generator will end up.

In terms of the dimensions that are being discretised, the R&A primal approach discretised only the storage and dispatch levels, a dual approach would discretise possible offers to be made and marginal fuel values, while the marginalistic approaches presented in this chapter discretise the storage and dispatch levels (in the same way as the primal approach) and the marginal fuel values (through the discretisation of the MCs in the pre-processing phase).

We have called these new approaches *Structured Marginalistic Dynamic Programming* methods, as they consider the particular structures apparent in this problem and apply the knowledge gained about how offers should change in various circumstances, developed through the analytic work described in earlier chapters. For example, we know that if the fixed MC level is increasing, then the offer must shift up and to the left[37], and thus there is no point considering an offer that would pass through the rest of the offering space.

6.5 Pre-Processing Phase

In this section, we present a flow diagram representing the PP phase algorithm, discuss the options that were chosen with respect to the approach of this phase, and identify the computational complexity of the algorithm.

6.5.1 The Pre-Processing Algorithm

Figure 6.6 presents the PP phase at a very high-level. The bulk of the computational complexity of the algorithm is found within the component labelled “Find Offer for MC

[37] As implied by Section 4.3, for simple types of market uncertainty, and in the remainder of Chapter 4 for more complex uncertainty.

level under this UMS". The details of this component are explained in Sections 6.5.3 and 6.5.4.

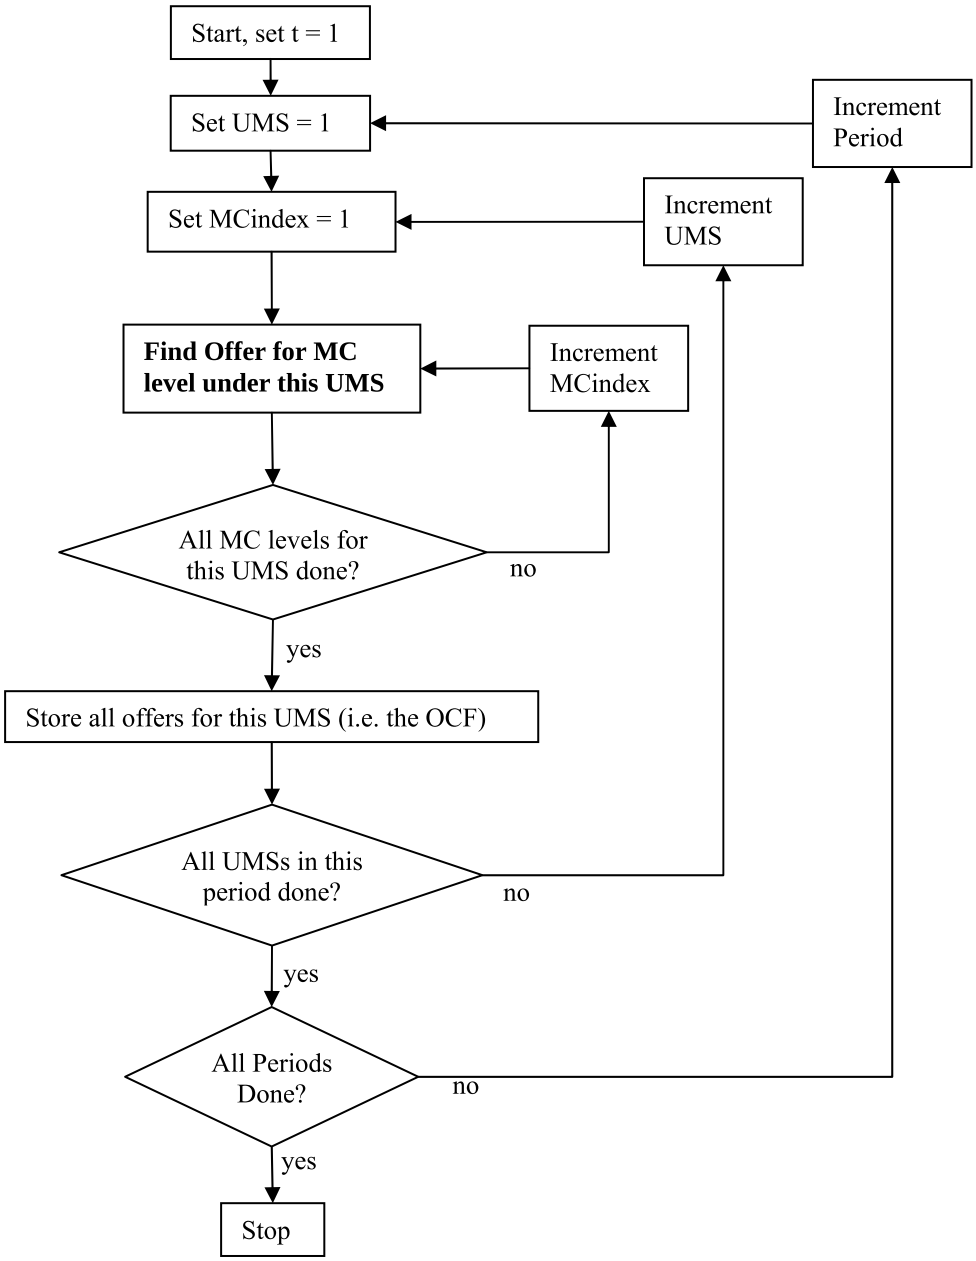

Figure 6.6 PP Phase Flow Diagram

6.5.2 Options for Finding Offers

There are two possible ways of going about the construction of the offers in the pre-processing phase:

1. Using DP, in a similar, but more structured and hence more computationally efficient manner to that used in the R&A algorithm, such as discussed in Section 3.4.1.
2. Using the analytic concepts for optimal offer construction developed throughout Chapter 4 of this thesis.

Both of these approaches have their merits, and would be appropriate under certain circumstances, where the most appropriate approach is determined purely by the method that would produce the results in the fastest time. In particular, it is fairly easy to show that under simple UMSs, the latter method would be the most efficient, while under complex UMSs, the DP approach would be appropriate.

Example 6.1 Simple UMS

Consider an example as shown in Figure 6.7, whereby the rest-of-market supply curve is known, but there is demand uncertainty. Therefore, we have an expected residual demand curve that is subject to horizontal uncertainty only. From Section 4.3.4, we know some parts of the pseudo-offer will definitely form part of the optimal offer, and these are very easy to define. We also know that there is a region of uncertainty that will be optimally overcome with a vertical step from the lower segment of the pseudo-offer to the upper, and we can use our optimality conditions to find the position of that step, as shown in Figure 6.7.

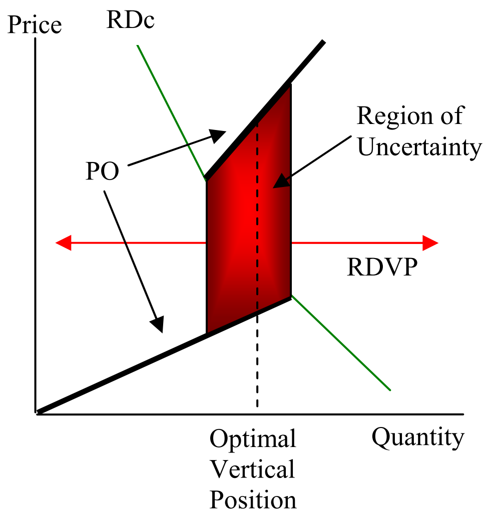

Figure 6.7 Simple UMS with MC=0

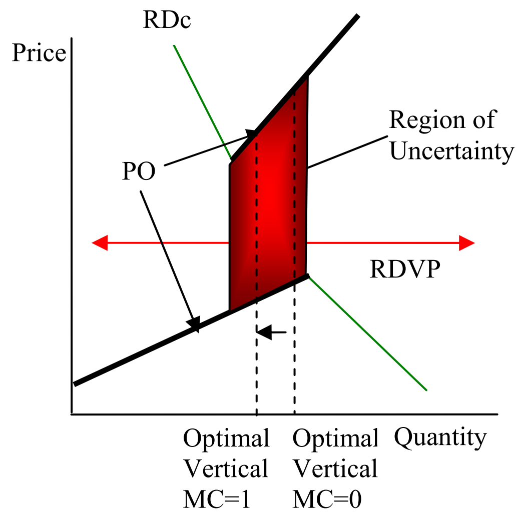

Figure 6.8 Simple UMS with MC=1

When we then incrementally increase our fixed marginal level, we know that there are again portions of the easy-to-find pseudo-offer that will form parts of the optimal offer, but we also know that the position of the vertical step must be to the left of where it previously occurred. This therefore gives us a very good starting point when searching for the new vertical position. This is shown in Figure 6.8. For this example, producing the OCF through this process would be much faster than forming and solving the low-level DP for each MC level.

Example 6.2 Complex UMS

Consider another example, this time where there is more complex uncertainty in the residual situation faced by the competitor. Here we consider that there are four possible linear residual demand curves that could be faced, forming a UMS as shown in Figure 6.9. Note that the stars indicate positions of local optimality along the contours of the MDF that is implied by the UMS.

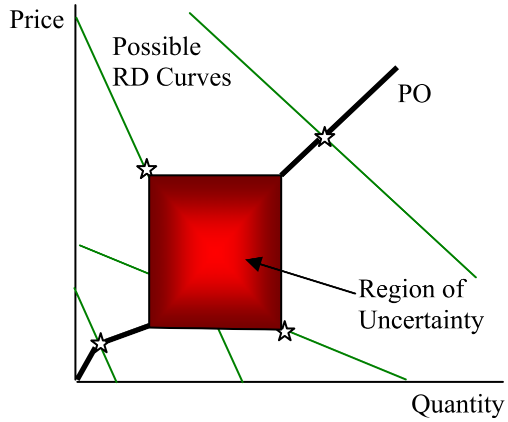

Figure 6.9 Complex UMS

We have shown in Section 4.3.1 that if monotonicity is not an issue, then the optimal offer would pass through the optimal position on each contour of the MDF[38]. If monotonicity is an issue, then the optimal offer will consist of vertical and horizontal segments through the region of uncertainty. In this case, we can see that monotonicity is an issue, and hence some sort of algorithm, or possibly heuristic would be needed to determine the optimal path through this region of uncertainty. As in the simple case, when we increase MC incrementally, we again know that the offer must shift up and to the left, but the variety of alternative movements make the process of finding the new optimal more complex than the first example. Although this is a very simple example of the more complicated type of case, we can see that when the region (or regions) of uncertainty become large, covering a greater portion of the offering space[39], the complexity of an analytic approach to pre-processing would increase greatly, and thus the DP approach would become more desirable.

[38] Note that when we are considering only discrete possible residual demand curves, the only points on the offer that are actually important are those where the offer intersects with the residual demand curves. The directions of improvement along linear contours of the MDF are still driven by the positions of the ALs, but the sloped sections of these attractor lines no longer need be contained within the solution.
[39] Because we are considering downward sloping residual demand curves, the non-zero expected vertex payoffs would occur (generally speaking) down the diagonal from (low quantity, high price) to (high quantity, low price), and hence would most likely lend themselves to creating larger regions of uncertainty than would likely result if the non-zero expected payoffs were more evenly spread out.

6.5.3 Basic Vertex-Traversing Algorithm

In addition to the benefits of removing the low-level DP calculations from the real-time calculations into a separate pre-processing phase, there are many other ways in which the general efficiency of these lower-level DP calculations can be improved on that presented in Rajaraman & Alvarado (2003). This section describes the basic DP vertex-traversing algorithm used, while Section 6.5.4 discusses some of the algorithmic improvements that have been applied, by recognising the structure of the particular problem.

The scenario for the basic vertex-traversing algorithm is defined by a (Q*P) grid with an expected payoff, $EP(q,p)$, defined on each of the QP vertices, where the objective is to find the path with the greatest total expected payoff from $(0,0)$ to $(Q_{\max},P_{\max})$ that is monotonically non-decreasing in both the quantity and price dimensions. This is done by first performing a backwards recursion, finding the optimal "Value-To-Go" (*VTG*) from each vertex to $(Q_{\max},P_{\max})$. This is a very straightforward calculation:

$$VTG(q,p) = EP(q,p) + \max(VTG(q+1,p), VTG(q,p+1))$$

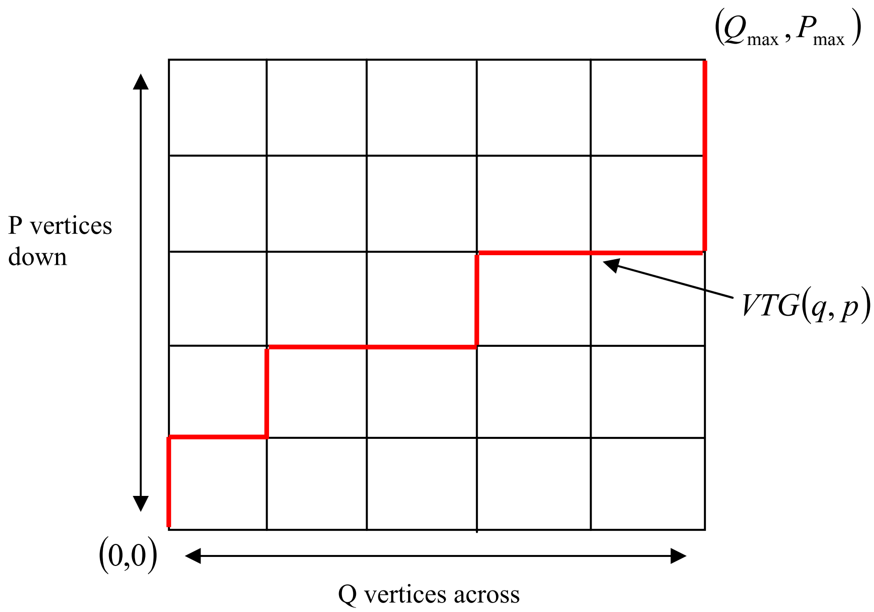

Figure 6.10 The Basic Vertex Traversing Algorithm

The offer is then traced forward from $(0,0)$, using the rule:

If $VTG(q, p) = EP(q, p) + VTG(q, p+1)$ then go up (increment p).

If $VTG(q, p) = EP(q, p) + VTG(q+1, p)$ then go right (increment q)

Note that in the R&A algorithm, the calculation of the Value-To-Go is different from this, and unnecessarily complicated. Rather than just restricting the path of the offer to edges of the grid, the R&A algorithm allows angled offer segments, by determining its Value-To-Go with the equation:

$$VTG(q, p) = EP(q, p) + \max(VTG(q, p+1), VTG(q+1, p+1), \ldots, VTG(Q_{\max}, p+1))$$

This method covers a very similar range of potential solutions (the differences are discussed in Section 6.7.1), but is significantly less efficient than the standard approach described above. If we were to apply the approach above to the R&A algorithm, then the order of computational complexity would be reduced from $trp^2d^3 + tpdr^2$ to $trp^2d^2 + tpdr^2$.

6.5.4 Problem-Specific Vertex-Traversing Algorithm

In this particular offer construction problem, there is a specific structure that lends itself to improved computational methods over those required for the basic vertex-traversing algorithm. The expected payoffs on the vertices are non-zero only on the vertices that residual demand curves pass through (as indicated by stars in Figure 7.6).

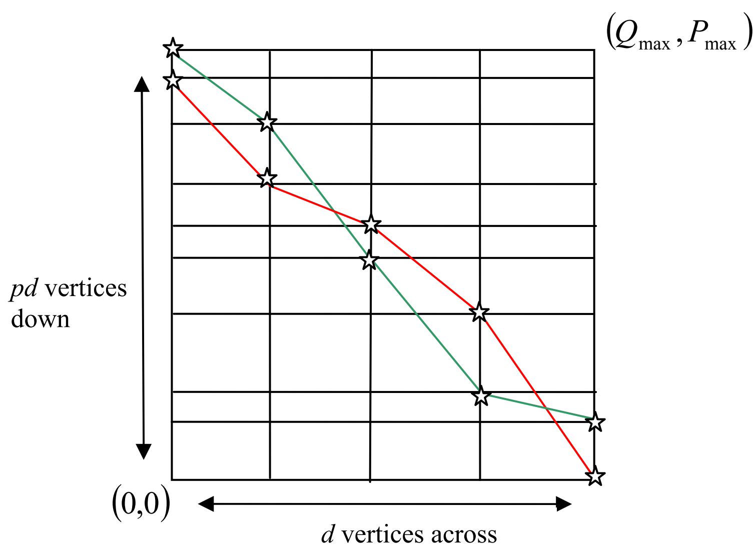

Figure 6.11 The Problem-Specific Vertex Traversing Algorithm

Therefore, if we have *p* residual demand curves and *d* possible dispatch levels, then we will have a grid of vertices *d* wide by *pd* tall, and hence a total of pd^2 vertices for which to potentially calculate the expected payoffs and values-to-go. In our model, we need to repeat this Vertex-Traversing Algorithm *M* times with the same set of residual demand curves, where *M* is the number of different MOC levels that we are considering.

The remainder of this section describes improvements over the basic vertex-traversing algorithm that we have implemented, in addition to improvements over its implementation in the R&A algorithm. Most of these improvements relate to the calculation of the VTG matrix, as this is the section of the algorithm to which most time is devoted in the algorithm. The following is a list of six different, but complimentary, improvements.

1. We know that there is only a single non-zero payoff in each row of the grid, as each of these rows corresponds to one of the prices on one of the RD curves. Hence the calculation of the expected payoffs can be limited to only these vertices, immediately reducing the calculation effort for expected payoffs to (1/d) of its original size. Note the calculation required is:

$$EP(q,p) = (pq - TotalCost(q)) * \Pr(associated\ RD\ occuring)$$

2. Rather than repeatedly calculating the revenue portion (*pq*) of the expected vertex payoff for each of the *M* different MOC levels, we can calculate it once outside this loop of the algorithm.

3. Rather than having a rule based forward recursion to trace out the optimal offer, the optimal direction of movement from each vertex is recorded in the calculation of the Value-To-Go stage. This substantially reduces the computational complexity of the offer finding process.

4. From the analytic work in the earlier chapters of this thesis, and general logic, we know that as the fixed MOC level rises, the optimal offer cannot move below (down and to the right of) the offer for the previous MOC. Therefore, as we successively move up the fixed MOC levels, we keep track of the lower bound price that should still be considered at each dispatch level. We do not calculate any expected payoffs that occur below this lower bound.

5. This previous modification leaves some rows with no non-zero expected payoffs. Therefore, we keep track of these rows and jump over them entirely when calculating the Values-To-Go and when tracing out the optimal offer.

6. An improvement has been made to the Value-To-Go calculations by recognising that the matrix of expected payoffs is very sparse. In particular, that there is only a single non-zero payoff in each price row of the matrix. Using this fact, we can simply copy many of the Value-To-Gos from either the right or above, rather than recalculating them. The process is:

   a. Work back from the last quantity level that is not below our *lower bound*, copy the VTGs from the row above until we reach the non-zero expected payoff.
   b. Do the normal VTG calculation for this vertex of the matrix.
   c. Keep working backwards, copying the VTG to the right while it is greater than the VTG above.
   d. Once, the VTG above is greater than the VTG to the right, simply copy all VTGs from above, back to the lowest quantity.

In Section 6.7, we will discuss a further addition to the algorithm, whereby payoffs on the edge of the grid are considered, in addition to payoffs on the vertices. When this addition is included, the first six improvements described here are unaffected, but the benefits of improvement six is lost to a large extent because the matrix of edges with positive payoffs is no-where near as sparse as the matrix of vertices with positive payoffs.

6.5.5 Computational Complexity of Pre-Processing Phase

We define the following additional parameter for use in expressing the computational complexity of the PP phase:

M = Number of MC levels at which to find offers in the OCF

From the source code for the PP phase, we can show that the computational complexity can be approximately summarised by Equation 6.2.

$$7tp + 12tdp + 0.667td^2 p^2 + 10tp^2 Md + 6tpMCd + 8tpM + 0.875tp^2 Md^2$$

Equation 6.2 Computational Complexity of the PP Phase in Full

Simplifying this down to the terms of the largest orders of magnitude, we can see that the algorithm is of order of complexity $tp^2 Md^2$.

Note that a decision must be made in the PP phase, as to the level of discretisation in MC levels, or in other words, the number of different MC levels for which to produce a pre-processed offer. Observe that the PP phase complexity is approximately linear in this parameter. In a preliminary trial, we have observed that anywhere between 50% and 85% of pre-processed fixed MC offers have been used in the RT phase. Of course, as the level of discretisation increases (and hence we have more MC based offers produced), the percentage decreases. In addition, if a pre-processed offer is used at all, then generally, most of it will be relevant (as opposed to just a small segment of it). This is because as we slide down the reservoir storage levels, that particular MC becomes appropriate at different segments of the dispatch range. This issue is discussed further in Section 6.6.5, with respect to the effect on the expected payoff of this decision.

6.6 Real-Time Phase

In this section, we present a flow diagram representing both RT phase algorithms, discuss the options that were chosen with respect to the approach of this phase, prove the consistency of this algorithm in constructing concave value curves (decreasing MV curves) and identify the computational complexity of the algorithm.

6.6.1 The Real-Time Algorithms

Figure 6.12 presents the RT phase algorithms at a very high-level. The bulk of the computational complexity of the RT phase of the Value Curve approach is found within the component labelled "Value Curve Approach: Find Offer Value and add it to the Value Curve for this Market and Reservoir State", while the bulk of the computational complexity for the Direct MV Curve approach is located in the corresponding component labelled "Direct MV Curve Approach: Find Expected MV and add it to the MV Curve for this Market and Reservoir State".

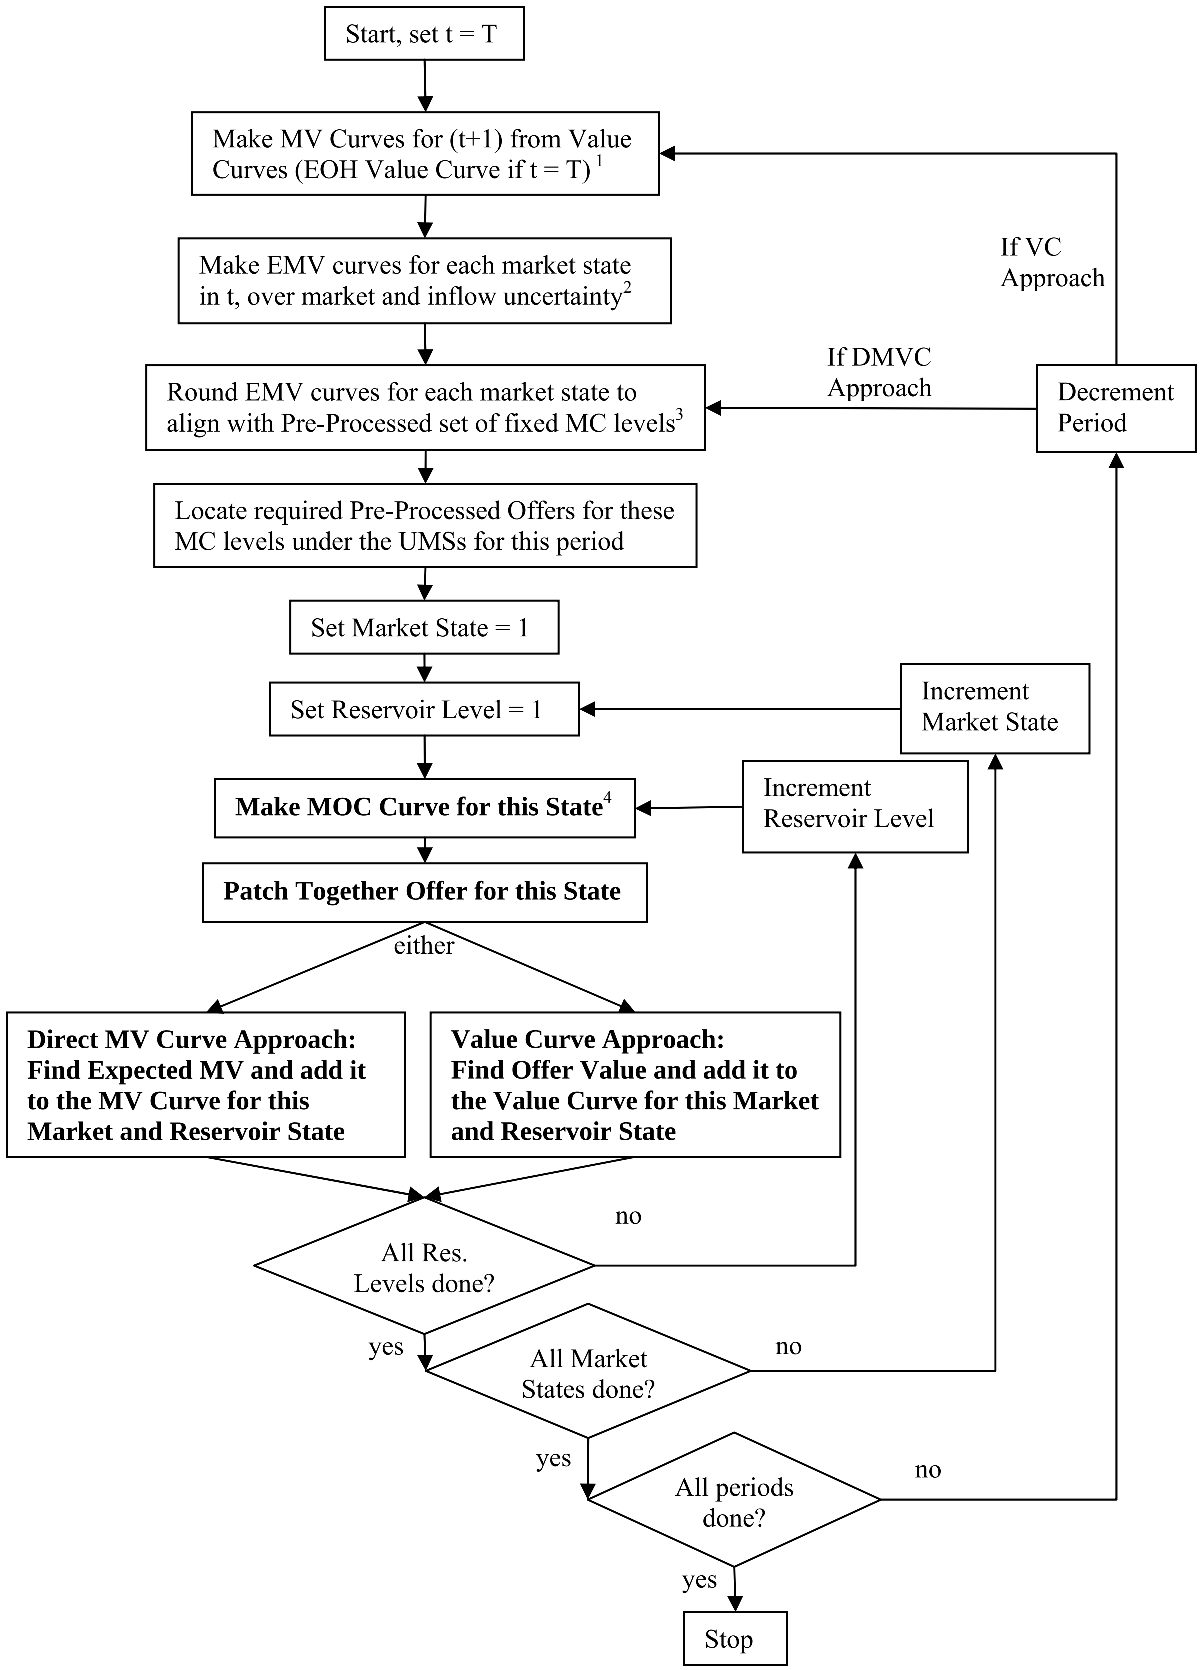

Figure 6.12 RT Phase Flow Diagram (both approaches)

The remainder of this section is dedicated to explaining certain components of this algorithm in greater detail. Each of the numbered notes corresponds to a component labelled with a superscript in Figure 6.12.

1. Make MV Curves

In order to produce monotonically non-decreasing offers, our algorithm needs to produce monotonically non-decreasing MOC curves. In order to achieve this, the Value Curve needs to be concave (i.e. the MV curve must be decreasing). Logically speaking, we would expect this to be the case: as the reservoir level gets higher and higher, the MV of increased storage should fall (this is proven in Section 6.6.4). However, the discretisation of the fixed MC levels and the discretisation of the grid on which the offers must traverse the edges can result in some minor inconsistencies in this respect. Additionally, when we limit our model to quantities and prices that must be positive, we have an additional effect that can produce a convex kink in the Value Curve. For low reservoir levels it may be that the MOC curve is too high to warrant an offer that would lead to any dispatch (and hence the expected payoff will be zero). As the reservoir level increases, there may then come a point where it becomes optimal to provide an offer that could lead to a positive dispatch (and hence would have a non-zero expected payoff). Figure 6.13 shows that this would result in a clear convex kink in the Value Curve.

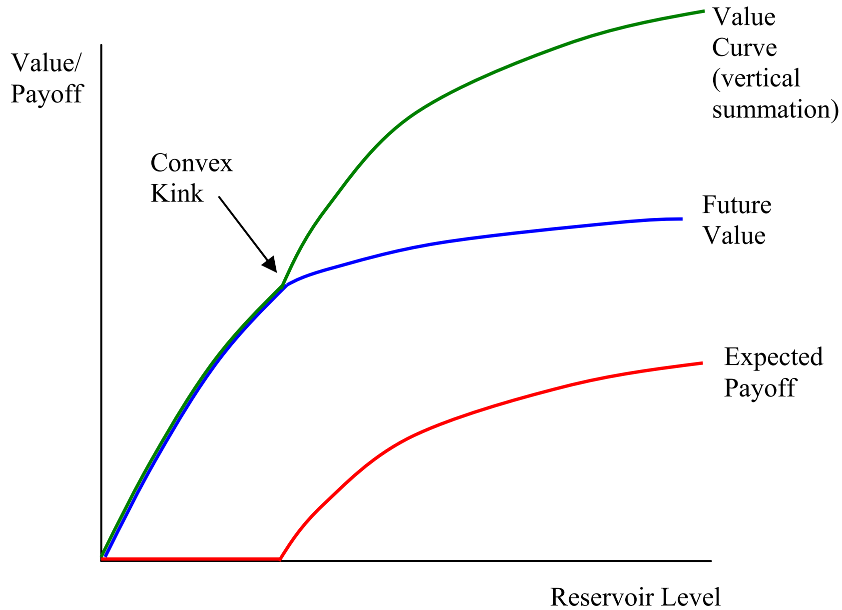

Figure 6.13 Convex Kink in the Value Curve

In order to overcome these non-concavities, the algorithm:

a. Produces the MV curve from the Value curves by taking the derivative
b. Smoothes the MV curve so that it is non-increasing

The smoothing procedure is demonstrated by the red line in Figure 6.14, and involves an adjustment to the marginal value curve over the range in which the MV is not monotonically non-increasing.

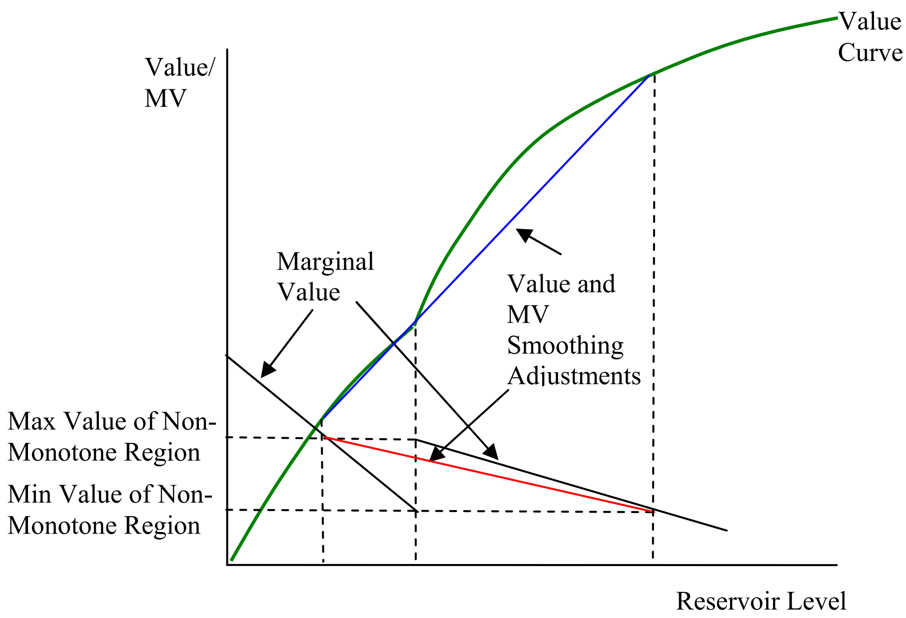

Figure 6.14 Convex Kink in the Value Curve

Note that the adjustments are made to the MV curve rather than the value curve, as this was found to be more computationally efficient.

2. Make EMV Curves

These are vertically weighted averages of the Marginal Value curves for the following period, and they provide the Marginal Opportunity Cost (MOC) curves for generation in this period.

3. Rounding EMV Curves

EMV curves are rounded up to the nearest fixed MOC level, as this provides more conservative result and hence leads to a more conservative use of fuel.

Note that offers have been pre-processed for a range of fixed MOC levels and parts of these pre-processed offers are patched together in real-time to determine the optimal offer under any specific Rounded MOC curve. This raises the question of whether we should interpolate between the pre-processed offers, rather than round the MOC curve. Due to the discrete nature of the problem, we find that in the pre-processing stage, the optimal offer will stay the same for a variety of MC levels, and then will suddenly jump up to a new position (by-passing various non-optimal alternatives in-between). Therefore, to interpolate between the offers would often not actually give a different result to just rounding the MOC, and on other occasions it could give a worse result than would be achieved by rounding the MOC either way. In addition, to interpolate between these offers would greatly increase the necessary calculations performed in the RT phase. Therefore, the best option is to simply pre-process offers for more fixed MC levels, particularly around the ranges where MOC is likely to sit, especially since the computational time for the RT phase is unaffected by the number of fixed MOC levels that have been pre-processed.

4. Make MOC Curve

This MOC curve is constructed either from the rounded EMV curve for the period, or by shuffling along the MC curve from the previous reservoir level. Note that this means the EMV curve only needs to be produced once for each market state, as the MOC curves for all reservoir levels can be derived from it.

6.6.2 Constructing the MV Curves under the Value Curve and Direct MV Curve Approaches

This section considers the core component of the two two-phase approaches, the construction of the MV curves as the algorithms work backwards through the horizon, on which offering decisions are based. This is the component that differentiates the two approaches. These differences are indicated in Figure 6.12. Note that the explanations given could be modified easily to account for inflow uncertainty, but this is ignored here for clarity.

To explain these two approaches, we define the following notation:

$V_{t,r}$	Value curve at the end of period *t*, given that residual demand curve index *r* occurred in the *previous* period (*t*-1)
$V_{t,r}(R_t)$	Value of storage at the end of period *t*, at reservoir level R_t, and given previous RD curve index *r*
$MV_{t,r}$	MV curve at the end of period *t*, given that residual demand curve index *r* occurred within this period
$MV_{t,r}(R_t)$	MV of storage at the end of period *t*, at reservoir level R_t, and given that residual demand curve index *r* occurred within this period
$EMV_{t,r}$	Expected MV curve at the end of period *t*, given that residual demand curve index *r* occurred in the *previous* period (*t*-1)
R_t	Reservoir level at beginning of period *t*
$\theta_t^r(R_t)$	The optimal offer provided to the market in period *t* from reservoir level R_t, given previous residual demand curve index outcome *r*
$Q_t^p(\theta_t^r(R_t))$	Release in period *t*, given residual demand curve outcome index *p* in this period, and offer $\theta_t^r(R_t)$ (which is based on residual demand curve index *r* having occurred in the *previous* period)
$P_t^p(\theta_t^r(R_t))$	Price in period *t*, given residual demand curve outcome index *p* in this period, and offer $\theta_t^r(R_t)$
P	Total number of possible market outcomes per period
$pr_{r,p}$	Probability of residual demand curve outcome index p_t, given residual demand curve outcome index *r* in the previous period.
$\text{infl}_{t,i}$	Level of inflow in period *t* under possible inflow index *i*

Value Curve Approach

To construct the $MV_{t-1,r}$ curves from the $MV_{t,r}$ curves under the VC approach, the following steps are followed:

i. Construct the $EMV_{t,r}$ curves from the $MV_{t,r}$ curves (one for each market outcome). These are vertically weighted averages, based on the probabilities in the Markov chain between the two periods.

ii. Construct the $V_{t,r}$ curves, one for each previous RD state

 For each reservoir level,

 a. Construct the MOC curve from the appropriate $EMV_{t,r}$ curve
 b. Determine the optimal offer
 c. Find the expected release under all possible market outcomes (the intercept of the offer and each of the RD curves)
 d. Determine the expected current and future value from this point for all possible RD curves (weighted by their probability of occurring under the current UMS).
 - Current value is revenue in the current period from this offer and given the market outcome.
 - Future value is determined from the point that you would end up on the following period's value curve, under the release level for this market outcome, given the optimal offer, and for the inflow level that occurs.

iii. Construct the $MVS_{t-1,r}$ curves by taking the derivative of the $V_{t,r}$ curves.

Steps (i) to (iii) are demonstrated in Figure 6.15, while step (d) is shown in Figure 6.16 (along with a numerical example for a single point).

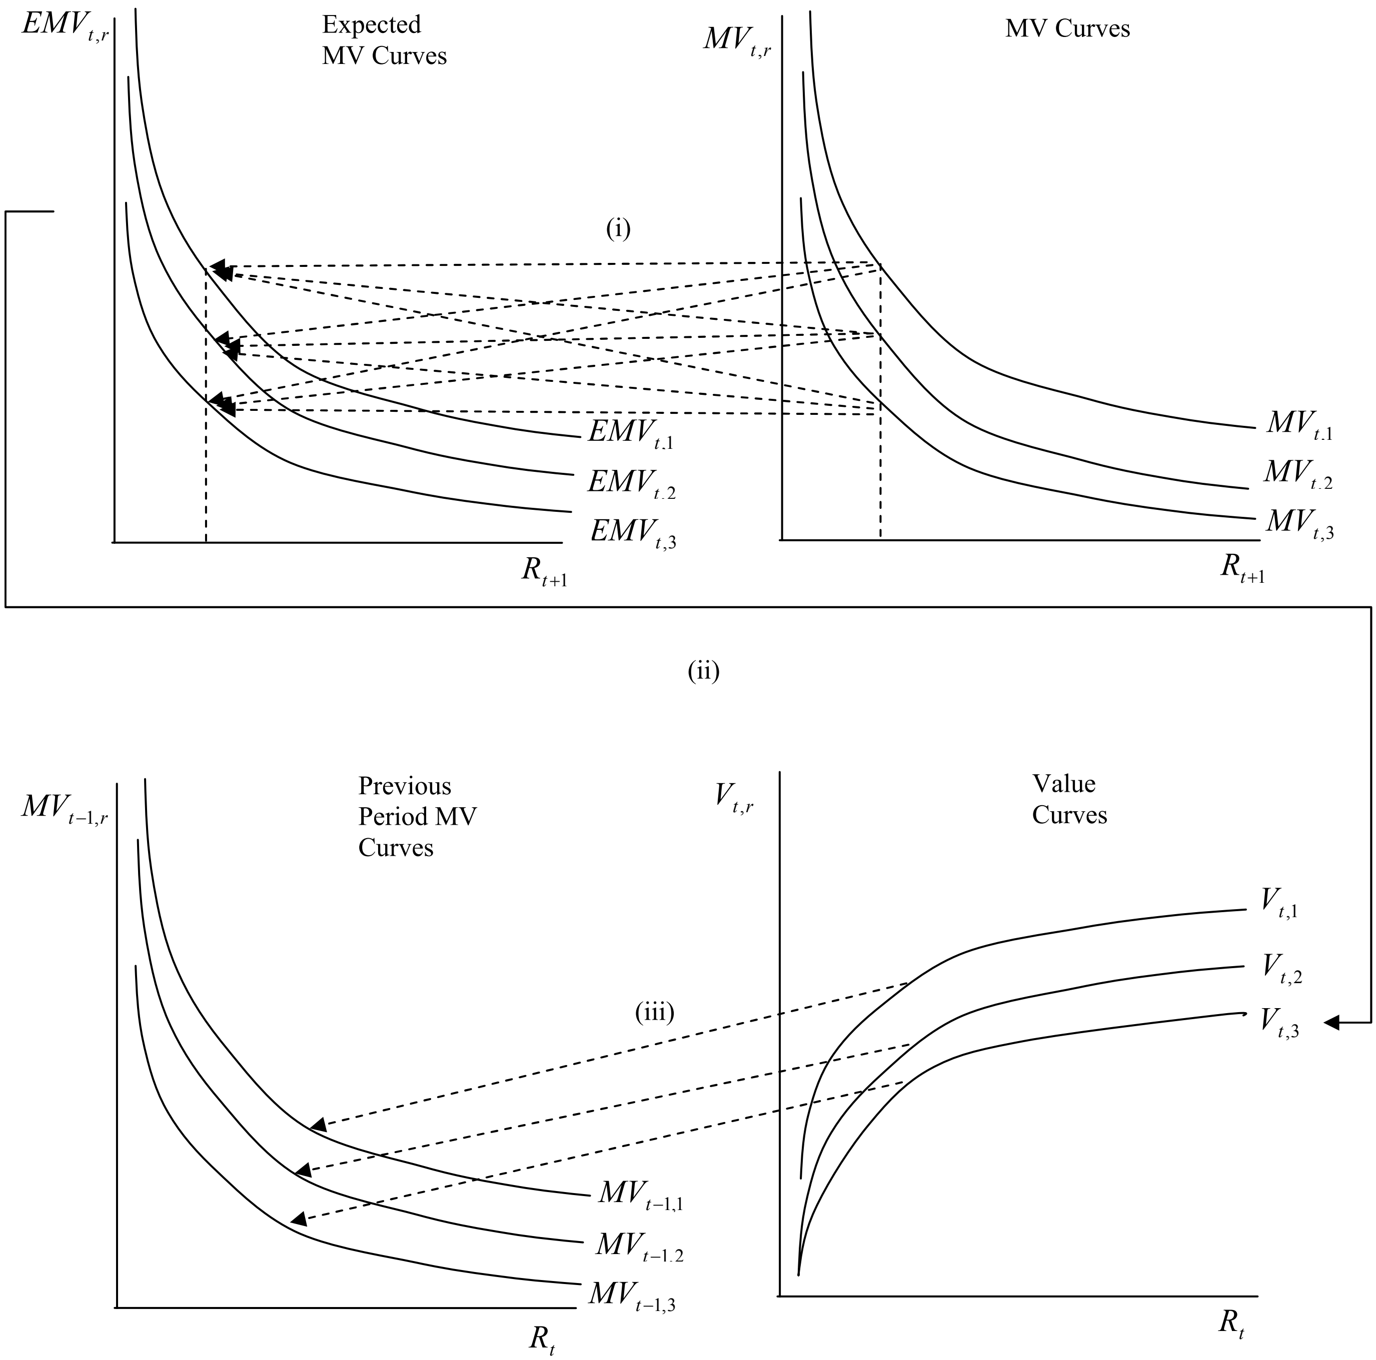

Figure 6.15 Producing the MV Curves under the Value Curve Approach

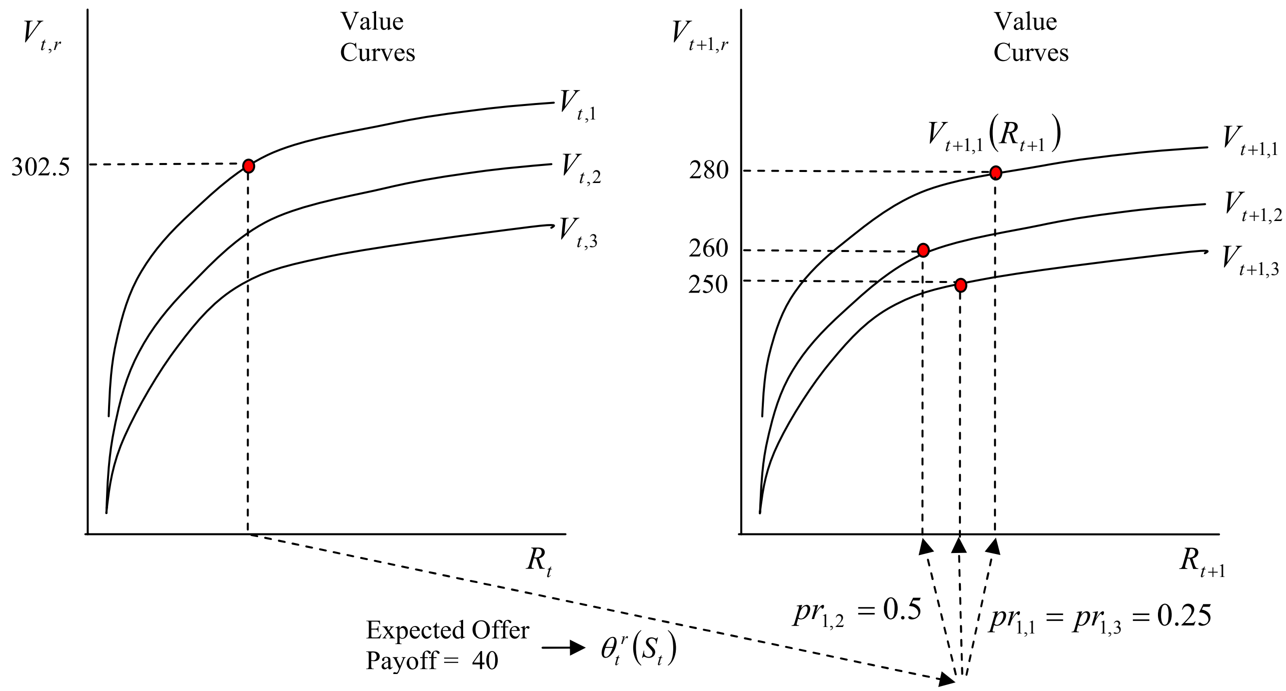

Figure 6.16 Calculating the Value Curve

In step (ii), for a given offer, if market outcome *p* and inflow index *i* occur, the new storage level can be defined mathematically as:

$$R_{t+1} = R_t - Q_t^p\left(\theta_t^r\left(R_t\right)\right) + \text{infl}_{t,i}$$

The value at the current reservoir level R_t under a previous RD curve p_{t-1} can be defined as:

$$V_{t,r}\left(R_t\right) = \sum_{p=1}^{P} pr_{r,p}\left[Q_t^p\left(\theta_t^r\left(S_t\right)\right) * P_t^p\left(\theta_t^r\left(S_t\right)\right) + V_{t+1,r}\left(R_{t+1}\right)\right]$$

Direct Marginal Value Curve Approach

The MV_{t-1} curves are constructed from the MV_t curves in a similar way under the DMVC approach, but there is no payoff considered from within the period, and so the expected MV is just the weighted average of the marginal values at the reservoir levels at which we could possibly end up, based on the selected offer and the current UMS. This is a much simpler process than the VC approach. Appendix B demonstrates the basic theory of these approaches, under the simpler scenario of inflow (rather than market) uncertainty and correlation.

To construct the MV_{t-1} curves from the MV_t curves under the DMVC approach, the following steps are followed:

i. Construct the EMV_t curves from the MV_t curves (one for each market outcome). These are vertically weighted averages, based on the probabilities in the Markov chain between the two periods.

ii. Construct the MV_t curves, one for each previous RD state

    For each reservoir level,

    a. Construct the MOC curve from the appropriate EMV_t curve
    b. Determine the optimal offer
    c. Find the expected release under all possible market outcomes (the intercept of the offer and each of the RD curves)
    d. Determine the expected future MV from this point for all possible RD curves (weighted by their probability of occurring under the current UMS). The future MV is determined from the point that you would end up on the following period's value curve, under the release level for this market outcome, given the optimal offer, and for the inflow level that occurs.

Steps (i) and (ii) are demonstrated in Figure 6.17, while step (d) is shown in Figure 6.18 (along with a numerical example for a single point).

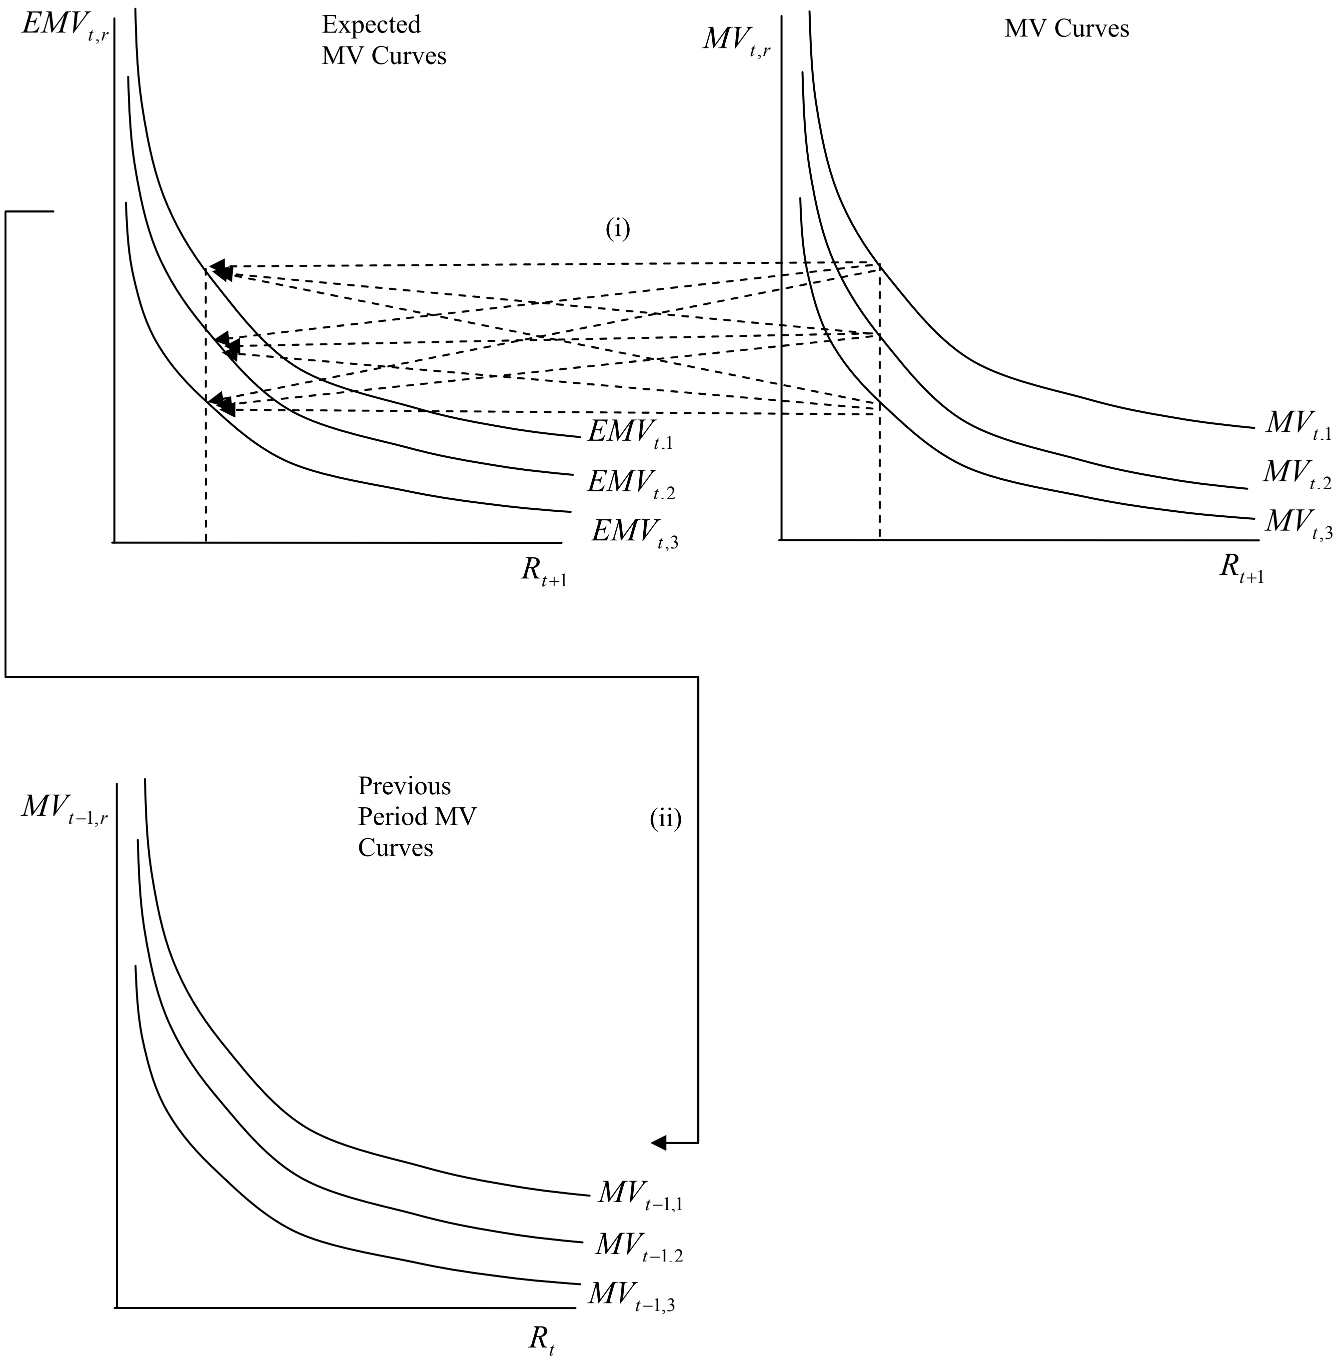

Figure 6.17 Producing the MV Curves under the Direct MV Curve Approach

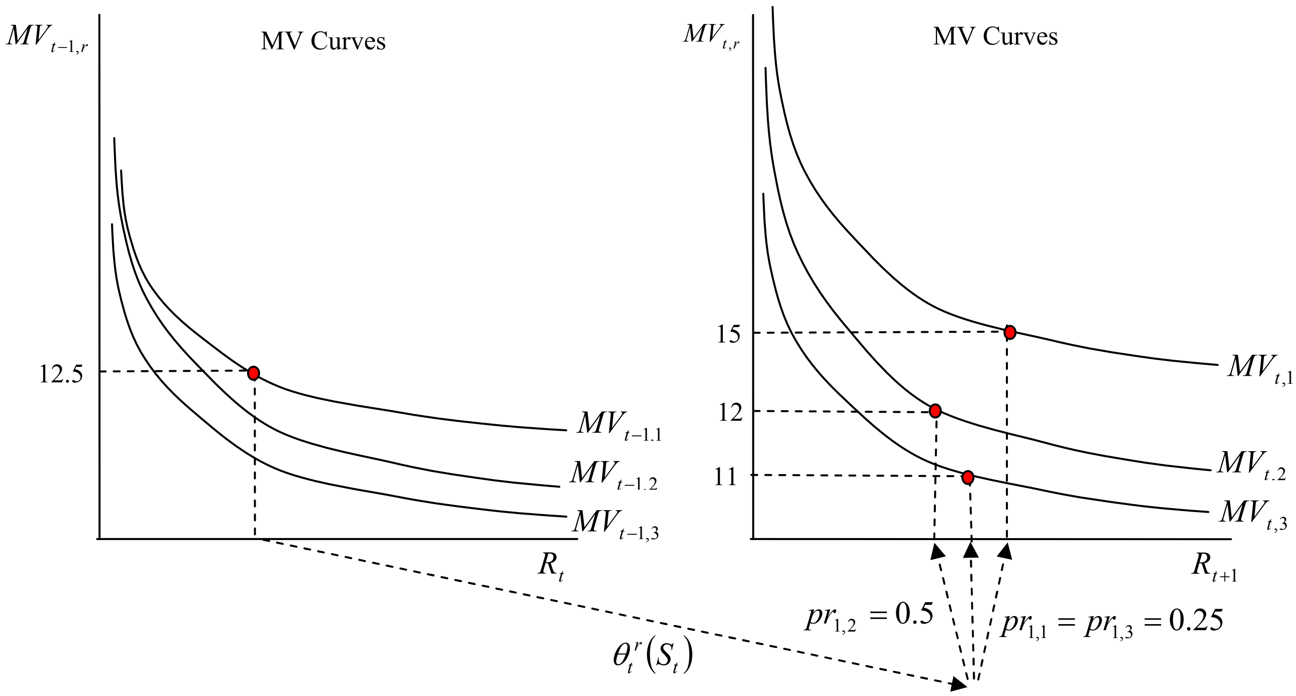

Figure 6.18 Calculating the MV Curve

For a given offer, if market outcome p occurs, the new storage level is defined in exactly the same way as for the VC approach. However, in the recursion, the MV at the current reservoir level R_t under a previous RD curve p_{t-1} can be defined mathematically as:

$$MV_{t-1,r}(R_t) = \sum_{p=1}^{P} pr_{r,p}\left[MV_{t,r}(R_{t+1})\right]$$

This shows that the value curve can be bypassed in the construction of the MV curves as we move backwards through the planning horizon.

Note that for the algorithm to be completely consistent, we would actually consider a separate MOC curve for each RD curve, rather than a weighted aggregate one, because the MOC of release depends on which MV curve we will end up on in the following period (which is dependent on the RD curve that occurs in this period, thus creating a circular problem). However, in Appendix C we show that basing the offers on a MOC curve that is produced from a vertically weighted average MV curve is both necessary (due to computational complexity of the alternatives) and acceptable.

6.6.3 Equivalence of the VC and DMVC Approaches

In this section we will prove that, in the continuous case, the VC and DMVC approaches will always produce the same offers, the same release levels, and the same MV curves and are therefore equivalent approaches.

Proof 6.1

Assume that the value and MV curves for the end of period t correspond to one another (that is, the MV curve is the derivative of the value curve). Under the VC approach, we construct a MV curve by taking the (discretised) derivative of the value curve. Therefore, the MV curves for the two approaches are identical. As explained in Section 6.6.2, the MOC curves on which the offers are based are determined from these MV curves, and as such, the offers (and thus dispatch) for the two approaches in period t will also be identical.

Now let us consider the construction of the MV curves as the respective DPs move back through the horizon. Start by dealing with the simplest case where there is only one possible RD curve in period t. We ignore inflows here, as they provide a fixed shift up in reservoir level for both approaches, which simply leads to a change in the MOC curves on which the offers are based and thus does not affect the results presented.

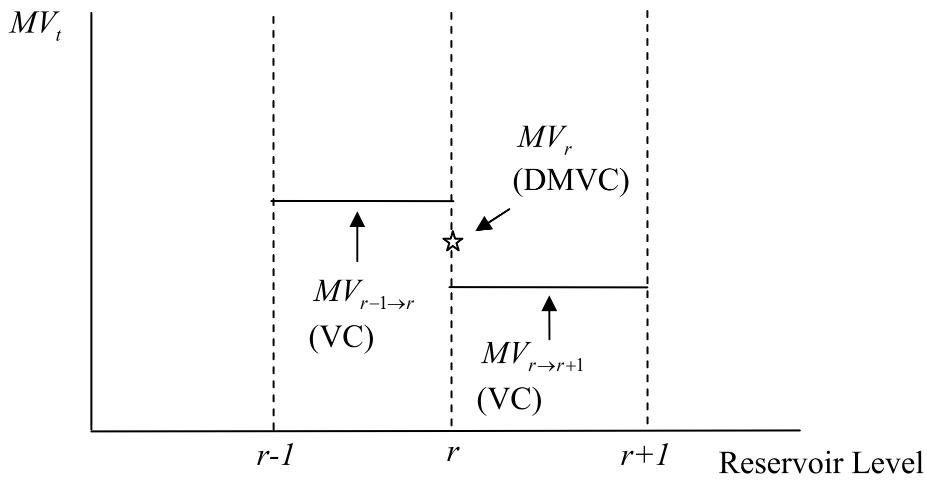

Figure 6.19 Different MV Points under the VC and DMVC Approaches

As illustrated in Figure 6.19, the DMVC approach calculates a MV at each reservoir level breakpoint, and the VC approach calculates a MV between consecutive reservoir level breakpoints (it calculates the *value curve* at the breakpoints). As such, we cannot compare the MV found from the two approaches directly at a single point. We must show that MV_r falls between $MV_{r-1\to r}$ and $MV_{r\to r+1}$, and that the difference between the terms tends towards zero as the width of the reservoir intervals tends towards zero. To show this we define the following terms:

MV_r = MV at reservoir level *r* under the DMVC approach

$MV_{r-1\to r}$ = MV on segment between reservoir levels *r-1* and *r* under the VC approach

V_r = Value curve at reservoir level *r*

q_r = Quantity dispatched as a result of the offer from reservoir level *r*

FV_{r-q_r} = Future Value of End-of-Period (EOP) reservoir level $r-q_r$

PV_{q_r} = Present Value of dispatch of quantity q_r

$MFV_{r\to r+1}$ = Marginal Future Value between reservoir levels *r* and *r+1*

$MPV_{r \to r+1}$ = Marginal Present Value between reservoir levels *r* and *r+1*

G = Gain of MV from calculating the segment below *r*

L = Loss of MV from calculating the segment above *r*

The first requirement defined for this proof is:

$$MV_{r \to r+1} \leq MV_r \leq MV_{r-1 \to r}$$

Or, equivalently:

$$V_{r+1} - V_r \leq MV_r \leq V_r - V_{r-1}$$

$$FV_{r+1-q_{r+1}} - FV_{r-q_r} + PV_{q_{r+1}} - PV_{q_r} \leq MV_r \leq FV_{r-q_r} - FV_{r-1-q_{r-1}} + PV_{q_r} - PV_{q_{r-1}}$$

$$MFV_{r \to r+1} + MPV_{r \to r+1} \leq MV_r \leq MFV_{r-1 \to r} + MPV_{r-1 \to r}$$

In other words, the MV produced by the VC approach on the sections *below* and *above* reservoir level *r* must be *less than* or *more than* the MV produced by the DMVC approach *at* reservoir level *r*, respectively.

The second requirement of the proof is that these inequalities to tend to equalities as the discretisation in the reservoir levels becomes finer.

Figure 6.20 shows an example situation with a single linear RD curve where the optimal dispatch quantity (offer) for both approaches is the same, and is at the point of intersection between the marginal revenue (MR) and MOC curves (with the price given by the RD curve at this quantity). The marginal value from the DMVC approach for reservoir level *r* is given by this intersection (MV_r).

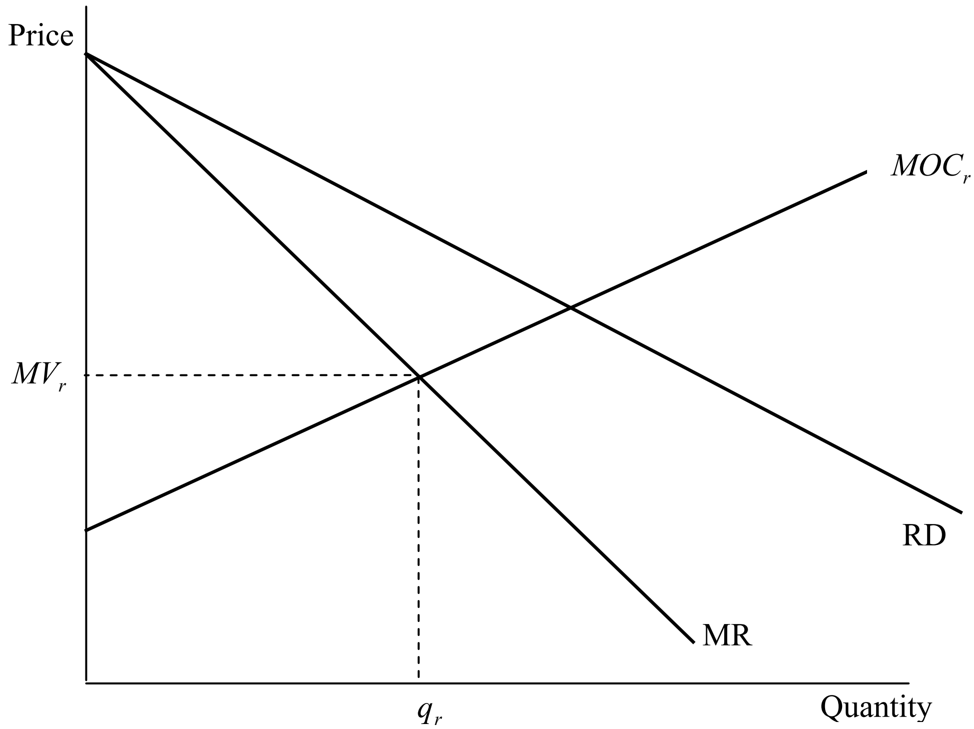

Figure 6.20 Optimal Dispatch Point under Both Approaches

Figure 6.21 demonstrates the change in the present value component as a result of the marginal increase in reservoir level from *r* to *r+1*. The additional value is the area under the MR curve from the dispatch level of q_r to q_{r+1}.

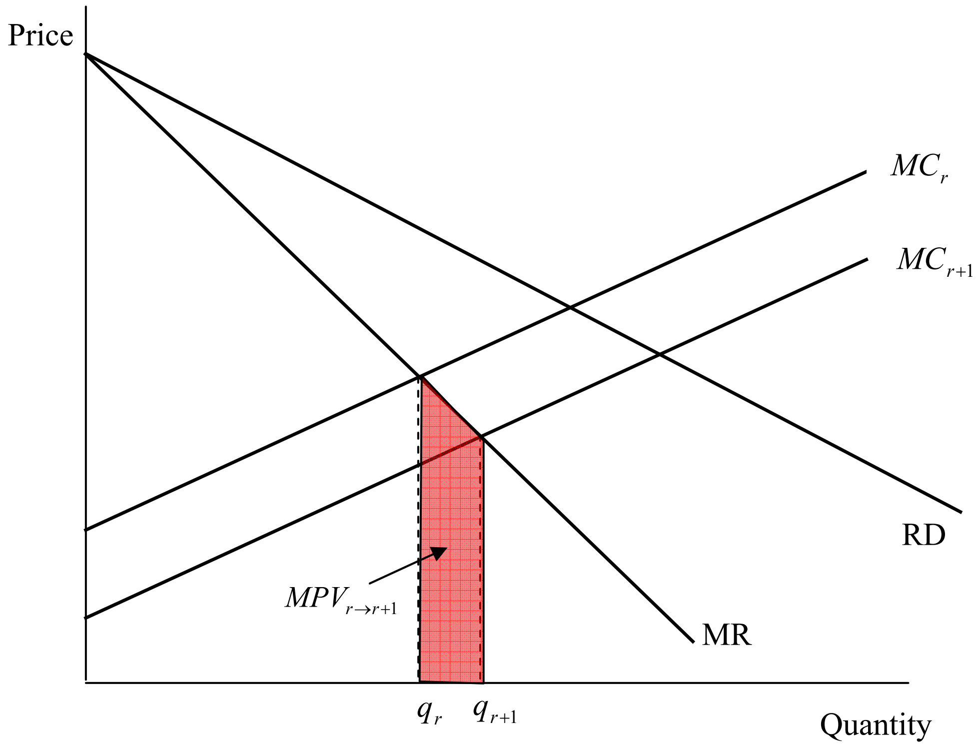

Figure 6.21 Value Curve Approach MV Calculation - Present Value Component

Figure 6.22 demonstrates the change in the future value component as a result of the marginal increase in reservoir level from *r* to *r+1*. The additional value is the area under the MV_{t+1} curve between the reservoir levels of $r - q_r$ and $(r+1) - q_{r+1}$.

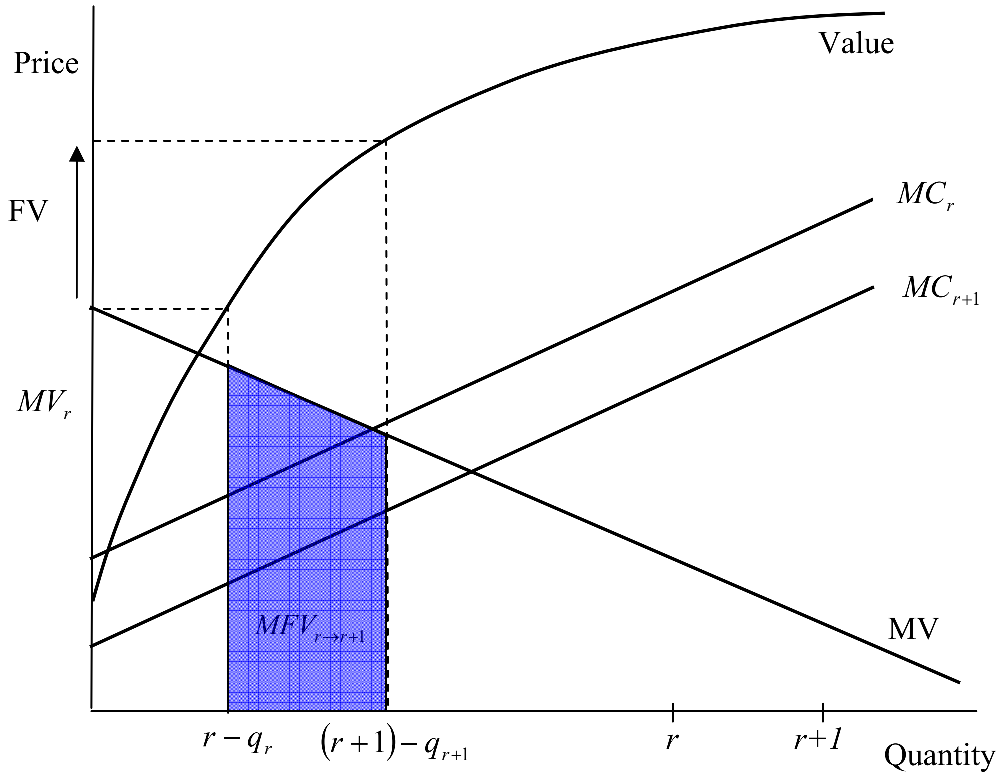

Figure 6.22 Value Curve Approach MV Calculation - Future Value Component

If we consider this in terms of a MOC curve based on a Beginning-of-Period (BOP) reservoir level of *r+1*, then:

$$(a) = (r+1) - ((r+1) - q_{r+1}) = q_{r+1}, \text{ and}$$
$$(b) = (r+1) - (r - q_r) = q_r + 1$$

These points are shown on Figure 6.23, along with regions corresponding to increased current and future value. We can see that the total width of additional payoff on this diagram is:

$$((q_r + 1) - q_{r+1}) + (q_{r+1} - q_r) = 1$$

Because the maximum point of the shaded area is at the level of MV_r and the total width of the area is 1, then we know that:

$$MFV_{r\to r+1} + MPV_{r\to r+1} + L = MV_r$$

$$MV_{r\to r+1} + L = MV_r$$

$$MV_{r\to r+1} \leq MV_r$$

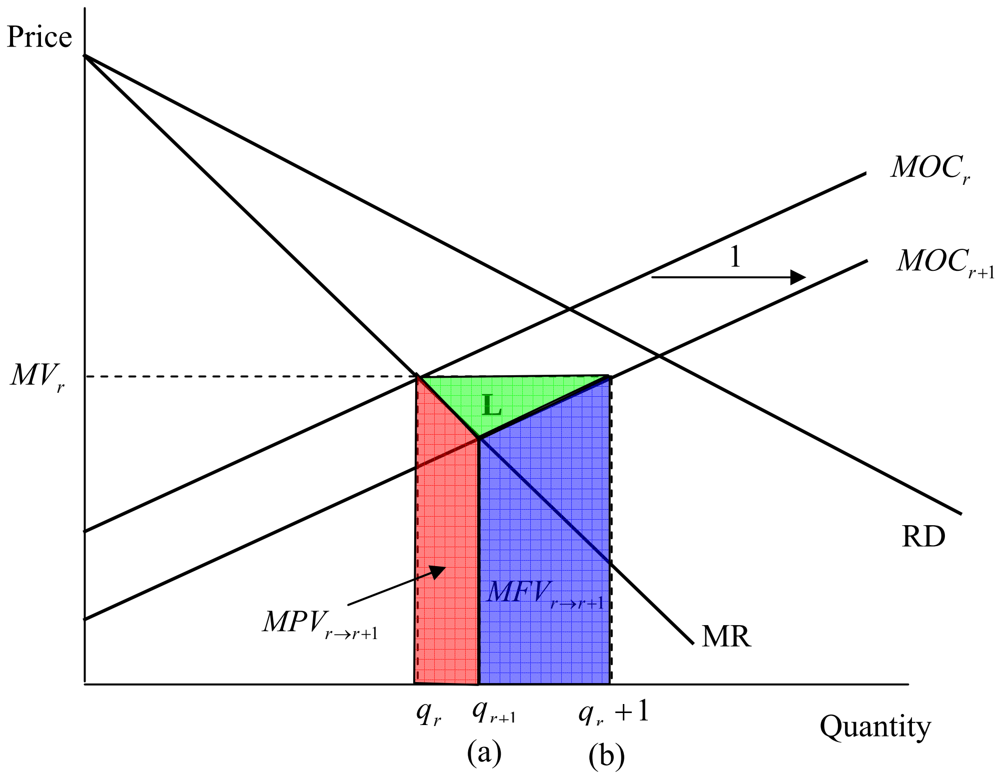

Figure 6.23 MV Comparison - Higher Reservoir Level

Figure 6.24 shows a similar result, comparing MV_r with $MV_{r-1\to r}$. We can see that:

$$MFV_{r-1\to r} + MPV_{r-1\to r} = MV_r + G$$

$$MV_{r-1\to r} = MV_r + G$$

$$MV_{r-1\to r} \geq MV_r$$

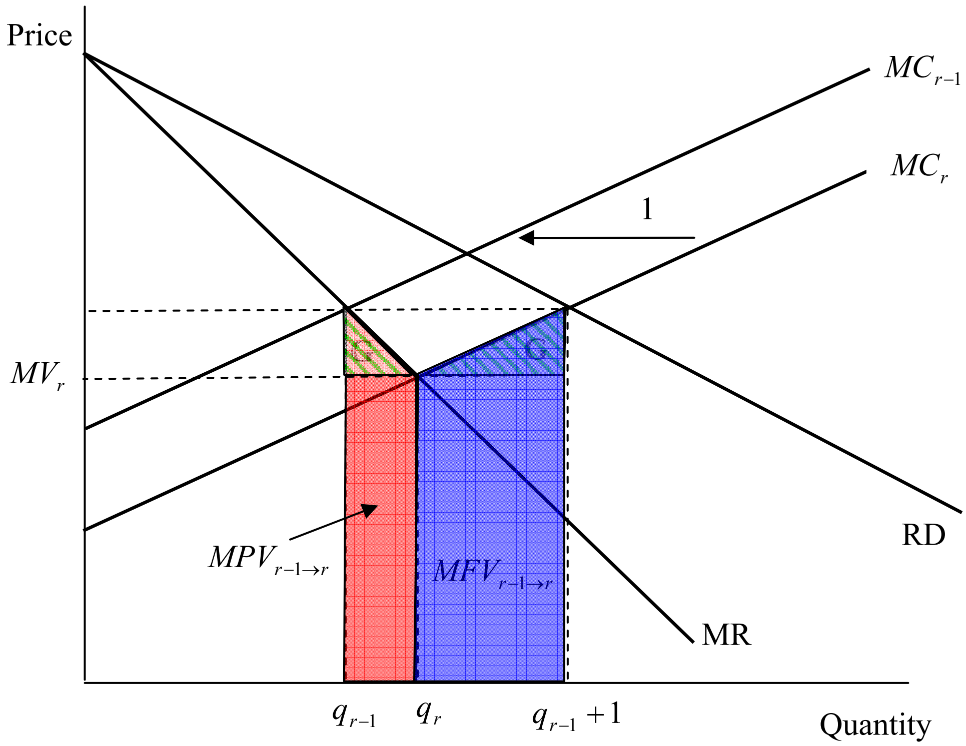

Figure 6.24 MV Comparison - Lower Reservoir Level

It is clear from Figure 6.23 and Figure 6.24 that as the discretisation of the reservoir levels becomes finer (less shift in the MC curve at each incremental move), the MV for the segments above and below a particular reservoir level point found from the VC approach become progressively closer to the MV found at that reservoir level point from the DMVC approach.

Therefore, we can see that under a scenario with a single RD curve, the VC approach and the DMVC approach are equivalent – they produce the same offers, the same dispatch level, and the same MV curve for the previous period, on which earlier offers will be based.

Q.E.D.

Throughout this thesis we consider scenarios with multiple RD curves under consideration. This situation is the same as that presented here, except that the offer is constructed to optimally trade-off between alternative residual demand curves, and the analysis demonstrated above would be performed separately for each possible market outcome, where the weighted average of these cases would give a consistent result.

6.6.4 Proof of Algorithm Convexity

In Section 6.6.3, we have shown that the VC and DMVC approaches are equivalent, and in the continuous case should produce the same results. Therefore, in order to prove that both algorithms will continue to naturally produce monotonically non-increasing MV curves (and thus monotonically non-decreasing offers) as the DP recursion moves backwards through the horizon, this just needs to be shown for one of the algorithms. Here, we demonstrate this for the DMVC, as the proof is simpler.

Recall that all offers pass along a grid of edges, and thus for each single unit of generation, our decision is essentially at what price to offer that tranch of width 1MW, based on the MOC on that tranch and determined by the pre-processed OCF for the relevant UMS. As with Section 6.6.3, we ignore inflows as the effect is again just to shift the MOC curve, independently of starting reservoir level or any other factor.

Proof 6.2

For this proof we define the following terms:

$Q_{\max}$ = Maximum capacity of the generating unit

$MOC_{q,r}$ = Marginal opportunity cost for q^{th} unit, given reservoir level r

$OP_q\left(MOC_{q,r}\right)$ = Offer Price for q^{th} unit of capacity, given the MOC at the q^{th} unit, given a reservoir level of r

$d_{r,RD}$ = Dispatch level under the optimal offer from reservoir level r and under RD curve outcome RD

1. Begin with a set of monotonically non-increasing MV curves at the end of period t, with respect to reservoir level.

2. Construct the set of Expected MV curves, also for the end of period t, by taking weighted averages, as described in Section 6.6.2. These new curves must also be monotonically non-increasing, as taking a vertically weighted average of n monotonically non-increasing curves will produce a monotonically non-increasing curve, regardless of the weights.

3. As reservoir level increases, the MOC curve used to construct the offer moves horizontally to the right (shown in Figure 6.25). Specifically, when r increases by 1MW from r_{old} to r_{new}:

For q = 1 to $Q_{\max}$

$$MOC_{q,r_{new}} = MOC_{q-1,r_{old}}$$

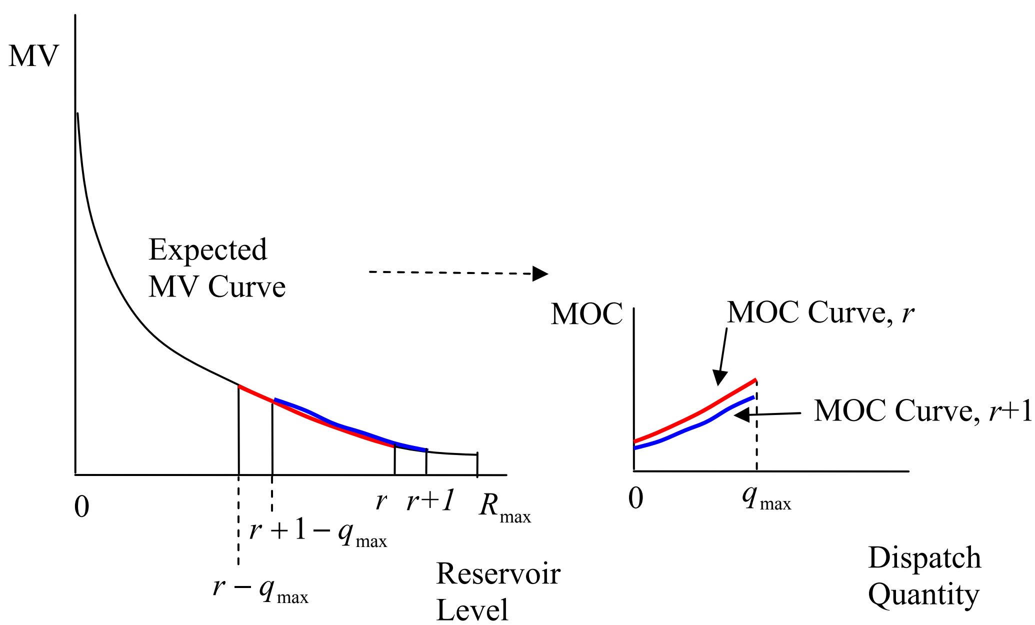

Figure 6.25 Constructing the MOC Curve

4. As a result, the offer:
    a. Becomes more generous, but
    b. For a given *q*, will not drop below $OP_{q-1}\left(MOC_{q-1,r_{old}}\right)$ because the offer must be at least as high in this *q* as it was for the same MOC in the previous *q* level (as pre-processed offers for fixed MC levels are monotone). This is demonstrated in Figure 6.26.

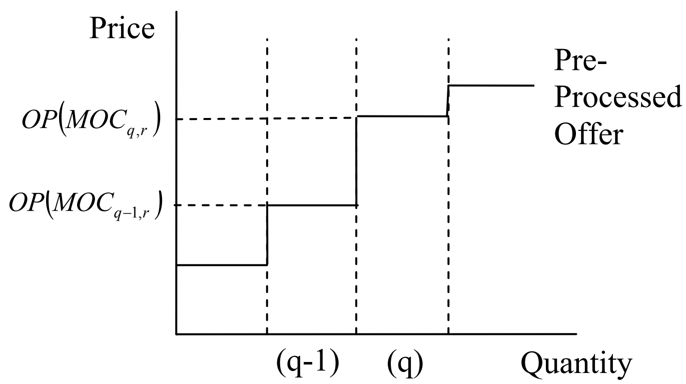

Figure 6.26 Monotone Pre-Processed Offer for Fixed MC Level

Condition b　　　Condition a

For q = 1 to $Q_{\max}$

$$\underbrace{OP\left(MOC_{q-1,r_{old}}\right) \leq}_{\text{Condition b}} \underbrace{OP\left(MOC_{q,r_{new}}\right) \leq OP\left(MOC_{q,r_{old}}\right)}_{\text{Condition a}}$$

In other words, when the reservoir level increases, giving the fall in MOC curve shown in red in Figure 6.27, the new offer will be somewhere in the highlighted range.

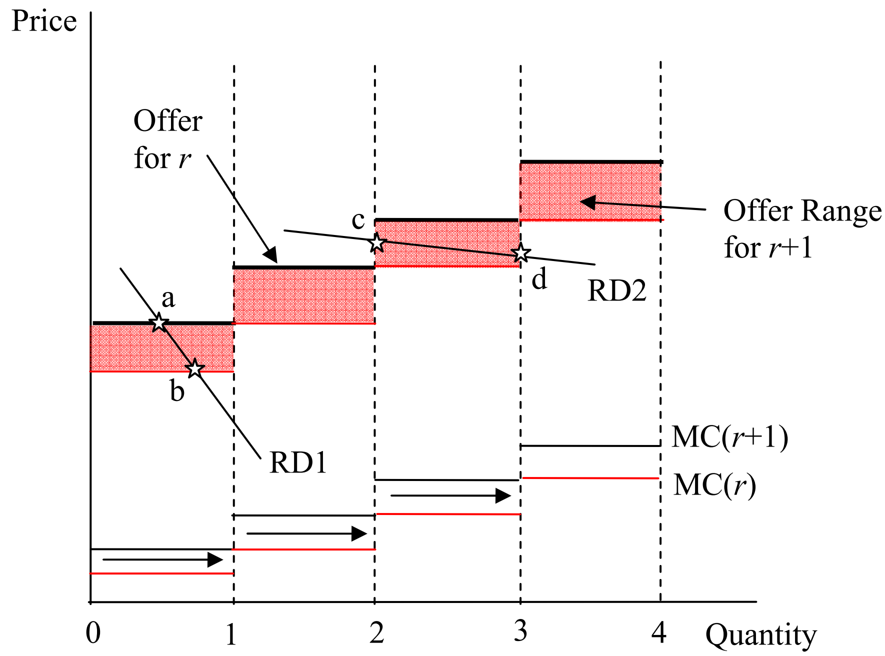

Figure 6.27 Shift in Offer as Reservoir Level Increases

5. Figure 6.27 tells us that under a monotonically non-increasing RD curve, *RD*, the following must be the case:

$$d_{r+1,RD} \leq d_{r,RD} + 1$$

In other words, there is no way that the dispatch level can increase by more than one unit of generation when the reservoir level increases by 1. For example, assume that the offer moves as far as it feasibly can in Figure 6.27, to the bottom of the highlighted range. Under RD1, dispatch increases by less than 1 unit, from a-b, while under RD2 (a very flat curve), dispatch increases by the maximum possible amount, that being exactly 1 unit (from c-d).

6. If we rearrange the equation above and add r to both sides, we get:

$$
\begin{aligned}
&-(d_{r,RD}+1) \le -d_{r+1,RD} \\
&r-(d_{r,RD}+1) \le r-d_{r+1,RD} \\
&r-d_{r,RD} \le (r+1)-d_{r+1,RD}
\end{aligned}
$$

This means that when the BOP reservoir level increases, the EOP reservoir level must also increase.

7. Now consider how the MV in the current period is calculated for market state p_2, and reservoir levels r and (r+1), under the DMVC approach.

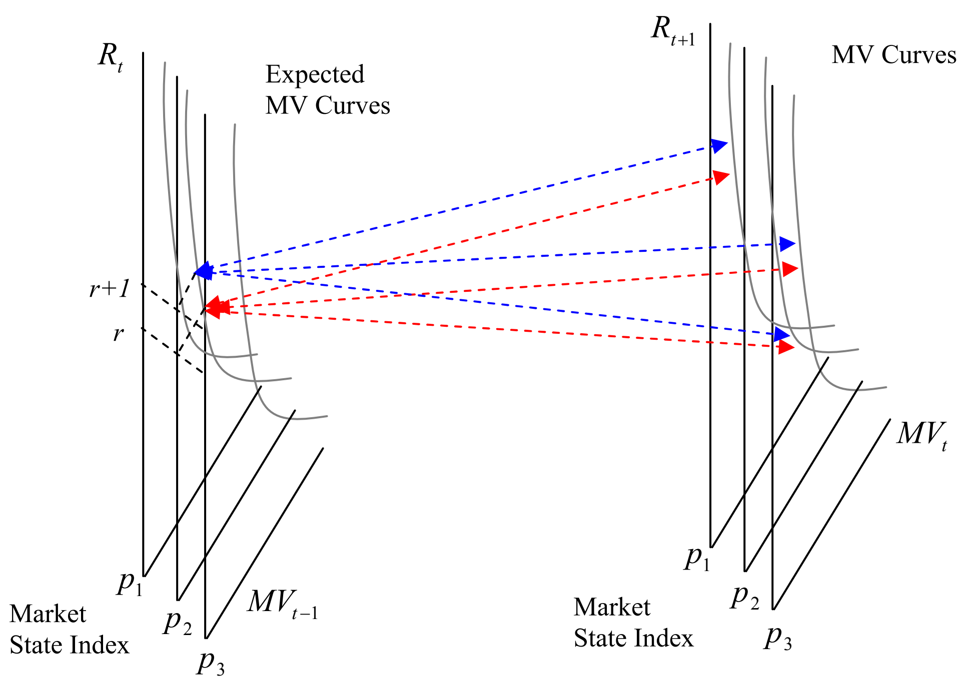

Figure 6.28 MV Calculation under the DMVC Approach

Figure 6.28 shows that to find the MV at this particular point, we take a weighted average of the MV levels that could be reached under the possible RD curve outcomes, given the optimal offer from the appropriate reservoir level.

In step 6, we showed that if the BOP reservoir level increases, then the EOP reservoir level that results must also increase. Given that the MV curves from which we are taking our weighted values are downward sloping, this means that under all RD curve outcomes, the corresponding MV must not increase when we calculate the BOP MV under a higher reservoir level. Hence, each of the MV curves in the period *(t+1)* must be monotonically non-increasing.

8. So, if we begin with a set of monotonically non-increasing MV curves at the end of period t, all the offers within period t will be monotone, and will lead to a set of monotonically non-increasing MV curves at the end of period *(t-1)*.

Q.E.D.

6.6.5 Computational Complexity of Real-Time Phase

From the source code for the RT phase based on the Value Curve and Direct MV Curve approaches, we can show that the computational complexities can be approximately summarised by Equation 6.3 and Equation 6.4.

$$0.1tp\left(400r + 20ri + 230pr + 140 + 65M + 28Md + 10p + 20d + 130rd\right)$$

Equation 6.3 Computational Complexity of the RT Phase Value Curve Approach in Full

$$tp\left(16r + 2ri + 11pr + 14 + 6.5M + 2.8Md + p + 2d + 13rd\right)$$

Equation 6.4 Computational Complexity of the RT Phase Direct MV Curve Approach in Full

Simplifying these down to the terms of the largest orders of magnitude, we can see that both algorithms are of order of complexity $tpr\left(p + d\right) + tpMd$.

When compared to the simplified form of Equation 6.1, this shows us that the computational complexity of the RT phase for both the Value Curve and Direct MV Curve approaches will increase at a considerably slower rate than the R&A algorithm (which had complexity $trp^2d^3 + tpdr^2$) as the values of the various parameters which define the problem size increase. In other words, as the scenarios considered increase in size and detail, the two-phase approaches will become increasingly faster in comparison to the R&A approach.

Recall from Section 6.5.5 that the order of computational complexity of the PP phase, which is performed just once per horizon, is tp^2Md^2. By comparing computational complexity with the R&A algorithm, we can observe that as the values of the various parameters which define the problem size increase, the computational time required will increase at a significantly slower rate than the R&A algorithm.

Note that the number of fixed MC levels at which pre-processing occurs does not greatly affect the computational complexity of the RT phase algorithm. Of course, as the discretisation of the MC levels becomes finer, we would expect the solution quality to improve. Although testing has shown this to be true to a certain extent, we have found that the optimal offers and payoffs in the RT phase are quite robust to the number of MC levels considered. Therefore, we should not be too concerned about the number of MC levels that are solved for in the PP phase (whose computational complexity is linear in this parameter).

6.7 Expected Solution Quality Comparison: Two-Phase and R&A Approaches

In Section 6.2, it was noted that the Two-Phase and R&A algorithms considered in this chapter all deal with the same market complexities and issues, and as such, we would expect that the solution quality from all three approaches would be very similar. This section explores the reasons for slight deviations that we expect to observe between the actual solution qualities produced by the three approaches.

6.7.1 Two-Phase Approaches versus R&A Approach

There are three differences between the Two-Phase and R&A algorithms which will mean that the outcomes achieved are not identical.

The first two of these have been discussed earlier in this chapter, and favour the R&A algorithm very slightly. These were the choice of factors/dimensions that were to be discretised (Section 6.4.1), and the fact that the construction of the MOC curves in the Two-Phase approaches comes from a weighted average rather than a perfect representation (in the Value Curve approach, which we have chosen to programme, as described in Section 6.6.2).

The final difference weighs the advantage heavily on the side of the Two-Phase algorithm, especially when the shape of the RD curves considered is "steppy" and conducive to market power exploitation, such as those found in the real world. The remainder of this section explains this difference.

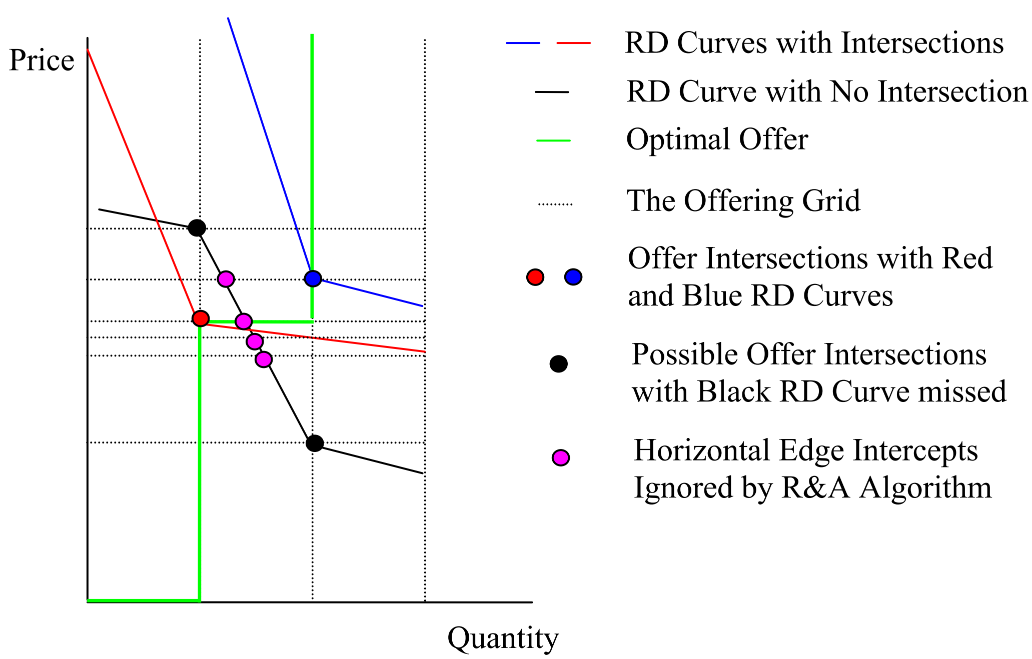

Figure 6.29 Inaccuracy in R&A Algorithm

The problem lies in the fact that the R&A algorithm considers only the payoffs that occur on the vertices of the grid, as shown in Figure 6.29. The example here demonstrates a case where the expected values on the vertices indicated by the blue and red dots are very high, and as such, the optimal offer should pass through these points. Under such a case, there is no way that an intersection with a vertex of the black RD curve can be recognised. This is because the offer must pass through this RD curve at one of the purple dots, on a horizontal edge of the offering grid. By ignoring this intercept, not only is the additional expected payoff within the current period omitted, but so is the future value associated with being dispatched on this ignored RD curve.

For example, Figure 6.30 shows a case where the solid pink line represents the R&A algorithm offer, while the dotted pink line represents the deviation from this for the VC approach offer. Quite clearly, the only difference between the two offers is the point at which they intercept RD curve 2. Given the rest of the offer beyond this point, the R&A algorithm considers only the possibility of intersecting this RD curve at point b (with a payoff + future value of $0.1(2*26.14+6424)=647.63$), as point a would require shifting around the remainder of the offer to maintain monotonicity[40]. On the other hand, the VC approach also considers the possibility of intersecting this RD curve at (amongst others) point c (with a payoff + future value of $0.1(1.168*220+6475)=673.20$). In other words, expected total payoff associated with this RD curve has increased by around $25.56, or almost 4%. by considering payoffs on the horizontal edges of the offering grid, in addition to those on the vertices.

[40] Note that 0.1 is the probability of this RD curve occurring under the given UMS

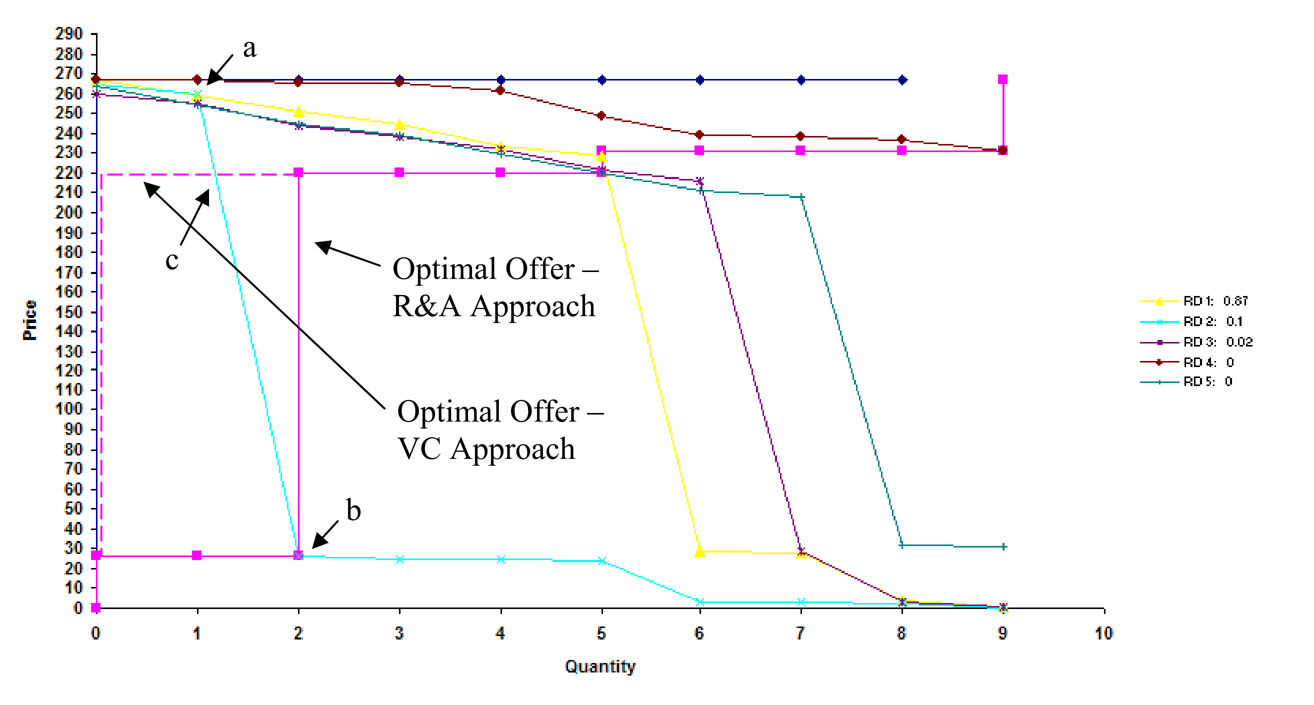

Figure 6.30. Problem with Ignoring Horizontal Intercepts

6.7.2 Value Curve Approach versus Direct Marginal Value Curve Approach

We expect that the VC approach will perform marginally better than the DMVC approach with respect to the optimality of the solutions, because of the way that the interpolation is performed, when we choose to find the value or MV at only a portion of reservoir levels.

Let us assume that the general shape of the value curves is increasing and concave, giving a MV curve that is decreasing and convex, as shown in Figure 6.31.

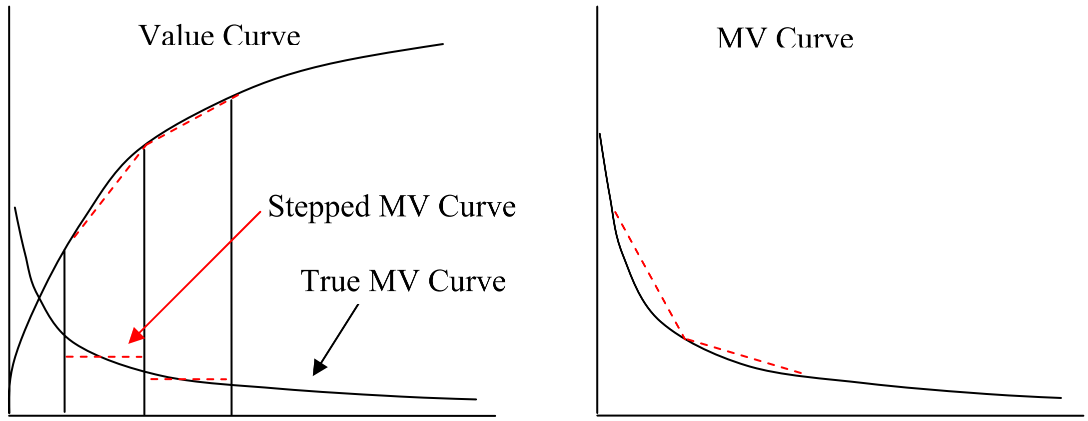

Figure 6.31. Standard Shape of Value Curve and Marginal Value Curve

Start by considering the MV curve. The offering decisions are based on the marginal value of each incremental unit of fuel. Therefore, if we have not explicitly calculated the marginal value of each of these units, then we must perform a *linear* interpolation on the MV curve. As we can see from the red lines on the right diagram of Figure 6.31, this will always bias the MV curve upwards, and hence lead to more restrictive offers than are optimal. On the other hand, if we use the Value Curve approach, the linear interpolation occurs on the Value Curve, as shown on the left hand side of Figure 6.31. It is apparent that this will produce a stepped MV curve, where each step is at a level somewhere in-between the maximum and minimum true marginal values over that range. With this approach there is no clear bias causing us to offer in either too conservative or too aggressive a manner in the long run. As such, we expect that the VC approach (which performs interpolation on the value curve) to produce slightly better results than the DMVC approach (which performs interpolation on the MV curve).

On the other hand, there is potential that the MV curve approach may be slightly more computationally efficient, as it avoids the need to construct the value curves, in constructing the MV curves for a given period directly from those for the following period.

6.8 Summary and Conclusions

In this chapter, we have described two new algorithms for the production and updating of an optimal offering strategy over a planning horizon for either a single fuel-constrained generation unit or a set of parallel units. The key complications that these algorithms consider are two of the important intertemporal considerations: stochastic market correlation between periods and limited fuel resources. The algorithm incorporates knowledge gained earlier in this thesis about the way that optimal offers should be formed and how they should change as a result of outside influences. In particular, it has capitalised on the idea of Marginal Cost Patching, introduced in Chapter 5.

We have proven that both algorithms will consistently produce concave value curves (decreasing MV curves), which is necessary for the construction of monotonically non-decreasing offers under the two-phase approach presented.

We have also compared, on a theoretical basis, the two new algorithms to an algorithm presented in Rajaraman & Alvarado (2003), which considers the same market complexities, but uses a very different approach. We have determined that the payoffs from the two new algorithms should be slightly higher for most RD curve types, and significantly higher under one specific (and realistic) type. We have also determined that, in terms of computational complexity, the two new algorithms are significantly superior to the algorithm found in Rajaraman & Alvarado (2003). These two hypotheses are tested in Chapter 8.

The following chapter considers incorporating a branching tree structure to the Markov chain system in an effort to further simplify the two-phase approaches and more closely reflect reality.

Chapter 7

DECISION ANALYSIS DYNAMIC PROGRAMMING FRAMEWORK

7.1 Introduction

This chapter describes what appears to be a new solution technique that combines traditional stochastic dynamic programming methods used to this point in the thesis with a decision analysis or branching tree type structure. This structure implies that the overall state of any system of interest at any stage is dependent on the various events that have occurred to this point. The combination of these events defines the *Macro-State* that the system is in at that time, and the optimal decision would therefore depend on the value of this macro-state. Note that these macro-states contain different information (the general, high-level state of the system) from that implied by a regular state (which in the algorithms presented in this thesis is generally defined by the previous residual demand curve outcome and the reservoir level).

Such an approach has many applications. We begin this chapter by describing the technique generically and describing its possible application to two different electricity generating problems that are not directly related to the offering problem addressed in this thesis. We then consider this new method in the context of the problem faced in this thesis, that being the construction of optimal offering strategies for a planning horizon. By applying this method to our problem of interest, we recognise that there are two types of correlation that may be occurring throughout the planning horizon. The first type is the direct correlation between the market states (or RD curves) that was included in the algorithm already presented, in the form of a Markov chain. The second type is the correlation between the overall macro-states. Ladurantaye, Gendreau, & Potvin (2006) presents an approach that is similar to that described in this chapter, but the branches used in their approach are purely of the probabilistic variety that are described in Section 7.2.2, and thus do not consider decision branches as described in Section 7.2.3. In addition, that approach considers deterministic pricing outcomes along each multi-period branch, rather than the stochastic and correlated RD curve outcomes that we consider here. For a fuller study and description of applied decision analysis, refer to Bunn (1984).

We show how applying this structure to an optimisation method enables us to better represent real-world scenarios that are themselves of this form, and as a consequence the computational complexity of a dynamic program can be significantly reduced. Finally, we show how it can be used to consider the effect on a generator’s offering behaviour where many events and phenomena, such as the reaction of rival generators to the generator’s own offering strategy, need to be taken into account.

7.2 Decision and Probability Tree Structure

In this section, we will discuss the combination of decision analysis with stochastic dynamic programming in a generic sense, and provide two brief examples of its possible application to problems other than that of short-term offer strategy optimisation primarily addressed in this thesis.

In a traditional stochastic dynamic program, a single Markov chain is considered to represent the transition probabilities between possible system states between consecutive stages. In combining decision analysis with dynamic programming, we incorporate a branching-type structure to represent the decision maker being faced with entirely different groups of possible system states, as a result of various events that could occur within the problem scenario between stages. We refer to each of these sets of possible system states as a macro-state. The events are represented by nodes in the decision tree and can be one of two distinct types.

Chance Nodes:

'Nature' can make decisions that are beyond the control of the decision maker, leading to chance nodes within the decision analysis branching structure. As explained in Bunn (1984) and demonstrated in Figure 7.1, chance nodes (or *outcome* nodes, as the text refers to them) are conventionally represented by circles and indicate that the set of subsequent nodes connected to this circle will be reached according to some probabilistic process over which the decision maker has no control.

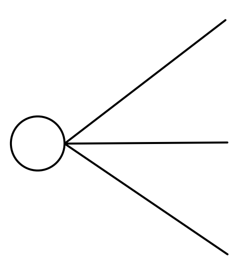

Figure 7.1 A Chance Node (Source: pg 182, Bunn (1984))

Decision Nodes:

The decision maker can make internal decisions, leading to decision nodes within the decision analysis branching structure. Bunn (1984) explains that decision nodes are conventionally represented by a square box, and indicate that subsequent nodes connected to this box can be reached according to a deterministic choice on the part of the decision maker at this point.

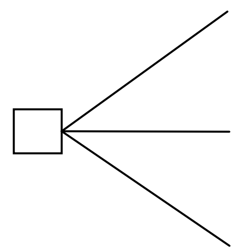

Figure 7.2 A Decision Node (Source: pg 182, Bunn (1984))

First let us consider the two techniques of dynamic programming and decision analysis separately.

7.2.1 A Simple Dynamic Program Based on a Markov Chain

Consider a single Markov chain, such as that represented in Figure 7.3, which describes the transition probabilities between four possible system states between two consecutive stages, for a given problem.

Next Stage System State

Current Stage System State	1	2	3	4
1	0.32	0.48	0.18	0.02
2	**0.24**	**0.56**	**0.04**	**0.16**
3	0.08	0.02	0.45	0.45
4	0.04	0.06	0.18	0.72

Markov Chain (Transition Probabilities)

Figure 7.3 Markov Chain Structure for System States

Assume that a decision maker must make a certain decision for the given stage based on the initial system state, before observing the probabilistic outcome of the final system state, where the payoff for this decision is dependant on the final system state. Therefore, we define the following notation:

i,j =	System state indices
NS_t =	Total number of possible system states in stage t
$pr_{i,j}$ =	Probability of transition from system state i in stage t to system state j in stage $(t+1)$ (the values in Figure 7.3)
$EV_{i,t}$ =	Expected value from the beginning of stage t onwards, associated with being in system state i
$D_{i,t}$ =	Decision in stage t, given that the system started the stage in system state i ($D_{i,t}$ * is the optimal decision)
$EP_{j,t}(D_{i,t})$ =	Expected payoff to the decision maker within stage t for a decision $D_{i,t}$, given that the system will end the stage in system state j

In this situation, the recursion equation can be defined as:

$$EV_{i,t} = \max_{D_{i,t}} \left(\sum_{j=1}^{NS_{t+1}} pm_{i,j} \left[EP_{j,t}\left(D_{i,t}\right) + EV_{j,t+1} \right] \right)$$

For example, the expected value formula if the system began stage t in system state 2 would be:

$$EV_{2,t} = \max_{D_{2,t}} \begin{pmatrix} 0.24\left[EP_{1,t}\left(D_{2,t}\right) + EV_{1,t+1}\right] + 0.56\left[EP_{2,t}\left(D_{2,t}\right) + EV_{2,t+1}\right] + \\ 0.04\left[EP_{3,t}\left(D_{2,t}\right) + EV_{3,t+1}\right] + 0.16\left[EP_{4,t}\left(D_{2,t}\right) + EV_{4,t+1}\right] \end{pmatrix}$$

Note that the relevant probabilities appear in bold in Figure 7.3. We continue this convention throughout the examples in this chapter.

7.2.2 A Simple Decision Analysis Tree Based on a Markov Chain – Chance Nodes

Taking the set of system states presented in Figure 7.3, let us now assume that system states 1 and 2 in each stage can only occur under a specific *macro-state* that the world is in at that given time, and conversely, states 3 and 4 can only occur under the alternative value of that macro-state. We will also assume that there is a higher level Markov chain on these macro-states, as presented in Figure 7.4.

		Next Period Macro-State	
		1	2
Current Period Macro-State	1	0.8	0.2
	2	0.0	1.0

Figure 7.4 Markov Chain Structure for Macro-States

This information could be reflected by chance nodes in a decision tree, combined with a set of smaller Markov chains, as shown in Figure 7.5. This structure implies that we have two levels of uncertainty; uncertainty with respect to the macro-state (observed first, at the chance nodes), and uncertainty with respect to the system state (observed second, given the macro-state observation).

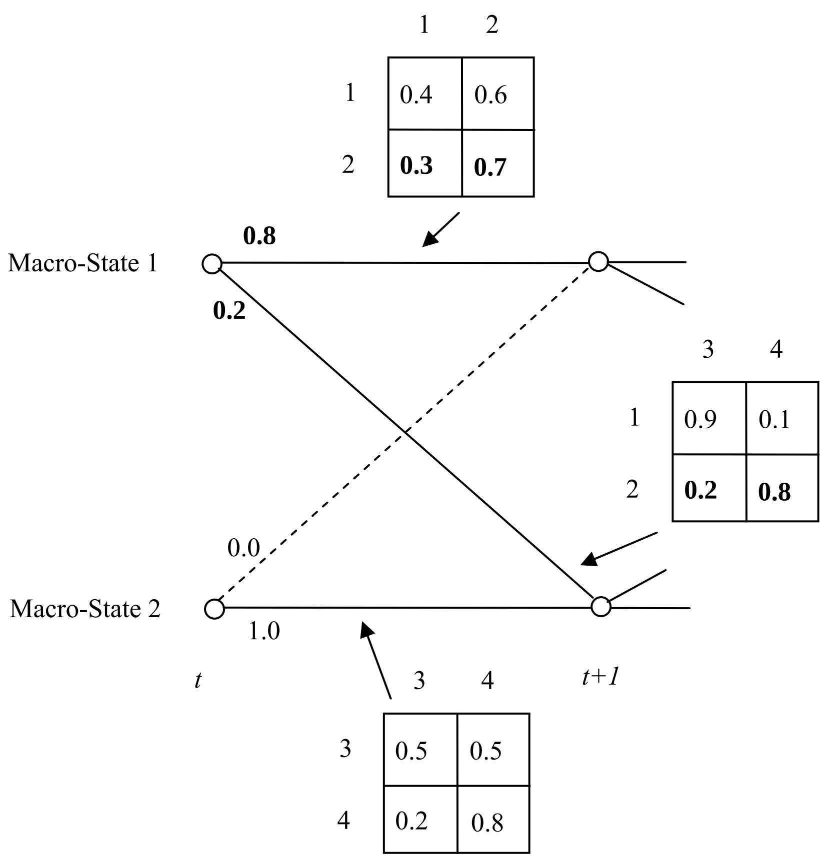

Figure 7.5 Branching Tree Representation of a Markov Chain with Macro-States

Given that this example considers only chance nodes, there is a simple, probabilistic formula to calculate the expected value from each node onwards, moving backwards across the decision tree towards the source node. In order to define this formula, we present the following additional notation:

$k,l =$ Macro-state indices

$NM_t =$ Number of possible macro-states in stage t

$m_{k,l} =$ Probability of transition from macro-state k in stage t to macro-state l in stage (t+1) (the values in Figure 7.4)

$NSM_t =$ Number of possible system states associated with each macro-state in stage t (for simplicity we will assume that there are the same number of system states associated with each macro-state within a single stage, i.e. $nsm_t = ns_t / nm_t$)

$pm_{i,j,k,l} =$ Probability of transition from system state i in stage t to system state j in stage (t+1), given that there has been a transition from macro-state k to macro-state l between these stages (the values in the Markov chains in Figure 7.5)[41]

$EV_{i,k,t} =$ Expected value from the beginning of stage t onwards, associated with being in system state i, which is associated with macro-state k at the beginning of stage t

The expected value is calculated using:

$$EV_{i,k,t} = \sum_{l=1}^{NM_{t+1}} m_{k,l} \left[\sum_{j \in l} pm_{i,j,k,l} EV_{j,l,t+1} \right] \text{ for } m_{k,l} \neq 0 \text{ terms}$$

Consider calculating the expected value of being in system state 2, which is associated with macro-state 1 in stage t, i.e. $EV_{2,1,t}$, for the example in Figure 7.5.

$$\begin{aligned} EV_{2,1,t} &= \sum_{l=1}^{2} m_{1,l} \left[\sum_{j \in l} pm_{2,j,1,l} EV_{j,l,t+1} \right] \\ &= 0.8\left[0.3EV_{1,1,t+1} + 0.7EV_{2,1,t+1}\right] + 0.2\left[0.2EV_{3,2,t+1} + 0.8EV_{4,2,t+1}\right] \\ &= 0.24EV_{1,1,t+1} + 0.56EV_{2,1,t+1} + 0.04EV_{3,2,t+1} + 0.16EV_{4,2,t+1} \end{aligned}$$

41 Note that the macro-state indices k and l are technically redundant here, as they are implied by the market states i and j. However, we include them for clarity regarding the current macro-state.

Note that these final probabilities defined above are the same as those in the second row of the original non-branched Markov chain presented in Figure 7.3. The reason for this is that in the special case where all macro-states can be reached from all other macro-states, the two methods are equivalent to one another, despite the fact that they use different structures to represent the stochasticity in the system being modelled. Specifically, the probabilities are related through the formula:

$$pr_{i,j} = m_{k,l} pm_{i,j,k,l} \text{ for all } i \in k, j \in l$$

However, it is reasonable to assume that some macro-states transitions may not be possible between certain periods, as is the case for the transition from macro-state 2 to macro-state 1 in our example (infeasible transitions are represented by a probability of zero). For this node, the expected value calculation, assuming system state 3, simplifies to:

$$EV_{3,2,t} = 0.5EV_{3,2,t+1} + 0.5EV_{4,2,t+1}$$

Section 7.2.5 explains more about these types of structures and the computational benefits of infeasible transitions such as this one.

7.2.3 A Simple Decision Analysis Tree Based on a Markov Chain – Decision Nodes

Again let us take the set of system states presented in Figure 7.3 assuming that system states 1 and 2 can occur under macro-state 1, while system states 3 and 4 can occur under macro-state 2. However, now let us assume that rather than a higher level Markov chain defining stochastic transition between the macro-states, these transitions are based on an internal decision by the decision maker. In other words, the decision maker chooses which branch to traverse.

This information could be reflected by decision nodes in a decision tree, combined with a set of smaller Markov chains to represent the system state uncertainty, as shown in Figure 7.6.

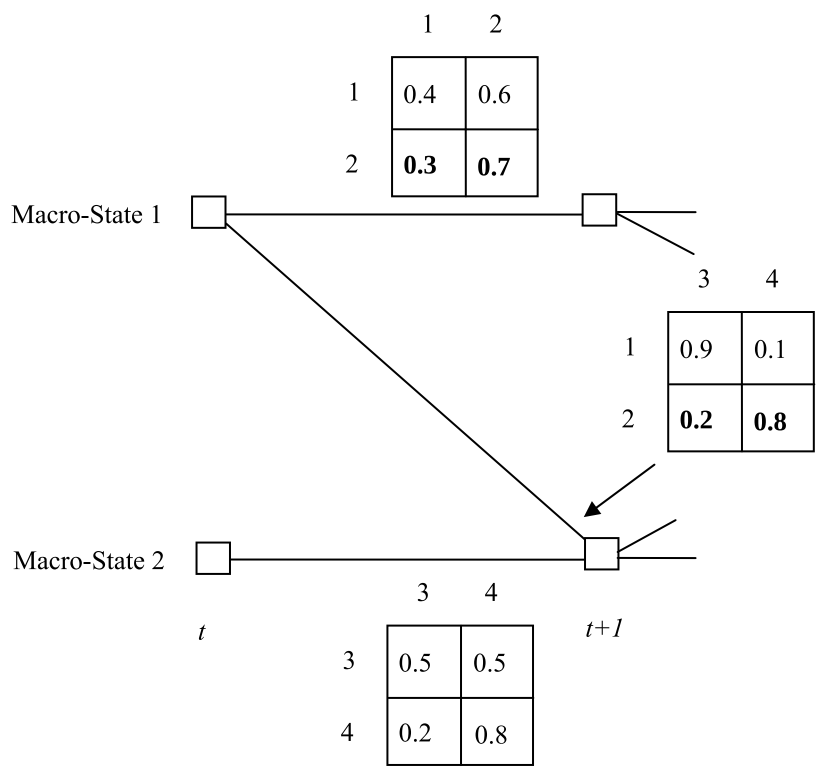

Figure 7.6 Branching Tree Representation of a Markov Chain with Macro-States

The expected value from any decision node onwards is evaluated in a similar general formula to that for a chance node, but rather than probabilistically determining the expected value over the possible branches, the expected value is the maximum value out of these branches. Therefore, the expected value formula is:

$$EV_{i,k,t} = \max_{l}\left(\sum_{j\in l} pm_{i,j,k,l} EV_{j,l,t+1}\right)$$

Note though, it is quite likely that there may be a cost associated with switching between macro-states. In this case, let us define:

$TC_{k,l}$ = Transition cost from macro-state k in stage t to macro-state l (for k=l, $TC_{k,l} = 0$)

Therefore, the expected value formula becomes:

$$EV_{i,k,t} = \max_{l}\left(\sum_{j\in l} pm_{i,j,k,l} EV_{j,l,t+1} - TC_{k,l}\right)$$

For simplicity though, we will leave this term out of the remaining equations in this chapter, as this has no significant impact on the complexity of the algorithm.

Considering the same example as before, the expected value of being in system state 2, which is associated with macro-state 1 in stage t can be calculated as follows:

$$\begin{aligned} EV_{2,1,t} &= \max_{l}\left(\sum_{j\in l} pm_{2,j,1,l} EV_{j,l,t+1}\right) \\ &= \max\{(0.3EV_{1,1,t+1} + 0.7EV_{2,1,t+1}),(0.2EV_{3,2,t+1} + 0.8EV_{4,2,t+1})\} \end{aligned}$$

7.2.4 Combining DP and DA: Chance and Decision Nodes

In Section 7.2.1, we described a simple DP based on a Markov chain, where a decision was made within the stage, in order to maximise expected benefit to the decision maker, given uncertainty with respect to the system state at the end of the stage. In Sections 7.2.2 and 7.2.3, we described simple decision trees, where at the beginning of each stage, either a probabilistic event was resolved, or a decision was made, that determined the macro-state that the system would be in at the following stage. For the remainder of this chapter, we will consider the combination of these two concepts in a new decision analysis and dynamic programming (DADP) framework.

For chance nodes in the branching tree the decision making and uncertainty resolution process in each stage is shown in Figure 7.7.

Resolve Macro-State Uncertainty → Make Decision subject to System State Uncertainty → Resolve System State Uncertainty and thus determine value of the decision

Figure 7.7 Decision Making and Uncertainty Process under Chance Nodes

For decision nodes in the branching tree, there is a two-level decision making process, and the process in each stage is shown in Figure 7.8.

Make Macro-State Decision → Make Decision subject to System State Uncertainty → Resolve System State Uncertainty and thus determine value of the decision

Figure 7.8 Decision Making and Uncertainty Process under Decision Nodes

For both of these node types, the operational decision is made after the high-level/macro-state transition has been resolved, once the decision maker knows the general realm of the system states that may occur, while the value of that decision to the decision maker remains dependent on the specific system state that occurs.

This information is demonstrated in Figure 7.9.

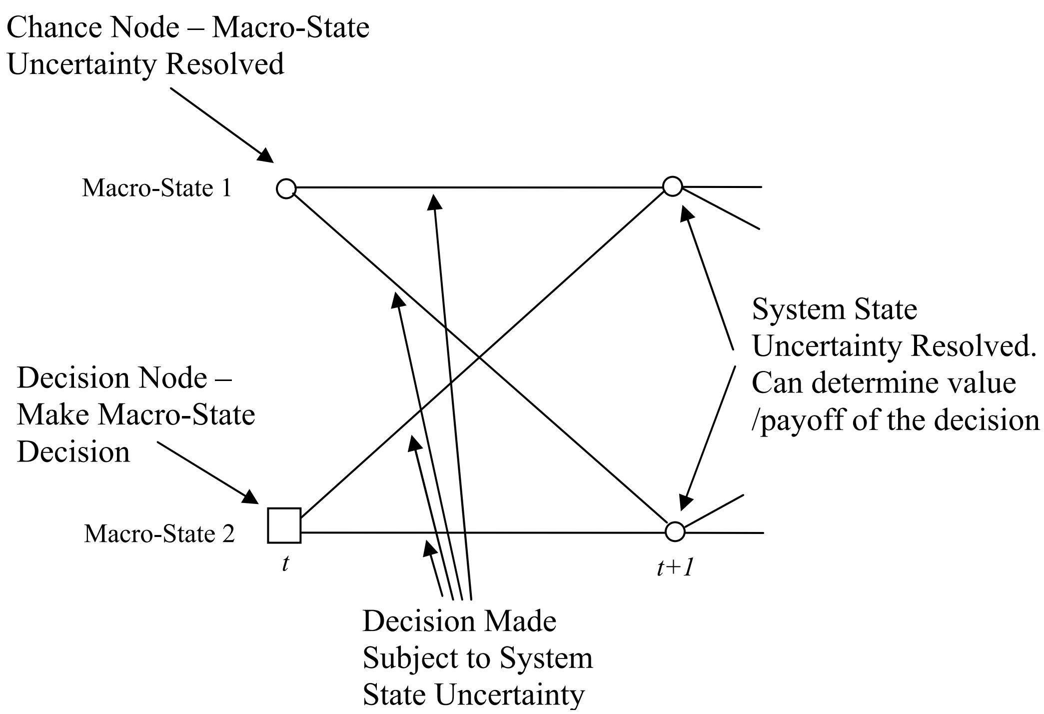

Figure 7.9 Process for each Stage

In order to define the recursion equation for the DADP framework, we present the following additional notation:

$D_{i,k,l,t}$ = Decision in stage t, given that the system started the stage in system state i (associated with macro-state k) and given that the decision maker has observed or caused a shift to macro-state l ($D_{i,k,l,t}$* is the optimal decision)

$EP_{j,t}\left(D_{i,k,l,t}\right)$ = Expected payoff to the decision maker within stage t for a decision $D_{i,k,l,t}$, given that the system will end the stage in system state j

The expected values and payoffs are demonstrated in Figure 7.10, for both a chance and a decision node.

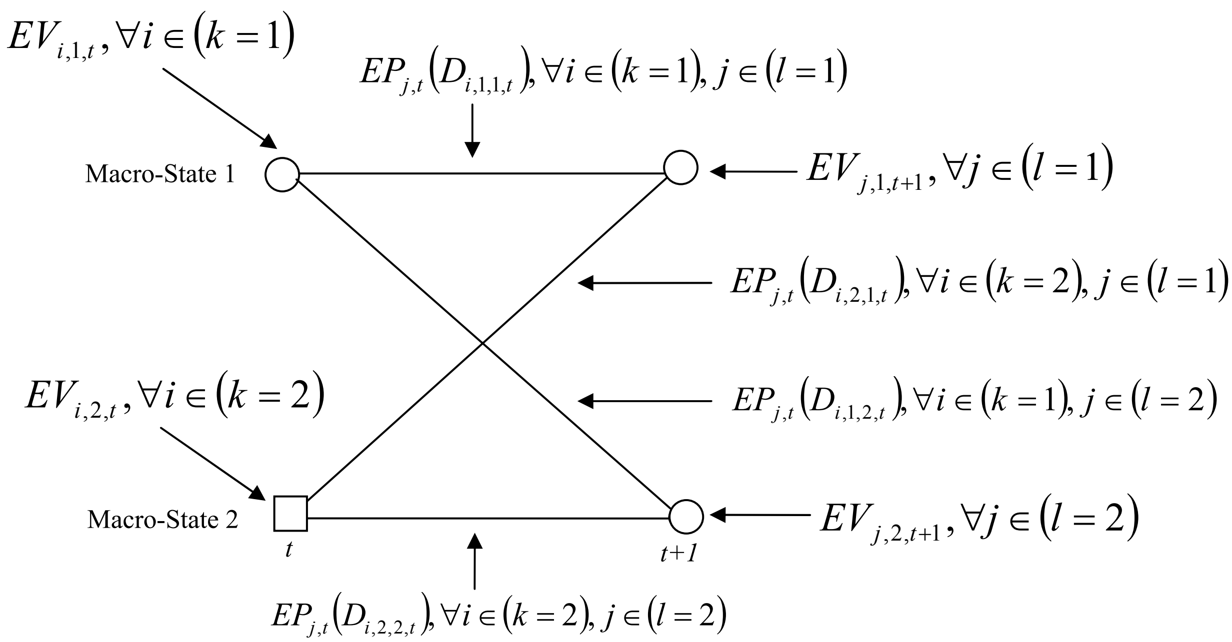

Figure 7.10 Values and Expected Values

For the chance node, the expected value is calculated for the recursion as follows:

$$EV_{i,k,t} = \sum_{l=1}^{NM_{t+1}} m_{k,l} \left[\max_{D_{i,k,l,t}} \left(\sum_{j \in l} pm_{i,j,k,l} \left[EP_{j,t}\left(D_{i,k,l,t}\right) + EV_{j,l,t+1} \right] \right) \right]$$

In other words, this is a weighted value of the expected payoffs and future values that the decision maker would receive under the optimal decisions along each possible branch.

Alternatively, for the decision node, the recursion equation is defined as:

$$EV_{i,k,t} = \max_{l} \left(\max_{D_{i,k,l,t}} \left(\sum_{j \in l} pm_{i,j,k,l} \left[EP_{j,t}\left(D_{i,k,l,t}\right) + EV_{j,l,t+1} \right] \right) \right)$$

In other words, this selects the branch that has the greatest expected payoff and future value associated with its optimal decision.

Consider the same simple example as in Section 7.2.2, but within this DADP framework. We can calculate the expected value of being in system state 2, which is associated with macro-state 1 in period t, i.e. $EV_{2,1,t}$ as follows:

$$
\begin{aligned}
EV_{2,1,t} &= \sum_{l=1}^{2} m_{1,l}\left[\max_{D_{2,1,l,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{2,t}\left(D_{2,1,l,t}\right)+EV_{j,l,t+1}\right]\right)\right] \\
&= 0.8\left[0.3\left[EP_{1,t}\left(D_{2,1,1,t}*\right)+EV_{1,1,t+1}\right]+0.7\left[EP_{2,t}\left(D_{2,1,1,t}*\right)+EV_{2,1,t+1}\right]\right]+ \\
&\quad 0.2\left[0.2\left[EP_{3,t}\left(D_{2,1,2,t}*\right)+EV_{3,2,t+1}\right]+0.8\left[EP_{4,t}\left(D_{2,1,2,t}*\right)+EV_{4,2,t+1}\right]\right] \\
&= 0.24\left[EP_{1,t}\left(D_{2,1,1,t}*\right)+EV_{1,1,t+1}\right]+0.56\left[EP_{2,t}\left(D_{2,1,1,t}*\right)+EV_{2,1,t+1}\right]+ \\
&\quad 0.04\left[EP_{3,t}\left(D_{2,1,2,t}*\right)+EV_{3,2,t+1}\right]+0.16\left[EP_{4,t}\left(D_{2,1,2,t}*\right)+EV_{4,2,t+1}\right]
\end{aligned}
$$

Alternatively, consider the expected value of being in system state 3, associated with macro-state 2 in period t. There is a decision node associated with this macro-state, and as such, given an optimal macro-state transition $l*$, the expected value is defined as:

$$
\begin{aligned}
EV_{i,k,t} &= \max_{l\in\{1,2\}}\left(\max_{D_{i,k,l,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(D_{i,k,l,t}\right)+EV_{j,l,t+1}\right]\right)\right) \\
&= \max\begin{pmatrix} 0.8\left[EP_{1,t}\left(D_{3,2,1,t}*\right)+EV_{1,1,t+1}\right]+0.2\left[EP_{2,t}\left(D_{3,2,1,t}*\right)+EV_{2,1,t+1}\right], \\ 0.5\left[EP_{3,t}\left(D_{3,2,2,t}*\right)+EV_{3,2,t+1}\right]+0.5\left[EP_{4,t}\left(D_{3,2,2,t}*\right)+EV_{4,2,t+1}\right] \end{pmatrix}
\end{aligned}
$$

There are significant advantages to structuring the stochasticity and decision making process with the DADP framework, not least of which is the potentially more accurate representation of the form of the stochasticity and decisions that exists in the real system that is being modelled. These advantages are dependent on the particular system to which the technique is being applied and are explained more fully for the generator offer optimisation problem addressed in this thesis, in Section 7.5.

7.2.5 Possible DADP Node and Transition Structures

In this section, we will present a small set of possible structures that the nodes and transitions between macro-states could take within the DADP structure. Specifically, the aim of this section is to highlight the differences, both practical and computational, between macro-states that are *captivating* for a given number of stages and those that are not. In addition, we will differentiate between the concepts of captivating and *absorbing* macro-states.

In Markov chain theory, an absorbing sub-chain is defined in Norris (1997) as a closed class, from which there is no escape. In our context, this would therefore mean that once an absorbing macro-state is entered, the system could never leave that macro-state. This could be the case if a decision maker would make a decision, or an external event could occur, from which there is no reversion. In the problems that we address in this thesis, the Markov chains that we apply are non-stationary between different periods[42], and thus it is also useful to define a captivating macro-state in a given period to be a macro-state from which it is not possible to escape in that period. A captivating macro-state would represent a situation where, if the system goes into a certain macro-state, then it will remain in that macro-state for a number of periods. Note it is possible that a macro-state may be absorbing over a short horizon, but only captivating over a longer horizon. As such, the diagrams presented in this section may appear to be absorbing, as they do not converge to a single macro-state in the final period, however, the diagrams show only a subset of the entire planning horizon.

Let us begin by noting that Figure 7.5 and Figure 7.6 represented a branching tree with an underlying Markov chain that represented macro-state 2 as being captivating, and macro-state 1 as not captivating between the given stages.

[42] Note that this concept of varying Markov chains is similar to the area of Periodic Markov Chains, except that we do not limit our method to a periodically repeating set of probabilities.

In Figure 7.11, we present a two macro-state structure over four periods, in which it is possible to reach both macro-states from any beginning macro-state in all periods.

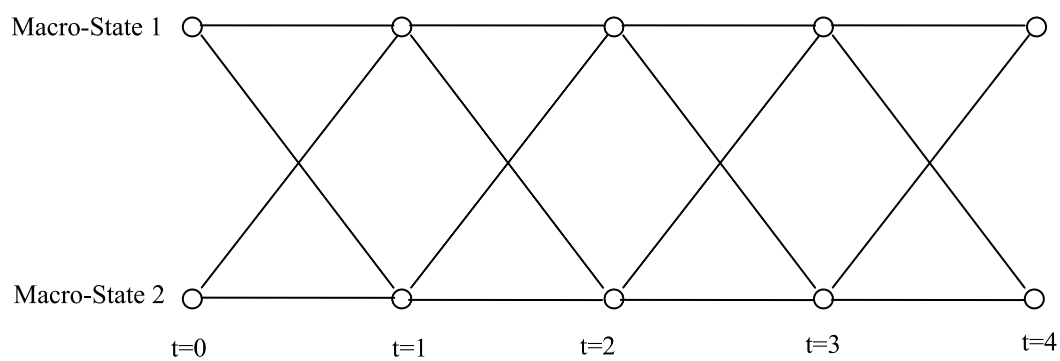

Figure 7.11 Non- Captivating Macro-States

Alternatively, it is possible that the macro-states of the system might be captivating in a particular period, meaning that it is not possible to shift to an alternative macro-state once a particular macro-state is observed. This situation might occur if a system is in a particular macro-state up to, say, $t = 1$, at which time it may or may not switch to a different macro-state, with a certain likelihood. Figure 7.12 illustrates this scenario, where the system may not return to the original macro-state at any point in the shown portion of the planning horizon.

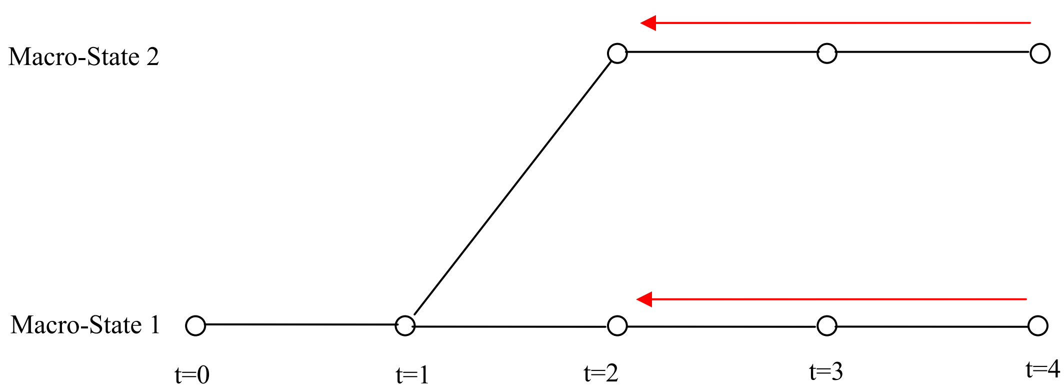

Figure 7.12 Captivating Macro-States

It is under structures like this where the major benefits of the DADP framework are evident. The lower level DP is performed backwards along each branch over multiple periods (as indicated by the red arrows in Figure 7.12), until the chance node is reached, at which point the values (or value curves, depending on the application) are combined probabilistically (or based on a maximum if it is a decision node).

Figure 7.13 represents a scenario with three possible macro-states, demonstrating a combination of captivating and non-captivating macro-states in various time periods over the short horizon. Observe that macro-state 3 is captivating within the portion of the horizon shown, macro-state 2 is captivating only in t = 1 and 2, and macro-state 1 is captivating only in t = 1.

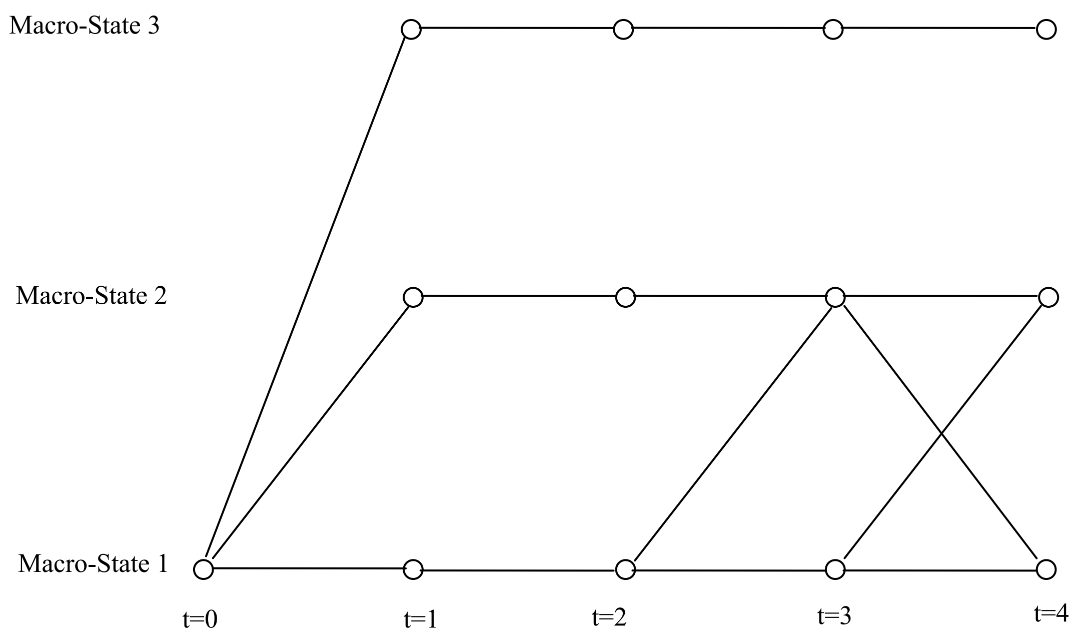

Figure 7.13 Captivating versus Non- Captivating Macro-States

Let us consider how the mathematics simplifies when we have two captivating (or absorbing) macro-states. Figure 7.14 demonstrates the branching tree and set of small Markov chains in this case.

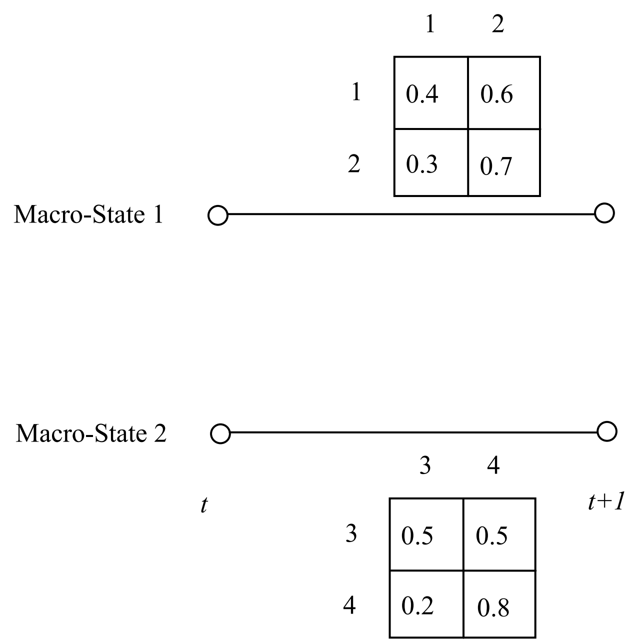

Figure 7.14 Branching Tree Representation of a Markov Chain with Captivating Macro-States

The corresponding overall Markov chain is demonstrated in Figure 7.15.

Next Period System State

Current Period System State	1	2	3	4
1	0.4	0.6	0	0
2	0.3	0.7	0	0
3	0	0	0.5	0.5
4	0	0	0.2	0.8

Markov Chain (Transition Probabilities)

Figure 7.15 Markov Chain for System States with Transient Macro-States

Observe that by using the branching structure under this captivating macro-states scenario, we automatically ignore the irrelevant regions of the overall Markov chain, which would otherwise have been included in the analysis. In other words, the expected value recursion for a macro-state that is captivating in a given period is simplified to:

$$EV_{i,k,t} = \max_{D_{i,k,k,t}} \left(\sum_{j \in k} pm_{i,j,k,k} \left[EP_{j,t} \left(D_{i,k,k,t} \right) + EV_{j,k,t+1} \right] \right)$$

This recursion equation assumes that the system will remain in the existing macro-state, and as such, the outer maximisation (relating to the choice of macro-state) has been removed. As such, the computational complexity of the expected value calculation is substantially reduced.

7.3 Possible Applications of the DADP Framework Methodology

The possible applications of the DADP framework presented in this chapter are very wide-ranging, both within dynamic programming applications in the field of energy market modelling and outside this field. In fact, the technique could be applied to any situation in which you have alternative dynamic programs that are valid under certain assumptions, and you do not yet know which one of those assumptions is valid.

In this section, we will consider just a few of these applications. In the following sections, we will consider the application of this framework to the problem of interest in this thesis, the optimal short-term offer strategy of an electricity generator, in much more detail.

7.3.1 Application 1: Long-Term Hydro Reservoir Management

Firstly, within the area of energy market modelling, this new framework could be applied to long-term (annual, for example) reservoir management, where the events that determine the macro-state of the system are based around longer time-frames. Consider one such macro-state, relating to whether the year that is being modelled is an El Niño year, which will affect hydrological reservoir inflows and thus has the potential to substantially change the operating strategy of a generator. Let us assume a simplified version of this model, where:

- Each stage is one month
- There is a value curve for water at each node defined over the reservoir level (the states of the DP) – as opposed to a single value at each node, as presented in Section 7.2. These would be calculated by the dynamic program.
- Within each month, there will be many decisions made, and many possible system states within each macro-state.
- There are three approaching months (Feb – April) in which we may or may not discover whether it is an El Niño year.

The branching diagram for the macro-states (ignoring the decisions and the system states) is shown in Figure 7.16, assuming that by the end of those three months, we will know one way or the other if it is such a year.

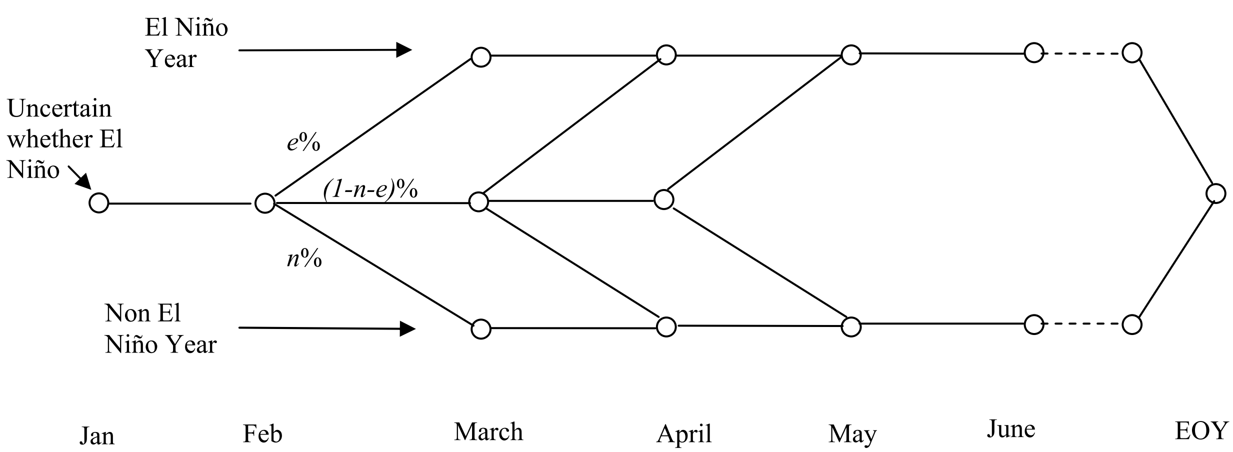

Figure 7.16 Branch Diagram for Long-Term Reservoir Management - El Niño Year

We can see that the El Niño macro-state is reflected by an captivating macro-state system, implying that once a year has been recognised as being or not being an El Niño year, it will not switch to the alternative. It will only revert to the common uncertain state at the end of the year/horizon (EOY). In February and March, there are expected probabilities that the year will be either recognised as an El Niño year (*e*%), recognised as not being an El Niño year (*n*%), and that we won't be able to tell at this point whether or not it is an El Niño year (*1-n-e*%). April is similar, but we assume that we will know by this time whether or not it is an El Niño year. As the DP works backwards through the horizon down each of the branches, the water value curves are combined probabilistically at each state to produce expectations over the likelihood of this weather phenomenon occurring.

If there is a major event such as the recognition of an El Niño year that could occur over the planning horizon in a case like this, it is beneficial to model the system in this manner, as it provides a more accurate representation of the problem situation and the cause of the uncertainties. For the example provided, compared with ignoring this structure, the number of nodes considered has effectively increased from 6 (one for each period – not shown) to 12 (one each for January and February, three each for March and April, and two each for May and June) over this 6 period horizon. This might imply that that although we are gaining a more accurate representation of the system, that the computational effort will increase substantially. However, it would be reasonable to assume that we are able to consider a smaller set of system states associated with each macro-state node under this branched representation, reducing the computational effort per node and as such, the overall computational effort may not be harmed, or in fact, may even be reduced. This is discussed further in Section 7.7.2.

7.3.2 Application 2: Managing a Gas Generator

Consider the owner gas-powered electricity generator, who must manage the amount of gas purchased and supplied along a pipeline, as well as the pressure of the gas in the pipeline, while simultaneously offering its generation into the electricity market. Both the input and output rates of gas are controllable by the decision maker, as it trades off storage costs and limitations with minimum/maximum rates of gas flow, and the value of saving gas to enable greater electricity generation in a later period.

Let us assume a simple version of this model, where:

- Each stage is one week
- Within each week, there will be an electricity generation decision made, and a stochastic distribution of price that may be received for that generation.
- At the start of each month (weeks 3 and 7 in Figure 7.17), the owner must decide whether to enter into a contract with an electricity market demand-side participant to sell a certain quantity of electricity at a fixed price each period for the remainder of the horizon. Note that for the first three weeks of the horizon shown there is no contract in place, and hence only a single macro-state.

The branching diagram for the macro-states (ignoring the lower-level operational decisions and the system states) is shown in Figure 7.17. The DP algorithm will work backwards through the horizon producing a value curve for gas availability at each node.

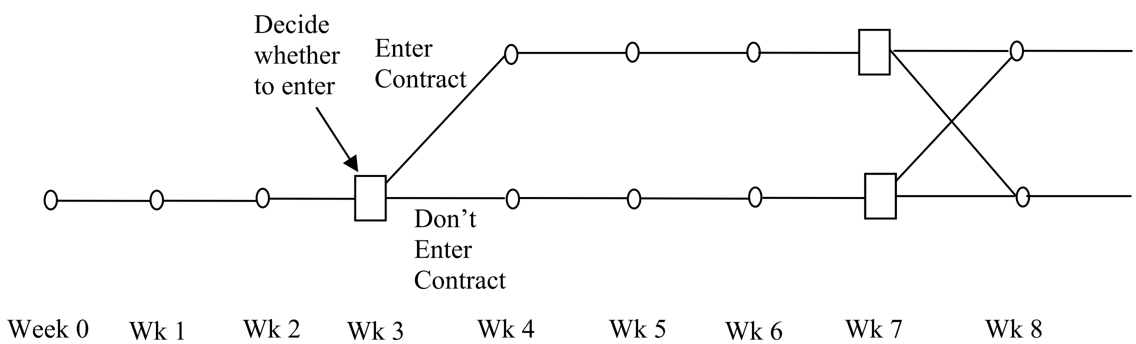

Figure 7.17 Branch Diagram for Gas Pipeline Management Problem – Entering a Contract

This structure could easily be extended to consider the possibility of an electricity market price shock at some point after the contract decision, which would affect the revenue from any non-contracted sale of electricity (use of gas). There are various possible causes of such a price shock, such as the outage of a major rival generation unit. Based on an x% chance that an event will occur to cause the electricity market prices will be significantly higher (a different macro-state) than normal from day 3 onwards, Figure 7.18 shows the appropriate branch diagram, for a subset of the horizon.

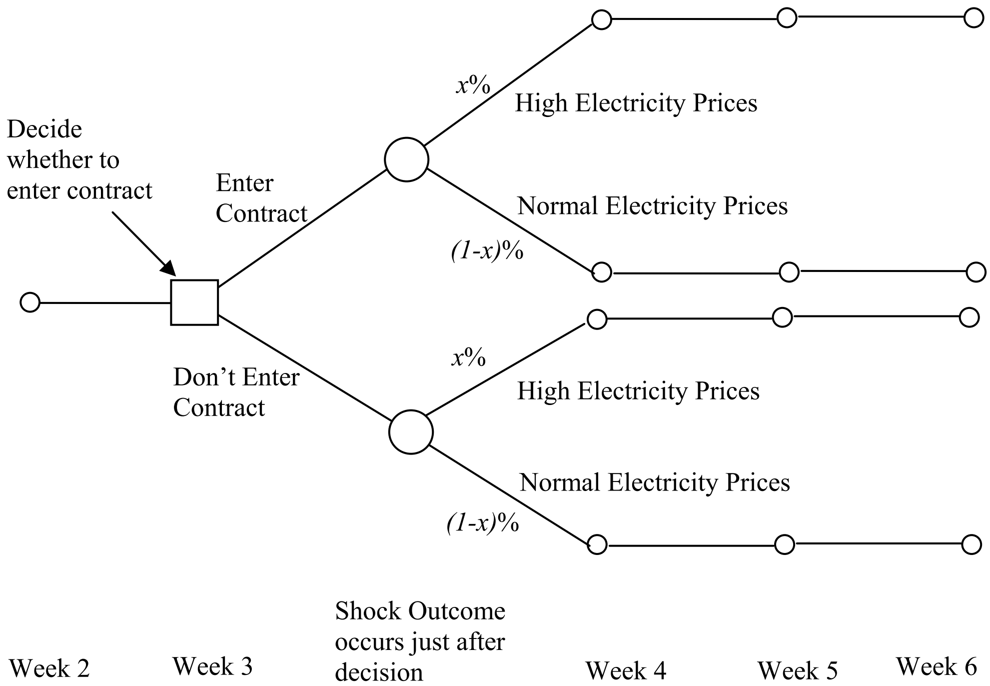

Figure 7.18 Expanded Branch Diagram for Gas Pipeline Management Problem

7.3.3 Application 3: Thermal Generation Unit Commitment Problem

The final application that we will discuss very briefly here relates to a thermal generation unit commitment problem under a deregulated market structure. The decision nodes would in each period relate to the decision of whether to switch the unit state (from on to off, or vice versa). The macro-states would therefore be the state of the unit, and could either take a very simple on/off form, or they could be more complex, and consider minimum up and down times, and shut-down and start-up phases. In many of the states of the more complex alternative, there would be no decision for the operator to make, as the behaviour/decision is implied by the state of the unit. In this problem, the DP states could represent fuel reservoir level (if it is a fuel-constrained thermal station), and/or previous price level or RD curve (if market correlation is being considered). Chapter 9 expands on this idea significantly, and presents the formulation for this problem.

7.4 Application of the DADP Framework to the Electricity Offer Optimisation Problem

Let us now consider the application of the DADP framework to the offer optimisation problem faced by hydro generators in an electricity market, and addressed in this thesis. In Section 6.6, the RT algorithm that was presented can be thought of as representing the simple situation where there is only one macro-state in each period, but many system states associated with this macro-state, representing all the possible residual demand curves that could possibly occur within the period. Typically, the real-world can be approximately represented by this type of branched structure, and thus, the system states (or RD curves) can be combined into groups that could only occur under certain macro-states. As such, there is value in applying the DADP framework to this problem. In the remainder of this chapter, we will demonstrate how this framework could be applied, and describe the benefits of such an implementation.

This section begins by briefly summarising the structure of the problem under a single macro-state structure, equivalent to the non-DADP framework presented in Section 6.6. The section goes on to consider how chance nodes and then decision nodes could be added on to this basic structure.

7.4.1 Single Macro-State Structure

The probabilities in the Markov chain on the left of Figure 7.19 represent the probability that if RD curve A has occurred in period t, then RD curve B will occur in period $(t+1)$, assuming just four possible RD curves per period. The UMS that the generator would face in period $(t+1)$, given that they had observed RD curve 3 in period t is shown on the right of this figure.

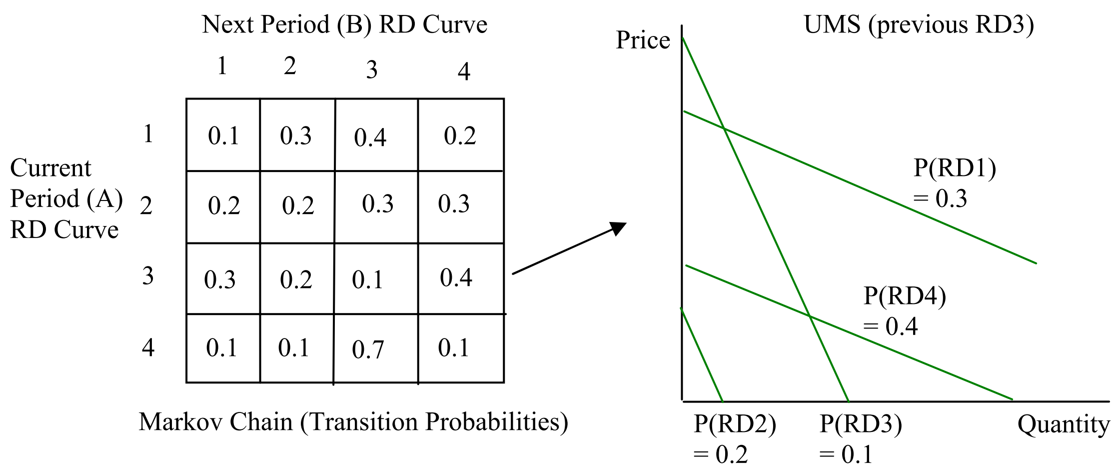

Figure 7.19 Markov Chain Interpretation

The overall process of the RT phase is shown in Figure 7.20, where there is a single Markov chain between the possible RD curves in each consecutive period, and the algorithm works backwards through the horizon (as indicated by the red arrow) determining the optimal decision (offer) under each combination of market and reservoir states, at each of these periods.

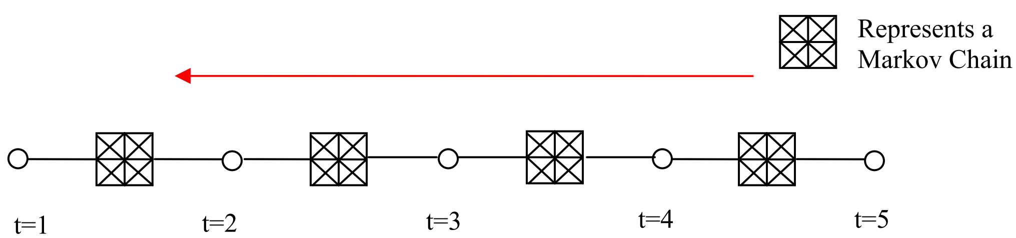

Figure 7.20 Current Market Uncertainty Structure

Under the DADP framework proposed in this chapter, the algorithm still works backwards through the horizon, but there is a separate Markov chain along each branch connecting macro-states in consecutive periods, as was demonstrated in Figure 7.5. The following two sub-sections of this chapter deal with chance and decision nodes, respectively.

7.4.2 Chance Nodes

Chance nodes within the offer optimisation problem scenario could be used to represent a wide range of possible events in the marketplace over which the generator of interest has no control. These could include:

- Weather Events[43]
- Interconnector binding status
- Rival breakdowns
- Initiation of rival offering strategies (aggressive or defensive, or maybe as a result of a change in the plant operator). These may be placed at the start of strategic periods, such as the start or end of peak and off-peak phases.
- General or rival-specific reservoir level indicators
- Demand shocks/realisation of errors in forecasted demand

Related to the last point, one possible idea for sourcing the RD curves to be used in each period under each probability branch macro-state is to base them purely on uncertainty in the rest-of-market supply curves, while assuming that the demand component of RD is fixed. This assumption is reasonable in a short-term model such as this, as good demand forecasting models are available in the literature (as reviewed, for example, in Fatai, Oxley, & Scrimgeour (2003)). If there were some potential external shock that could affect these forecasts significantly, then this could be covered through the use of the branch structure.

Figure 7.21 demonstrates a very simple example of the new structure, where in period 3 there is a chance that the HVDC link (an interconnector with a major effect on the residual situation faced by this generator) will become binding for the remainder of the short horizon (a captivating macro-state). Note that only a subset of the planning horizon

[43] While weather events will directly affect demand, they may also affect contractual positions of a generator. In other words, through their retail-side interests, they may be contracted to provide all electricity demanded by their customers at a fixed price. However, the level of this demand is stochastic, and highly subject to weather events.

is shown, and these macro-states are captivating, rather than absorbing. Also, we make no assumption with respect to the end-of-horizon value curves; they may either be different under each macro-state, or the same.

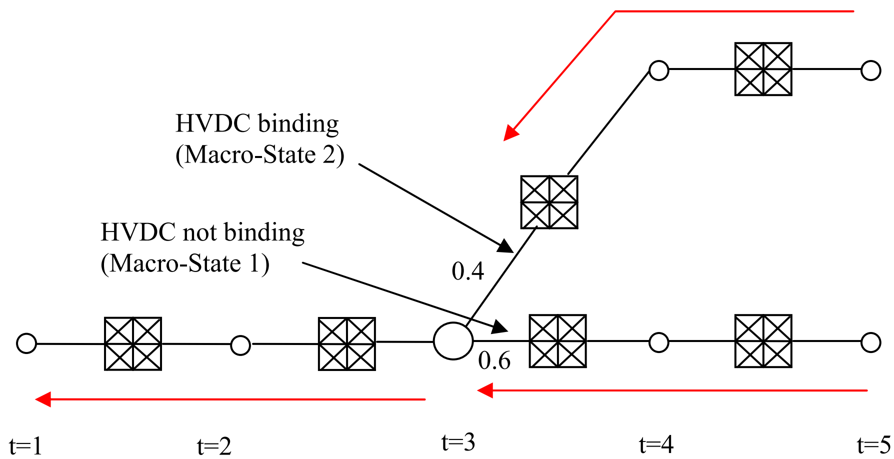

Figure 7.21 Probability Node Example

Mathematically, the key adjustment that needs to be made to the generic description of the DADP framework for chance nodes, in order to apply this framework to the offering problem, is to recognise that at each node, there is an expected value *curve* defining the total value of water for the remaining horizon from each level in the reservoir, as opposed to just a single expected value at these points. As such, the notation must be modified and expanded slightly.

$r =$	Reservoir level index, from 0 to $R_{\max}$, where $R_{\max}$ is the maximum reservoir level index.
$\theta_{i,k,l,r,t} =$	Offer in period *t*, given a starting reservoir level index *r* and given that the previous RD curve index was *i* (associated with macro-state *k*) and given that the decision maker has observed (chance node) or caused (decision node) a shift to macro-state *l* ($\theta^*_{i,k,l,r,t}$) is the optimal offer)
$Q_j\left(\theta_{i,k,l,r,t}\right) =$	Quantity dispatch that results from this offer, given that RD curve index *j* occurs

$EP_{j,t}\left(\theta_{i,k,l,r,t}\right) =$ Expected payoff to the generator within period t for the offer $\theta_{i,k,l,r,t}$, given that RD curve index j occurs

$EV_{i,k,r,t} =$ Expected value from the beginning of period t onwards, given previous RD curve index i, which is associated with macro-state k, and given a starting reservoir level of r

$EMV_{i,k,r,t} =$ Expected marginal value at the beginning of period t, given previous RD curve index i, which is associated with macro-state k, from a reservoir level of r. It is the expected marginal value curves that are used to construct the marginal opportunity cost (MOC) curves on which offers for the previous period are based.

$\text{infl}_t =$ Expected hydro inflow in period t [44]

The expected value for a given reservoir level at a particular node and given a previous RD curve (system state) is calculated using Equation 7.1.

$$EV_{i,k,r,t} = \sum_{l=1}^{NM_{t+1}} m_{k,l}\left[\max_{\theta_{i,k,l,r,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r,t}\right) + EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r,t}\right)+\text{infl}_t,R_{\max}\right),t+1}\right]\right)\right]$$

Equation 7.1 Expected Value Calculation for a Given State under a Chance Node

The expected marginal value curve is then calculated as a simple difference between consecutive points on the expected value curve, as shown in Equation 7.2.

[44] Note that it is straightforward to extend this to model stochastic inflows, rather than a simple expected inflow level for each period. The results presented in Chapter 8 are based on a model that considers stochastic inflows.

$$
\begin{aligned}
EMV_{i,k,r,t} = &\sum_{l=1}^{NM_{t+1}} m_{k,l}\left[\max_{\theta_{i,k,l,r+1,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r+1,t}\right)+EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r+1,t}\right)+\text{infl}_t,R_{\max}\right),t+1}\right]\right)\right]- \\
&\sum_{l=1}^{NM_{t+1}} m_{k,l}\left[\max_{\theta_{i,k,l,r,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r,t}\right)+EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r,t}\right)+\text{infl}_t,R_{\max}\right),t+1}\right]\right)\right] \\
= &\sum_{l=1}^{NM_{t+1}} m_{k,l}\left[\begin{array}{l}\max_{\theta_{i,k,l,r+1,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r+1,t}\right)+EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r+1,t}\right)+\text{infl}_t,R_{\max}\right),t+1}\right]\right)- \\ \max_{\theta_{i,k,l,r,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r,t}\right)+EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r,t}\right)+\text{infl}_t,R_{\max}\right),t+1}\right]\right)\end{array}\right]
\end{aligned}
$$

Equation 7.2 Expected Marginal Value Calculation under a Chance Node

Note that in practice, a *value* curve is produced first for all macro-states based on expected current and future values for the optimal offer from each reservoir level. A *marginal value* curve is then produced, calculated as the derivative of this value curve. The expected marginal value curve is calculated finally as a simple weighted average of these marginal value curves, based on the $m_{k,l}$ probabilities.

7.4.3 Decision Nodes

Decision nodes within the offer optimisation problem scenario could be used to represent a wide range of possible decisions over which the generator of interest has control, other than the specific offers themselves (which are the second level of decisions, made after the branch decision has been resolved). These could include:

- Unit start-ups or shut-downs
- Contracting Decisions
- Decisions regarding other offering units under the generator’s control
- General offering strategy decisions (e.g. aggressive versus conservative, etc)

With regard to this last point, this implies that decision nodes can effectively be used to model the expected reactions of rivals to changes in this generator’s behaviour. In other words, if we shift our strategy, then this changes our expectations about the supply curves that other generators may provide to the market and thus our expectations on the possible RD curves that we may face. A separate branch could be created, associated with each of the alternative set of expected rival reactions.

Figure 7.22 demonstrates a very simple example of the new structure, where there is a decision to be made in periods 3, 4 and 5 regarding whether the generator should embark on an aggressive (high-priced) or defensive (low-price) offering strategy, from the existing moderate strategy (neither are captivating macro-states – the generator is free to switch to the other strategy at any time[45]). The decisions would be defined by the generator itself and probably correspond to regions of the offering space through which offers under the given strategy are permitted to pass. For each state (as defined by reservoir level and previous RD curve) in period 3, the algorithm compares:

1) The payoff of the optimal aggressive offer within this period plus the expected future value down the upper branch, with
2) The payoff of the optimal defensive offer within this period plus the expected future value down the lower branch

The higher of these two values implies the decision to be made from this state, and provides the value from that point. Note that in this example, the macro-states are absorbing at the end of the horizon, and thus have a potentially different end-of-horizon value curve associated with each.

[45] Note that this example essentially assumes that rivals have no extended memory of previous behaviour, as we consider there to be no difference in a generator's market expectations once they reach a particular macro-state, based on the macro-state path that they followed to reach that point.

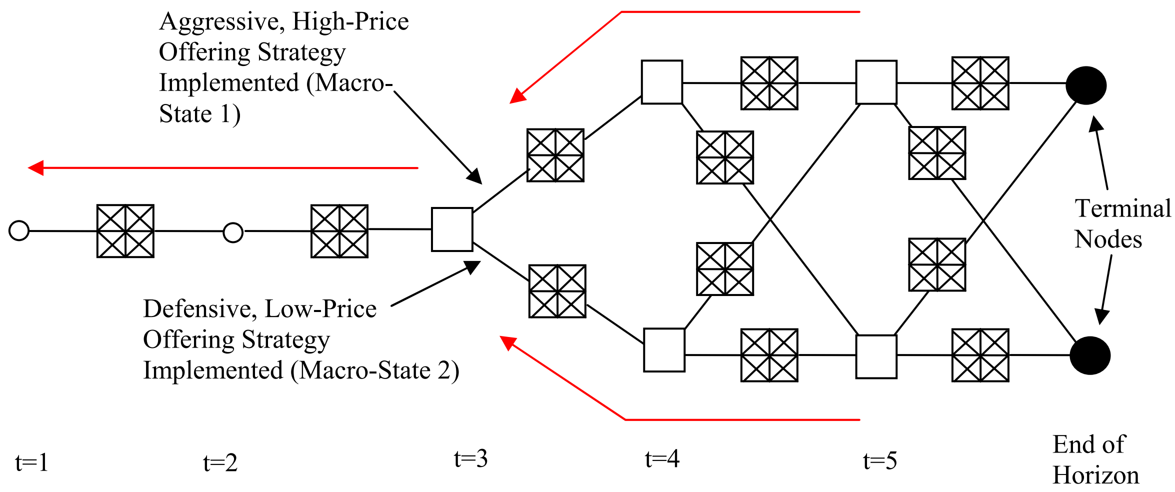

Figure 7.22 Decision Node Example

As with chance nodes, the key adjustment that needs to be made to the generic description of the DADP framework for decision nodes is to recognise that we have a value curve at each node, rather than just a single value. The new expected value for a given reservoir level at a particular node and given a previous RD curve (system state) is calculated using:

$$EV_{i,k,t} = \max_{l}\left(\max_{\theta_{i,k,l,r,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r,t}\right) + EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r,t}\right)+\mathrm{infl}_t,R_{\max}\right),t+1}\right]\right)\right)$$

This equation recognises that the decision the generator makes regarding macro-state transition will depend on both the reservoir level and the previous RD curve observed.

Note that in practice, a *value* curve is produced first for all macro-states based on expected current and future values for the optimal offer from each reservoir level. The expected value curve defined above is then produced at the source node of these macro-states, by taking the maximum value down each of the possible macro-states that could be chosen. Finally an expected marginal value curve is produced as a simple difference between consecutive points on the expected value curve, as shown in Equation 7.3.

$$EMV_{i,k,r,t} = \max_{l}\left[\max_{\theta_{i,k,l,r+1,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r+1,t}\right)+EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r+1,t}\right)+\mathrm{infl}_t, R_{\max}\right),t+1}\right]\right)\right]-$$
$$\max_{l}\left[\max_{\theta_{i,k,l,r,t}}\left(\sum_{j\in l} pm_{i,j,k,l}\left[EP_{j,t}\left(\theta_{i,k,l,r,t}\right)+EV_{j,l,\min\left(r-Q_j\left(\theta_{i,k,l,r,t}\right)+\mathrm{infl}_t, R_{\max}\right),t+1}\right]\right)\right]$$

Equation 7.3 Expected Marginal Value Calculation under a Chance Node

It is this expected marginal value curve from which the marginal opportunity cost curves are produced for the previous period, on which offers are based.

7.5 Comparison of the DADP Framework with Alternative Approaches

In this section, we consider the alternatives with respect to representing and dealing with a system that inherently contains a branching type structure and can thus be identified at any given point in time as being in one of a limited number of macro-states. Clearly, the two basic options are to consider this branched structure explicitly, or to ignore it and apply an approximation. The DADP approach that we have suggested in this chapter considers the structure explicitly, and it is important to recognise that this explicit representation could also be achieved through the use of a (possibly very large) Markovian Decision Process (MDP) matrix, covering all periods in the horizon. The main conceptual differences between this Structured MDP, and a basic MDP representation would be that the latter assumes that transfer probabilities and decision rules are constant over time, and that the system loop through the same set of possible states indefinitely. This means that once the decisions for all states have been determined, so that combined with the uncertain system outcomes, they define transition probabilities between any possible pair of system state, the representation simplifies from a MDP to a (possibly very large) stationary Markov Chain in which all states are recurrent. These stationary/recurrent assumptions are not made with the DADP representation, and the corresponding MDP, and eventual optimal Markov Chain, would be non-stationary, and include both transient and (possibly) recurrent states and sub-chains, along with what we have termed “captivating” sub-chains. Techniques exist to analyse Markov Chains with such structures, albeit in a somewhat ad hoc fashion, and one could imagine applying such techniques, iteratively, to determine the optimal decision strategy for the MDP, and hence the optimal Markov Chain representation. But the point is that, at least if they are to be efficient, such techniques need to account for and exploit the specific structure of the MDP and, in the limit, the hybrid DADP technique we have proposed is simply the logical expression of such an efficiently structured technique.

Thus, there is no point discussing a "comparison" between our DADP approach and a large structured MDP representation. The former is merely a practical way of solving the latter.[46] The real issue is whether there is anything to be gained by representing and analysing whatever natural structure may exist in this way, or more exactly whether the gains to be made by attempting such a representation can be justified in terms of the increased effort, or whether, perhaps better value might be delivered by, say, adopting a more fine-grained, but unstructured, representation. Thus, for example, an Unstructured stationary Markov Chain could be constructed to represent the uncertain system state transitions over the horizon, considering the same overall set of possible system (micro-) states as the Structured MDP, but ignoring any macro-state structure which may exist. And such a representation could be refined by increasing the number of states modelled, and/or allowing transition probabilities to vary over time. Such an approach would be consistent with both the algorithm developed in Chapter 6 and the R&A algorithm.

In this section we will compare this type of unstructured MDP approach with the structured MDP/DADP framework set forward in this chapter, with respect to the representation of the real-world structure, data construction and data management issues, and the complexity of the algorithms needed to represent these approaches.

[46] Note that, in all of our examples, the MDP is not stationary, because load levels, and hence generation, vary over time. The underlying Markov Chain describing market state transitions has been assumed to be stationary, though. But, while this simplifies the notation and reduces the data requirements, neither the R&A algorithm, nor any of ours, actually require this assumption.

Representation of Real-World Structure and Data Construction

As described above, a Structured MDP representation would be capable of the same structural representation of the system as could be achieved by the DADP Framework. On the other hand, by ignoring the specific structure, an Unstructured MDP would result in a different, and more detailed data collection and construction process than the Structured MDP.

There are two possible ways in which this data construction could occur: bottom-up or top-down. A bottom-up approach would take information about all possible macro-states of the market and determine what the possible RD curves might look like in each of these macro-states. A top-down approach would use historical RD curve data to forecast an overall set of possible RD curves for each period of the horizon. If the former approach were applied, then the RD curves would automatically be grouped according to macro-states, and thus no extra effort is required in order to get them into this form. However, if the latter approach is applied, then further effort is needed to break these RD curves down into groups associated with each macro-state, using a method such as that presented in Hernaez et al. (2004). In this case, a more accurate real-world representation would come at the cost of increased data construction effort. With respect to estimation of transfer probabilities, the number of probabilities to be estimated has the potential to be substantially lower under the Structured MDP and DADP Framework approaches, compared with the Unstructured MDP, given the same total number of possible RD curves. The reason for this is that the structured methods eliminate consideration of a large number of possible RD curve transfers between consecutive periods.

It is our belief that, if the real-world can reasonably be represented by a branched structure, then the quality of the solution achieved can, in principle, be improved by recognising this structure, as in the Structured DADP Framework. However, if the underlying branched structure is weak or non-existent, then the unstructured algorithm presented in Chapter 6 may well be sufficient. Therefore, the decision as to the choice of approach in any given scenario will be dependent on the strength of the underlying branched structure.

Data Management

We have noted earlier that, while the underlying Markov Chain describing market state transitions has been assumed to be stationary, neither the R&A algorithm, nor any of ours, actually require this assumption. Thus there is no reason why those algorithms can not cope with situations where probabilities change with time, per se. Conversely, we have also noted that the Markov chain probabilities should not change in response to actual dispatch results, as these outcomes are already built into the Markov chain. However, it is also possible that the Markov chain should change in response to changes in expected market supply and demand, if and when they occur. These changes may be exogenous changes that can be observed by a generator (such as break-downs, etc), or they may be changes to our expectations of events occurring as a result of observing pre-dispatch market outcomes.

Under the DADP framework, chance nodes and macro-states could be used to represent the probabilities of events such as these occurring. When new probabilistic information becomes available, then this would involve changing just the probabilities in the high-level macro-state transition Markov chain (the $m_{k,l}$ terms), as shown in Figure 7.4, and hence the probability of any particular lower-level Markov chain occurring. But this would not necessarily imply a need to change the probabilities on the lower-level system state (or RD curve) transition Markov chains, per se. Computationally, the effect of this would be simply to change the calculation of the EV curve at each node, as explained in Section 7.4.2. As such, this would be very simple to implement, whether the new probabilities are produced by an outside decision (manual input from the user), or using some sort of Bayesian updating process.

On the other hand, if an Unstructured SDP were used, then the branch probabilities of the underlying system would be built into the system state transition probability matrices, and thus it is likely that many elements of these matrices would need to be changed. This would mean that the UMSs on which the offers are based would change, and so a whole new set of OCFs would need to be produced. Computationally, this is far more expensive than the effect on the same problem considered under the Structured MDP or the DADP framework.

Also, as we move forward through the planning horizon, we are likely to be able to update our expectations of macro-state transition probabilities as a result of pre-dispatch outcomes and other information learned from the market far more easily under the Structured MDP or DADP framework approaches than under an Unstructured MDP approach. No simple analogous solution to this event would be available under an Unstructured SDP approach. In the limit, in the case of an event occurring which eliminates the whole set of system states associated with a given macro-state, then under the Structured MDP or DADP Framework approaches, this could be dealt with very simply by eliminating that set of system states from consideration.

Algorithmic Complexity

As we progress through the planning horizon, we will actually be able to observe the occurrence (or non-occurrence) of the events to which the chance nodes relate, and will know the decisions that we have made at past decision node. We have already noted the way in which this may simplify the task of updating data, but it also means that if the generator finds itself in a macro-state that has no way of transitioning to a particular macro-state in some future period, then there is no need to consider the decisions that would be made under this unreachable macro-state. Under a static end-of-horizon point, this would mean that the computational time of the RT phase would reduce dramatically as time progresses, for the DADP algorithm, but not for the non-branched algorithm (which, as explained earlier, is unable to disregard this set of system states). A static end-of-horizon is a reasonable assumption in electricity markets such as that of New Zealand, as the final period for which offers must be provided, changes only once every 24 hours. Of course, if a rolling horizon was used, then in each period, a whole new set of macro-states would open up for the new period considered, and therefore this computational benefit would be eliminated.

The remainder of this section compares the likely computational efficiency of the DADP Framework, with the non-branched algorithm presented in Chapter 6 (which is equivalent to the Unstructured MDP approach).

Under the DADP framework, the decisions at each node consider only the RD curves that could occur under the associated macro-state, which reflects the overall position that the market is in at that point. Under the non-branched algorithm these macro-states are not distinguished, and hence all the possible RD curves under each of these macro-states must be considered simultaneously, rather than separately. This simplification means that computational time increases much more slowly under the DADP algorithm as the total number of RD curves (or system states) considered increases, in both phases of the two-phase algorithms. Alternatively stated, for the same computational resource, this method is able to consider a more detailed representation of the possible residual scenarios that might be faced. More detailed computational complexity comparisons are presented in Sections 7.6.2 and 7.7.2.

7.6 Offer Optimisation Problem: Pre-Processing Phase under DADP Framework

In this section, we present a flow diagram representing the PP Phase algorithm under the DADP framework and identify the computational complexity of the algorithm.

7.6.1 The Pre-Processing Phase Algorithm under Branch Structure

Figure 7.23 presents the PP Phase at a very high-level, where the additional components of the algorithm are highlighted. In the non-branched algorithm, we defined a term p as the total number of possible RD curves per period. Assuming that under the DADP framework, all macro-states have the same number of possible RD curves, we define:

b = Number of macro-states per period

P = Number of RD curves per macro-state

Note that:

$$p = bP$$

Therefore, although there is now an extra loop to the algorithm, the total number of times that we iterate through the core components is unchanged. It is the computational complexity of those components themselves that will determine the change in computational complexity of the branched algorithm as a whole. Note that the PP phase is the same regardless of whether the DADP structure contains decision or chance nodes.

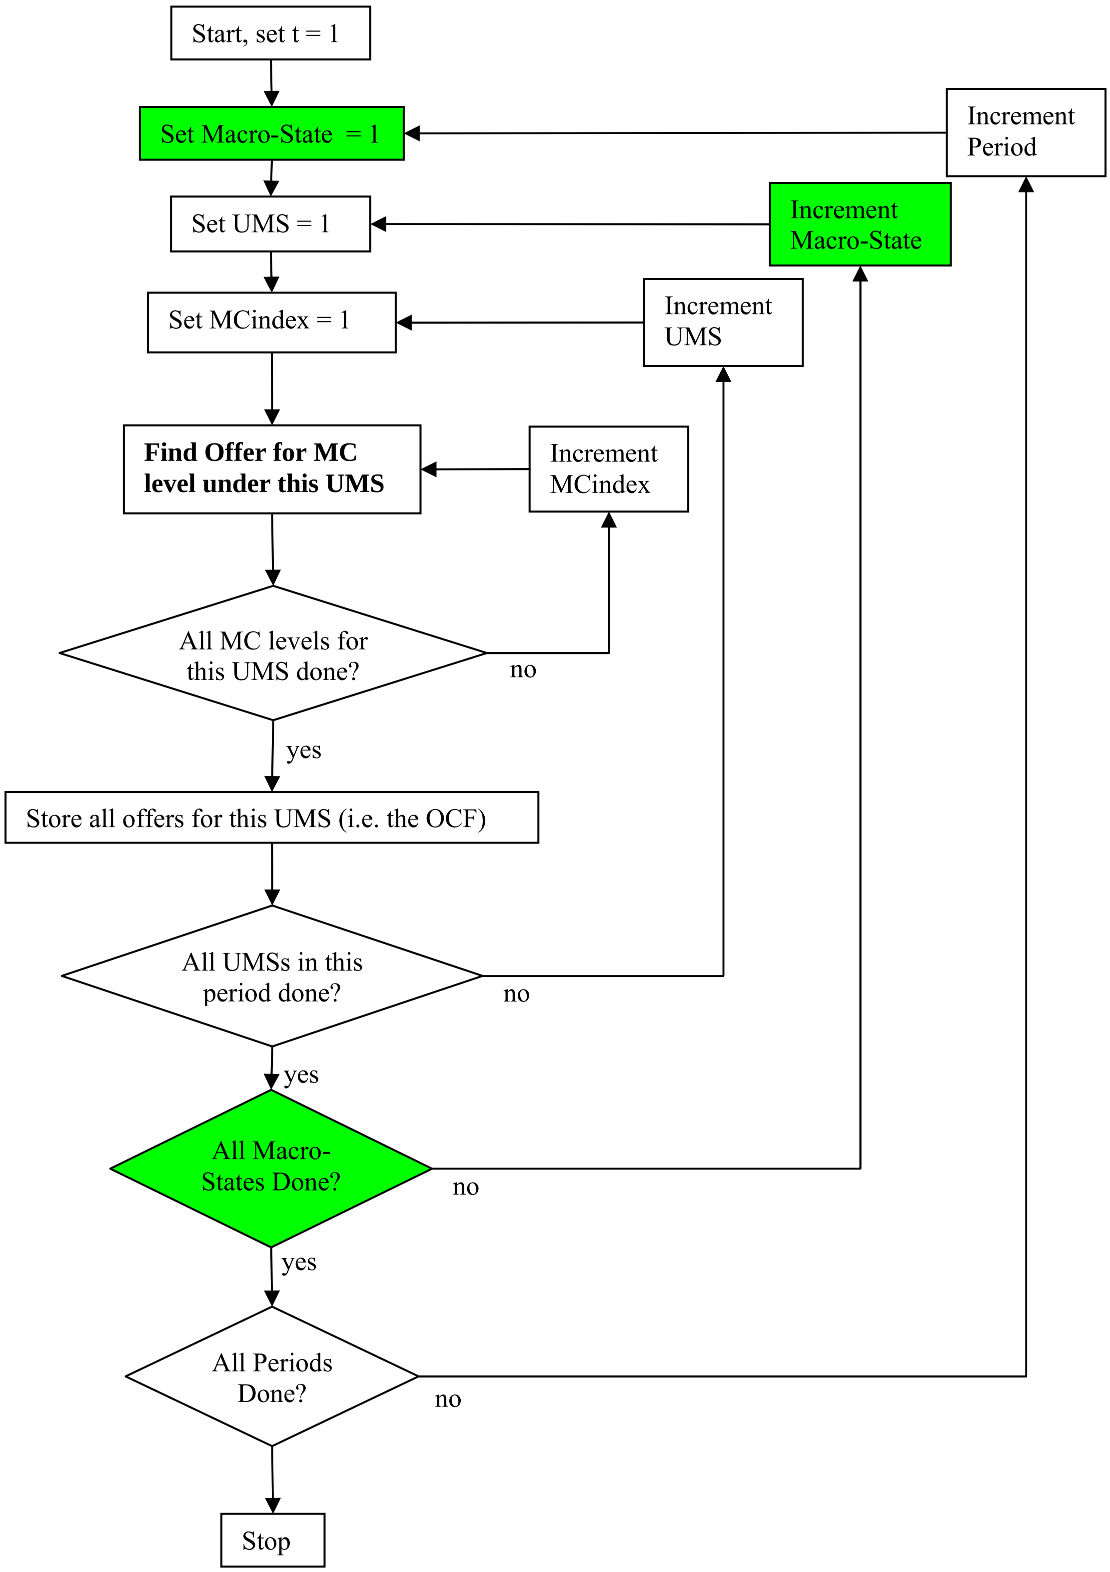

Figure 7.23 PP Phase Flow Diagram

7.6.2 Computational Complexity of the Pre-Processing Phase under Branch Structure

From the source code for the PP Phase under the DADP framework, we can show that the new computational complexity can be approximately summarised by Equation 7.4.

$$11tb + 8Ptb + 14Ptbd + 2tbP^2 + 0.667tbP^2d^2 + 13P^2tbMd + 6PtbMd + 8PtbM + 0.875P^2tbMd^2$$

Equation 7.4 Computational Complexity of the Branched PP Phase in Full

Simplifying this down to the terms of the largest orders of magnitude, we can see that the algorithm is of order of the complexity shown in Equation 7.5.

$$tbP^2Md^2$$

Equation 7.5 Order of Computational Complexity for Branched PP Phase

Substitute $p = bP$ to give Equation 7.6.

$$tbP^2Md^2 = \frac{tp^2Md^2}{b}$$

Equation 7.6 Comparable Computational Complexity for Branched PP Phase

In Section 7.5, we discussed the fact that if the real world can be represented by a branched structure, we can choose to recognise this structure (through the DADP Framework), or ignore it, and apply an approximation, as represented by the non-branched algorithm presented in Chapter 6. Let us now compare the computational requirements of the two alternative approaches. Recall the simplified version of Equation 6.2 from Chapter 6:

$$tp^2Md^2$$

Equation 6.2 Order of Computational Complexity for PP Phase

Note that the order of complexity for the PP Phase algorithm under the DADP framework is $\frac{1}{b}$ that of the non-branched algorithm[47]. Consider a simple numerical example with 24 periods, 30 fixed MC levels, 20 possible dispatch levels, starting with a single branch of 10 RD curves and a computational time of x, and then increasing to consider 40 RD curves (either split equally over four branches, or grouped all together for the branched and non-branched algorithms respectively). From the equations above we would expect the computational time for the non-branched algorithm would increase to $16x$, while the computational time for the branched algorithm should increase to only $4x$.

In other words, if we hold the number of RD curves under each macro-state constant, the computational time for the PP phase under the DADP framework will increase only linearly in the number of macro-states that we have. However, under the non-branched algorithm, computational time would increase quadratically with respect to the total number of RD curves.

7.7 Offer Optimisation Problem Real-Time Phase under DADP Framework

In this section, we present a flow diagram representing the RT Phase algorithm under the DADP framework and identify the computational complexity of the algorithm.

[47] Note that this factor is also achieved if the full complexity equations are compared, as the coefficients of the terms in the two equations are very similar.

7.7.1 The Real-Time Phase Algorithm under Branch Structure

Figure 7.24 presents the RT Phase algorithms at a very high-level, where the additional components of the algorithm are highlighted (assuming the Value Curve approach presented in Section 6.6.1, rather than the Direct MV Curve approach – where the adjustments to the algorithm would be the same as presented in Figure 6.11).

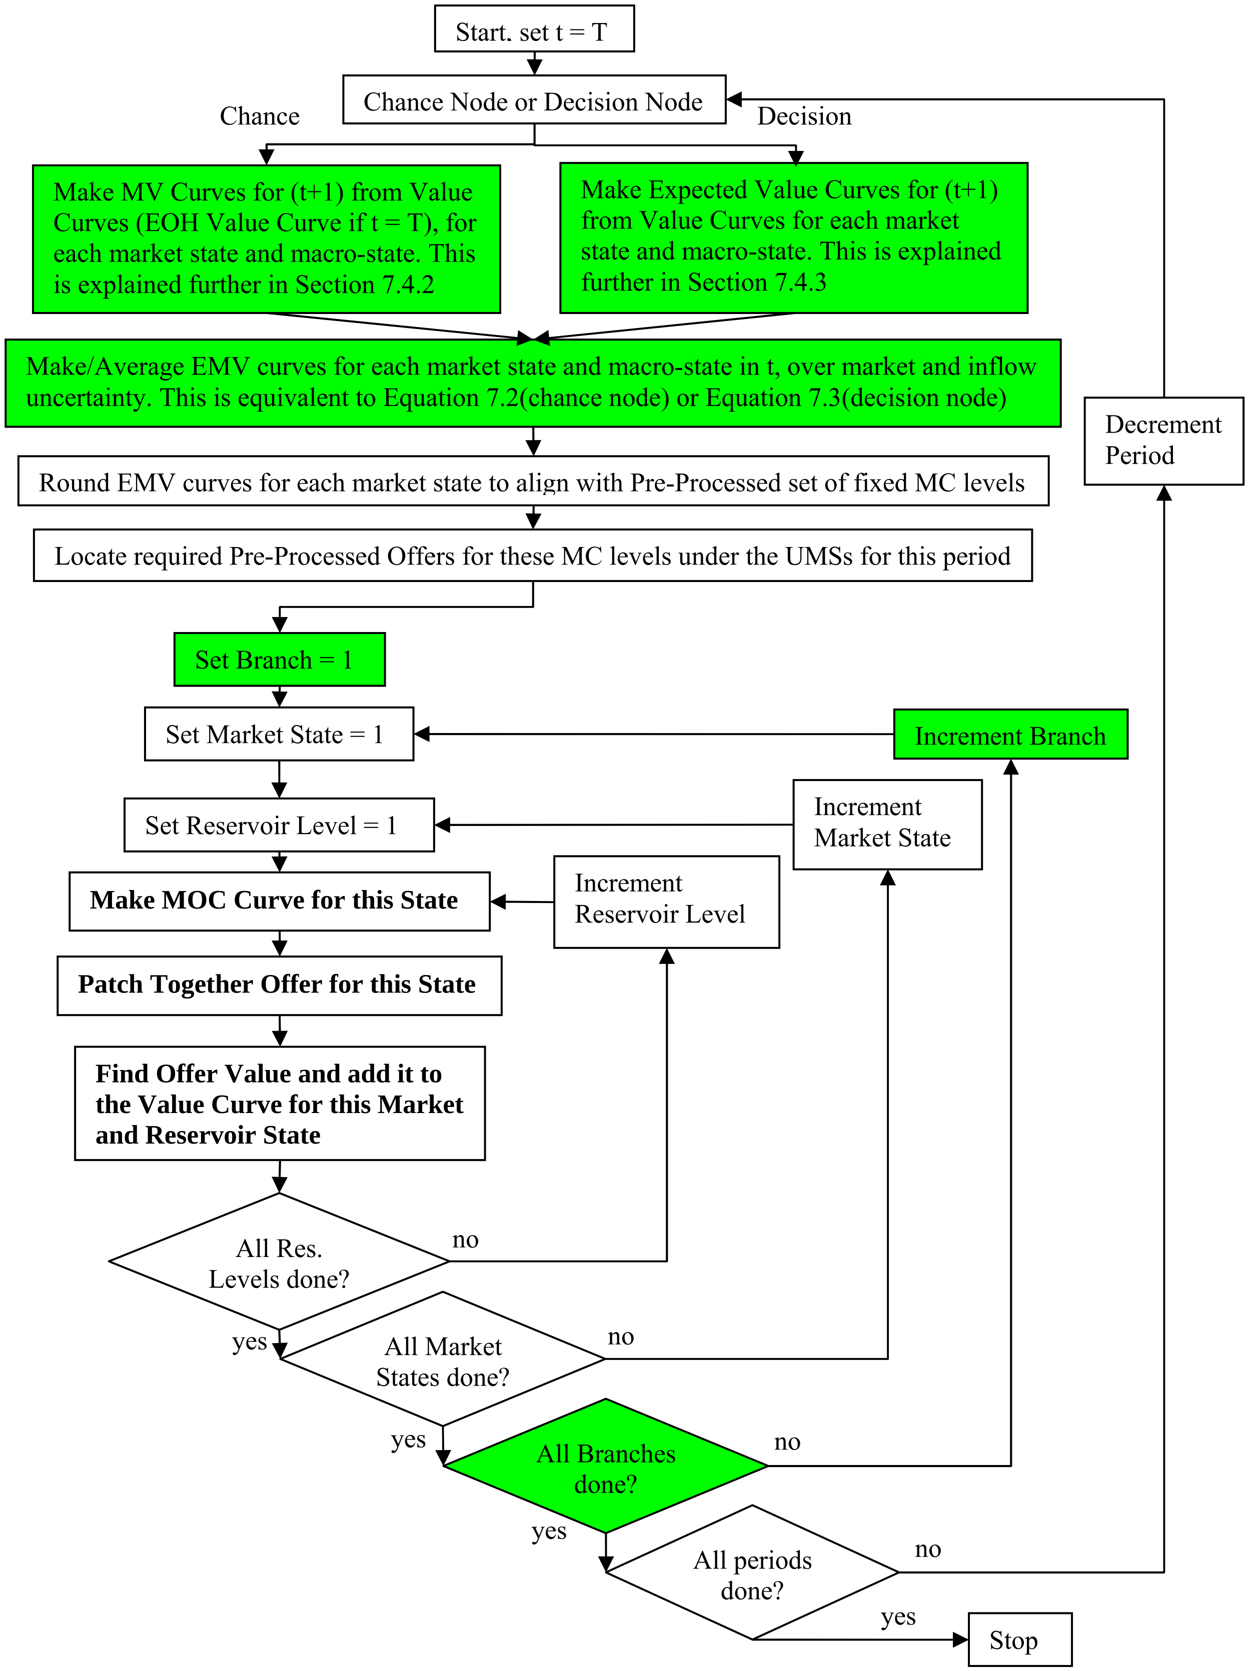

Figure 7.24 RT Phase Flow Diagram (Value Curve Approach)

7.7.2 Computational Complexity of the Real-Time Phase under Branch Structure

From the source code for the RT Phase for the DADP framework (assuming chance nodes), we can show that the new computational complexity for the Value Curve and Direct MV curve approaches can be approximately summarised by Equation 7.7 and Equation 7.8, respectively.

$$0.1tb\begin{bmatrix} 100+160P+10b+420Pr+40Pd+30P^2+10bP+230P^2r \\ +65PM+30Pbr+20Pr\inf+28PMd+130Prd \end{bmatrix}$$

Equation 7.7 Computational Complexity of the Branched RT Phase Value Curve Approach in Full

$$0.1tb\begin{bmatrix} 100+160P+10b+180Pr+40Pd+30P^2+10bP+110P^2r \\ +65PM+30Pbr+20Pri+28PMd+130Prd \end{bmatrix}$$

Equation 7.8 Computational Complexity of the Branched RT Phase Direct MV Curve Approach in Full

Simplifying these down to the terms of the largest orders of magnitude, we can see that both algorithms are of order of complexity shown in Equation 7.9.

$$tbP^2r+tb^2Pr+tbPMd+tbPrd$$

Equation 7.9 Order of Computational Complexity for Branched RT Phase

Substitute $p = bP$ to give Equation 7.10.

$$tpr(P+b+d)+tpMd$$

Equation 7.10 Comparable Computational Complexity for Branched RT Phase

Recall again that in Section 7.5, we discussed the fact that if the real world can be represented by a branched structure, we can choose to recognise this structure (through the DADP Framework), or ignore it, and apply an approximation, as represented by the non-branched algorithm presented in Chapter 6. We can compare the computational requirements for the RT phase of the two alternative approaches. Recall the simplified version of Equation 6.3 from Chapter 6.

$$tpr(p+d)+tpMd$$

Equation 6.3 Order of Computational Complexity for RT Phase Value Curve Approach

Note that the orders of complexity for the RT phase under the DADP framework and under the non-branched approach are quite similar. The only difference (other than coefficients and minor terms) is that one of the components of the complexity order expression has simplified from p to $(P+b)$. Generally speaking, $(P+b) \leq p = bP$. For example, if there were four macro-states of ten RD curves each, then p = 40, while $(P+b)$ = 14.

As we discussed in Section 6.6, the majority of time spent in the RT Phase can be attributed to the process of finding the value of the proposed offers. Under each UMS, for each reservoir level, the intersection between the proposed offer and each of the possible RD curves for that UMS is found. If there are non-zero transition probabilities between all possible macro-states between two consecutive periods, this means that the number of intersections that need to be found will be unchanged. However, if there are some zero transition probabilities, this means that we do not need to consider a UMS along the associated branch, thus reducing computational effort. This was demonstrated in more detail in Section 7.2.5.

7.8 Summary and Conclusions

In this chapter, we have presented a new approach to dynamic programming, which incorporates a decision analysis or branching tree type structure to the possible macro-states that a system can be in at any point in time. We have presented this DADP framework in a very general manner and suggested that it has a wide range of possible applications, both within the electricity market and beyond. We have then presented the application of this framework to the main issue of this thesis, electricity market offer optimisation strategies. The key benefit of this new technique is that it potentially enables much more accurate representation of the uncertainty and decision making processes that occur in the real world. In addition, with respect to the application in the offer optimisation problem, we have shown that when there are only chance nodes, there are additional computational benefits compared with an alternate approximate algorithm, as presented in Chapter 6.

The next chapter presents results of experiments on both two-phase algorithms (value curve and direct MV curve), with chance nodes, compared with the R&A algorithm. We consider only chance nodes for these comparisons because the R&A algorithm cannot deal with the complexity related to the future impact of current decisions.

Chapter 8

EXPERIMENTAL DESIGN AND RESULTS

8.1 Introduction

In the preceding chapters of this thesis, we have proposed two two-phase electricity generator offer construction algorithms and hypothesised that they will produce similar results to the model presented in Rajaraman & Alvarado (2003) with respect to expected payoffs over the planning horizon, but with substantially reduced computational requirements. The reason for these hypotheses was that they are essentially three different approaches to the same problem, but the algorithms have vastly different complexities. In this chapter we describe an experimental design to test these hypotheses and to investigate the value of including the various complexities considered within these models, and then present the results of the experiments.

8.2 Experimental Design

There is a need for a formal experimental design to test the hypotheses on the relative quality and efficiency of the various algorithms considered in this chapter, for two reasons:

1) We want to explore how the values of various parameters that describe a particular problem scenario's size or nature affect the two measures of algorithm performance (algorithm speed and solution quality).
2) The conclusions need to be justified statistically.

Here, we present a design based on suggestions in Rardin & Uzsoy (2001). Although the paper is targeted at experimental design for heuristic optimisation algorithms, the techniques are equally applicable to the DP optimisation approach presented in this thesis. The terminology used throughout this chapter is that proposed in Rardin & Uzsoy (2001).

8.2.1 Problem Parameters

There are seven parameters that describe the size and nature of a particular problem scenario. These are the:

1) Number of periods in the horizon
2) Number of reservoir levels
3) Number of possible dispatch levels
4) Number of residual demand curves per period
5) Number of possible inflow levels in each period

6) Type of transition probability structure to have (i.e. whether the Markov chain should be highly or loosely correlated).
7) Type of residual demand curve structure to have (i.e. nominal data relating to the type of shapes that the residual demand curves have).

The first five parameters are numeric and thus describe the size of the problem, whereas the last two are qualitative and describe the nature of the correlation and RD curve structure used in the problem.

8.2.2 Outline of Experimental Design

Unfortunately, there is no library of problems available for testing the performance of the algorithms presented in this thesis. As such, we must design a set of problems with various possible characteristics, in order to cover a broad range of the *problem domain*. Therefore, we use the following definition:

Instance: a given set of values for the seven model parameters.

An *instance*, however, does not fully define a problem 'scenario', as the transition probabilities between macro-states, the exact form and shape of the residual demand curves, the Markov chains indicating correlation between market outcomes in consecutive periods, and the end-of-horizon value curve are randomly constructed within the general restrictions implied by the values of the last two parameters. The specific 'scenarios' tested are therefore referred to as *replicates*. The definition we will use is:

Replicate: an actual set of model inputs defining all facets of the scenario faced. We could produce infinite replicates for a given instance if we desired, based on different random number seeds.

Our specific design is the commonly used *full-factorial* approach, where all algorithms are run on *replicates* representing all *instances*. In particular, we submit the *same* replicates to all algorithms, utilising an approach known as *blocking*.

Number of Replicates

Rardin & Uzsoy (2001) states that at least three replicates sharing each parameter combination are needed to get a reasonable idea of algorithm robustness. We have chosen to test five replicates for each instance, where possible[48], in order to trade off between increased confidence in algorithm robustness and computational resource restrictions.

Model Process

For a single replicate and single algorithm implementation, successive trials can produce quite different outcomes because of the impact of random choices along the way. As such, several *runs* need to be made, with different random number seeds controlling the evolution of computation, to get a sense of the robustness of the procedure (Rardin & Uzsoy (2001)).

To evaluate a given algorithm for a replicate we carry out the following:

1) An optimisation is performed, which determines the optimal offer to provide for each situation throughout the horizon (market states, reservoir level state, etc).
2) A simulation is performed, consisting of 1000 runs[49], where each run is a randomly created series of market and inflow outcomes for the planning horizon. Again using *blocking*, the same series of outcomes is applied to each of the

[48] Except for two of the algorithms tested, where we used just a single replicate. The reasoning behind this is explained in Section 8.4 and Section 8.5.
[49] We have chosen 1000 runs as it is sufficiently large such that the percentage deviations from the optimal payoff under perfect information for each of the iterations forms an approximately normal distribution.

algorithms, and the expected value for each algorithm under the given *replicate* is the mean of the payoffs achieved under the 1000 runs.

Simulation

Rather than the real-time dynamic program expected payoff, a simulation is used to measure the solution quality performance of the various algorithms, for two reasons.

1) Simulating these results gives us a distribution of payoffs that could occur under various demand and inflow scenarios, rather than just a single value, and hence we are able to analyse the results in a more in-depth manner. Such a distribution is an important feature of risk management within a generation company, as it enables them to identify worst case scenarios, variation of possible outcomes, etc.
2) We want to be able to determine the relative optimality of our results for different modelling assumptions. Because of the stochastic nature of the results, we need to have actual scenarios of market and inflow outcomes in order to determine optimal values and hence relative optimality.

Of course, the expected value from the dynamic program is equal to the long-run expected value from the simulation.

Algorithm Performance Measures

There are two measures of algorithm performance that must be considered. These are:

1) Algorithm speed
2) Solution quality

<u>Measures of Algorithm Speed</u>

The main comparisons made throughout this chapter on algorithm speed will be based on the solving time required in the Real-Time phases of the various algorithms proposed. This is the fairest comparison, as it is what determines whether a particular algorithm could be used in a real-life generator offering situation. The average Pre-Processing phase times[50] will also be reported in this chapter but, as explained in Section 6.4, these need only be performed once for the whole horizon (rather than each period), and can be completed in advance.

Note that all model runs and simulations are performed on an Intel Pentium 4 3.20GHz machine, with 1 GB of RAM.

Measures of Solution Quality

As stated earlier, the value taken initially to represent the performance of a particular algorithm on a particular replicate is the mean of the expected payoffs over 1000 simulation runs. As each of these simulation runs represents a deterministic set of market and inflow outcomes over the horizon, the optimal payoff that could have been achieved under each of these paths under this perfect information (or hindsight) scenario can be determined. In order to measure a solution quality we therefore use the mean percentage difference between the algorithm performance and the optimal payoff. This is a preferable measure for two reasons:

1) It standardises results, and as such, overall results will not be affected by the *magnitude* of the profit that can be achieved under each replicate.
2) It provides an absolute measure of performance, enabling a comparison with an optimal solution, rather than just a relative measure of performance against alternative algorithms.

[50] Note that the PP phase is identical for both of the two-phase approaches.

Algorithm Analysis and Comparison Methods

In order to determine the statistical significance of various parameters (including the algorithm used) on the solution quality, we have followed a two-stage process, using the statistical package MINITAB. The two stagers are:

1) Performed a General Linear Model (GLM) ANOVA on the datasets, including the approach and all input parameters individually, in addition to interaction variables for the combination of each of the input parameters with the approach[51].
2) Performed Bonferroni Multiple Comparison Significance Tests and produced Bonferroni Confidence Intervals on all individual and interaction ANOVA variables which appeared as significant in the GLM ANOVA.

When the GLM ANOVA has identified a factor (such as the approach used) as significant, we can not be sure[52] whether this is caused by:

a) A single value of that factor producing a significantly different performance measure. For example, computational time being significantly less under a particular technique A than under the other techniques B and C, which are not statistically significantly different from one another.

Or

b) Multiple values of that factor producing a significantly different performance measure. For example, computational time being significantly less under a particular technique A than under another technique B, which, which in turn is significantly less than the computational time under the remaining technique C.

[51] Note that for each run of the GLM ANOVA in this chapter, I removed the insignificant factors and re-solved the model, but there was no change to the significance of the remaining factors.
[52] Harraway (1997)

We have performed a GLM ANOVA rather than a Balanced ANOVA because under a GLM ANOVA, it is possible to perform multiple comparison tests (such as Tukey or Bonferroni), which are not possible under the Balanced Model. We need these multiple comparison tests in order to distinguish between the two causes of significant factors (a and b) identified above, while ensuring that the probability of incorrectly rejecting the null hypothesis for *any* of the comparisons in the family does not exceed α (Harraway (1997)).

We have chosen to use Bonferroni Multiple Comparison Tests (as opposed to alternatives such as Tukey, etc) in order to produce the most conservative results regarding the significance of variables. Conservatism, in this context, means that the true family-wise error rate is less than the stated one. We produce both Bonferroni Significance Tests and Confidence Intervals because the significance tests easily show whether a factor (or an interaction) is statistically significant, while the confidence intervals give a better indication of whether a factor is *practically significant* (Rardin & Uzsoy (2001)).

To determine statistical significance in the computational time differences between the Value Curve approach and the various alternative offering approaches, we have used simple paired comparison t-tests. This is because ANOVA ignores the fact that the samples are paired, and thus overestimates the variance in the datasets.

Note that a level of significance of $\alpha = 0.05$ is assumed throughout the analysis in this chapter.

8.2.3 Instances Used For Simulations

As listed in Section 8.2.1, there are seven parameters that define the size and nature of a particular problem scenario, or instance. In order to cover a broad range of the problem domain, we have selected two values for each of the first 6 parameters, and five values for the RD curve type parameter, and produced instances based on every possible combination of these parameters ($5 \times 2^6 = 320$ instances). As stated earlier, five replicates were then created randomly based on each of these instances, producing an overall total of 1600 replicates on which the models could be tested and compared, covering a large area of the problem domain. The values of these parameters that were used are shown in Table 8.1.

Parameter	**Levels**
1. Number of Periods	20, 40
2. Number of Reservoir Levels	100, 300 (0-99 and 0-299 MWh)
3. Number of Dispatch Levels	20, 40 (0-19 and 0-39 MW/period)
4. Number of Residual Demand Curves per Period	25 (over 5 macro-states), 100 (over 10 macro-states)
5. Number of Possible Inflow levels per Period	1, 5 (both with an expected inflow level of 4 MWh/period)
6. Transition Correlation type	Loose, Strong
7. RD Curve Structure type	Single Similar, Single Different, Branch Different, Naturally Monotone, Stepped

Table 8.1 Parameter Values Used

Note the following definitions of the RD Curve Structure types:

Single Similar: All RD curves under all macro-states are of the same general shape, that being slightly convex.

Single Different: Within each macro-state, the RD curves are randomly selected to be either 1) Slightly convex, 2) Strongly convex, 3) Concave, 4) Linear

Branch Different: Each macro-state has a different type of RD curve, selected from 1) Convex, 2) Concave, 3) Linear and Steep, 4) Linear and Flat, 5) Stepped

Naturally Monotone: The possible RD curves within each macro-state run parallel to one another within the range of minimum and maximum possible prices. In Section 4.3, it was shown that under such a structure, the optimal offer was able to be theoretically determined and would be naturally monotone. As such, we would expect that under this RD structure, we could achieve very close to optimal results, despite the stochasticity in the market outcomes.

Stepped: The possible RD curves under all macro-states take a stepped form, which recognises that if you hold back generation, then in some cases, it can lead to a large jump in the market clearing price. It also recognises the convex shape that the market supply curve is generally considered to take. Such recognition of market power makes this the most realistic kind of RD curve to use.

In setting the parameter values of Table 8.1, many factors were taken into account, such as the ability to:

1) Produce realistically-sized and scaled instances across the problem domain.
2) Produce instances that demonstrate a range of market dominance levels.
3) Produce instances that are able to be solved in realistic times (for testing purposes) under all algorithms on the computers available (P4, 3.2GHz).
4) Produce instances that are able to be solved and simulated, given the memory restrictions of the computers available (1GB RAM).
5) Produce instances that represent a variety of types of offering patterns from rival generators.

These considerations of scaling and reality of instances were dealt with in the following ways:

1. Number of Periods

In the New Zealand market, the number of future periods that a generator is able to provide offers for at any given time ranges from 24 to 72 periods, depending on the time of day. For simplicity, we have used rounded values of 20 and 40 periods. We have chosen relatively short horizons to enable us to solve instances with greater values for the other parameters in the time that we have had available. The results will not be significantly different from those that we would find over longer horizons, for two reasons:

- We know from the computational complexity formulations in Section 6.6.5 that solving time will increase linearly as the number of periods increases.
- We know that the possible market outcomes towards the end of the horizon don't have as much of an effect on our behaviour now, as more impending periods do, because we are much more uncertain about these market outcomes in the relatively distant future[53,54].

2, 3. Number of Reservoir Levels and Dispatch Levels

Due to computer memory restrictions and the large computational time required under the R&A approach, we chose to set our parameter values for the number of reservoir levels and number of dispatch levels to reflect small to medium sized (up to 40MW) generating units (hydro and coal) and reservoirs. Note, however, that the Value Curve approach developed in this thesis is capable of handling much bigger scenarios than those considered in the experimental design. This is demonstrated in Section 8.8.

The reservoir levels that we have chosen are set such that, allowing for inflows, a generator running at full capacity could empty its small to medium sized reservoir

[53] Note that the level of certainty that we have about these outcomes decreases gradually as the horizon extends, and as such there is no specific cut-off point after which we don't need to consider possible outcomes. It is a subjective decision.
[54] Note also that an alternative set of horizon lengths might have been 24 and 48 periods, enabling us to cover one and two price cycles, respectively, as demonstrated in Figure 8.1.

within 12 to 24 hours. This is comparable to reservoirs found in New Zealand, such as those on the Waikato River, operated by Mighty River Power. Note that, most importantly, the expected inflows and reservoir sizes were set (in most cases) to create scenarios where sufficient fuel would be available to enable offering patterns to emerge that would produce non-zero release levels in most periods. However, a small number of the cases do consider the extreme situations of water shortage or surplus water, as would periodically be observed in a real hydro system. The focus is not on these types of situations, though, as the short term optimal offering strategy would generally be a straightforward decision in these cases (i.e. reduce or increase release as much as possible).

4. Number of Residual Demand Curves per Period

The number of RD curves to consider per period, and the number of macro-states on which to place them, would be completely up to a generator, who is using the model, to determine. This decision would be based on the data available to the generator, and their belief regarding the reliability of that data to predict future outcomes. We selected the values of 25 and 100 RD curves per period, as we believe that this represents a range of possible numbers of market scenarios a generating firm could estimate, without performing in-depth research into the topic.

The instances that we use all reflect an underlying world in which the branched macro-state structure of information and macro-states exists, as presented in Section 7.2. Because the R&A algorithm ignores this information, we include in Section 8.4 a comparison between that algorithm and the VC algorithm presented in Section 6.6, which ignores the branched structure. This provides the most direct comparison between the two approaches, while our concurrent analysis of the VC algorithm under the DADP framework with branching and macro-states shows how capitalising on such an underlying real-world structure can further improve the performance of the model.

5. Number of Possible Inflow levels per Period

Inflow levels are likely to be well known or easy to predict over the short-term horizon that is considered by this offering model. As such, we have considered only the case of inflow certainty (with a value of 4 MW/period) and the case of five possible inflow levels (with values of 1, 2, 4, 6 and 7 MW/period). The issue of the expected inflow levels is addressed in the reservoir size section above.

6. Transition Correlation type

As with the number of RD curves per period, the level of correlation between market outcomes in consecutive periods is a highly subjective measure, and should be set based on the data available to a generator using this model and according to their beliefs about the future. As such, we have considered two possible broadly-defined levels of correlation, those being “weak” (*stability* factor 1) and “strong” (stability factor 5). The probabilities and RD curves are structured such that we consider a shift between RD curves with similar indices to be more likely than a shift between RD curves with highly different indices. A stability factor of 1 implies that transitions are almost completely random, while a stability factor of 5 implies that a RD curve with the same index in the following period is approximately 5 times more likely to occur than a RD curve with an index with a difference of one (which in turn is approximately 5 times more likely than a RD curve with an index with a difference of two, and so on).

7. RD Curve Structure type

The five RD curve structure types that we have considered under this experimental design cover a broad range of possibilities that could occur, with respect to offering patterns of rival generators and demand fluctuations. From observing market offers, we believe the most likely shape of the RD curves faced by a generator to be either convex or stepped, reflecting the fact that most of the time, the price will be relatively low, but is subject to sharp price spikes which can be created and exploited by generators with market power. These shapes are reflected in the five RD curve structure types described above, which also account for other less-likely possibilities. Note that stepped RD curves produce the highest levels of market power, followed by convex RD curves. At the other extreme, flat linear RD curves provide the least market power (as holding back generation has little effect on the market clearing price).

An additional factor built into the replicates was to account for the relative *peakiness* of periods throughout the horizon. This was done by providing a weighting factor to each period, such that the height and position of the RD curves reflect whether the period is a peak or an off-peak period. These weighting factors are shown for 40 periods, in Figure 8.1.

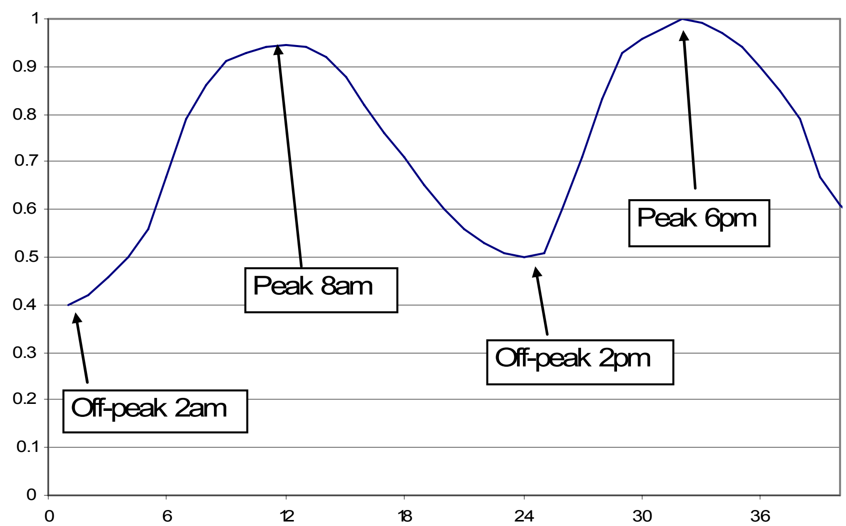

Figure 8.1 Peakiness Factors for 40 Periods

The actual price ranges used for the RD curves was taken from approximate price ranges for electricity at the three reference nodes of Haywards, Benmore, and Otahuhu in New Zealand, from the beginning of February 2000 to the end of June 2003[55].

The final consideration with respect to the data construction that needed to be considered was with respect to the reservoir cycle as the price (or RD curve) level cycles

[55] http://www.comitfree.co.nz/fta/ftapage.main, 31/8/03

through peak and off-peak periods in the short-term. We would expect that the reservoir level would be approximately a quarter-phase out of sync with price levels. In other words, the reservoir level should reach its peak level as the price is in the shoulder-peak zone heading into the peak period (e.g. period 6 in Figure 8.2), and it should reach its lowest levels as the price is in the shoulder-peak zone heading out of the peak period (e.g. period 18). The reason for this is that the generator would desire to maximise release between these two periods, when price is at its highest, and to minimise release in other periods, when price is at its lowest.

For all instances, we have assumed that the generator is starting in a relatively off-peak period, where the reservoir level is relatively near to full (as shown in Figure 8.2 – note that we have assumed that the reservoir begins completely full for simplicity). Under both the 20 and 40 period instances, the figure shows that the horizon is ending on shoulder peak periods, and as such we would expect the reservoir level to be near its lowest point in the cycle. As such, the end-of-horizon value curves (or marginal value curves) have been constructed to ensure that this reality is reflected in the test instances, and thus that the net result over both horizon lengths is that the reservoir level is being run down from a high level to a low level.

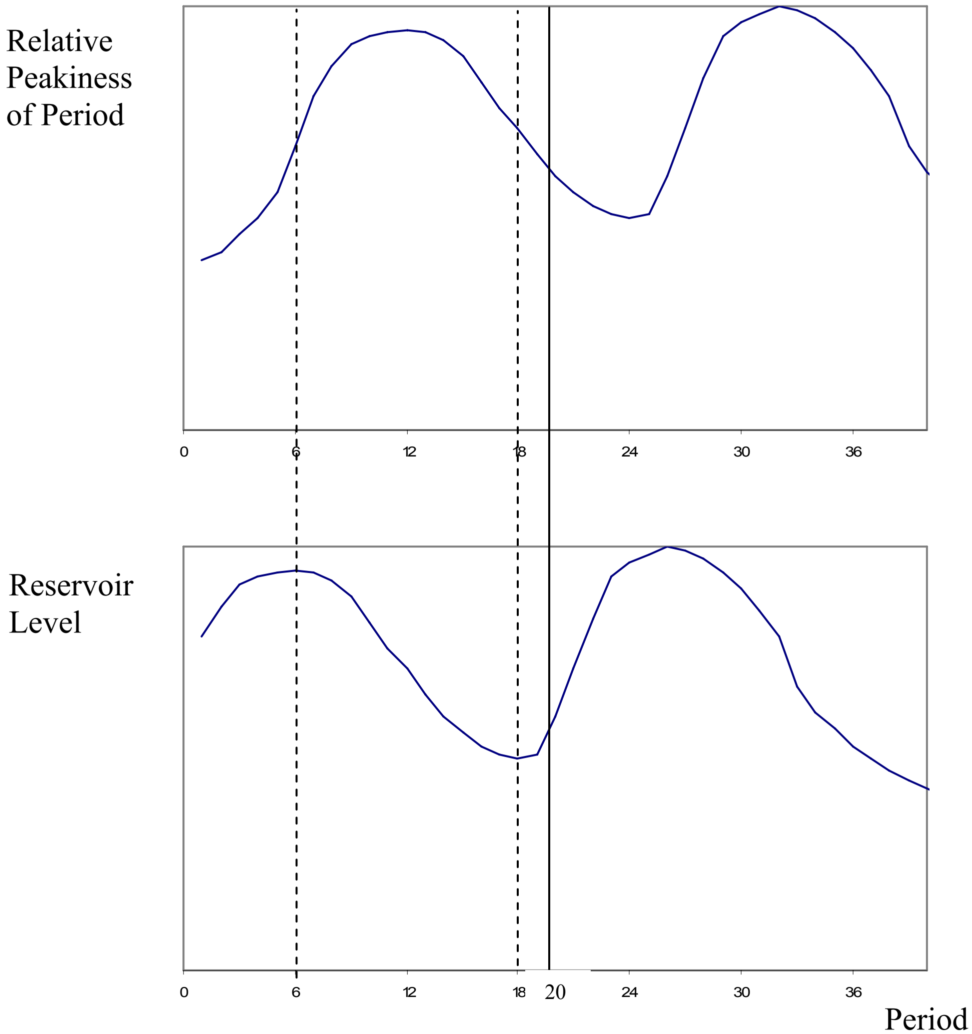

Figure 8.2 Reservoir Cycle Factors for 40 Periods

Note that many of the model inputs described in this section have been produced randomly, based on the various parameters defined, or arbitrarily. The construction and prediction of realistic market-data based data is beyond the scope of this thesis. The reason for this exclusion is two-fold. Firstly, for much of this information, a generation company would be far better placed to construct realistic data. Secondly, other researchers have performed work on predicting and grouping together information such as RD curves and macro-state probabilities (although for much different purposes). For example, Hernaez et al. (2004) uses clustering and artificial neural network techniques to group together historical RD curve types, providing a good basis for prediction of future possible RD curves.

8.3 Simulation Model

For a given replicate, the simulation model begins by producing 1000 schedules of possible residual demand and inflow scenarios for the horizon, based on the probabilistic information available. The model then takes the set of optimal offers (i.e. offers for all possible states throughout the horizon) for the given algorithm being considered, and simulates forward through the horizon, to obtain 1000 point estimates of the profit that would be achieved under the given scenario (or of percentage deviation from the optimal profit). Specifically, the logical process and order of events simulated is as follows:

a) Start out with knowledge of the MDF (uncertain market scenario) that we expect for the coming period (most likely based on the RD curve that occurred in the previous period, and on the macro-state that we are currently in)
b) Identify the optimal offer for this period based on that MDF and the current reservoir level.
c) Observe the RD outcome for this period, the inflow for the period, and the outcome of the event that defines which macro-state we will be in for the next period. The RD, combined with the offer determines our dispatch level and combined with the observed inflow will provide our new reservoir level.
d) Identify the MDF that we are expecting in the next period, based on the macro-state we are in and the previous RD that occurred.
e) Repeat b-d for all periods until the end of the horizon

Figure 8.3 demonstrates this process, showing where steps a-d listed above are applied. Note that under the R&A algorithm and the VC algorithm without branching, this process is simplified to ignore the different possible branches or macro-states.

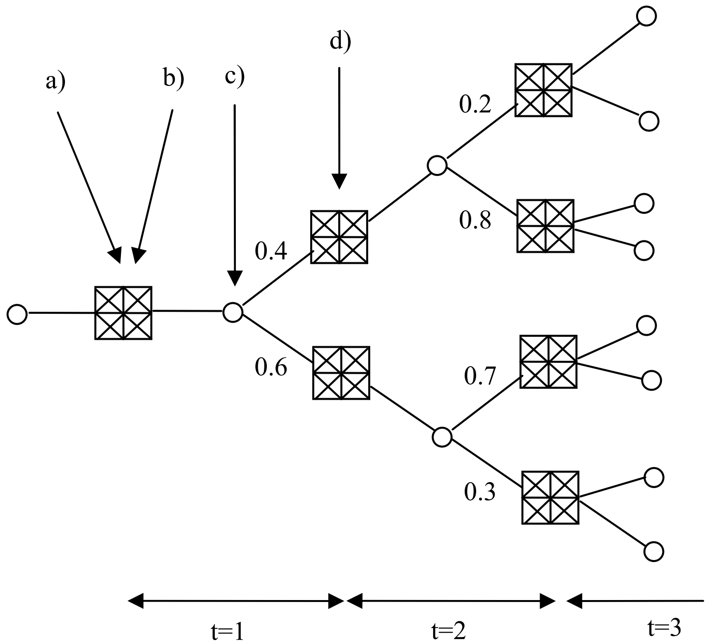

Figure 8.3 Simulated Event Process

8.4 Comparison between Full Complexity Algorithms without Branching

This section of the thesis compares the results of the simulation from two algorithms considering the same complexities and with the same modelling assumptions. They are:

1. The Rajaraman and Alvarado Offer Construction Algorithm
2. The Value Curve Approach Offer Construction Algorithm ignoring branched, macro-state structure

Both these algorithms make the modelling assumption that there is no underlying branched, macro-state structure to the real-world. The first algorithm is simply that described in Rajaraman & Alvarado (2003), while the second is the Two-Phase Value Curve approach without branching, described in Section 6.6.

The results used in this section relate to only a single replicate for each of the 320 instances. To solve all 1600 replicates under these algorithms would have been impractical, due to their long computational times, but the 320 replicates actually used still provide good coverage of the problem domain and is more than sufficient to prove the results presented. All relevant MINITAB output regarding computational times on the two models being compared can be found in Appendix D, while solution quality information can be found in Appendix E.

8.4.1 Algorithm Speed Comparison

A comparison between the overall mean computational times for the two models under the 320 replicates described above is shown in Table 8.2.

Model	Mean Computational Time (seconds)
R&A Approach[56]	80,324
Value Curve Approach (not branched)	99
PP Phase[57] (not branched)	4578

Table 8.2 Mean Computational Times

We can see that Real-Time phase of the Value Curve approach without branching requires far less computational time than the R&A Model, although the PP phase for this approach is still relatively slow. Specifically, on average, using the Value Curve approach saved 99.7% of the Real-Time computational time required by the R&A approach. In other words, the computational time for the R&A approach is on average 554.3 times that required for the Value Curve approach without branching, as demonstrated in Figure 8.4[58].

[56] Note that the R&A algorithm had been programmed as it is reported in Rajaraman & Alvarado (2003). This is to enable the best possible comparison between the approaches presented in this thesis and that presented in the referenced paper.
[57] Note that this PP Phase is required to be performed once in the horizon for the VC (branched and non-branched) and DMVC approaches, but not the R&A approach.
[58] Note that under all 320 replicates tested, the computational time is smaller for the Value Curve approach than it is for the Direct MV Curve approach, which in turn is less than that for the R&A approach.

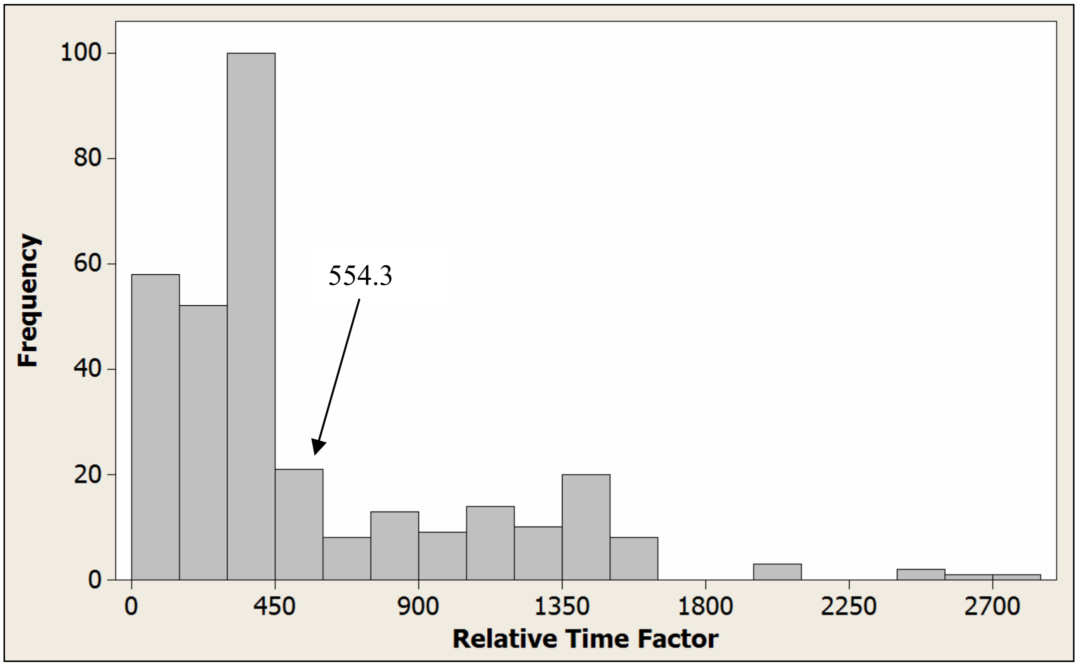

Figure 8.4 Histogram of Relative Computational Time for R&A Approach vs VC Approach

Recall that the instances being considered here correspond to small to medium sized generating units and reservoirs. For larger scenarios, the computational time savings would be even more substantial, as the computational complexity of the R&A algorithm increases at a much faster rate than the Value Curve approach without branching. For example, the largest instance took 851,506 seconds (9 days, 20.5 hours) to solve under the R&A approach, but just 538 seconds under the Value Curve approach; a saving of 99.94%. In a half hour offer based system, such as that found in New Zealand, this computational time for the Value Curve approach is fast enough to be used in real time. However, the computational time of almost 10 days under the R&A approach is clearly not practical.

Table 8.3 presents a two-way break-down of mean computational times by approach and parameter value. The shaded regions represent factors that *do not* have a statistically and practically significant effect on the computational times (approach held constant), while the un-shaded regions represent factors that *are* both statistically and practically significant. The statistical significance of these factors has been shown using paired comparison tests, and is presented in Appendix D. The shading practice used in this table is repeated in the equivalent tables in Sections 8.5, 8.6 and 8.7 of this chapter.

Mean Computational Times (seconds)	Approach	
	Value Curve (without branching)	R&A
Number of Periods		
20	60	53709
40	138	106939
Number of Reservoir Levels		
100	59	38303
300	138	122344
Number of Dispatch Levels		
20	72	25439
40	126	135208
Number of RD Curves		
25	30	9391
100	168	151256
Number of Inflows		
1	101	76548
5	97	84100
Spread Type		
Weak	102	80108
Strong	96	80539
RD Type		
Single Similar	98	82744
Single Different	89	87745
Branch Different	95	75595
Naturally Monotone	89	74906
Stepped	124	80628
Overall Mean	99	80324
Overall Worst	933	883565
Overall Best	7	882

Table 8.3 Mean Computational Times broken down by Factor and Approach

To summarise, Table 8.3 contains many pieces of statistically and practically significant information, which confirm our expectations of computational complexity presented in Section 6.4. For example:

- Computational time appears linear in number of periods for both approaches (Figure 8.5a).
- Computational time appears to be less than linear in number of reservoirs for the Value Curve approach, but more than linear for the R&A approach (Figure 8.5b).
- Computational time is approximately linear in the number of dispatch levels under the Value Curve approach, but substantially more than linear under the R&A approach (Figure 8.5c).
- Computational time is less than quadratic in the number of RD curves per period under the Value Curve approach, but approximately quadratic under the R&A approach (Figure 8.5d).
- The difference between the overall mean computational times for the two techniques is statistically significant.

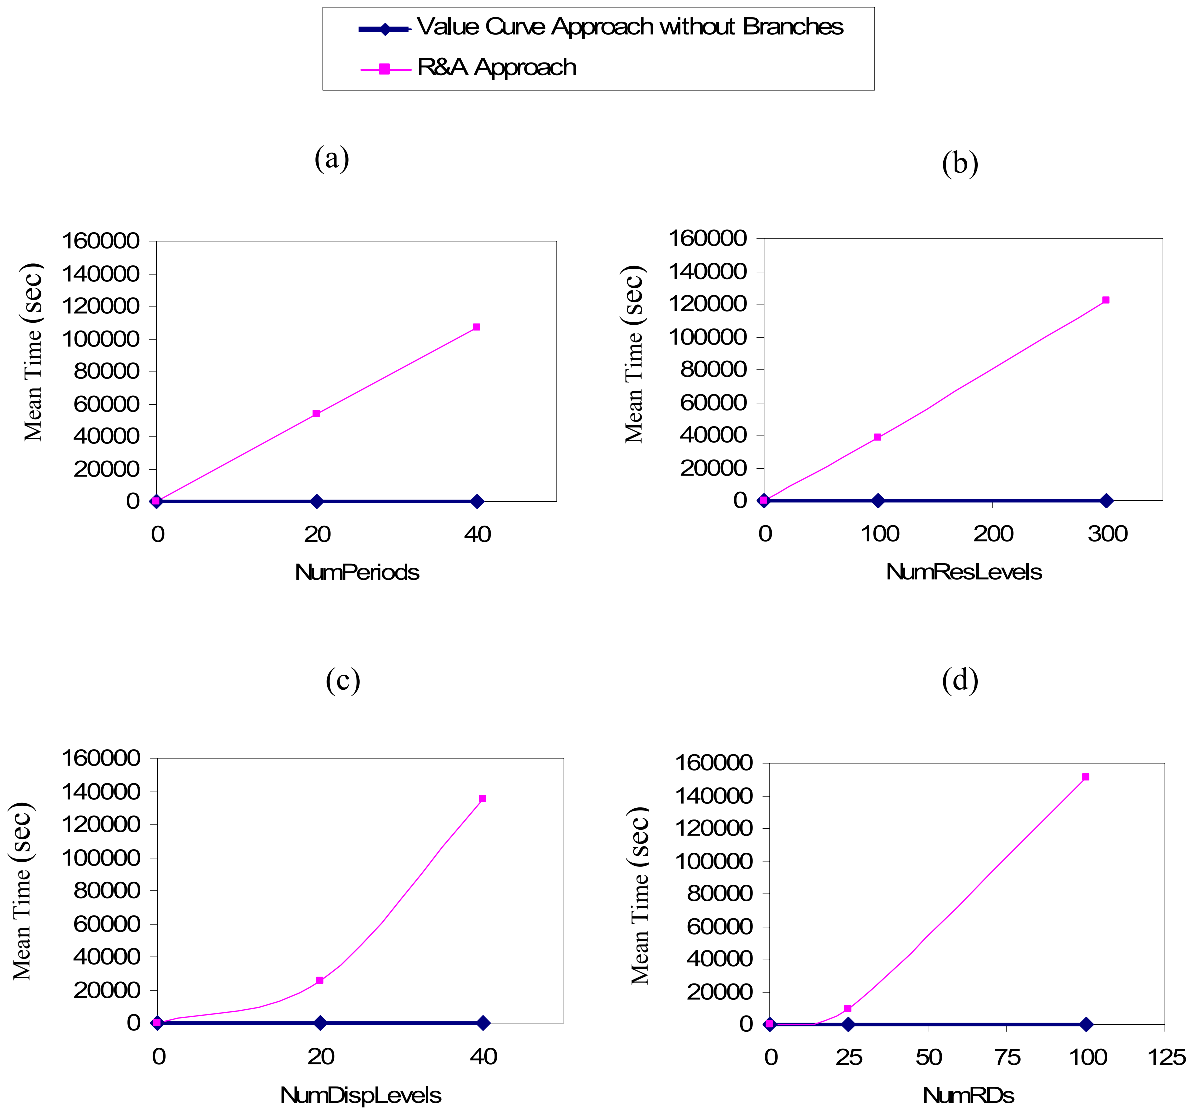

Figure 8.5 Effect on Mean Computational Time of the Parameters under each Model

8.4.2 Solution Quality Comparison

As reported in Section 6.7, we would expect that the two-phase approaches developed in this thesis would perform moderately better than the R&A approach with respect to mean percentage error in payoff under some RD curve structures, and significantly better under others. A comparison between these mean errors in payoff for the R&A and Value Curve approach without branching under the 320 replicates described earlier is shown in Table 8.4.

Model	Mean Percentage Error (%)
Value Curve Approach without Branching	1.681
R&A Approach	4.983

Table 8.4 Mean Percentage Error Comparison: R&A vs Value Curve Approach without Branching

On an aggregate basis, we can see that the Value Curve approach without branching performs substantially better than the R&A approach. Table 8.5 presents a two-way breakdown of mean percentage error by approach and parameter value. Un-shaded regions again represent factors that are both statistically and practically significant, for the given approach. Note that for RD type, where there are five possible parameter values, summary figures with matching groups in brackets are not statistically different from one another. For example, under the R&A approach, the mean percentage error from “Single Different” and “Naturally Monotone” are not different from one another at a statistically significant level (because they are both in group ‘a’, but the mean percentage error for “Branch Different” RD curves is statistically different (and worse) from both of these RD types (as it is in group ‘b’). The shading and grouping practice used in this table is repeated in the equivalent tables in Sections 8.5, 8.6 and 8.7 of this chapter. The statistical significance of the un-shaded factors has been shown using a general linear ANOVA model along with Bonferroni comparison tests, and is presented in Appendix E.

Mean Percentage Error	Approach	
	Value Curve without Branching	R & A
Number of Periods		
20	1.789	5.125
40	1.573	4.841
Number of Reservoir Levels		
100	1.952	5.581
300	1.410	4.386
Number of Dispatch Levels		
20	1.896	6.300
40	1.466	3.666
Number of RD Curves		
25	1.653	4.580
100	1.709	5.386
Number of Inflows		
1	1.524	4.813
5	1.838	5.153
Spread Type		
Weak	1.784	5.953
Strong	1.578	4.013
RD Type		
Single Similar	0.782 (a)	1.077 (a)
Single Different	0.939 (a, b)	1.734 (a)
Branch Different	2.074 (b)	5.376 (b)
Naturally Monotone	1.157 (a, b)	1.248 (a)
Stepped	3.453 (c)	15.480 (c)
Overall Mean	1.681	4.983
Overall Worst[59]	6.856	27.601
Overall Best	0.000	0.098

Table 8.5 Mean Percentage Errors broken down by Factor and Approach

[59] Note that in just 23 or the 320 replicates, the DMVC approach performed better than the VC approach, with a maximum of only 0.03% closer to the optimal payoff. In the remaining 297 replicates, the VC approach performed better, with a maximum of 1.43% closer to the optimal payoff.

To summarise, Table 8.5 reveals the following interesting pieces of statistically and practically significant information:

Number of Reservoir Levels and Number of Dispatch Levels

Under the R&A approach, the mean percentage error is significantly less for instances with 300 reservoir levels, as opposed to those with 100 reservoir levels. This is because the instances with the lower available fuel over the horizon are more constrained and, as with any mathematical program, a tighter constraint will cause the solution quality to be the same or worse. For the same reason, the mean percentage error is significantly less for instances with 40 dispatch levels, as opposed to those with 20 dispatch levels, under the R&A approach. There is no statistically significant difference in mean percentage error under the Value Curve Approach without branching.

Spread Type

Under the R&A approach, the mean percentage error is lower (better) by almost 2% under strong correlation between market outcomes between periods, compared with weak correlation. This is because under a strong correlation, the generator has greater knowledge of what is going to happen in the future from any given current state, and as such can make better informed generation offers.

Residual Demand Curve Type

Figure 8.6 summarises the mean percentage error information in Table 8.5, broken down by RD curve structure. We can see that the performance is quite similar under all RD curve structures except "Branch Different" and in particular under "Stepped", which performs significantly worse under both models. As explained in Section 8.2.3, this is the most realistic type of scenario. In particular, we note that the R&A approach performs especially poorly under this structure. The reason for this is that given the way the R&A algorithm was presented, it is possible for some payoffs to be overlooked, particularly when RD curves have steep, stepped segments, as explained in Section 6.7.1.

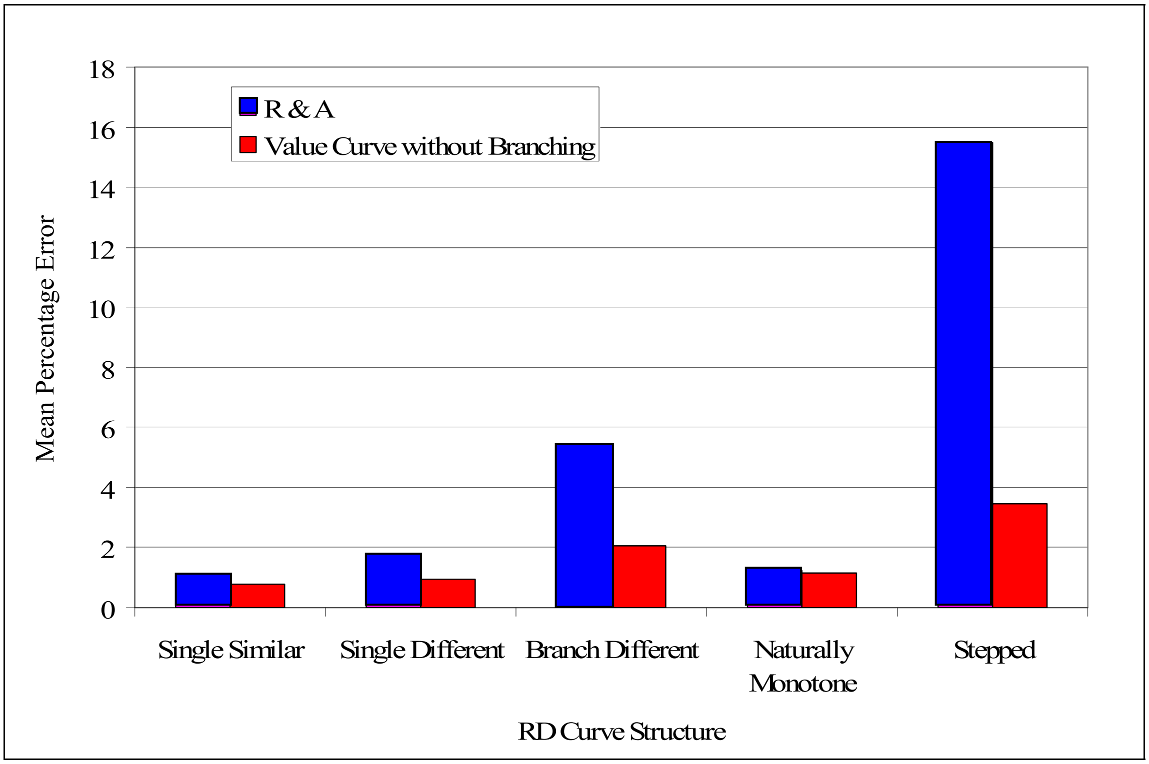

Figure 8.6 Mean Percentage Errors broken down by RD Curve Structure

Finally, note also that the difference between the overall mean percentage error values of 1.681% for the Value Curve approach without Branching and 4.983% for the R&A approach, presented in Table 8.5 is statistically significant.

8.4.3 Summary of Full Complexity Algorithm Comparisons without Branching

In this section of the chapter we have compared two approaches that consider the same complexities of the problem scenario and make the same assumption regarding the data not taking a branched form, but contain very different methods of dealing with them. Table 8.11 provides a summary of mean computational times and mean percentage errors for both approaches, over the 320 replicates described at the beginning of this section.

	Approach	Value Curve without Branching	R&A Approach
Computational Time	**Mean (seconds)**	99	80,324
	vs VC Approach without Branching	-	*Worse*
Algorithm Performance	**Mean Percentage Error**	1.681	4.983
	vs VC Approach without Branching	-	*Worse*
	Worst RD Curve Structure	Stepped	Stepped
	Mean Percentage Error under the Worst RD Curve Structure	3.453	15.480
	Worst Case	6.856	27.601

Table 8.6 Summary of Results for Full Complexity Algorithms

We have shown that the Value Curve approach without branching is significantly superior to the R&A approach in terms of the expected payoff that can be achieved, particularly in the case of stepped RD curves, which we believe to be the most realistic case. We have also shown that it is significantly superior to the R&A approach in terms of computational time. The extent of this saving is such that we go from having a model that could only be solved in real time for very small and simplistic examples, to a model that can now be solved in real time for large problems that incorporate great amounts of complexity. Quite clearly, the Two-Phase Value Curve approach without branching is a preferable algorithm to the R&A approach. The following section of this chapter will investigate the further gains that can be made by considering that the real world can be approximately represented with a branched structure and thus applying the DADP framework developed in Chapter 7.

Note also that another benefit of the two-phase approaches is that they provide MV and MC curves for all states throughout the horizon, hence providing less of a “black-box” solution than that provided by R&A.

8.5 Comparison between Full Complexity Algorithms with and without Branching

In this section we compare, with respect to both solution quality and computational time, the Value Curve approach without branching (proven to be the best of the full complexity non-branching algorithms from Section 8.4) to the branched, macro-state version of the two Two-Phase algorithms developed in Chapter 7. Therefore, the three models for comparison are:

1. The Value Curve Approach Offer Construction Algorithm ignoring branched structure
2. The Value Curve Approach Offer Construction Algorithm considering branched structure (DADP framework)
3. The Direct Marginal Value Curve Offer Construction Algorithm considering branched structure (DADP framework)

The first algorithm makes the modelling assumption that there is no underlying branched, macro-state structure to the real-world, while the final two algorithms recognise that such a structure is recognised in the instances, and capitalise on this added information appropriately. All three algorithms are variations on the Two-Phase approach developed in this thesis and described in the preceding chapters.

As in the previous section, the results used in this section relate to only a single replicate for each of the 320 instances. To solve all 1600 replicates under the DMVC approach would again have been impractical, due to the long computational times, while the 320 replicates actually used continue to provide good coverage of the problem domain and are more than sufficient to prove the results presented. All relevant MINITAB output regarding computational times on the three models being compared can be found in Appendix F, while solution quality information can be found in Appendix G.

8.5.1 Algorithm Speed Comparison

A comparison between the overall mean computational times for the three models under the 320 replicates described above is shown in Table 8.7.

Model	Mean Computational Time (seconds)
Value Curve Approach (not branched)	99
Value Curve Approach (branched)	74
Direct MV Curve Approach (branched)	999
PP Phase (not branched)	4578
PP Phase (branched)	235

Table 8.7 Mean Computational Times

We can see that the Real-Time phase of the Value Curve approach is substantially faster than the Direct MV Curve approach, whether or not branching is considered. Specifically, the computational time for the Direct MV Curve approach is on average 12.1 times that required for the Value Curve approach, as demonstrated in Figure 8.7. We recognise that the computational complexities presented in Section 7.7.2 implied that these computational times should have been similar, and we leave for further research investigating the cause of this anomaly, and thus whether the implementation of the Direct MV Curve approach could be improved.

Further, we can see that a 25% improvement in Real-Time computational requirements is achieved under the Value Curve approach by considering branching. The biggest effect of considering branching, though, is the 95% computational saving in the Pre-Processing phase of the algorithm.

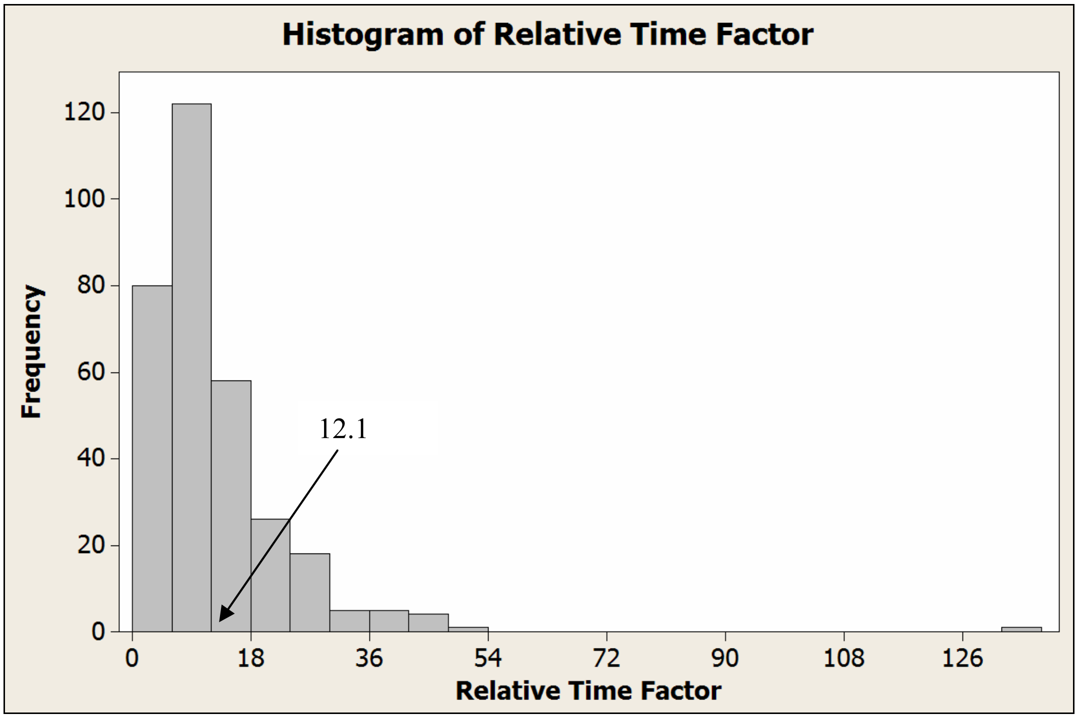

Figure 8.7 Histogram of Relative Computational Time for the Direct MV Curve Approach vs VC Approach

In Section 8.4, we noted that the largest instance took 538 seconds to solve the Real-Time phase when branching was ignored. When we consider branching under the Value Curve approach, this instance takes just 295 seconds to solve, a further computational saving of 45.17%.

Table 8.8 presents a two-way break-down of mean computational times by approach and parameter value. The statistical significance of the un-shaded factors has been shown using paired comparison t-tests, and is presented in Appendix F. Recall from Section 8.4.2 that the comparisons are made between the parameter values within each technique, and that the letters after the figures under the RD Type break-down indicate groups that are statistically significantly different from one another.

Mean Computational Times (seconds)	Approach		
	Value Curve without Branching	Value Curve with Branching	Direct MV Curve
Number of Periods			
20	60	50	663
40	138	99	1335
Number of Reservoir Levels			
100	59	42	449
300	138	106	1550
Number of Dispatch Levels			
20	72	54	1131
40	126	94	867
Number of RD Curves			
25	30	29	235
100	168	120	1764
Number of Inflows			
1	101	76	978
5	97	72	1021
Spread Type			
Weak	102	74	992
Strong	96	74	1007
RD Type			
Single Similar	98	72	971 (b)
Single Different	89	73	740 (a)
Branch Different	95	76	971 (b)
Naturally Monotone	89	74	743 (a)
Stepped	124	76	1571 (c)
Overall Mean	99	74	999
Overall Worst	933	390	7148
Overall Best	7	7	34

Table 8.8 Mean Computational Times broken down by Factor and Approach

The results presented in Table 8.8 reveal very similar trends over all three models with respect to changing parameter values, where the computational times are lowest under the Value Curve approach with branching, one third slower under the Value Curve approach without branching, and substantially slower still under the Direct MV Curve approach. One difference of note though, is that the Direct MV Curve approach performs significantly worse under the stepped RD type than under any other type. A selection of the results are demonstrated in Figure 8.8, showing simple comparisons between the computational times of the three approaches.

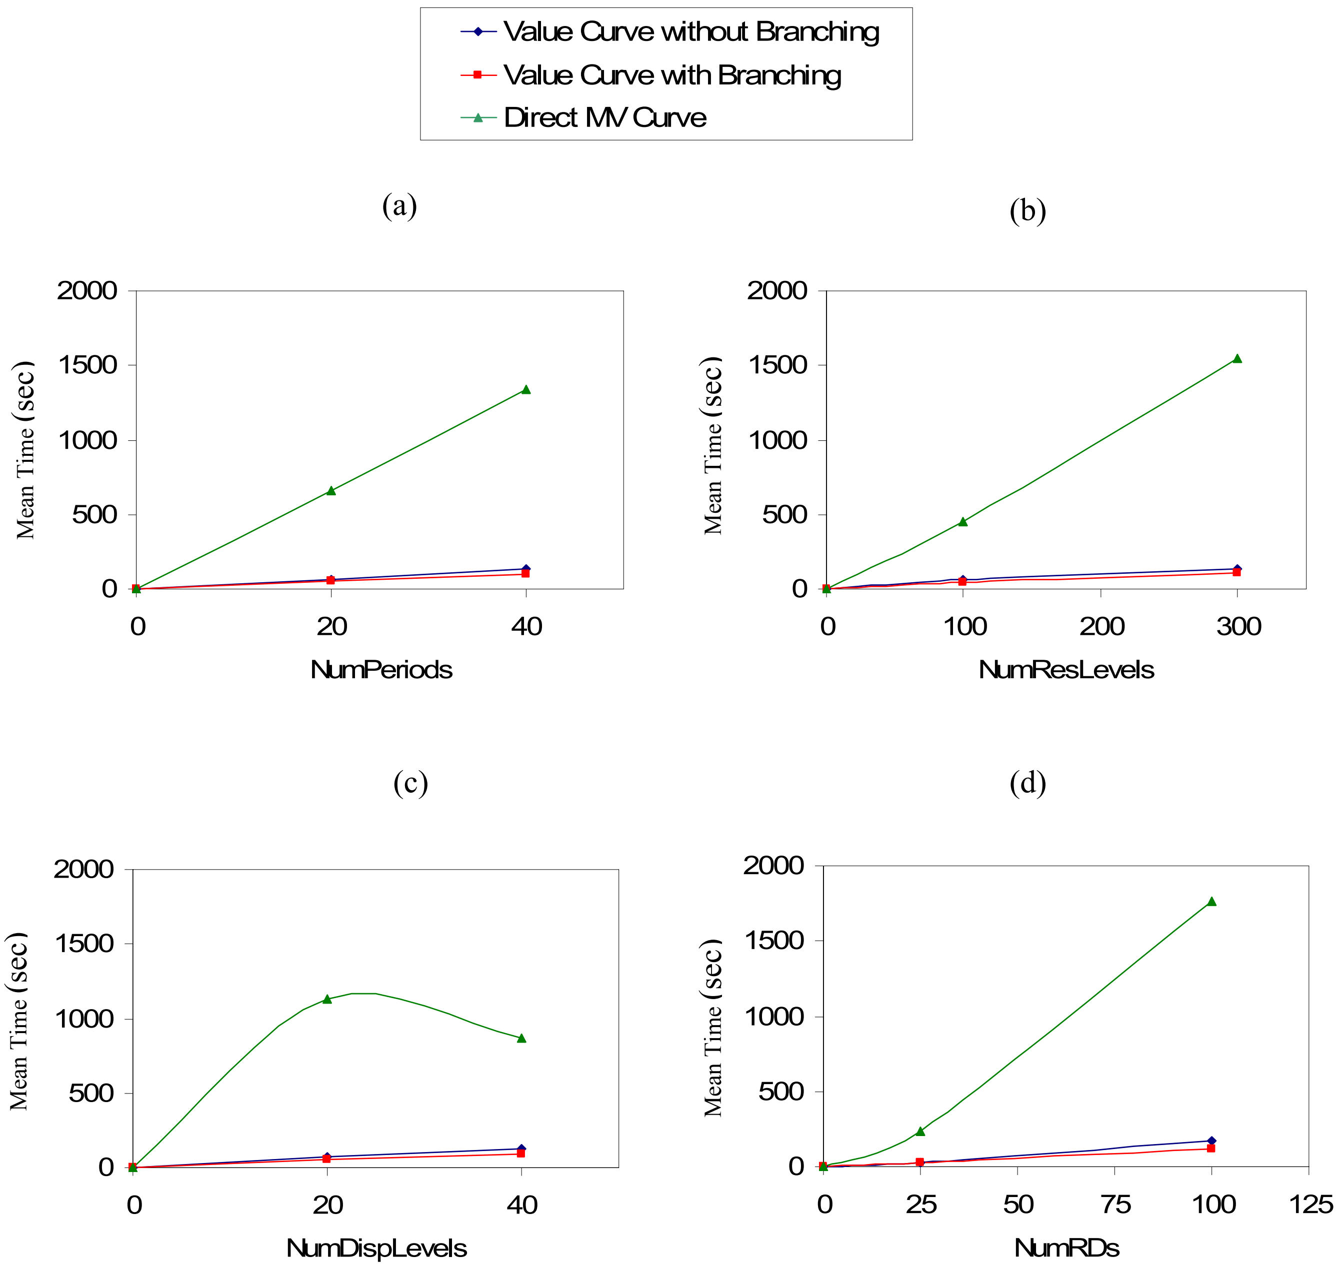

Figure 8.8 Effect on Mean Computational Time of the Parameters under each Model

Note also that the differences between the overall mean computational times for all three pairs of techniques are statistically significant.

8.5.2 Solution Quality Comparison

As reported in Section 7.5, we would expect that the two-phase approaches with branching would perform moderately better than the Value Curve approach without branching, with respect to mean percentage error in payoff. A comparison between these mean errors in payoff for the three models under the 320 replicates described earlier is shown in Table 8.9.

Model	**Mean Percentage Error (%)**
Value Curve Approach without Branching	1.681
Value Curve Approach with Branching	1.300
Direct MV Curve Approach	1.404

Table 8.9 Mean Percentage Error Comparison: R&A vs Two-Phase Models

On an aggregate basis, we can see that the Value Curve approach with branching performs marginally better than the Direct MV approach, which in turn performs marginally better than the Value Curve approach without branching.

Table 8.10 presents a two-way breakdown of mean percentage error by approach and parameter value. The statistical significance of the un-shaded factors has been shown using a general linear ANOVA model along with Bonferroni comparison tests, and is presented in Appendix G. Again, the significance is calculated between the parameters, within each technique, rather than between the techniques.

Mean Percentage Error	Approach		
	Value Curve without Branching	Value Curve with Branching	Direct MV Curve
Number of Periods			
20	1.789	1.313	1.417
40	1.573	1.287	1.390
Number of Reservoir Levels			
100	1.952	1.628	1.802
300	1.410	0.972	1.005
Number of Dispatch Levels			
20	1.896	1.338	1.405
40	1.466	1.262	1.402
Number of RD Curves			
25	1.653	1.252	1.365
100	1.709	1.348	1.443
Number of Inflows			
1	1.524	1.141	1.306
5	1.838	1.459	1.501
Spread Type			
Weak	1.784	1.368	1.442
Strong	1.578	1.232	1.365
RD Type			
Single Similar	0.782 (a)	0.565 (a)	0.671 (a)
Single Different	0.939 (a)	0.742 (a, b)	0.912 (a, b)
Branch Different	2.074 (b)	1.469 (c)	1.569 (c)
Naturally Monotone	1.157 (a)	1.153 (b, c)	1.245 (b, c)
Stepped	3.453 (c)	2.570 (d)	2.621 (d)
Overall Mean	1.681	1.300	1.404
Overall Worst[60]	6.856	5.601	5.598
Overall Best	0.000	0.001	0.0001

Table 8.10 Mean Percentage Errors broken down by Factor and Approach

[60] Note that in just 23 or the 320 replicates, the DMVC approach performed better than the VC approach, with a maximum of only 0.03% closer to the optimal payoff. In the remaining 297 replicates, the VC approach performed better, with a maximum of 1.43% closer to the optimal payoff.

Table 8.10 reveals that only the RD Curve Type has a statistically and practically significant effect on the mean percentage error under these three models.

Residual Demand Curve Type

Figure 8.9 summarises the mean percentage error information in Table 8.10, broken down by RD curve structure. We can see that the performance is quite similar under all RD curve structures. The worst structure under all three approaches, though, is "Stepped", which, as explained in Section 8.2.3, is the most realistic type of scenario. In addition, we can observe that the Value Curve approach with branching performs the best under all RD Curve types.

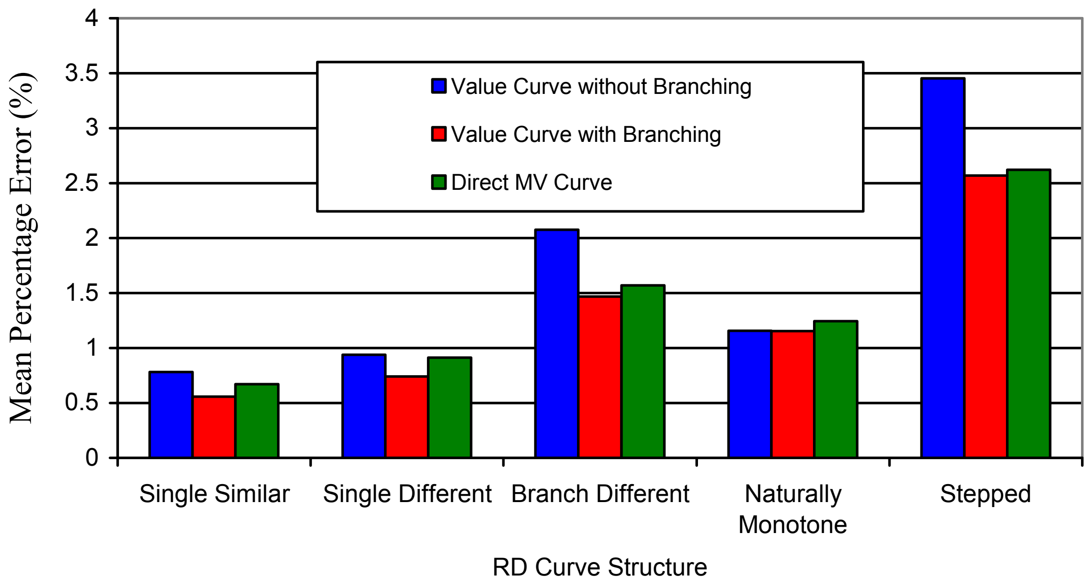

Figure 8.9 Mean Percentage Errors broken down by RD Curve Structure

Note also that the differences between the overall mean percentage error values for each pair of these three approaches is statistically significant.

8.5.3 Summary of Full Complexity Algorithm Comparisons with and without Branching

In this section of the chapter we have compared three approaches that consider the same complexities of the problem scenario, one of which ignores the underlying branched structure of the data, and two of which capitalise on this structure. Table 8.11 provides a summary of mean computational times and mean percentage errors for all three approaches, over the 320 replicates described at the beginning of this section.

	Approach	Value Curve without Branching	Value Curve with Branching	Direct MV Curve
Computational Time	**Mean (seconds)**	99	74	999
	vs VC Approach with Branching	*Worse*	-	*Worse*
Algorithm Performance	**Mean Percentage Error**	1.681	1.300	1.404
	vs VC Approach with Branching	*Worse*	-	*Worse*
	Worst RD Curve Structure	Stepped	Stepped	Stepped
	Mean Percentage Error under the Worst RD Curve Structure	3.453	2.570	2.621
	Worst Case	6.856	5.601	5.598

Table 8.11 Summary of Results for Full Complexity Algorithms

We have shown that although the Direct MV Curve approach is nice analytically, practically the Value Curve approach performs marginally better in terms of percentage payoff error (when they both consider branching) and substantially better in terms of computational requirements. As the Value Curve approach is also easier to interpret, we will not consider the Direct MV Curve Approach in the remainder of this thesis.

We have also shown that the Value Curve approach with branching is significantly superior to the same approach without branching with respect to both computational time and mean percentage error.

Given these comparisons, the Two-Phase Value Curve approach with branching is by far the best of the three algorithms discussed in this section when the real world can reasonably be approximated with a branched structure, and as such will be the only model considered in the remainder of the thesis.

8.6 Comparison between the Full Complexity Algorithm and the Simplified Algorithms

In this section we compare, with respect to both solution quality and computational time, the Value Curve approach with branching (proven to be the best of the full complexity algorithms from Section 8.5, and hereafter referred to simply as the Value Curve approach) to various alternative simplified algorithms. These algorithms are simplified in various ways, either changing the allowable form of offers, or by reducing the amount of available information that we actually consider when determining the choice of offer. The purpose of this is to establish if there is value in considering these various complexities, or alternatively, if the simplified approaches will provide solutions of similar quality in less computational time.

Specifically, we consider five different modelling simplifications.

1. Quantity-Based Offer Simplification

This simplification requires that the offers provided to the market by the generator are quantity-based (similar to a Cournot-style game), rather than sculpted (as assumed by the three models considered in Section 8.4). In other words, the generator decides the quantity that they will be dispatched from any given system state, and the price is determined by the stochastic outcome of the market scenario. With such a simplification, the Two-Phase approach with offer patching can still be applied. The Pre-Processing Phase is greatly simplified, and thus far quicker to apply, while we would expect the Real-Time solving requirements to be very similar to that for the sculpted offer Two-Phase Approach.

2. Price-Based Offer Simplification

This simplification requires that the offers provided to the market by the generator are price-based (similar to a Bertrand-style game, but with a maximum quantity level). In other words, the generator decides the price that they will offer their maximum generation output to the market, and the quantity that they are actually dispatched is determined by the stochastic outcome of the market scenario.

With such a simplification, the Two-Phase approach with offer patching can no longer be applied. The reason for this is that through patching, we would not end up with a single-price offer. As such, the Price-Based Offer approach is really a different model to the Value Curve approach entirely, where all calculations are performed in a single phase (as they are for the R&A approach) at the expense of computational efficiency. Therefore, despite reduced complexity associated with this simplification, we would expect the real-time solving requirements to increase over the VC approach.

3. Naïve Correlations Simplification

This simplification ignores the correlation between market scenarios in consecutive periods, in addition to ignoring the branching structure of the problem along with the associated branch probabilities. The probabilities of each RD occurring in each period is therefore established in advance, and is defined to be independent of any other market events or scenarios that occur[61]. The result of this simplification is that the state-space of

[61] To produce these RD curve probabilities, the first step is to work through from the start of the horizon and establish the probability of being in each branch at each period. We then produce the average probabilities for each RD on a given branch, calculated as shown in the table below. Finally, we multiply these RD probabilities by the probability of reaching that branch. This is shown in the table below.

Transition Probabilities		RD			
		1	2	3	4
Previous RD	1	0.2	0.1	0.3	0.4
	2	0.3	0.2	0.2	0.3
	3	0.1	0.5	0.4	0
	4	0.3	0.1	0.1	0.5
Averages		0.225	0.225	0.250	0.300

the DP model becomes one-dimensional (as there is no longer a need to consider the previous RD curve observed). A simplified Two-Phase Model can be applied with this simplification, where we would expect the solving times in both phases to be significantly reduced as a result of the reduced number of dimensions in the DP.

4. Naïve Water Simplification

This simplification is to completely ignore the reservoir level when determining offers. In Chapter 3 we demonstrated how this would adversely affect our outcomes in simple deterministic cases, but we extend this here to show how our model can be used to demonstrate that the effect is just as bad in more complex, stochastic scenarios. The effect of this simplification is that we consider the marginal cost to be zero over the entire feasible generating range, regardless of the reservoir level state. If we were to run this in a DP, we would not need to store the Value or MV curves at any point, as the offers would be based simply on the uncertain market scenarios. For the purposes of this study, we have not actually produced a separate DP model for this scenario. Rather, we have simply applied the Pre-Processed offers for a marginal cost of zero, at all states and times, and as such, the solving time required is zero.

We know from simple mathematical programming theory that the optimal expected payoff achievable in any problem situation must either stay the same or decrease if additional constraints are added to the scenario. This is exactly what is occurring if we require the offers to be purely quantity-based or price-based, rather than sculpted. We know this because each of these feasible offer sets form a subset of the feasible sculpted offers, as illustrated by Figure 8.10.

5. Perfect Competitor Simplification

The final simplification is to assume that the generator considers themselves to be a perfect competitor, and thus unable to affect the market clearing price. As such, the optimal offer is simply to provide electricity at the generator's marginal cost, which is determined by the marginal value curve.

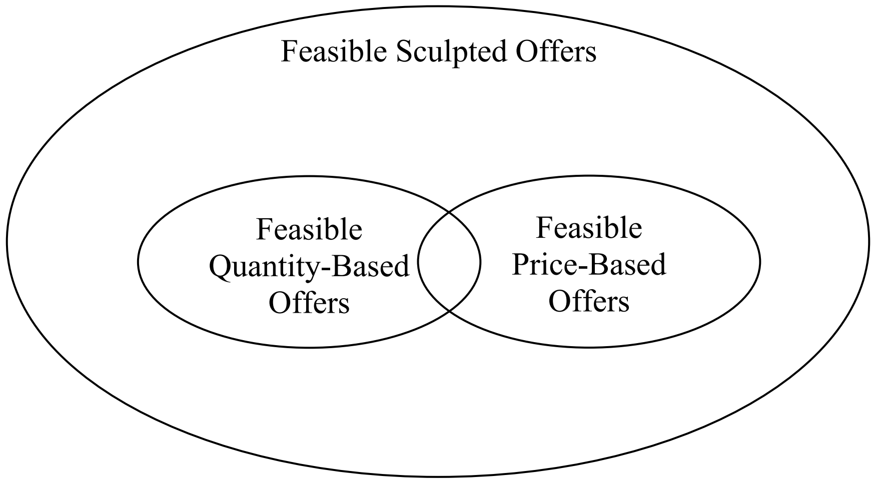

Figure 8.10 Feasible Offer Sets

Assuming that the information available to be used is accurate, we can expect that when a decision does not use all available information, the expected payoff will again be lower (at least in the long-run). Both the Naïve Correlations and Naïve Water simplifications are based on such a reduction in information.

Note that all of these simplified approaches, other than the Naïve Correlation approach, assume that the real-world can reasonably be represented by a branched structure, and thus apply this structure.

The results used throughout this section relate to all 1600 replicates described in Section 8.2.3. All relevant MINITAB output regarding computational times for the five approaches (the four described above, plus the Value Curve approach) can be found in Appendix H, while output regarding solution quality is located in Appendix I.

8.6.1 Algorithm Speed Comparison

A comparison between the overall mean computational times for the five models over the 1600 replicates described in Section 8.2.3 is shown in Table 8.12.

Model	**Mean Computational Time (seconds)**
Value Curve Approach	52
Quantity-Based Offers	52
Price-Based Offers	410
Naïve Correlation Approach	1.4
Naïve Water Approach	0
Perfectly Competitive Approach	50
PP Phase – Value Curve Approach	131
PP Phase – Quantity-Based Offers	56
PP Phase – Naïve Correlation Approach	17

Table 8.12 Mean Computational Times

We can see that the mean computational times for the Quantity-Based Offers and Perfectly Competitive approaches are very similar to that for the Value Curve approach, and any differences would certainly not be practically significant[62]. However, the mean computational time for the Naïve Correlation approach is extremely small (due to the reduced dimension of the DP), at just 3% of that required for the Value Curve approach, while the mean computational time for the Price-Based Offers is substantially larger, at almost eight times that required for the Value Curve approach. Table 8.13 presents a two-way break-down of average computation times by approach and parameter value. We can see that under the factors that do present both a practical and a statistical significance (un-shaded), the effects are very similar under all approaches, and thus the

[62] With respect to spread of computational times, note that the standard deviations for the three approaches are similar, and the worst case scenario was far worse under the Quantity-Based Offers and Perfectly Competitive approaches, than under the Value Curve approach (as presented in Table 8.13).

analysis would be very similar to that presented in Section 8.5.1, and so is not repeated here.

Mean Computation Times (seconds)	Approach					
	Value Curve	Quantity-Based Offers	Price-Based Offers	Naïve Correlation Approach	Naïve Water Approach	Perfectly Competitive Approach
Number of Periods						
20	38	38	318	1	0	36
40	66	66	501	2	0	65
Number of Reservoir Levels						
100	33	31	211	1	0	32
300	72	73	608	2	0	69
Number of Dispatch Levels						
20	37	37	329	1	0	37
40	68	68	490	2	0	64
Number of RD Curves						
25	30	30	196	1	0	27
100	74	75	623	2	0	73
Number of Inflows						
1	51	51	413	1	0	47
5	53	53	406	1	0	53
Spread Type						
Weak	52	52	409	1	0	51
Strong	52	53	410	1	0	50
RD Type						
Single Similar	51	51	466	1	0	53
Single Different	51	50	442	1	0	51
Branch Different	54	55	416	1	0	54
Naturally Monotone	52	52	421	1	0	47
Stepped	53	53	303	1	0	47
Overall Mean	52	52	410	1	0	50
Overall Worst	390	633	9992	11	0	959
Overall Best	6	5	31	0	0	5

Table 8.13 Mean Computational Times broken down by Factor and Approach

Note that the differences between the overall mean computational times for all possible pairs of techniques (other than those between the Value Curve, Quantity-Based Offers, and Perfectly Competitive approaches) are statistically significant.

8.6.2 Solution Quality Comparison

Overall, we would expect that the Value Curve approach developed in this thesis would perform better than each of the simplified algorithms with respect to mean percentage error in payoff, due to the reasons regarding increased constraint or reduced information levels, explained earlier in this section. A comparison between these mean errors in payoff for the Value Curve approach and the five simplified algorithms under the 1600 replicates is shown in Table 8.14.

Model	Mean Percentage Error (%)
Value Curve Approach	1.186
Quantity-Based Offers	4.132
Price-Based Offers	2.304
Naïve Correlation Approach	3.080
Naïve Water Approach	36.009
Perfectly Competitive Approach	9.393

Table 8.14 Mean Percentage Error Comparison: Value Curve approach vs Simplified Algorithms

On an aggregate basis, we can see that the Value Curve approach performs substantially better than all of the simplified algorithms, as would be expected. All of these differences are statistically significant, according to the Bonferroni tests presented in the appendices.

Table 8.15 presents a two-way breakdown of mean percentage error by approach and parameter value, showing how the approaches perform under different groupings of instances. Un-shaded regions again represent factors that are both statistically and practically significant. The blue shading is used under the RD type parameter to signify approaches whose performance is not significantly different to the Value Curve approach, under each given RD curve type.

Mean Percentage Error	Approach					
	Value Curve	Quantity-Based Offers	Price-Based Offers	Naïve Correlation Approach	Naïve Water Approach	Perfectly Competitive Approach
Number of Periods						
20	1.109	3.494	2.246	2.640	36.174	7.717
40	1.262	4.771	2.362	3.520	35.843	11.070
Number of Reservoir Levels						
100	1.588	5.317	3.099	4.032	35.636	11.814
300	0.784	2.948	1.509	2.128	36.381	6.973
Number of Dispatch Levels						
20	1.129	4.320	2.053	3.224	31.379	10.484
40	1.242	3.944	2.555	2.935	40.638	8.303
Number of RD Curves						
25	1.194	4.117	2.307	2.963	36.069	9.507
100	1.177	4.148	2.301	3.197	35.948	9.280
Number of Inflows						
1	0.823	3.747	2.054	2.834	35.816	8.994
5	1.549	4.517	2.554	3.326	36.201	9.793
Spread Type						
Weak	1.117	5.443	2.214	2.910	36.501	8.905
Strong	1.254	2.822	2.394	3.250	35.516	9.882
RD Type						
Single Similar	0.688	0.825 (c)	0.973 (a,b)	1.090 (c)	47.627 (b)	5.848 (d)
Single Different	0.993	3.284 (b)	2.611 (b)	1.783 (b,c)	45.551 (c)	2.461 (e)
Branch Different	1.825	2.431 (b)	2.313 (b,c)	3.146 (b)	51.396 (a)	7.543 (c)
Naturally Monotone	1.686	7.790 (a)	4.801 (a)	1.768 (b,c)	31.843 (d)	11.387 (b)
Stepped	0.737	6.333 (a)	0.822 (c)	7.614 (a)	3.626 (e)	19.727 (a)
Overall Mean	1.186	4.132	2.304	3.080	36.009	9.393
Overall Worst	5.601	25.455	10.935	22.390	67.860	50.829
Overall Best	0.001	0.071	0.075	0.018	0.038	0.419

Table 8.15 Mean Percentage Errors broken down by Factor and Approach

Some of the significant information presented in Table 8.15 is consistent with that analysed in Section 8.5.2, and so we will not repeat the analysis here[63]. Below we will highlight some new insights from Table 8.15. Note that the differences between the overall mean percentage error values for each pair of these six approaches is statistically significant.

Number of Periods

Table 8.15 shows us that under the Value Curve, Price-Based Offer and Naïve Water approaches, there is no statistically significant difference between the mean percentage errors with respect to the number of periods in the horizon. For the remaining approaches, the mean percentage error is greater for the instances with a longer planning horizon. The reason for this is that with more periods, all else held constant, available generation capacity must be used more sparingly. This results in a higher marginal value of fuel available and hence the cost of non-optimality in the offering policy is greater.

Number of Reservoir Levels and Number of Dispatch Levels

Under all approaches other than the VC and Naïve Water, the mean percentage error is significantly higher in the replicates with only 100 reservoir levels. This is because under the scaling of these instances, the cases with the smaller reservoir require much greater conservation of fuel resources (i.e. a greater constraint), and as such it is more difficult to achieve a close-to-optimal result[64]. Likewise, regarding the number of dispatch levels, the instances that have a lower maximum generation level are more constrained and as such, a greater mean percentage error results.

[63] Note, however, that under the Value Curve approach with 1600 replicates shown here, there is no statistically significant difference between the mean percentage errors of the different RD types. This is different from the smaller sample of 320 replicates, as shown in Table 8.10. The 320 replicate sample was not large enough to recognise this lack of significant difference between the groups.
[64] Note that similarly, the percentage error will be higher under scenarios with lower expected inflows of the horizon.

Residual Demand Curve Type

Table 8.15 shows us that under all five RD curve types, the Value Curve approach performs at least as well if not better than all five other approaches, with respect to mean percentage error in payoff, consistent with our expectations from the beginning of this section (recall that differences with the VC approach that are not statistically significant are highlighted in blue). The worst mean percentage error for the Value Curve approach, of 1.825%, corresponds to the "Branch Different" RD curve type. While each of the other approaches work very well on certain RD curve types, each of them also works particularly poorly on other RD types. These relative performances are detailed below:

Quantity-Based Offers:

Although the Quantity-Based offer approach performs very well on the "Single Similar" RD curve type (with a mean percentage error of just 0.8255), it performs particularly poorly under "Naturally Monotone" RD curve types (mean percentage error of 7.790%) and under the "Stepped" RD curve type (mean percentage error of 6.333%), which we believe to be the most realistic RD curve type. The worst case percentage errors over the 1600 instances for these two RD curve types were 25.455% and 24.125% error respectively.

Price-Based Offers:

The Price-Based offer approach has a statistically significant difference in performance with the VC approach under only the "Single Different" and "Naturally Monotone" RD curve types. The worst performance was in the "Naturally Monotone" cases, with a mean percentage error of 4.801%, and the worst-case percentage error under this RD curve type was 10.935%.

Naïve Correlation Approach:

The Naïve Correlation approach performs particularly poorly under the “Stepped” RD curve type, with a mean percentage error of 7.614%. The worst-case percentage error under this RD curve type was 22.390%. Under all other RD curve types, there was no statistically significant difference in performance with the VC approach.

Naïve Water Approach:

Because the Naïve Water approach is so simplistic, it performs very poorly over most RD curve types, with an overall mean percentage error of 36.009%. The only RD curve type under which this approach performs well is the “Stepped” RD curves, with a mean percentage error of just 3.626%. The reason for this is that under the “Stepped” RD curve type, the optimal decision is generally to restrict generation output in order to produce higher market clearing prices. The effect of this is to naturally conserve fuel, reducing the marginal value of this fuel, and thus bringing it closer to the marginal value of zero implicit in the Naïve Water approach. The worst-case percentage error under all RD curve types for this approach was 67.860%.

Perfectly Competitive Approach:

The Perfectly Competitive approach has a relatively high overall mean percentage error, of 9.393%. The difference between mean percentage errors for the Value Curve approach and the Perfectly Competitive approach can be attributed to market power. The Perfectly Competitive approach performs particularly poorly under the “Stepped” RD curve type (mean percentage error of 19.727%). The reason for this high percentage is that under the “Stepped” RD curve type, the optimal behaviour is generally to restrict output, thus exerting the market power which the Perfectly Competitive approach assumes the generator is unable to exert, and thus ignores. The worst-case percentage error under this RD curve type was 50.829%.

Figure 8.11 summarises the mean percentage error information in Table 8.15, broken down by RD curve structure. It can clearly be seen that the Value Curve approach performs better, on average, than each of the other approaches, under all RD curve structures.

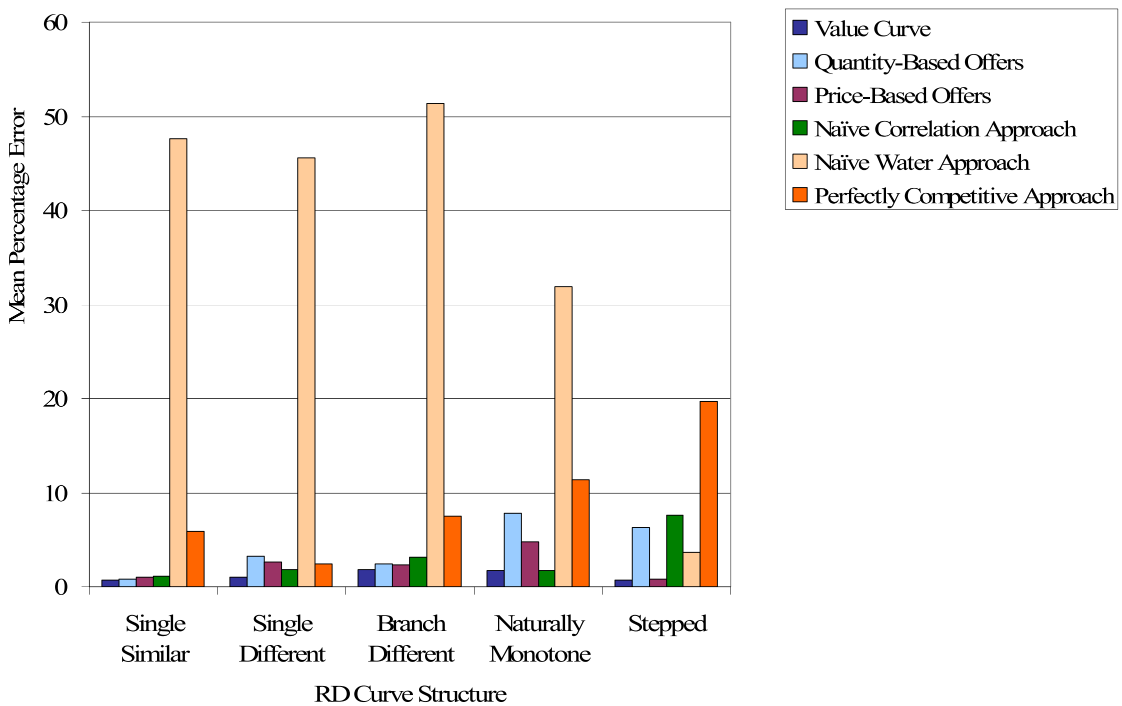

Figure 8.11 Mean Percentage Errors broken down by RD Curve Structure

8.6.2.1 Reasons for the Differences between Performance of Value Curve and Simplified Algorithms

In this section, we explore the reasons for the (sometimes substantial) differences between the quality of the solutions produced by the simplified algorithms considered above when compared to the Value Curve approach, using specific examples from the 1600 replicates on which the analysis has been performed.

Quantity-Based Offers:

We know that the Value Curve approach can trade off between the likelihood of each RD occurring, and thus create a sculpted offer that has relatively good outcomes under each RD curve that could occur in a given scenario. However, under the Quantity-Based Offer approach, we must determine the fixed quantity that will be dispatched, regardless of the RD curve that occurs. Therefore, under middling RD curve outcomes, the offer may perform acceptably, but under very high or very low RD curves, it is likely to do poorly.

Figure 8.12 demonstrates a possible set of offers for period 12. This replicate has a "Single Different" RD curve, and as such, there are a wide variety of RD curves positioned all over the offering space. The figure shows a scenario where the RD curve is relatively low. Point (a) indicates the optimal dispatch point in this period under perfect information, while points (b) and (c) indicate the dispatch points under the Value Curve and Quantity-Based Offer approaches respectively. Observe that the offer from the Value Curve approach results in a dispatch (and price) quite close to that which should be chosen under perfect information, while the dispatch under the Quantity-Based Offer approach is much greater (and price much lower).

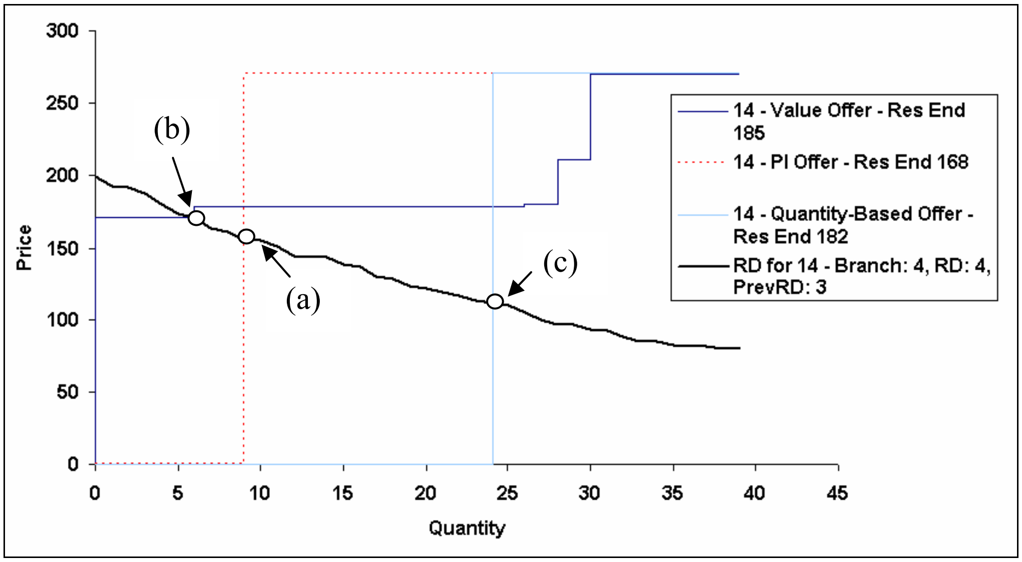

Figure 8.12 Offers for Value Curve and Quantity-Based Offer Approaches – Low RD

On the other hand, Figure 8.13 shows the scenario in period 14 of the same replicate, where the RD curve is relatively high. The offer from the Value Curve approach again results in dispatch (and price) very close to that which should be chosen under perfect information, while the dispatch under the Quantity-Based Offer approach is now much lower (and price much higher).

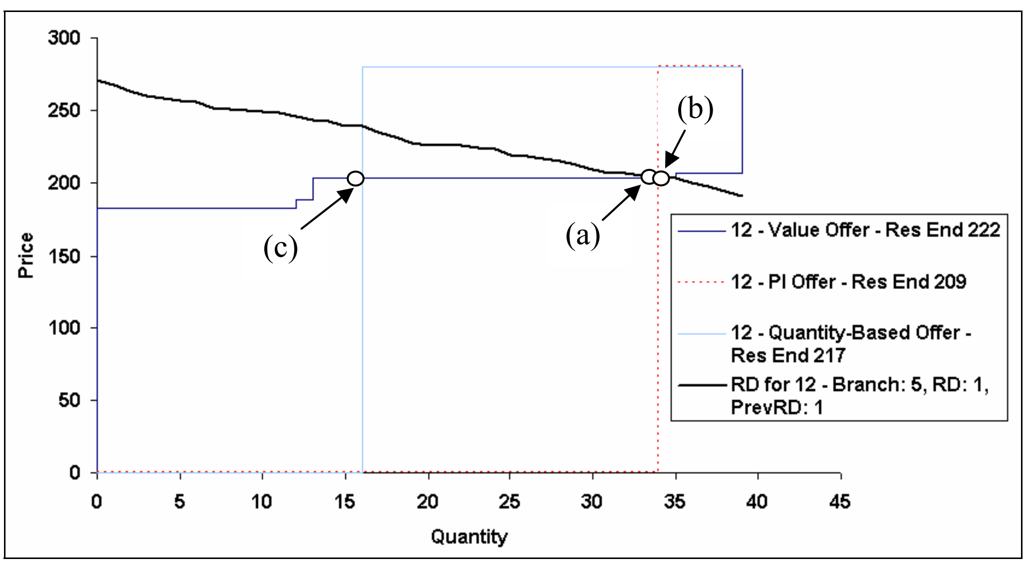

Figure 8.13 Offers for Value Curve and Quantity-Based Offer Approaches – High RD

Price-Based Offers:

As in the Quantity-Based Offer approach, offers under the Price-Based approach must trade-off between the likelihood of various RD curves occurring, and provide the best compromise price level to offer.

Figure 8.14 and Figure 8.15 demonstrate possible sets of offers for periods 16 and 8, under a new replicate, where the RD outcomes are low and high, respectively. Point (c) now indicates the dispatch point under the Price-Based Offer approach. Observe that the offer from the Value Curve approach yet again results in dispatch (and price) very close to that which should be chosen under perfect information, for both these periods. However, when using the Price-Based Offer approach, under a low RD, dispatch is much lower (and price much higher) than the desired point, and under a high RD, dispatch is much higher (and price much lower) than the desired point.

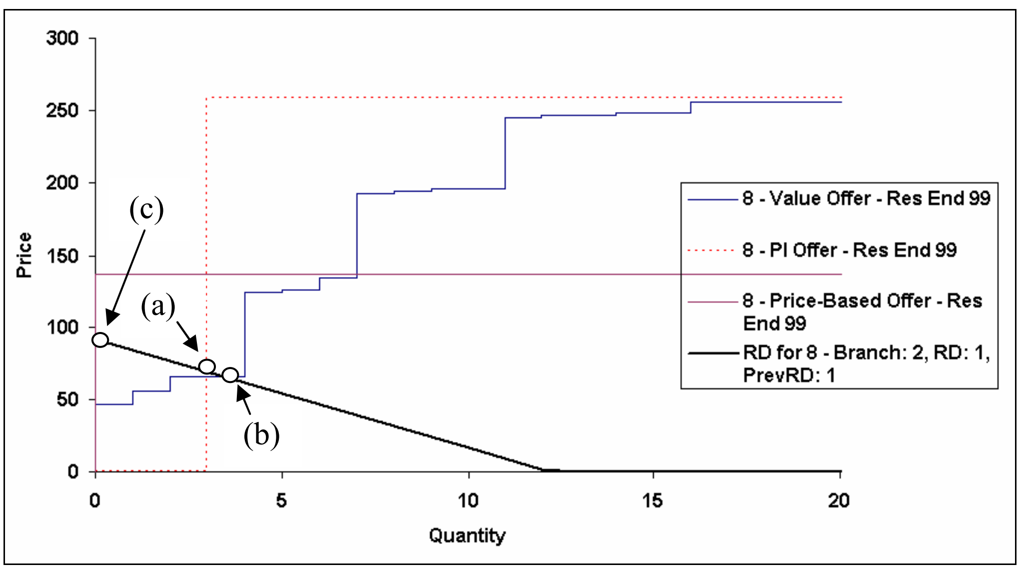

Figure 8.14 Offers for Value Curve and Price-Based Offer Approaches – Low RD

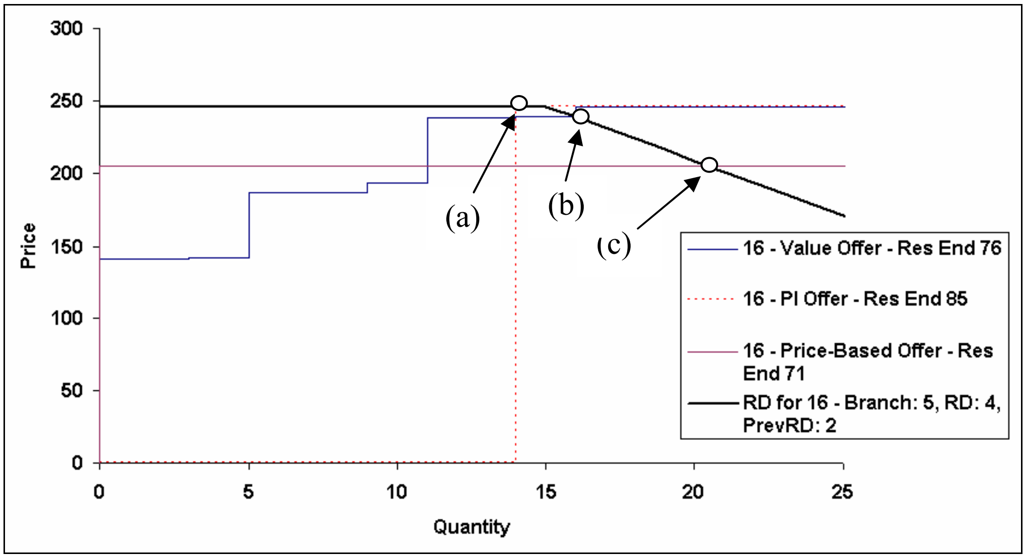

Figure 8.15 Offers for Value Curve and Price-Based Offer Approaches – High RD

Naïve Correlation Approach:

Under the Value Curve approach, when it comes time to make the offer on the coming period, we know which branch (or group of possible RD curves) we are on and so have a specific type (or subset) of RD curves to consider. On the other hand, if we use the Naïve Correlation approach, we do not know which branch we are on when it comes to making an offer for a period, and hence a much larger number of possible RD curves (or portion of the offering space) need to be considered[65].

One simple interpretation of considering correlations is to say that if we observe a low RD curve in one period, then it is very likely that we will observe a low RD curve in the following period as well. Therefore, we would be more inclined to provide a relatively restrictive offer in that following period. If, however, we did not consider these correlations, we would believe that a high RD would be more likely, and the offer would be adjusted accordingly.

[65] This would be even more of an issue if the approaches were modified to allow only limited tranches.

Figure 8.16 demonstrates a possible set of offers for period 10 under a new replicate. Points (c) now indicate the dispatch points under the Naïve Correlation approach. This replicate is of RD curve type “Stepped”, and thus provides the generator with opportunities to apply market power, if they have sufficient knowledge of the market conditions as to exploit it. Under perfect information and the Value Curve approach, generation capacity would be held back such that dispatch would occur before the large drop-off in the RD curve, thus raising the market clearing price substantially. However, as a result of ignoring correlations between market outcomes, the Naïve Correlation approach has not exploited this RD curve step, and his lead to dispatch (c) at the bottom of the step, with only marginally higher quantity, but a much lower price.

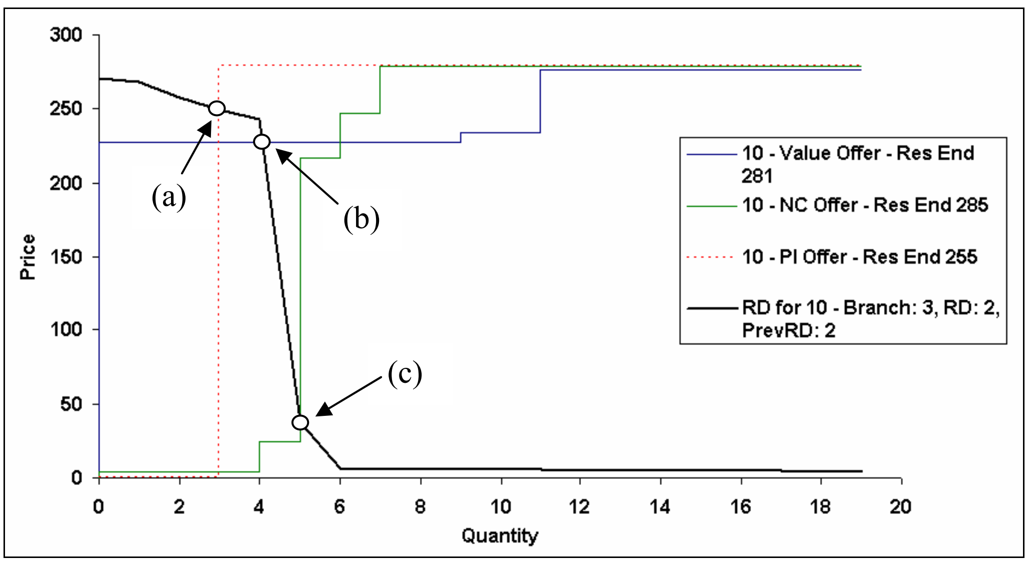

Figure 8.16 Offers for Value Curve and Naïve Correlation Approaches – Low RD

Figure 8.17 demonstrates the opposite scenario, where the actual RD curve does not step down in price until a very high dispatch level is reached. This behaviour is expected under perfect information and the Value Curve approach, and so high dispatch is exploited at high prices. Under the Naïve Correlation approach, we would not realise that we would be able to achieve a high price for such a high dispatch level, and so a more restrictive offer would be provided, leading to a much lower dispatch level (c).

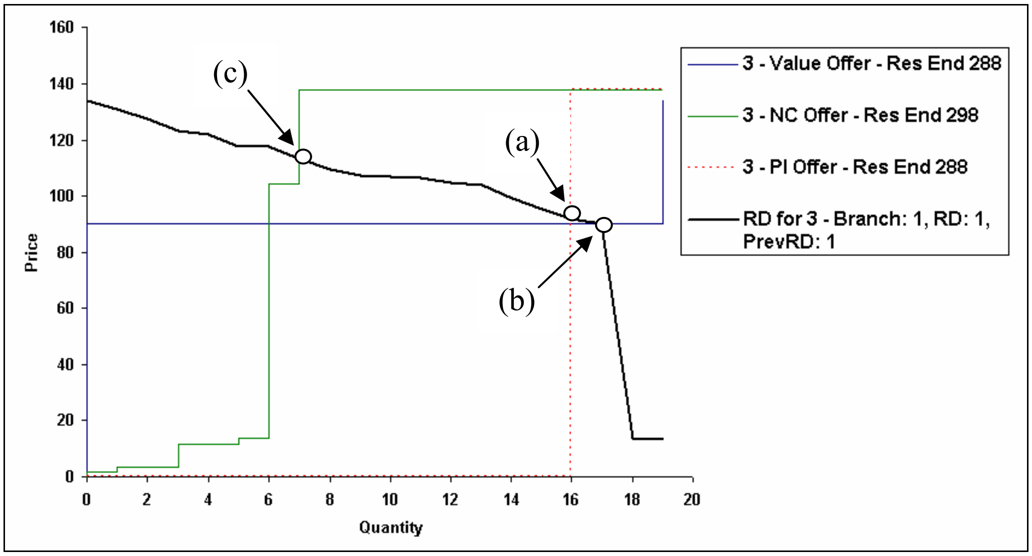

Figure 8.17 Offers for Value Curve and Naïve Correlation Approaches – High RD

Figure 8.18 shows another interesting scenario, in period 11 of a different replicate. Because of the flatness of the possible RD curves on this branch, the MV of dispatch remains high throughout the dispatch range, rather than falling off and becoming negative at higher reservoir levels. As a result, the dispatch under both Perfect Information and the Value Curve approach is around 15 MW/period. However, because the Naïve Correlation model considers all the RD curves on other branches, most of which have much more rapidly falling MV curves, it provides a much more restrictive offer than what is optimal, and hence leads to a much lower dispatch of only 11 MW/period.

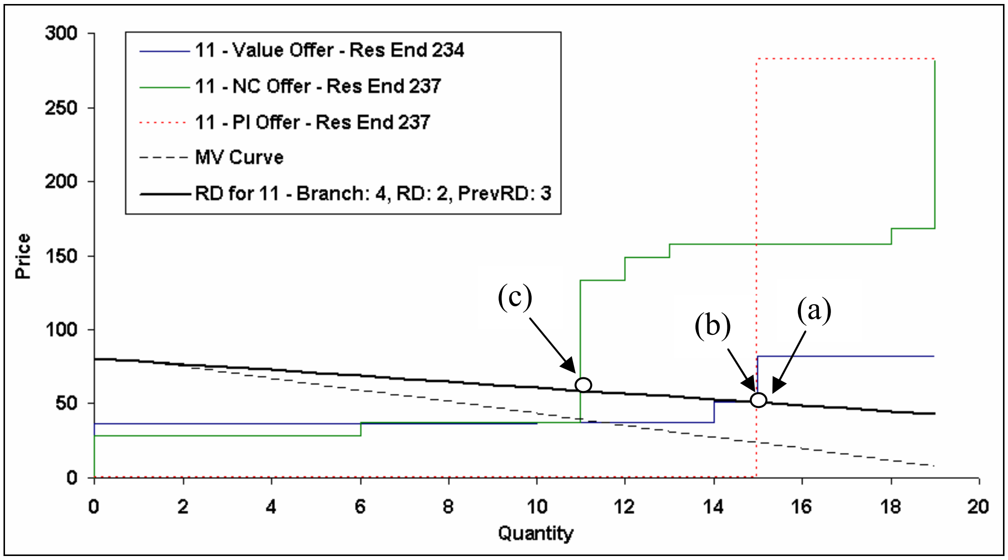

Figure 8.18 Offers for Value Curve and Naïve Correlation Approaches – Flat RD

Naïve Water Approach:

Under the Naïve Water approach, the offers provided to the market by the generator ignore the reservoir level, and as such assume a marginal cost of zero across the entire generating range. This will clearly lead to very generous offers, and thus high levels of dispatch. The result is that fuel resources will be run down very quickly, and little or none will be available to generate electricity later in the horizon. This will be particularly disadvantageous if all fuel is used up in off-peak periods, leaving the generator unable to benefit from higher electricity prices in the peak periods.

Figure 8.19 demonstrates this for period 3 (off-peak) of a different replicate, where (c) is now the dispatch point under the Naïve Water approach.

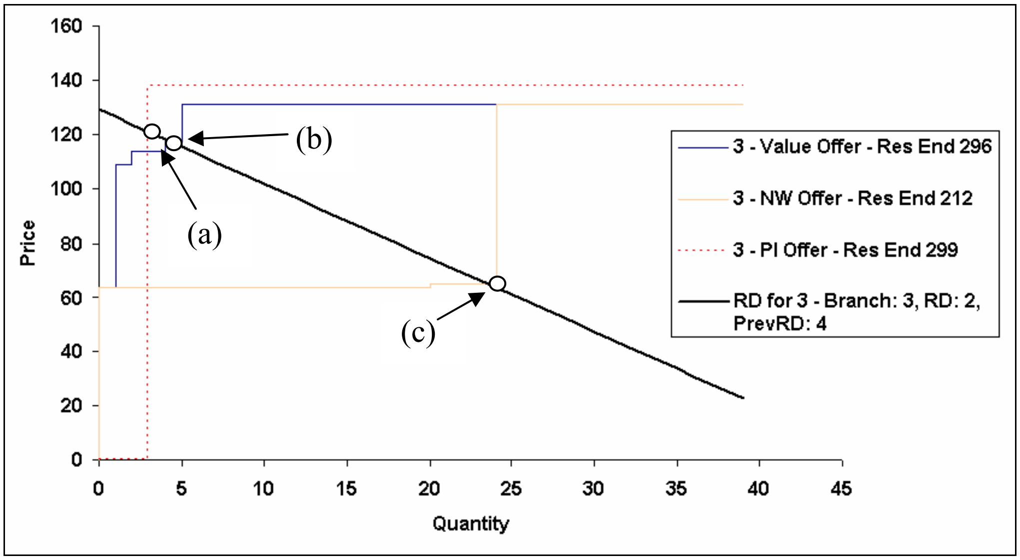

Figure 8.19 Offers for Value Curve and Naïve Water Approaches

Perfectly Competitive Approach:

Under the Perfectly Competitive approach, the generator does not believe it has any market power, and as such, its optimal offering decision is to offer at marginal cost. These offers are clearly more generous than those which would be provided by a generator that is exploiting its market power.

Figure 8.20 demonstrates this for period 7 of a different replicate, where (c) is now the dispatch point under the Perfectly Competitive approach. We can see that the offer is indeed more generous than the offer under the Value Curve approach, and as such dispatch is higher, and market clearing price lower.

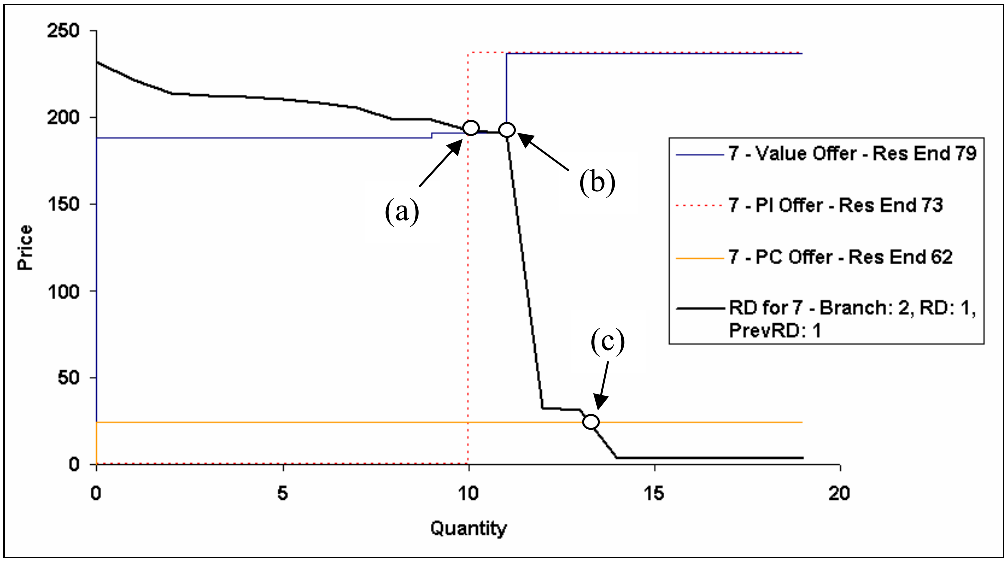

Figure 8.20 Offers for Value Curve and Perfect Competitor Approaches

8.6.3 Summary of Simplified Algorithm Comparisons

In this section of the chapter we have compared the best model from Section 8.3, the Value Curve approach, to various simplifications of this model. These simplifications have included changes to the allowable form of the offers and to the model's approach in dealing with market outcomes, water limitations, and market power. Table 8.16 provides a summary of mean computational times and mean percentage errors for all approaches considered in this section, over all 1600 replicates.

	Approach	Value Curve	Quantity-Based Offers	Price-Based Offers	Naïve Correlation Approach	Naïve Water Approach	Perfectly Competitive Approach
Comp. Time	Mean (seconds)	52	52	410	1	0	50
	vs VC Approach	-	Same	Worse	Better	Better	Same
Soln. Quality	Mean Percentage Error	1.186	4.132	2.304	3.080	36.009	9.393
	vs VC Approach	-	Worse	Worse	Worse	Worse	Worse
	Worst RD Curve Structure	Branch Different	Naturally Monotone	Naturally Monotone	Stepped	Branch Different	Stepped
	Mean Percentage Error under the Worst RD Curve Structure	1.825	7.790	4.801	7.614	51.396	19.727

Table 8.16 Summary of Results for Simplified Algorithms

We have shown that there is no practical difference between the computational time for the Value Curve, Quantity-Based offer, and Perfectly Competitive approaches. Because of its algorithm structure, the Price-Based offer approach actually requires significantly more computational time than the Value Curve approach, while the Naïve Water and Naïve Correlation approaches require zero and almost-zero solving time respectively.

With respect to mean percentage error, the performance of each of these simplified approaches is significantly worse than the Value Curve approach over the full set of replicates. Specifically, each of the simplified approaches performs especially poorly under at least one type of RD curve structure.

The summary in Table 8.16, combined with the more detailed information in Table 8.13 and Table 8.15, implies that the only simplified approaches worth considering as an alternative to the full complexity Value Curve approach are the Naïve Correlation and Naïve Water approaches with their negligible computational times and mean percentage errors of 7.614% and 3.626% under the "Stepped" RD curve type, respectively. The choice between the Value Curve, Naïve Correlation and Naïve Water approaches would depend on the computational time available and the RD curve structure, but we believe that the Value Curve approach, with its mean computational time of 52 seconds, is preferable, as this time is still realistic for use in practical applications, and it avoids the expense of greater deviation from optimality associated with the Naïve Correlation approach (which is particularly large under the most realistic RD curve structure).

8.7 Interpolating the Value Curve under the Value Curve Approach

In this section, we consider the possibility of interpolating the value curve under the Value Curve approach. As explained in Section 6.6, for each period and market state, the Value Curve approach currently determines the optimal offer to provide, and thus the total expected current and future value of being at every possible reservoir level (up to 300 of them in the replicates used in this chapter). Here we propose a simplification of this, whereby these offers and values are determined for just a subset of evenly-spaced reservoir levels, and estimates are produced for the values in between, via simple linear interpolation. Figure 8.21 illustrates this idea, under the assumption of a concave value curve. We can see that the value curve is underestimated using this approach, but there is no consistent bias to the marginal value curve, on which offering decisions are based. As such, we would expect to have a deterioration in solution quality as the resolution of the reservoir levels considered decreases, but it is unlikely that the deterioration will be too large. Here, we consider the interpolation options of solving for every 2^{nd}, 5^{th} or 10^{th} reservoir level.

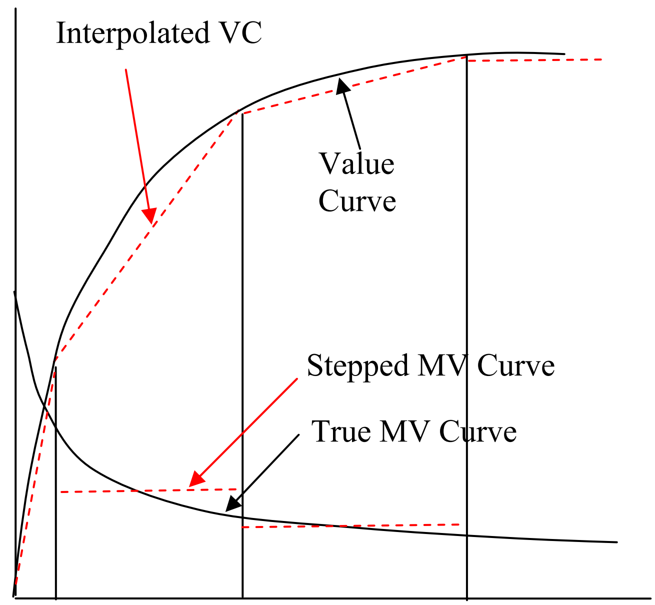

Figure 8.21 Interpolated Value and Marginal Values Curves

The results used throughout this section relate to all 1600 replicates described in Section 8.2.3. All relevant MINITAB output regarding computational times for the interpolation levels can be found in Appendix J, while output regarding solution quality is located in Appendix K.

8.7.1 Algorithm Speed Comparison

A comparison between the overall mean computational times for the original Value Curve approach and the three levels of interpolation models over the 1600 replicates described in Section 8.2.3 is shown in Table 8.17, and plotted in Figure 8.22. Note that Table 8.17 contains a mean computational time for using the interpolation approach solving at all reservoir levels. The difference between this mean time of 67 seconds and the mean time of the original Value Curve approach of 52 seconds can be attributed to the computational overhead associated with applying the value curve interpolation code. For the analysis in this section, we will use the original times, which provide the fairest representation of the computational time achievable with no interpolation.

Model	Mean Computational Time (seconds)
Original Value Curve Approach	52
Interpolation – Every 2nd	42
Interpolation – Every 5th	25
Interpolation – Every 10th	19
Interpolation – Every Level	67

Table 8.17 Mean Computational Times

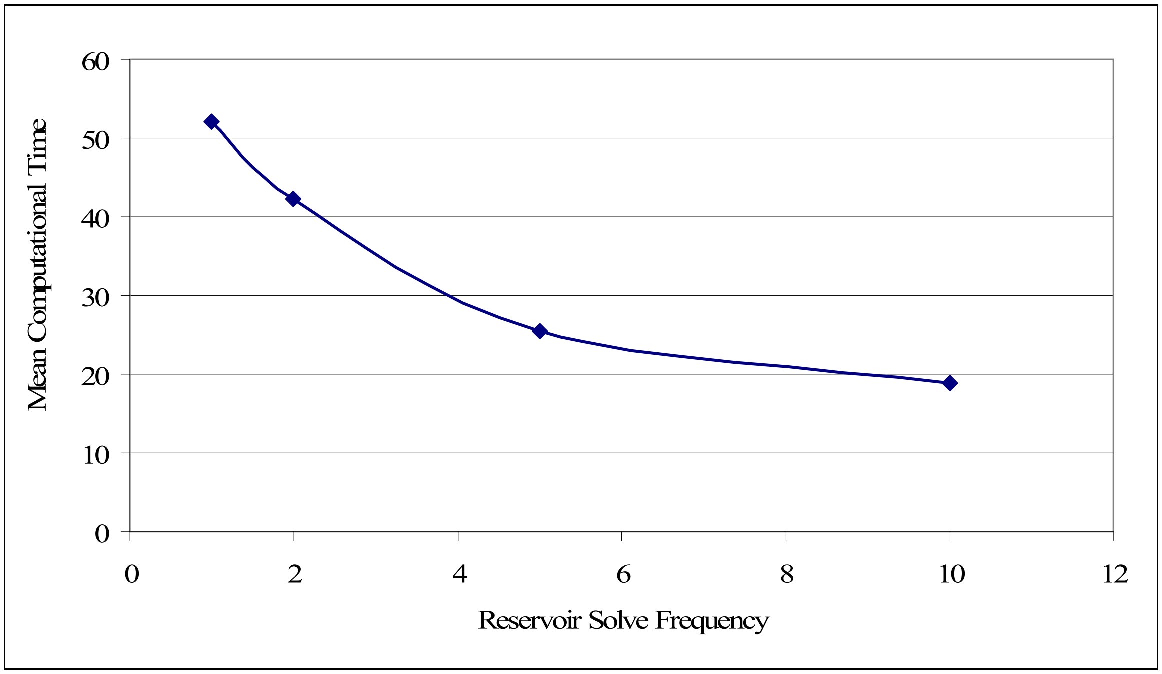

Figure 8.22 Mean Computational Times for Interpolation Frequencies

Appendix J shows that the computational time is falling by a statistically significant amount as we compute the actual value of the value curve at fewer and fewer points. However, we can see that the marginal gains are reducing as the level of interpolation increases. Table 8.18 shows us that the rate at which computational times increase with instance size is fairly independent of the interpolation level. In other words, the computational gains from the interpolation approaches are not lost as the instance size increases.

Mean Computation Times (seconds)	Approach			
	Value Curve	Interp 2	Interp 5	Interp 10
Number of Periods				
20	38	26	16	13
40	66	59	35	25
Number of Reservoir Levels				
100	33	23	15	12
300	72	62	36	26
Number of Dispatch Levels				
20	37	26	16	12
40	68	58	35	26
Number of RD Curves				
25	30	21	13	9
100	74	64	38	28
Number of Inflows				
1	51	40	24	17
5	53	44	26	20
Spread Type				
Weak	52	42	26	19
Strong	52	42	25	19
RD Type				
Single Similar	51	43	26	19
Single Different	51	41	25	18
Branch Different	54	42	27	19
Naturally Monotone	52	40	25	18
Stepped	53	45	25	20
Overall Mean	52	42	25	19

Table 8.18 Mean Computational Times broken down by Factor and Interpolation Frequency

8.7.2 Solution Quality Comparison

As explained above, we would expect that as the resolution of the value curve decreases, the mean percentage error of the approach should increase. A comparison of these mean errors for all interpolation levels, including the original Value Curve approach with no interpolation, is shown in Table 8.19, and plotted in Figure 8.23.

Model	Mean Percentage Error (%)
Original Value Curve Approach	1.186
Interpolation – Every 2^{nd}	1.242
Interpolation – Every 5^{th}	1.487
Interpolation – Every 10^{th}	1.867

Table 8.19 Mean Percentage Error Comparison: Interpolation

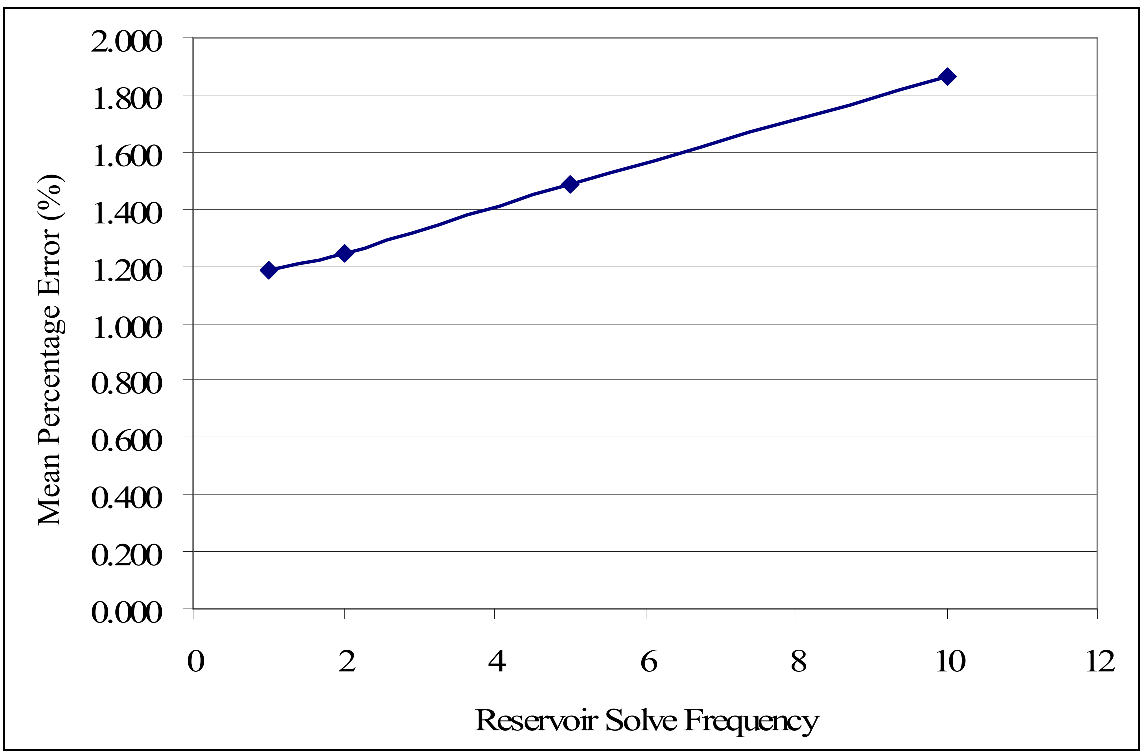

Figure 8.23 Mean Percentage Errors for Interpolation Frequencies

We can see that there appears to be an almost perfectly linear relationship between the level of interpolation and the mean percentage error that results, within the ranges considered. Figure 8.24 illustrates how the shapes of the distributions of the percentage errors change and increase with the interpolation level.

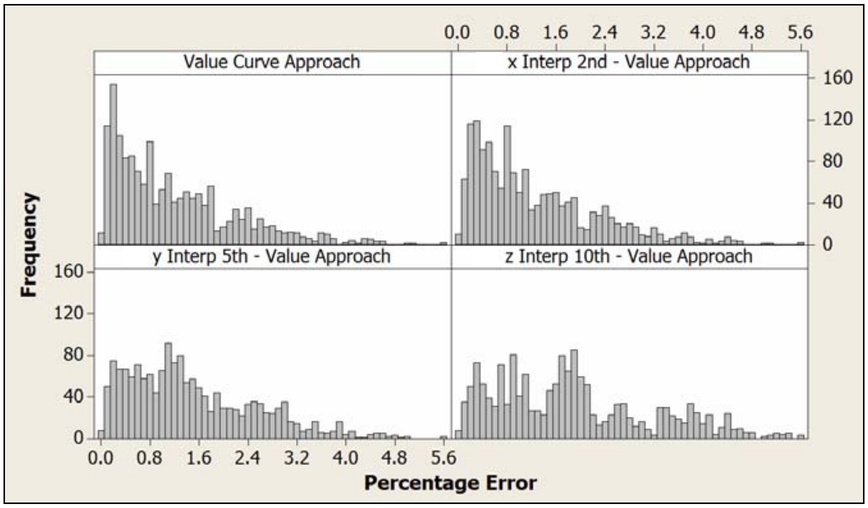

Figure 8.24 Percentage Error Distributions by Interpolation Level

Table 8.20 presents a two-way breakdown of the mean percentage error by interpolation frequency and parameter value.

Mean Percentage Error	Approach			
	Value Curve	Interp 2	Interp 5	Interp 10
Number of Periods				
20	1.109	1.171	1.424	1.827
40	1.262	1.314	1.550	1.907
Number of Reservoir Levels				
100	1.588	1.667	2.006	2.537
300	0.784	0.818	0.968	1.196
Number of Dispatch Levels				
20	1.129	1.188	1.429	1.797
40	1.242	1.297	1.546	1.937
Number of RD Curves				
25	1.194	1.250	1.494	1.872
100	1.177	1.235	1.480	1.862
Number of Inflows				
1	0.823	0.905	1.191	1.420
5	1.549	1.580	1.783	2.314
Spread Type				
Weak	1.117	1.175	1.424	1.814
Strong	1.254	1.310	1.550	1.920
RD Type				
Single Similar	0.688 (c)	0.789 (c)	1.153 (c)	1.737 (c)
Single Different	0.993 (b)	1.086 (b)	1.479 (b)	2.055 (b,c)
Branch Different	1.825 (a)	1.870 (a)	2.107 (a)	2.437 (a,b)
Naturally Monotone	1.686 (a)	1.726 (a)	1.918 (a)	2.252 (a)
Stepped	0.737 (c)	0.741 (c)	0.778 (d)	0.854 (c)
Overall Mean	1.186	1.242	1.487	1.867

Table 8.20 Mean Percentage Errors broken down by Factor and Interpolation Frequency

Some of the significant information presented in Table 8.20 is consistent with that analysed in Sections 8.5.2 and 8.6.2, and so we will not repeat that analysis here. Below we will highlight some new insights from Table 8.20.

Number of Reservoir Levels

We can see that under the instances with 100 reservoir levels, the mean percentage error increases by 59.7% (1.588% → 2.537%) when we compare the Value Curve approach with no interpolation to the case with an interpolation frequency of 10. However, under the instances with 300 reservoir levels, the mean percentage error increases by only 52.7% (0.784% → 1.196%) under the same comparison. In other words, as the size of the instances increases, with respect to number of reservoir levels considered, the cost of performing interpolation on the value curve falls.

Number of Inflows

As stated in Section 8.2.3, the *expected* hydro inflow of 4 MWh/period is the same in all instances, regardless of whether the number of possible inflow levels is one or five. As such, the difference in mean percentage error between these two sets of instances is caused purely by the uncertainty that results from the increased number of possible inflow levels. We can see that as the number of possible inflows is increased from one to five, the mean percentage error increases over half a percent under all levels of interpolation.

The reason for this increased percentage error is as follows. A change in the inflow level moves the marginal cost curve horizontally. It can easily be shown analytically that with a convex marginal cost curve (as we have in this problem and explained in Section 6.6), the vertically weighted average marginal cost curve will be above the MC curve corresponding to the expected value of inflows. In other words, the marginal cost curve is slightly higher in the uncertain case, and as such the resulting offer provided to the market would be slightly higher (more restrictive), and thus marginally less profitable.

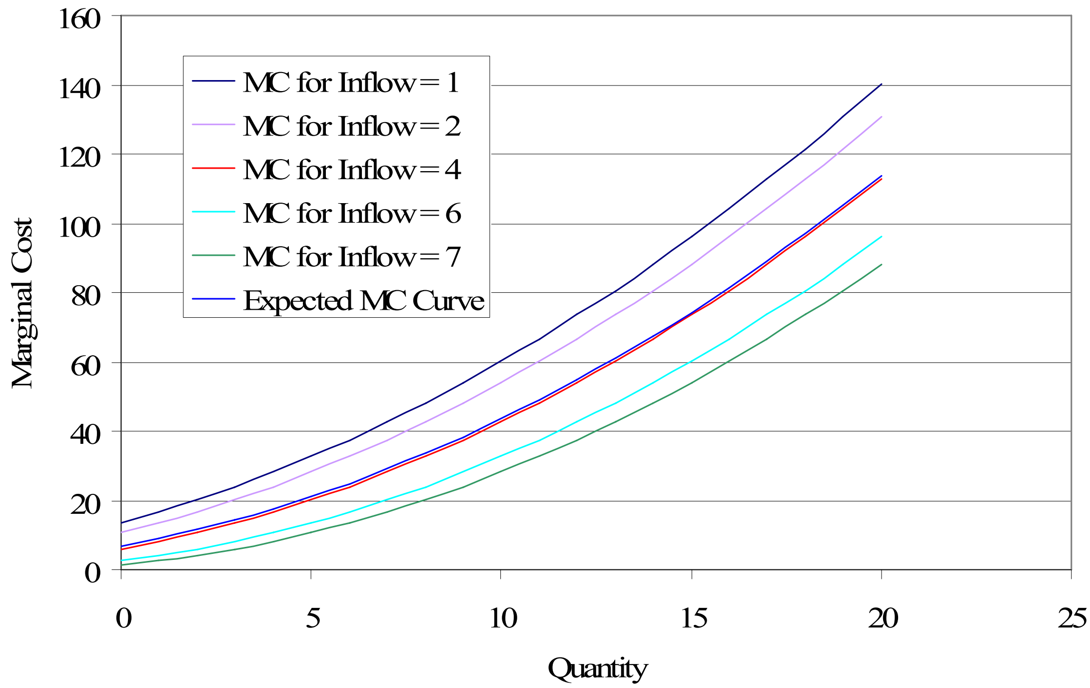

Figure 8.25. Marginal Cost Curves under Inflow Certainty and Uncertainty

Figure 8.25 illustrates this point comparing the expected marginal cost curve over the five possible inflow levels (in blue), and the marginal cost curve associated with the expected inflow level of 4 MWh/period (in red).

8.7.3 Summary of Value Curve Interpolation Comparisons

In this section, we have shown that when interpolation is used on the value curve, computational times reduce but at the expense of greater mean percentage errors. This trade-off is summarised in Table 8.21. Further, we have shown that as the resolution of the interpolation falls, the marginal reductions in computational time fall quickly (see Figure 8.22), while the marginal costs (increases in mean percentage error) are constant (Figure 8.23). This indicates that reducing the resolution beyond the frequency of 10 explored in this section would not produce much further computational time improvement, but the mean percentage error would get progressively worse.

	Approach	Value Curve	Interp 2	Interp 5	Interp 10
Computational Time	**Mean (seconds)**	52	42	25	19
	vs VC Approach	-	*Better*	*Better*	*Better*
Algorithm Performance	**Mean Percentage Error**	1.186	1.242	1.487	1.867
	vs VC Approach	-	*Worse*	*Worse*	*Worse*

Table 8.21 Summary of Results for Interpolation Levels

The best choice of interpolation level from the options above depends on the real-life scenario to which the model is being applied. If the computational time of the Value Curve approach with no interpolation is acceptable in the circumstances, there is no point interpolating the value curve to save time at the expense of a greater mean percentage error. On the other hand, if the scenario has high parameter values such that computational time becomes a factor, then the user should consider one of the interpolation options. Both interpolation levels 2 and 5 provide good trade-offs in this case, while the low computational time saving of an interpolation level of 10 (compared with 5) may not justify the increase to the mean percentage error.

Specifically, compared with not interpolating, the interpolation level of 2 leads to a 19% computational saving at the expense of 0.056% additional mean percentage error, while the interpolation level of 5 would result in a further computational saving of 40% (52% saving on no interpolation) at the expense of an additional 0.245% mean percentage error (0.301% expense compared to no interpolation). However, if the interpolation level of 10 were used, this would result in a further saving of 24% of the computational time, at the expense of 0.380% additional mean percentage error.

8.8 Value Curve Approach on Large Scenarios

In Section 8.2.3, two of the considerations in setting the parameter values related to time and computer memory constraints. In this section, we will consider the feasibility and potential of the Value Curve approach developed in this thesis to be applied to larger problems than those considered throughout this chapter. By doing this, we are recognising that:

- The restrictions placed on size of the instances used in this chapter due to computational time reflected the large solving times under the R&A and Direct MV Curve approaches, not the Value Curve approach.
- The main constraint on instance sizes used in this chapter due to computer memory restrictions are caused by large multi-dimensioned array variables in the *simulation* phase of the process, rather than the optimal offer determination phase. This simulation phase is needed in this thesis to evaluate the relative optimality of the various approaches, but would not be needed in the real-world in order to just produce the offers.

Note that instances proposed in this Section were run on an Intel Pentium D 3.20GHz machine, with 2GB of RAM available, as opposed to the 1GB considered in this chapter up to this point.

8.8.1 Some Larger Instances

Here, we consider four larger instances, solved only under the Value Curve approach, and based on the “Stepped” RD Curve type, which we believe to be the most realistic. The parameters sizes for each of these instances are listed in Table 8.22.

Parameter	Instance 1	Instance 2	Instance 3	Instance 4
Number of Periods	24	24	24	24
Number of Reservoir Levels	1400	700	500	300
Number of Dispatch Levels	40	80	60	35
Number of Residual Demand Curves per Period	25	25	50	200
Number of Possible Inflow levels per Period	1	1	1	1
Transition Correlation type	Medium	Medium	Medium	Medium
RD Curve Structure type	Stepped	Stepped	Stepped	Stepped

Table 8.22 Large Instance Parameter Values

Table 8.23 presents the results from these four instances, based on an interpolation level of 5.

Instance #	Percentage Error (%)	Computational Time[66] (seconds)
1	0.502	52
2	0.567	41
3	0.567	46
4	0.895	98

Table 8.23 Results for Large Instances

Quite clearly, despite the fairly large increases in instance size that these four cases represent, computational times remain feasible. In addition, the percentage errors remain very small.

[66] Note that computational time is approximately linear in the number of periods considered.

8.8.2 Potential Instance Sizes without Simulation

Through experimentation, I have been able to determine the approximate maximum sized instances that could be solved on a computer with 2GB of RAM available. The size limits are based on a combination of the sizes of various parameters, and as such certain trade-offs could be made to increase one parameter size at the expense of another. Table 8.24 shows two example combinations of parameters that are feasible within this available memory.

Parameter	Example 1	Example 2
Number of Periods	Unlimited	Unlimited
Number of Reservoir Levels	50,000	110,000
Number of Dispatch Levels	170	650
Number of RDs per Period	500 (10 branches of 50 RDs each)	400 (40 branches of 10 RDs each)
Number of Possible Inflow Levels per Period	Unlimited	Unlimited
other…		
Number of MC Levels that OCF is solved at	Unlimited	Unlimited

Table 8.24 Maximum Instance Size Examples

These examples show that the algorithm could quite clearly account for significantly larger instances than those considered throughout this thesis, and would certainly be able to be applied to a wide variety of generating units found in the markets of New Zealand, Australia, and around the world. Of course, with larger instances sizes comes a corresponding increase in computational time required, but in this chapter we have shown that the rate at which this increases with instance size is significantly less than under alternative algorithms. In real-world situations, the trade-off between instance size and computational time would need to be made by each company using the algorithm, depending on their specific needs.

Note that there is also some potential to exploit parallel processing of different components of the algorithm (or different branches) on separate CPUs in order to improve the size of the models that can be considered with respect to both memory limits and computational time available.

8.9 Conclusions

Figure 8.26 presents a summary of the mean computational times and mean percentage errors under all possible algorithms and simplifications considered in this chapter[67], as shown in Sections 8.4.3, 8.5.3, 8.6.3, and 8.7.3.

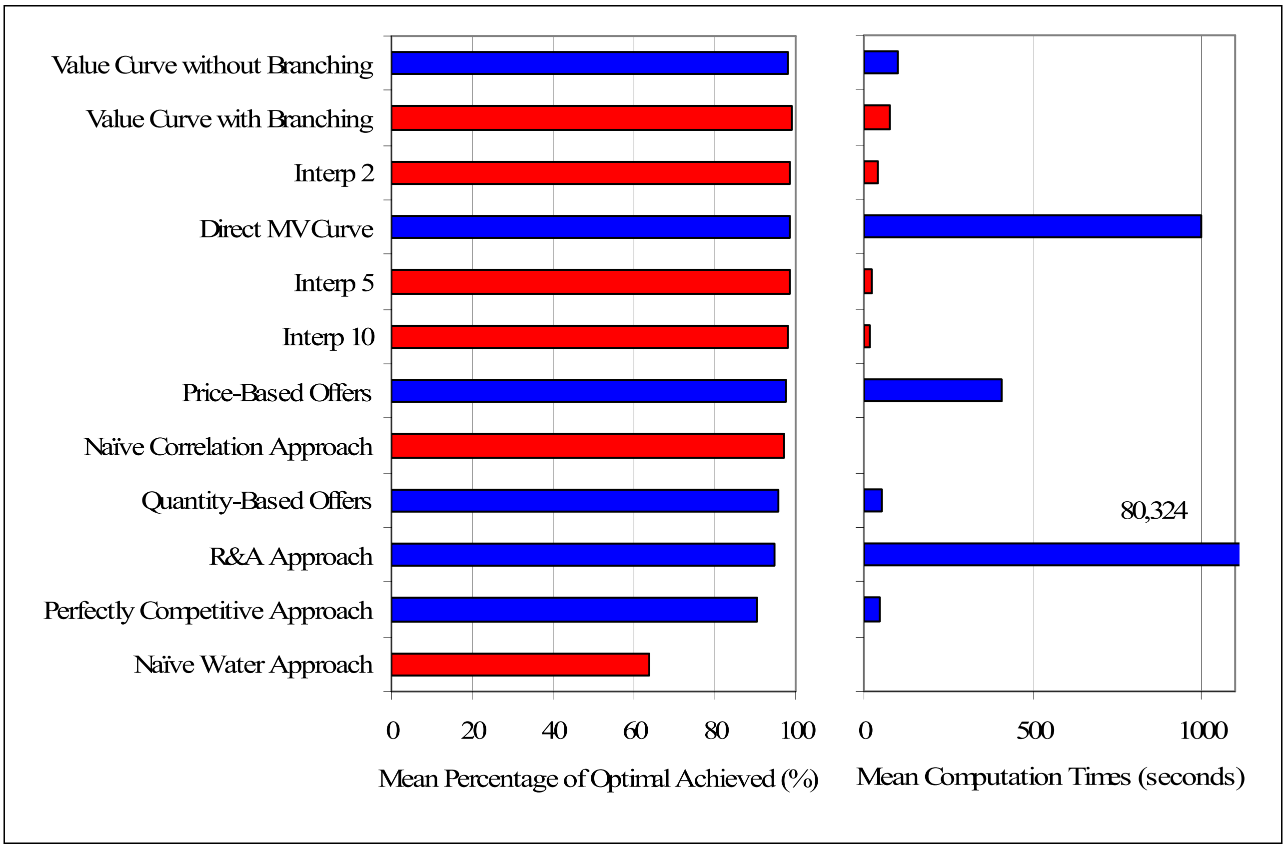

Figure 8.26 Summary of Performance Measures for all Algorithms

[67] Note that the values for the Direct MV curve and R&A approaches are based on the 320 replicates described in Section 8.4, while the rest are based on all 1600 replicates.

Those approaches highlighted in red are on the efficient frontier of the possible approaches, as demonstrated in Figure 8.27. In other words, for each of these algorithms, there is no alternative algorithm that is preferable under both performance measures (that *dominates* this algorithm).

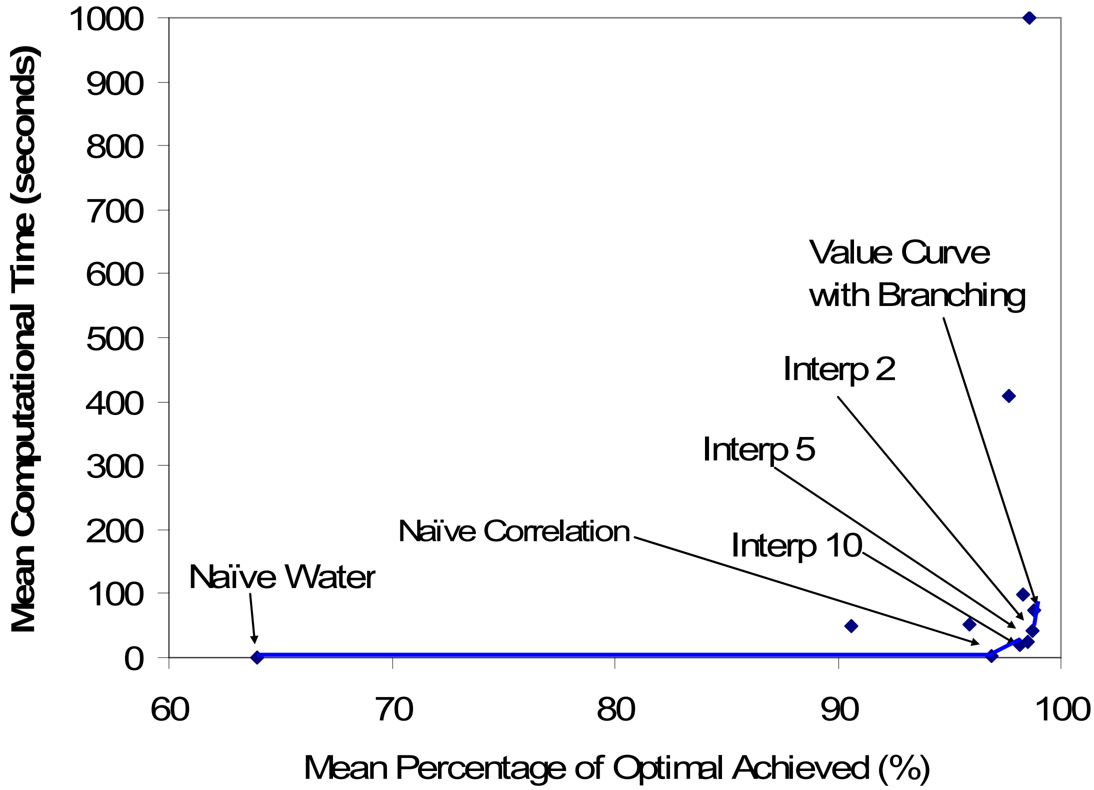

Figure 8.27 Efficient Frontier of Approaches

From this efficient frontier, we can conclude that there are only six algorithms worth considering, and the selection between them is dependent on the specific scenario that is being faced. Depending on the time available and the size of the instance being solved, the best choice of algorithm (from slowest with best mean percentage error to fastest with worst mean percentage error) will be either the:

1. Value Curve approach with branching and with no interpolation,
2. Value Curve approach with branching, solving only every 2nd reservoir level and interpolating the value curve between these points, and

3. Value Curve approach with branching, solving only every 5th reservoir level and interpolating the value curve between these points, and
4. Value Curve approach with branching, solving only every 10th reservoir level and interpolating the value curve between these points, and
5. Naïve Correlation Approach.
6. Naïve Water Approach

The detailed trade-offs between computational time and solution quality for these algorithms are presented in Table 8.25.

	Approach	Value Curve	Interp 2	Interp 5	Interp 10	Naïve Correlation	Naïve Water
Computational Time	**Mean (seconds)**	52	42	25	19	1	0
	vs VC Approach	-	*Better*	*Better*	*Better*	*Better*	*Better*
Algorithm Performance	**Mean Percentage Error**	1.186	1.242	1.487	1.867	3.080	36.009
	vs VC Approach	-	*Worse*	*Worse*	*Worse*	*Worse*	*Worse*
	Worst Case	5.601	5.601	5.599	5.602	22.390	51.396

Table 8.25 Summary of Results for the Best Algorithms

Finally, given sufficient time resource, we have shown that it is beneficial to consider all of the following aspects of the generator offering decision process:

- The underlying branched nature of the real-world
- Sculpted, forward-sloping offers (rather than simple quantity or price-based offers)
- Market outcome correlations
- Water/fuel limitations
- Market power.

Chapter 9

OPTIMAL OFFER CONSTRUCTION ALGORITHM UNDER UNIT OPERATING RULES

9.1 Introduction

In the preceding chapters, we have proposed two two-phase algorithms, and shown that one of them, the Value Curve approach performs particularly well, with respect to both computational time and solution quality. This algorithm considered many of the intertemporal complexities faced by a fuel-constrained generator, including market uncertainty and correlation, uncertain inflows, and fuel conservation. In this chapter, we extend the model to deal with a whole new set of intertemporal constraints, regarding the operating rules of generating units. The rules and conditions that we extend the model to include are:

1. Minimum feasible generation level
2. Start-up and shut-down costs
3. Minimum up and down times
4. Fixed start-up process
5. Ramp rate restrictions

The fixed start-up process defines the generation levels required over the *s* periods that it takes the particular generating unit to start up, while the minimum feasible generation level is the lowest feasible generation once the unit has completed its start-up process. Ramp rate restrictions place a limit on the difference between the levels of generation in consecutive periods.

In Section 9.2 we deal with the first four rules listed, by adding an extra dimension to the DP which defines the state of the unit at any given time (in other words, whether it is off, starting up or on). In Section 9.3 we deal with the ramp rate restrictions by adding a further dimension to the DP, to record the previous generation level (which will affect the range of generation quantities in which we are able to offer in the current period). In Section 9.4, we provide a summary of our findings in this chapter.

9.2 Incorporating Unit Rules

In this section, we consider the addition of the following complexities to our VC model:

1. Minimum feasible generation level
2. Start-up and shut-down costs
3. Minimum up and down times
4. Fixed start-up process

These complexities have not previously been considered together with the types of RD curve uncertainty and correlation that the VC model considers, and as such, the model that we develop in this section contains additional contributions in the offer optimisation field to those developed in earlier chapters.

9.2.1 Unit States Dimension - Simple

Let us start by ignoring the minimum up and down times and the start-up process. Figure 9.1 shows how at each (reservoir level, previous dispatch) state from the original VC algorithm, we can make a decision whether to

- Keep the unit off or turn it on (if it is currently off), or
- Keep the unit on or turn it off (if it is currently on)

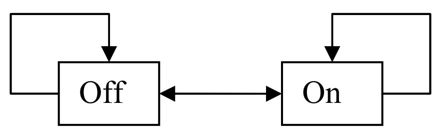

Figure 9.1 Simple Unit State Dimension

Therefore, we now have an extra dimension to the DP, making the state-space three-dimensional. The dimensions are now:

1. Previous market outcome
2. Reservoir level
3. Unit state

In order to describe the new recursion process, we define the following terms:

$V_{t,r}(R_{t,U}, U)$	Value of storage at the beginning of period *t*, at reservoir level R_t, given previous RD curve index *r*, and assuming state *U* at the end of period *t*
SU	Start-up cost (occurs at the moment a generator is switched on)

SD	Shut-down cost (occurs at the moment a generator is switched off)
$R_{t,U}$	Reservoir level at beginning of period *t*, given that the unit has been in state *U* in the previous period
r	Residual demand curve index that occurred in the previous period
p	Residual demand curve index that occurs in the current period
$\theta_t^r\left(R_{t,U}\right)$	The optimal offer provided to the market in period *t* from reservoir level R_t, given the previous market outcome *r* and previous state U, assuming that the unit is on and free to offer as it chooses above its minimum feasible generation level.
$R_{\max}$	Maximum reservoir level at any given time
$Q_t^p\left(\theta_t^r\left(R_{t,U}\right)\right)$	Release in period *t*, given residual demand curve outcome index *p* in this period, and offer $\theta_t^r\left(R_{t,U}\right)$ (which is based on residual demand curve index *r* having occurred in the *previous* period)
$P_t^p\left(\theta_t^r\left(R_{t,U}\right)\right)$	Price in period *t*, given residual demand curve outcome index *p* in this period, and offer $\theta_t^r\left(R_{t,U}\right)$
$pr_{r,p}$	Probability of residual demand curve outcome index p_t, given residual demand curve outcome index *r* in the previous period.
$\mathrm{infl}_{t,i}$	Level of inflow in period *t* under possible inflow index *i*

If the unit is operating in period *t*, then the new reservoir level is defined as:

$$R_{t+1,on} = \min\left(R_{t,U} - Q_t^p\left(\theta_t^r\left(R_{t,U}\right)\right) + \mathrm{infl}_{t,i}, R_{\max}\right) \text{ for } U \in \{on, off\}$$

If the unit is not operating in period *t*, then the new reservoir level is defined as:

$$R_{t+1,off} = \min\left(R_{t,U\in\{on,off\}} + \mathrm{infl}_{t,i}, R_{\max}\right)$$

Therefore, the value curves are constructed as follows:

$$V_{t,r}\left(R_{t,on}, on\right) = \max\left(stayon, turnoff\right)$$

$$= \max\begin{pmatrix} \sum_{p=1}^{P} pr_{r,p}\left[Q_t^p\left(\theta_t^r\left(R_{t,on}\right)\right) * P_t^p\left(\theta_t^r\left(R_{t,on}\right)\right) + V_{t+1,r}\left(R_{t+1,on}, on\right)\right], \\ V_{t+1,r}\left(R_{t+1,off}, off\right) - SD \end{pmatrix}$$

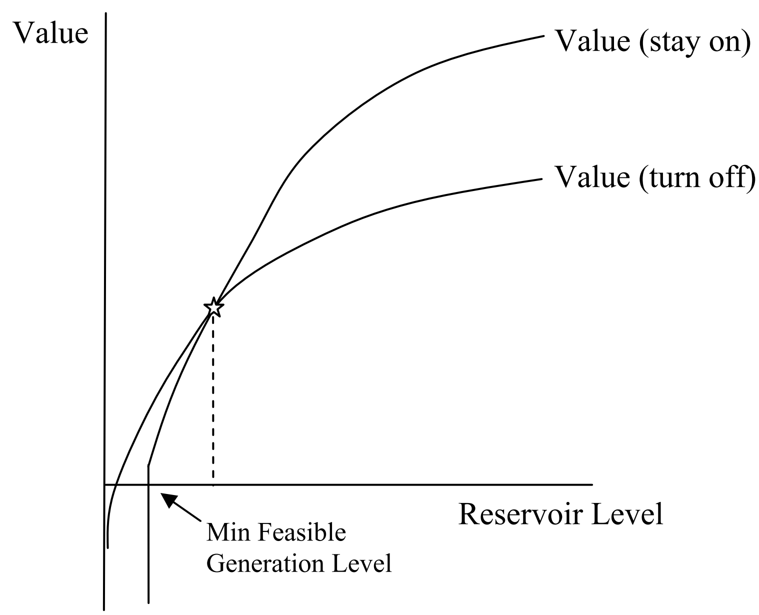

Figure 9.2 Constructing the Value Curve under Unit State Decisions – From On

Figure 9.2 shows an example of how as the beginning reservoir level increases, it starts to become relatively better to keep the unit on rather than switch it off. Note that the position of these two value curves will be dependant on the situation faced and whether or not dispatch in the current period is desired, and as such they will not necessarily cross one another within the reservoir range. Note also that it is infeasible to stay on if the reservoir level is less than the minimum feasible generation level, and as such the value associated with this option is assumed to be $-\infty$.

$$V_{t,r}\left(R_{t,off}, off\right) = \max\left(turnon, stayoff\right)$$

$$= \max\left(\begin{array}{l} \sum_{p=1}^{P} pr_{r,p}\left[Q_t^p\left(\theta_t^r\left(R_{t,off}\right)\right) * P_t^p\left(\theta_t^r\left(R_{t,off}\right)\right) + V_{t+1,p}\left(R_{t+1,on}, on\right)\right] - SU, \\ V_{t+1,p}\left(R_{t+1,off}, off\right) \end{array}\right)$$

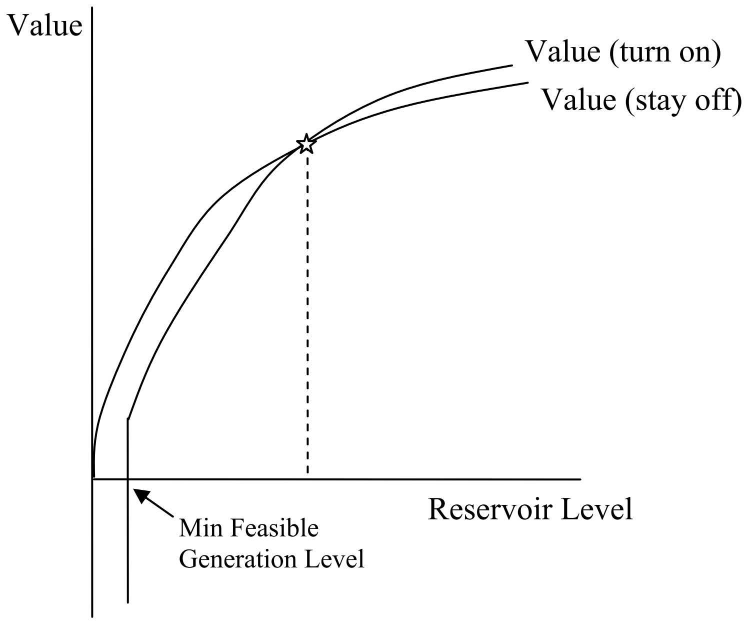

Figure 9.3 Constructing the Value Curve under Unit State Decisions – From Off

Figure 9.3 shows an example for when the unit starts switched off. It becomes relatively better to switch the unit on as the reservoir level increases, as is the case where the unit starts switched on. We note here though that the point of cross-over, where it becomes optimal to switch the unit on/keep it on would be at a higher reservoir level than had the unit started switched on. This is because, at all reservoir levels, the value associated with switching on is lower than the value for staying on, and the value for staying off is higher than the value for switching off, as a result of the *SU* and *SD* costs.

We know from Section 6.6 that if staying on is the only option, the MV curve would be monotonically non-increasing as long as the value curve for the following period is concave when the unit is *on*. We also know that if the only option is for the unit to stay off, this is equivalent to providing an offer from which dispatch will always be zero, and as such, the MV associated with this offer will also be monotonically non-increasing, so long as the value curve for the following period is concave when the unit is *off*. In addition, we know that the *SU* and *SD* costs would be constant terms in the value curves in these two cases, and thus do not affect the MV curves. An example situation using this information is shown in Figure 9.4. Again, these MV curves may or may not cross one another. Note that the MV for the unit being switched on is not shown for values less than the minimum feasible generation level, as this is not a feasible range in which to be switched on.

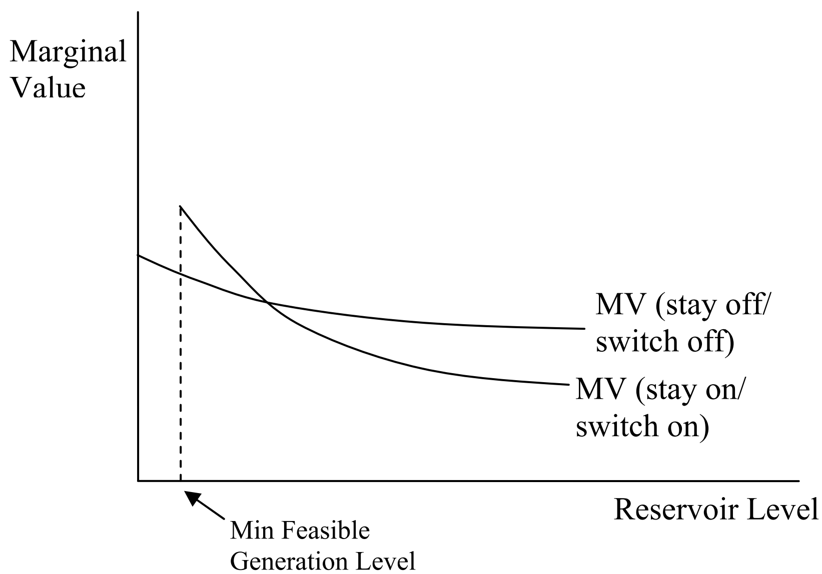

Figure 9.4 MV Curves under Pre-Determined Decision

Because these MV curves associated with each decision are decreasing, the associated value curves (presented in Figure 9.2 and Figure 9.3), must be concave. Figure 9.5 shows how the optimal decision changes from switching the unit off at r* to keeping it on (given the example presented in Figure 9.2), and that this produces a kink in the value curve when we are able to select the optimal option. This kink leads to a MV curve that is no longer monotonically non-increasing.

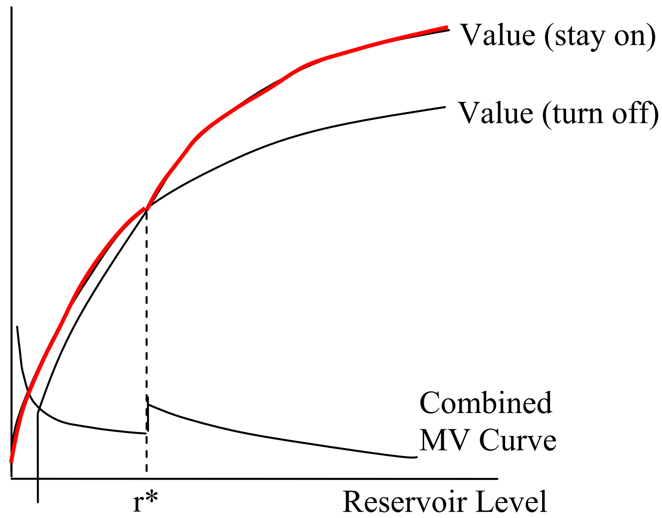

Figure 9.5 Combined MV Curve Construction

As mentioned above, this is only an example position of these respective value curves, but all possible combinations of positions will produce either this same result, with a single range of non-monotonicity in the MV curve, or will produce a MV curve that is naturally monotone.

We noted earlier in this thesis that in order to produce offers from the OCFs that are monotonically non-decreasing, the MV curves must be monotonically non-increasing. As such, a modification needs to be made to this curve, and this is demonstrated in Figure 9.6. Note that this modification is the same as was required for the value and marginal value curves of the basic algorithm, as described in Section 6.6.1.

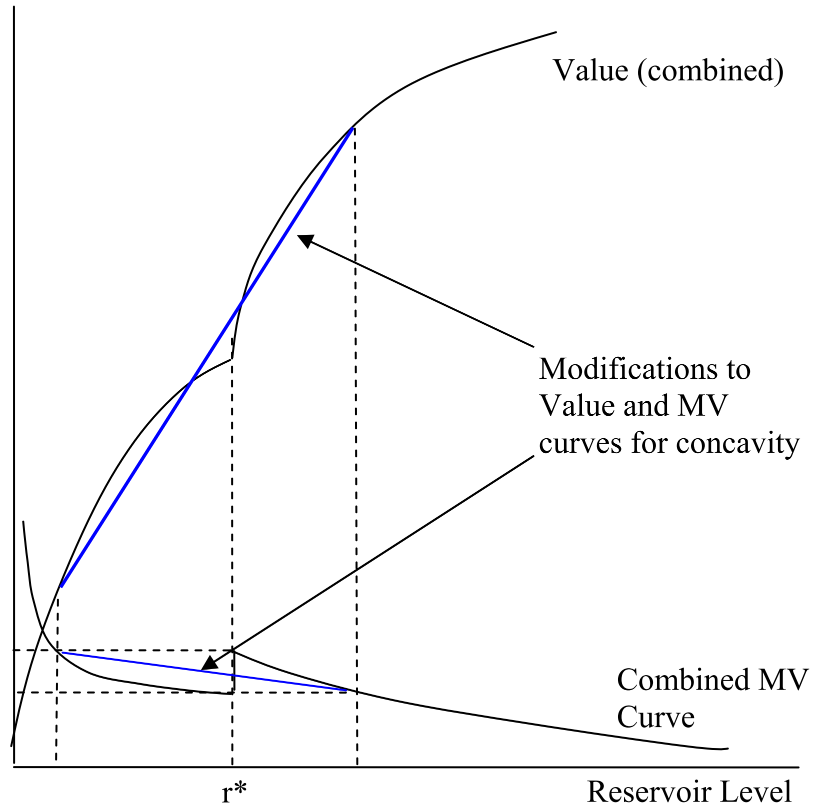

Figure 9.6 Combined MV Curve Construction

The effect of this change is a slight loss in accuracy compared to a purely primal approach. However, to achieve the computational gains from the marginalistic approach, this loss in necessary. The cause of the non-monotonicity is caused by the non-convex costs associated with unit commitment, and is well recognised in the literature, such as in O'Neill, Sotkiewicz, Hobbs, Rothkopf, & Stewart (2005).

9.2.2 Unit States Dimension – Including Minimum Up and Down Times and a Start-Up Process

We will now consider how the dynamic program described in Section 9.2.1 changes if we include minimum up and down times and the start-up process. Figure 9.7 shows an example full unit states problem with a minimum up time of two periods, a minimum down time of two periods, and a two period start-up phase. The green nodes indicate unit states where decisions are made at the start of the period, and correspond to the only two nodes found in the simple unit states example in Section 9.2.1. In the yellow nodes, an offer curve decision is made, but the change in unit state is pre-defined. Clearly there are quite a number of unit states from which no decision needs to be made at all, and thus as the complexity of these aspects of the unit state rules increases, the computational time of the algorithm will be relatively unaffected.

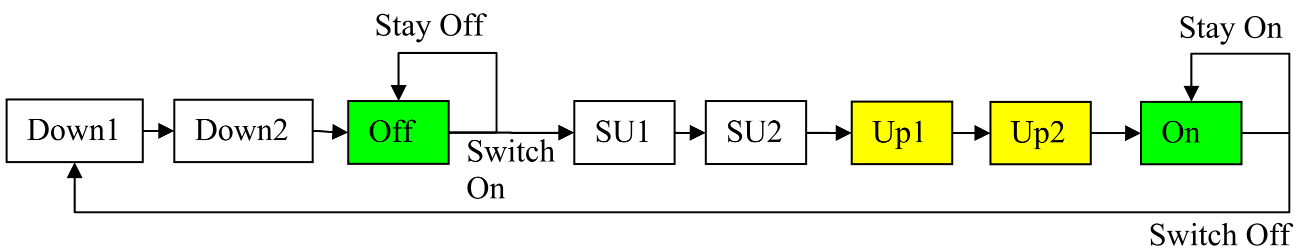

Figure 9.7 Full Unit State Dimension

In order to describe the complete recursion process, we must define the following additional term(s):

SUG_{ss}	Required generation level in phase *ss* of the start-up phase
DS	Number of “down” unit states/minimum down time
US	Number of “up” unit states/minimum up time
SS	Number of start-up states

The reservoir balance equations and DP recursion equations are different depending on which state the unit is in at the beginning of the period.

If the unit is in “Down” state ds at the start of period t ($ds \neq DS$), then the reservoir balance equation and value curves are constructed using:

$$R_{t+1,ds+1} = \min\left(R_{t,ds} + \mathrm{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,ds}, ds\right) = V_{t+1,p}\left(R_{t+1,ds+1}, ds+1\right)$$

If the unit is in the final "Down" state (*DS*) at the start of period *t*, then the reservoir balance equation and value curves are constructed using:

$$R_{t+1,Off} = \min\left(R_{t,DS} + \text{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,DS}, DS\right) = V_{t+1,p}\left(R_{t+1,off}, Off\right)$$

If the unit is in the "Off" state at the start of period *t*, then the reservoir balance equation and value curves (where ss_1 is the first start-up state) are constructed using:

$$R_{t+1,ss_1} = \min\left(R_{t,Off} - SUG_1 + \text{infl}_{t,i}, R_{\max}\right) \text{ or } R_{t+1,off} = \min\left(R_{t,Off} + \text{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,Off}, Off\right) = \max\left(\begin{array}{l}\sum_{p=1}^{P} pr_{r,p}\left[SUG_1 * P_t^p\left(SUG_1\right) + V_{t+1,p}\left(R_{t+1,ss_1}, on\right)\right] - SU, \\ V_{t+1,p}\left(R_{t+1,off}, Off\right)\end{array}\right)$$

If the unit is in "SU" state *ss* at the start of period *t* ($ss \neq SS$), then the reservoir balance equation and value curves are constructed using:

$$R_{t+1,ss+1} = \min\left(R_{t,ss} - SUG_{ss} + \text{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,ss}, ss\right) = V_{t+1,p}\left(R_{t+1,ss+1}, ss+1\right)$$

If the unit is in the final "SU" state at the start of period *t*, then the reservoir balance equation and value curves (where us_1 is the first up state) are constructed using:

$$R_{t+1,us_1} = \min\left(R_{t,SS} - SUG_{SS} + \text{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,SS}, SS\right) = V_{t+1,p}\left(R_{t+1,us_1}, Up_1\right)$$

If the unit is in one of the "Up" states at the start of period t ($us \neq US$), then the reservoir balance equation and value curves are constructed using:

$$R_{t+1,us+1} = \min\left(R_{t,us} - Q_t^p\left(\theta_t^r\left(R_{t,us}\right)\right) + \text{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,us}, us\right) = \sum_{p=1}^{P} pr_{r,p}\left[Q_t^p\left(\theta_t^r\left(R_{t,us}\right)\right) * P_t^p\left(\theta_t^r\left(R_{t,us}\right)\right) + V_{t+1,p}\left(R_{t+1,us+1}, us+1\right)\right]$$

If the unit is in the final "Up" state at the start of period t, then the reservoir balance equation and value curves are constructed using:

$$R_{t+1,on} = \min\left(R_{t,US} - Q_t^p\left(\theta_t^r\left(R_{t,US}\right)\right) + \text{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,US}, US\right) = \sum_{p_t=1}^{P} pr_{p_{t-1},p_t}\left[Q_t^p\left(\theta_t^r\left(R_{t,US}\right)\right) * P_t^p\left(\theta_t^r\left(R_{t,US}\right)\right) + V_{t+1,p}\left(R_{t+1,on}, on\right)\right]$$

If the unit is in the "On" state at the start of period t, then the reservoir balance equation and value curves are constructed using:

$$R_{t+1,on} = \min\left(R_{t,on} - Q_t^p\left(\theta_t^r\left(R_{,t,on}\right)\right) + \text{infl}_{t,i}, R_{\max}\right) \text{ or } R_{t+1,ds_1} = \min\left(R_{t,on} + \text{infl}_{t,i}, R_{\max}\right)$$

$$V_{t,r}\left(R_{t,On}, On\right) = \max\begin{pmatrix} \sum_{p=1}^{P} pr_{r,p}\left[Q_t^p\left(\theta_t^r\left(R_{t,on}\right)\right) * P_t^p\left(\theta_t^r\left(R_{t,on}\right)\right) + V_{t+1,p}\left(R_{t+1,on}, On\right)\right], \\ V_{t+1,p}\left(R_{t+1,ds_1}, ds_1\right) - SD \end{pmatrix}$$

9.2.3 Experimental Evidence of Offer Construction Theories

In Chapter 3 of this thesis, we demonstrated the desired dispatch behaviour of a generator that faces unit rules and deterministic scenarios. In this section we show, through the use of four examples, that these anticipated behaviours translate across to stochastic scenarios as well. Note that this section is not designed to prove that these results will always occur, but rather to show that the model is consistent with the offering behaviour logic under these sample cases.

The first example has one very high-price period in the middle of the off-peak section of the planning horizon, the second example has one relatively low-price period in the middle of the peak section of the planning horizon, while the third example has half the horizon with very low prices (or RD curves), and then suddenly jumps up to high prices for the remainder of the horizon. The final example considers two possible branches (high-price and low-price) that we may traverse over the horizon. The high-price branch produces much more extreme prices in the peak-period than the low-price branch. Note that all scenarios presented are relatively extreme cases, enabling us to demonstrate the effect of the unit rule constraints and operational costs clearly.

Example 9.2.3.1: Anomalous High-Price Period in the Off-Peak Range

In this first example, the general height of the RD curves, or the "peakiness" of each period is shown in Figure 9.8. Clearly, there is one anomalous high-price period (period 11) in the middle of the off-peak section of the planning horizon.

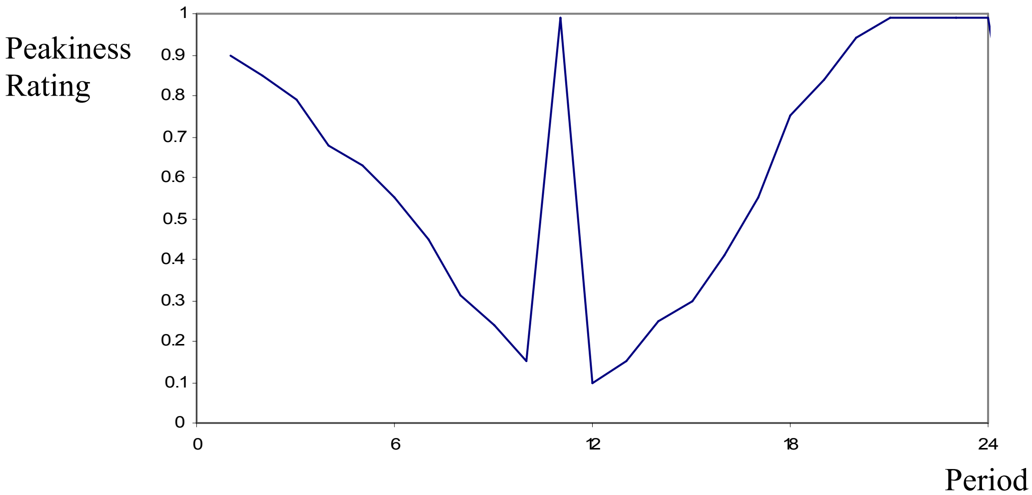

Figure 9.8 Peakiness Ratings for Example 9.2.3.1

To begin with, let us consider how the expected payoff over the 24-period horizon is affected by the imposition of unit rules, and the subsequent adjustment (tightening) of these rules, for a random set of RD curves. For this example, the expected payoff ignoring any unit rules is $35,540. If we were to consider a "base set" of rules (startup and shutdown cost of $0, minimum feasible generation level = 2, two period minimum up and down times, and a two-period start-up phase), then the expected payoff drops to $34,222.80. This information, and a summary of subsequent (and independent) rule adjustments is shown in Table 9.1.

Case	Expected Payoff
No rules	$35,540.95
Base set of rules	$34,222.80
Changes:	
1. Start-up and shut-down costs = $700	$31,110.79
2. Minimum up and down times = 3 periods	$34,038.43
3. Three period start-up phase	$33,922.31
4. Minimum feasible generation level = 4	$32,702.79
Combining all changes	$30,007.32

Table 9.1 Effect on Expected Payoff of Rules

Clearly the addition of constraints has reduced expected payoffs, simply confirming that the model is behaving as would be expected.

In order to get an idea of the structure of the value curves for the different unit states, consider Figure 9.9, which shows the value curves for all states at the beginning of period 22, in the lead up to some very high-price periods.

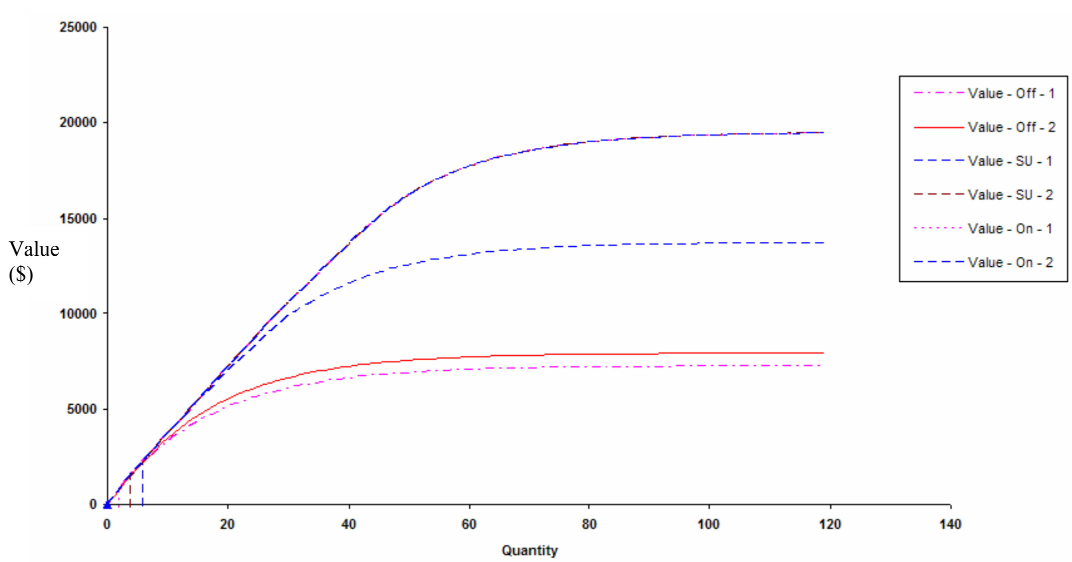

Figure 9.9 Value Curves for all Unit States at Beginning of Period 22

In this case, because of the approaching peak periods, the highest value curves are associated with the states where the unit is fully started up, and thus able to generate within the feasible bounds of generation. Next best are the start-up phase states, as they are on their way to being able to generate freely, while the off states have the lowest value curves (in fact, the value curves are almost identical for the two 'on' states and the final SU state). Intuitively, the value curve is lowest for the off state associated with the unit having just shut down, as this unit has to wait the longest before it is able to generate and take advantage of the forthcoming high-price periods.

Now let us consider the offers and value curves that occur around the anomalous high-price period (t=11). Figure 9.10 shows a subset of the optimal offers in period 11, ignoring any unit rules, and thus assuming the generator can be dispatched at any level, including zero. As it is a high-price period, the offers are relatively generous, leading to positive dispatch in all but the very lowest reservoir levels. On the other hand, Figure 9.11 shows the equivalent subset of offers if the generating unit is switched off at the beginning of period 11, considering the unit rules. In both of these diagrams of offers, and many others presented in the remainder of this chapter, there are many pieces of information in the legend, for each offer plotted. The first number gives the reservoir level to which the offer relates, while the second number is the DP value from the start of this period onwards, given this starting reservoir level. In addition, if the graph is associated with a situation considering unit rules, the note in the legend ends with an arrow followed by a binary variable, indicating whether the unit will optimally be chosen to stay off/turn off (0) or stay on/turn on (1) in this period, from the given state[68]. For the RD curve lines, the information given in the legend are simply the RD curve index for this period, and the probability of that RD occurring under the given UMS.

We can see from Figure 9.11 that the unit remains off, regardless of the reservoir level. This is because if it were to turn on at this point in time, it would be only in the first phase of the start-up process for this high-price period, and the minimum up-time rule

[68] The offers shown in these figures are the offers assuming that the unit is switched on and able to generate freely within its generation bounds: if this is not the case, then the form of the offer is not a decision, but rather clearly defined by the unit operating start-up rules.

would force the unit to then stay on (generating at least the minimum feasible generation level) for the following two very low-price periods.

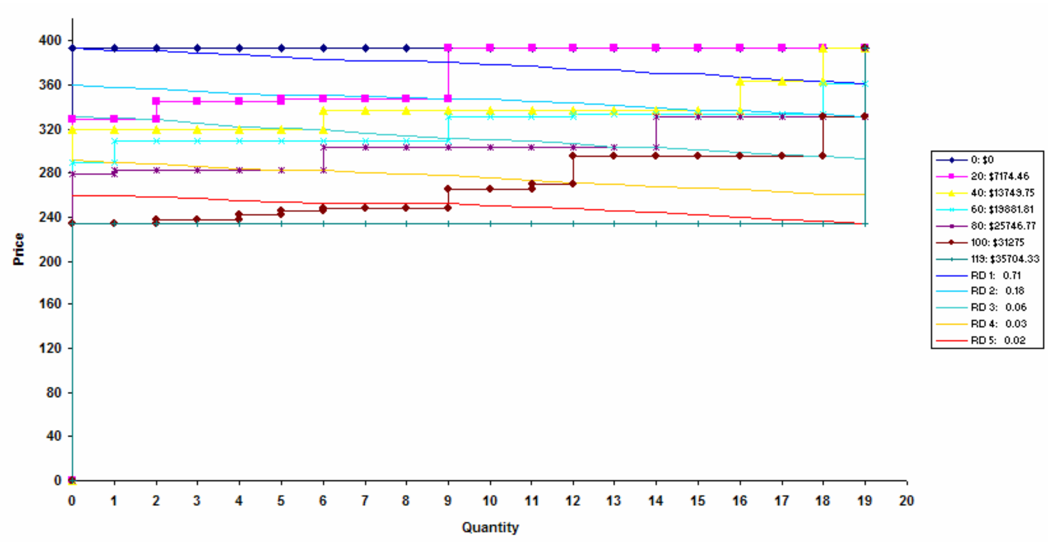

Figure 9.10 Offers for Period 11 Ignoring Unit Rules

Note that the terms in the legend indicate either the reservoir level to which the offer relates along with the expected value (current and future) from this point, or the RD curve index along with its probability of occurring. This is the same for all similar offer diagrams that appear in this chapter.

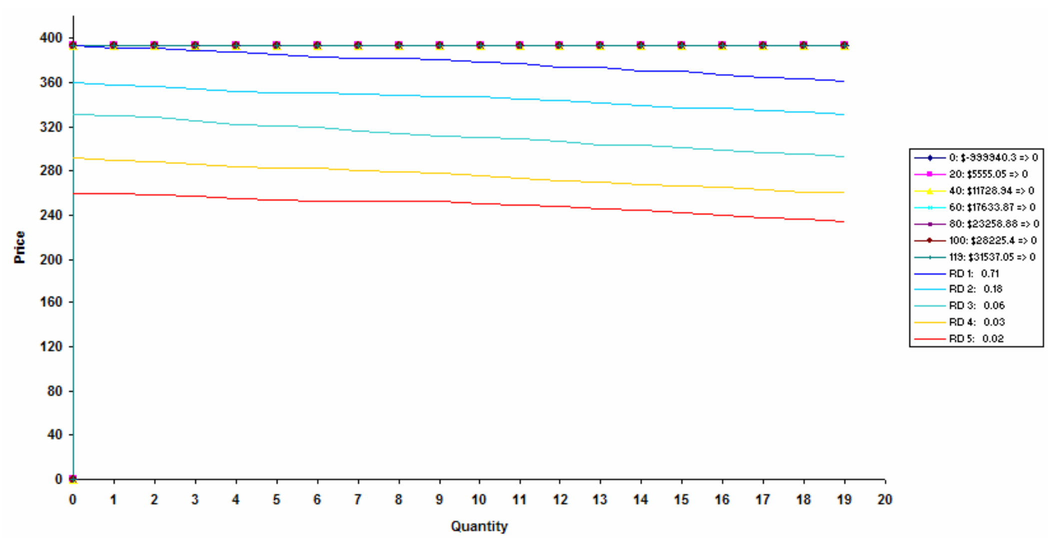

Figure 9.11 Offers for Period 11 Considering Unit Rules – from Off

Figure 9.12 shows the value curves for all possible unit states at the beginning of period 11. The general order of these curves, from highest to lowest is:

1. On (and free to shut down)

   On (and not free to shut down for one period)
2. In start-up phase (and not free to shut down for two periods)
3. Off (and free to start up)

   Off (and not free to start up for one period)
4. In start-up phase (and not free to shut-down for three periods)

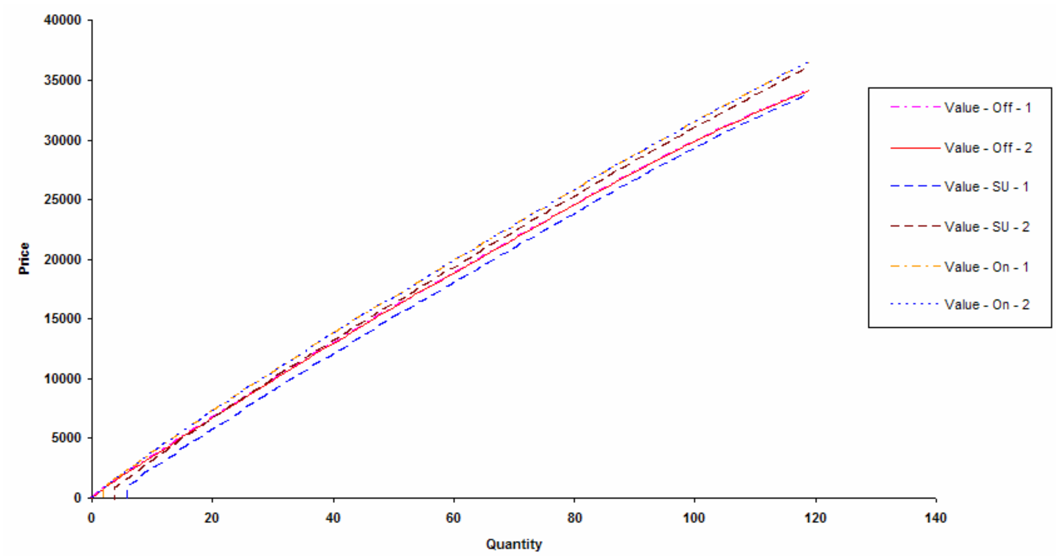

Figure 9.12 Value Curves for all Unit States at Beginning of Period 11

The reason for this order is that it is profitable to generate at a high level in this period, but costly to generate at all in the immediate following periods. The first two states listed enable the generator to capitalise on this high-price period, but then shut down immediately afterwards (as shown by the desired generating level of zero from all reservoir levels in Figure 9.13, which gives the offers assuming that the unit is on but free to switch off if it wishes). The final start-up phase state will mean that, due to the minimum up-time constraint, the generator will have to generate in one unprofitable period before being able to shut-down again. The lowest value curve is associated with the first start-up phase state, because not only is it not able to generate freely in the high-price period, it must stay on for at least two unprofitable periods beyond that point. The two states associated with the unit being off have a payoff slightly above this, even though they do not receive any benefit from the high-price period at all, they conserve fuel for use in later high-price periods, rather than wasting it on low-price periods as forced by the up-time constraint.

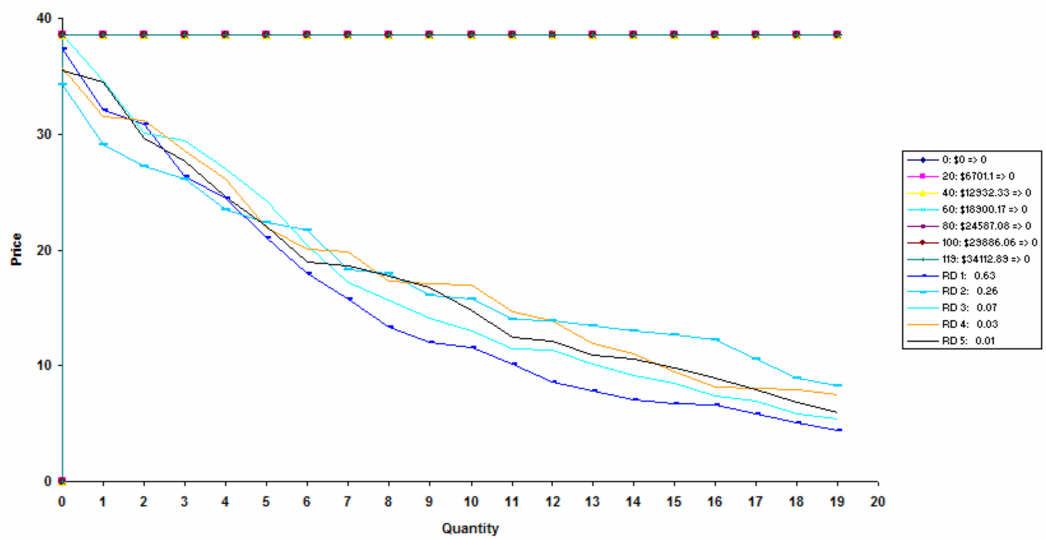

Figure 9.13 Offers for Period 12 Considering Unit Rules – from On

All this would imply that if the unit is switched off in the early periods of the horizon, but wants to take advantage of the high-price period 11, then it would need to begin its start-up phase earlier, in period 9, so that it is free to generate anywhere between the feasible generation bounds in period 11. Figure 9.14 shows that from relatively high reservoir levels in period 9, if the unit is switched off, then it is indeed optimal to start the unit up, proceeding through the start-up phases in periods 9 and 10, enabling a free choice of generation level in the high-price period 11.

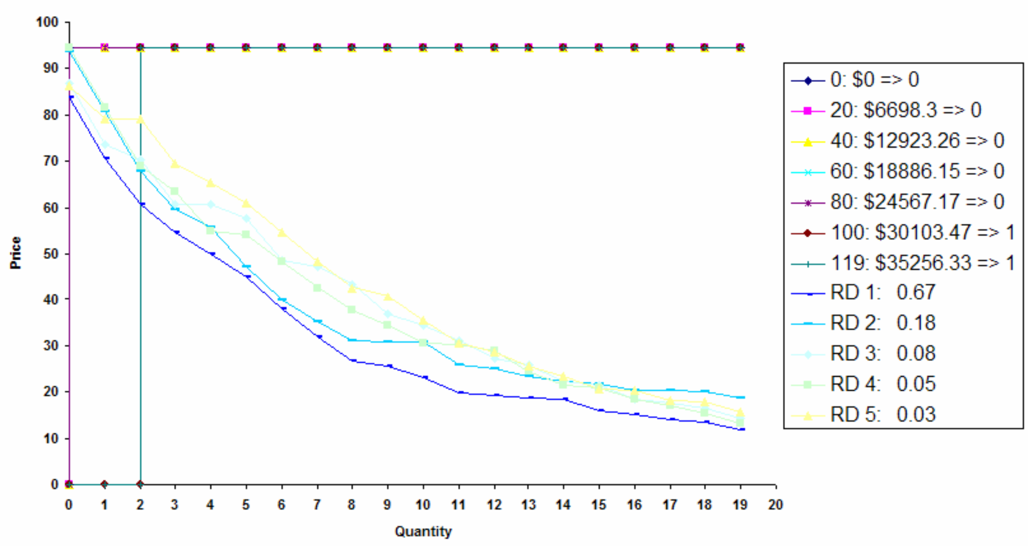

Figure 9.14 Offers for Period 9 Considering Unit Rules – from Off

It is interesting to observe how this decision in period 9 would change if the start-up and shut-down costs increased to $400 each (from the base case of $0). Clearly this discourages the short-term start-up planned to capitalise on the single high-price period, and Figure 9.15 shows that the reservoir level now needs to be higher in period 9 in order for it be optimal to start the unit up at this point.

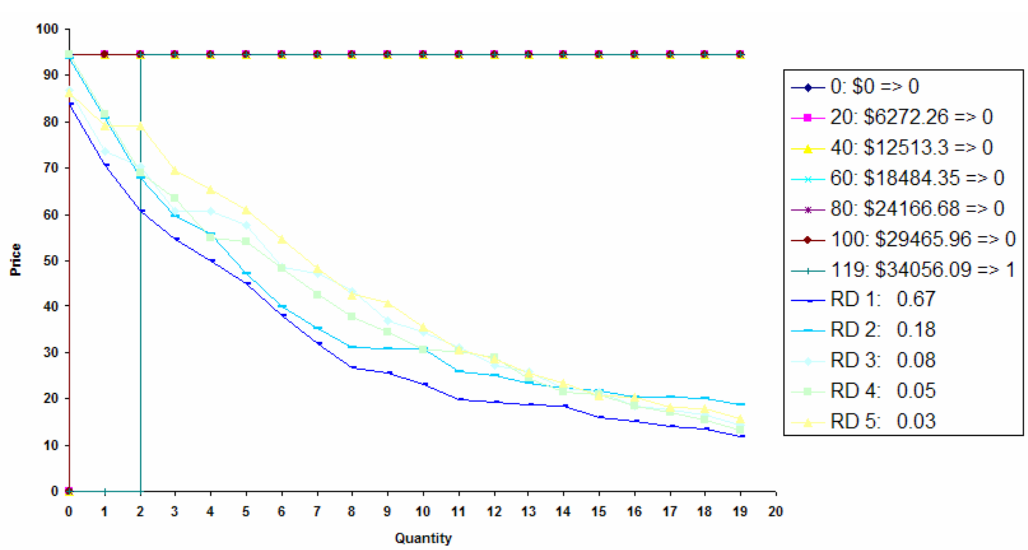

Figure 9.15 Offers for Period 9 Considering Unit Rules – from Off (Higher SU/SD Costs)

Example 9.2.3.2: Anomalous Low-Price Period in the Peak Range

In this second example, the general height of the RD curves, or the "peakiness" of each period is shown in Figure 9.16. Clearly, there is one anomalous low-price period (period 11) in the middle of the peak section of the planning horizon.

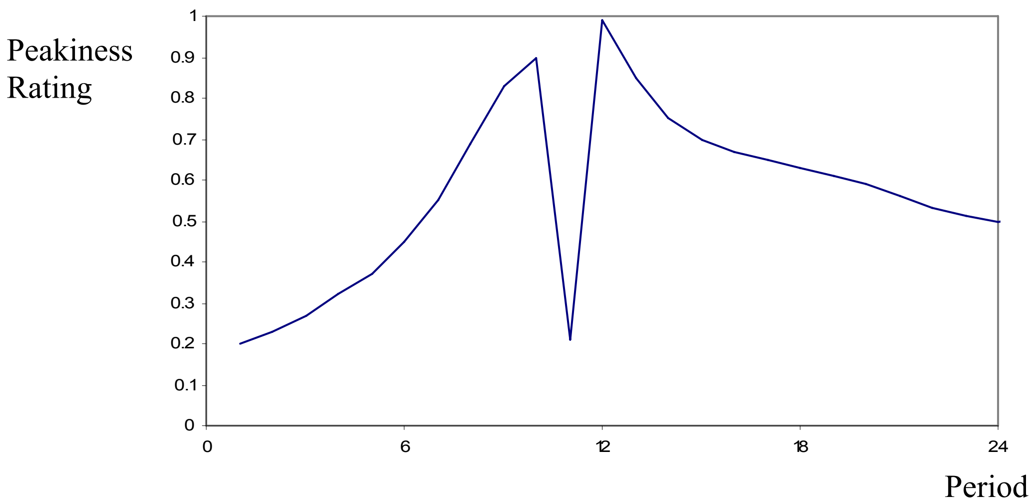

Figure 9.16 Peakiness Ratings for Example 9.2.3.2

Figure 9.17 shows the offers in this very low-price period. We can see that unless the reservoir is almost empty, then the optimal decision is to keep the unit switched on (albeit generating at the lowest feasible level) over this period, despite the large opportunity costs associated with generating at this point rather than later in the horizon. The reason is quite clearly that the unit rules would mean that shutting the unit down near would mean missing out on the potential for high prices in the periods that follow immediately after this one.

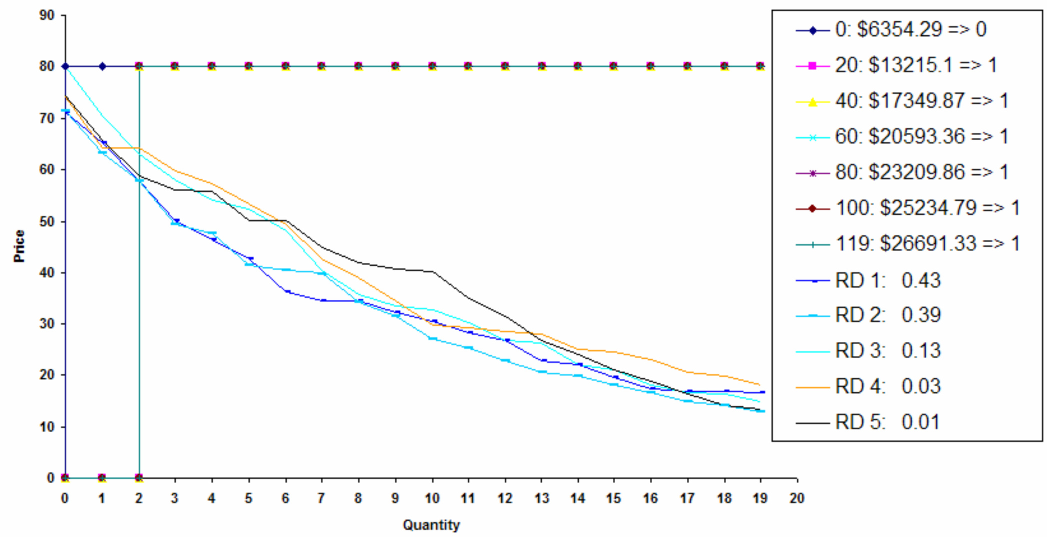

Figure 9.17 Offers for Period 11 Considering Unit Rules – from On

Example 9.2.3.3: Low-Price Periods Changing Suddenly to High-Price Periods

In this third example, the general height of the RD curves, or the “peakiness” of each period is shown in Figure 9.18. We can see that the first half of the horizon has very low prices (or RD curves), but then at period 14, the prices suddenly become very high, and remain that way for the rest of the horizon. A situation such as this may occur if a “macro-state” event was expected to occur at this point in time. A simple example might be that an interconnector becomes binding, isolating the pricing of the region in which the generator operates.

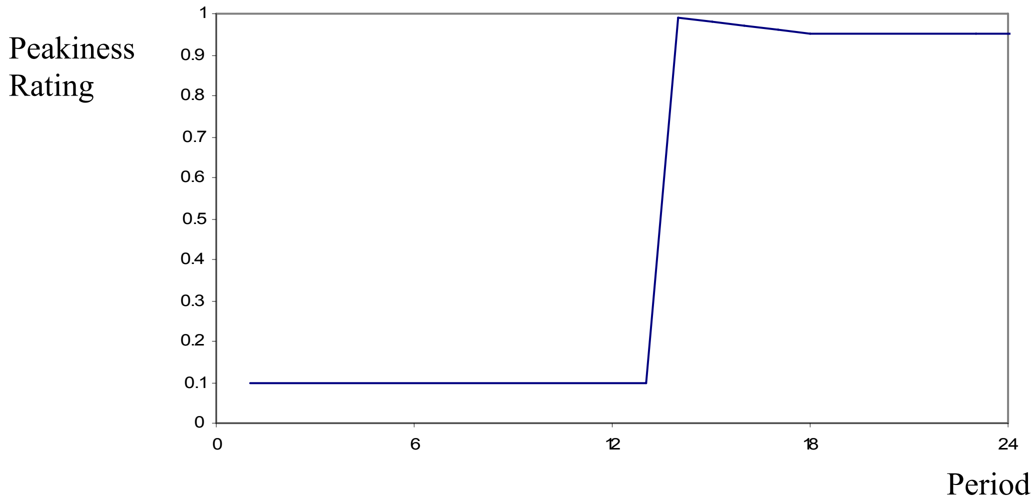

Figure 9.18 Peakiness Ratings for Example 9.2.3.3

This example is similar to Example 9.2.3.1, in that we are considering the matter of building up to generate in a high period, but the difference here is that we wish to keep the unit switched on beyond this point, as high prices continue. Given a two period start-up process, the unit must be switched on in period 12 if it is to be free to generate anywhere within its feasible bounds in the very high-price period 14. Figure 9.28 demonstrates that this switch is indeed the optimal decision in period 12, as long as the reservoir level is sufficiently high[69].

[69] note that the minimum feasible generation level in this example is 5 MW

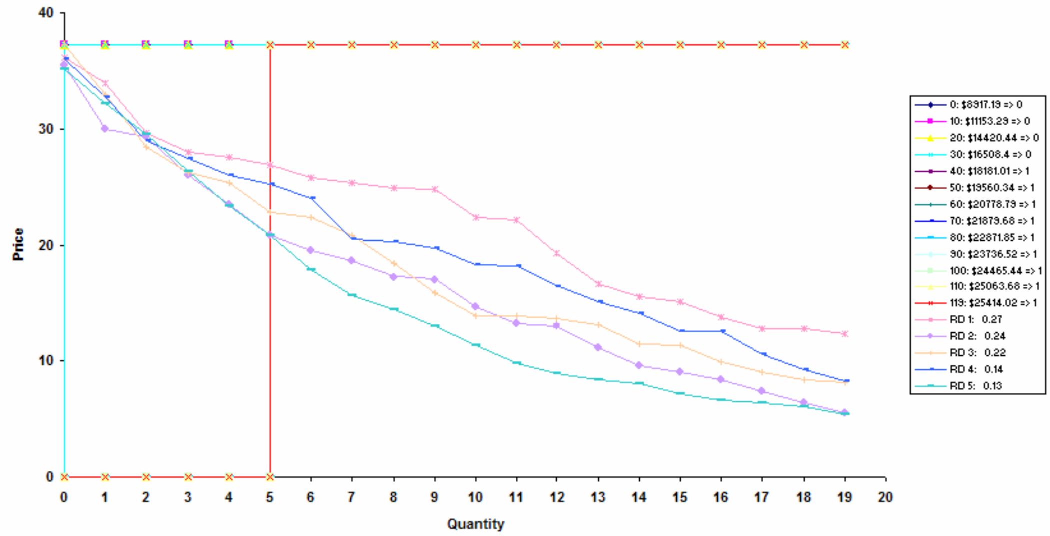

Figure 9.19 Offers for Period 12 Considering Unit Rules – from Off

In all periods prior to this it is only optimal to start the unit if the reservoir level is very high and the alternative is to spill the water/fuel that is arriving into the reservoir, as demonstrated for period 1 in Figure 9.20.

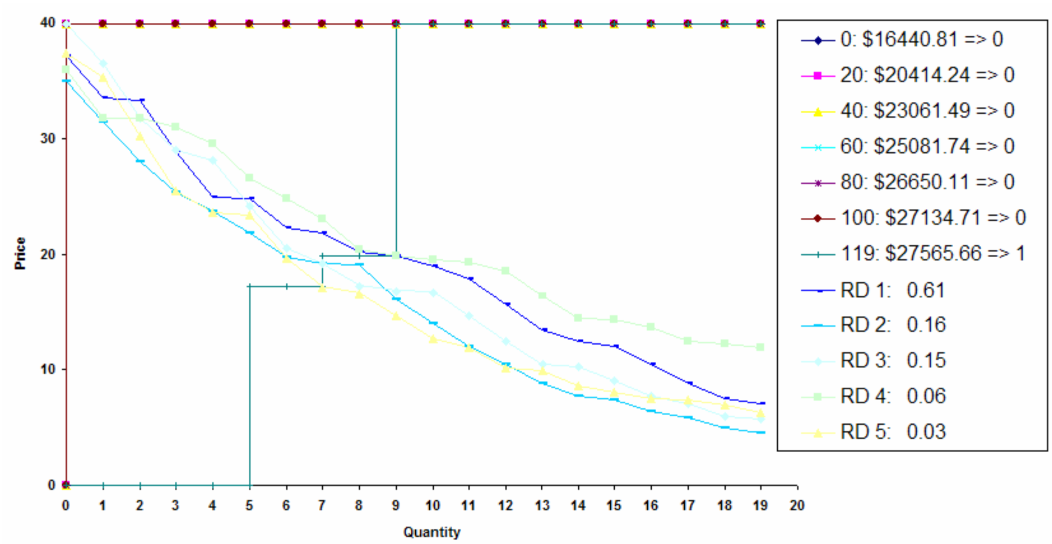

Figure 9.20 Offers for Period 1 Considering Unit Rules – from Off

Figure 9.27 demonstrates the value curves for all possible unit states at the beginning of period 13, one period before the very high-price period. The general order of these curves is very similar to that in Figure 9.12, the major difference being that the value associated with the unit only just having been switched off is now lower than that associated with a unit that has cooled down and has almost satisfied its minimum down-time requirement. The reason for this is that, as opposed to Example 9.2.3.1, at this point the generator is approaching a sustained set of high-priced periods, as opposed to just a one-off high-price period.

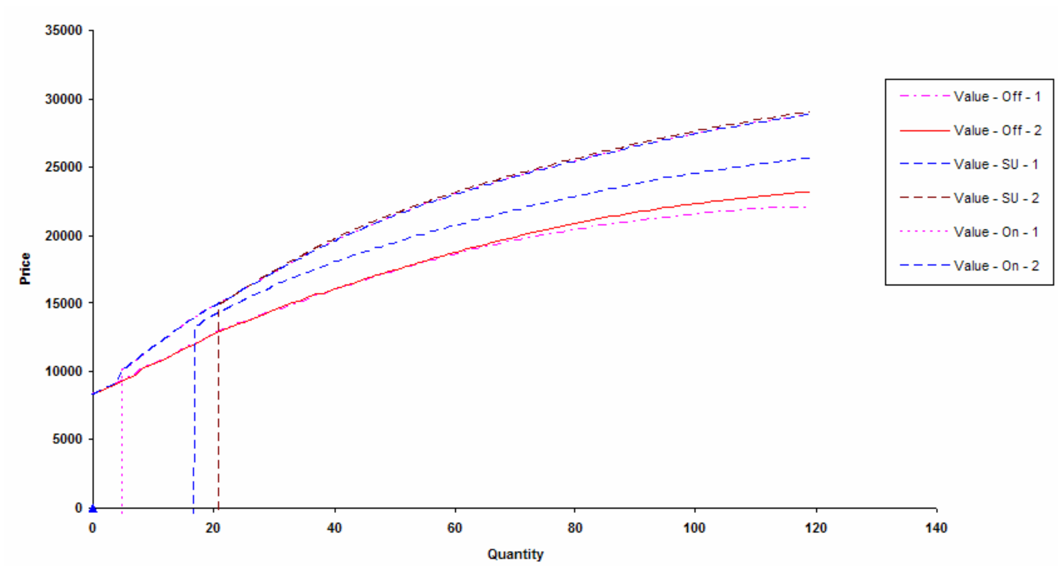

Figure 9.21 Value Curves for all Unit States at Beginning of Period 13

Example 9.2.3.4: Two Branches: One High and One Low-Priced

This final example does not relate to unit rules, but still illustrates an idea presented in Chapter 3, that offers could potentially be counter intuitively more restrictive under higher prices in off-peak periods, when there is high correlation between the prices (or RD curve levels) in consecutive periods. The reason for this is that it would be desirable to save fuel for the peak periods, when the price levels may be even more extreme. Here, we consider two possible macro-states that the market could be in over the day, where the transition probabilities are such that it is highly unlikely that the market will switch between the two states as the day progresses. Figure 9.27 demonstrates the "peakiness" of each period under the two possible states. The first state/branch corresponds to a high-price day, while the second state/branch corresponds to a low-price day, with relatively much lower extreme prices at the peak periods.

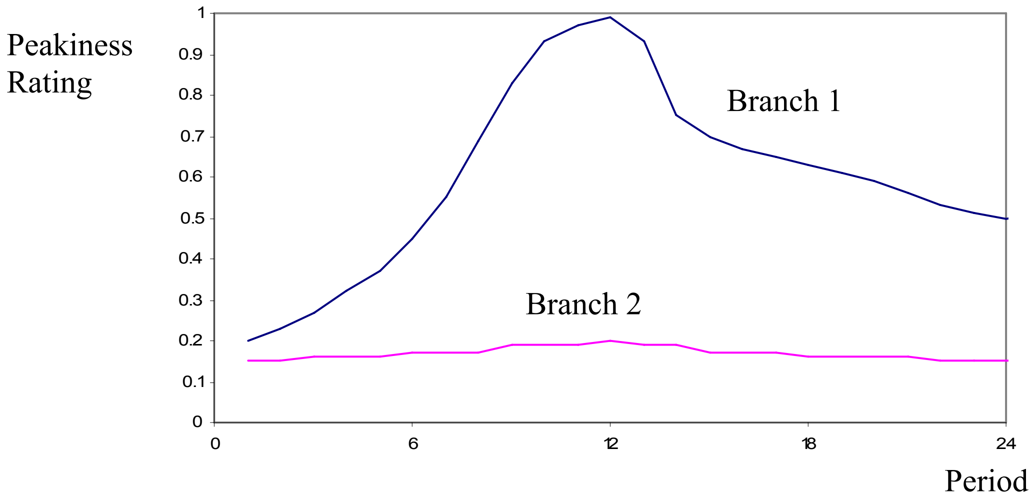

Figure 9.22 Peakiness Ratings for Example 9.2.3.4

Consider the offers in period 1 under the low and high-price branches, shown in Figure 9.24 and Figure 9.23 respectively. We can see that under the low-price branch, from the relatively high reservoir levels, dispatch becomes significantly positive, while there is no dispatch from any reservoir level under the high-price branch, despite its potential for far greater immediate profits.

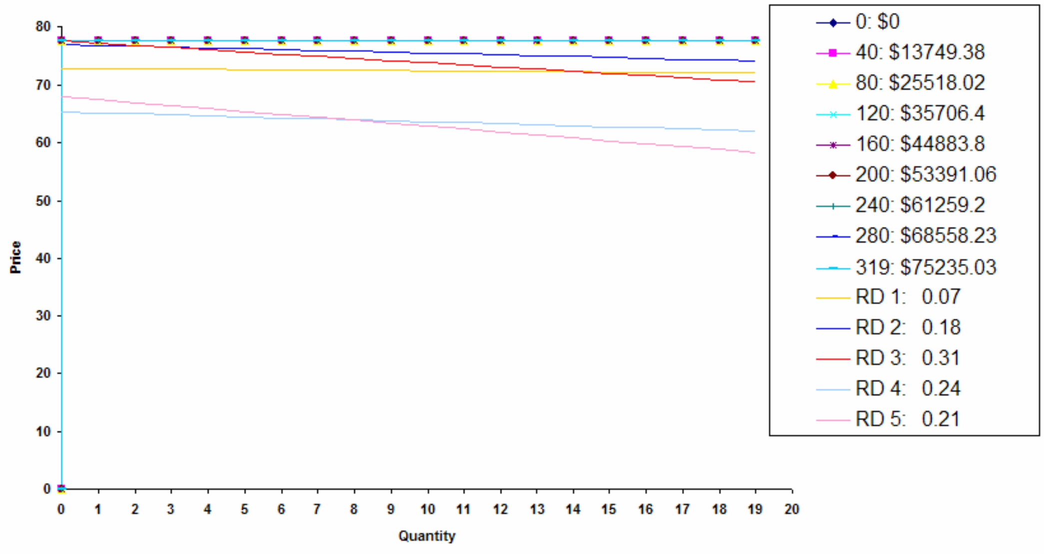

Figure 9.23 Offers for Period 1 on Branch 1

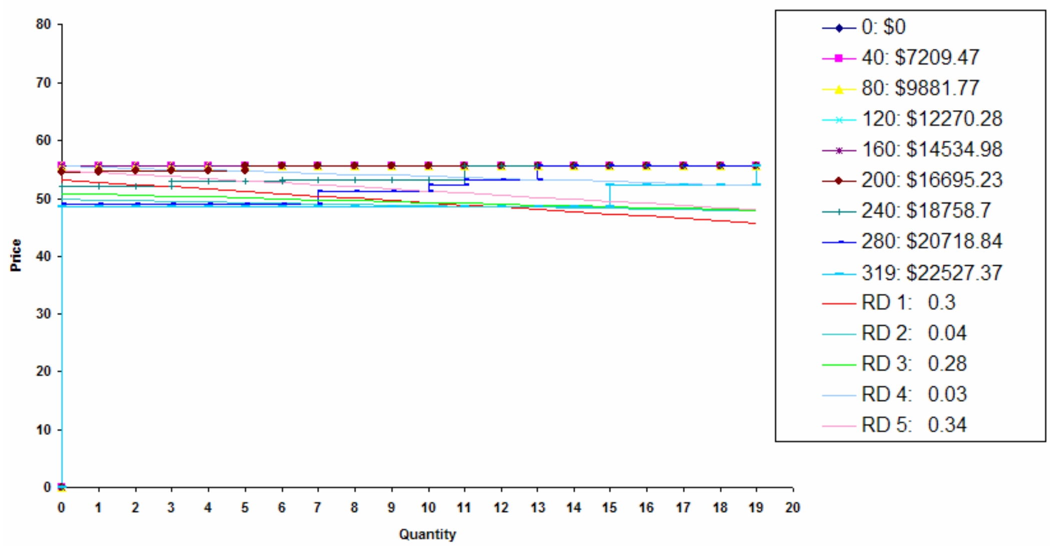

Figure 9.24 Offers for Period 1 on Branch 2

The reason for this is, of course, that the higher-priced branch scenario recognises that there will be potential for even greater profits if fuel is saved for the peak periods. This is reflected in the corresponding sets of MC curves, as shown in Figure 9.25 and Figure 9.26.

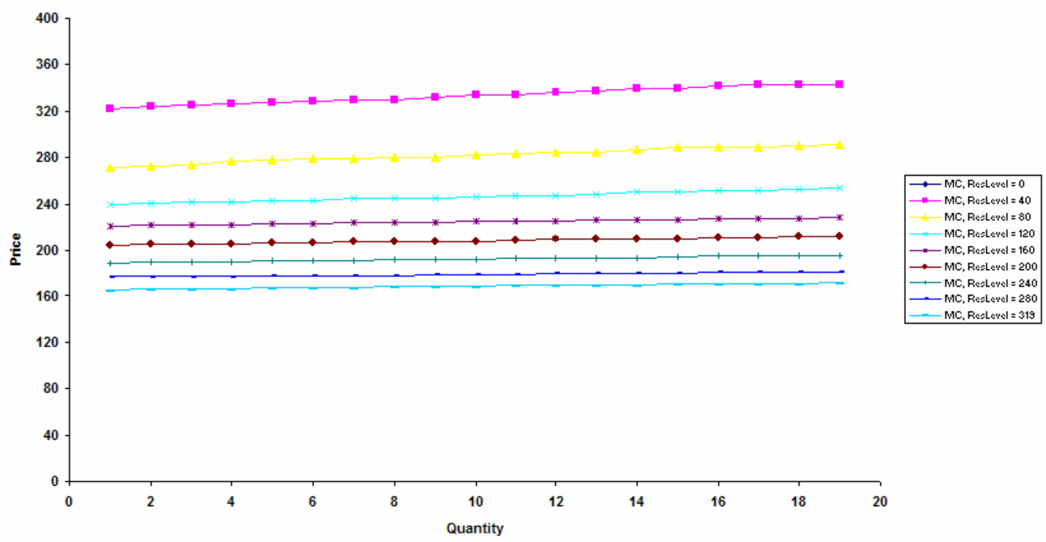

Figure 9.25 MC Curves for Period 1 on Branch 1

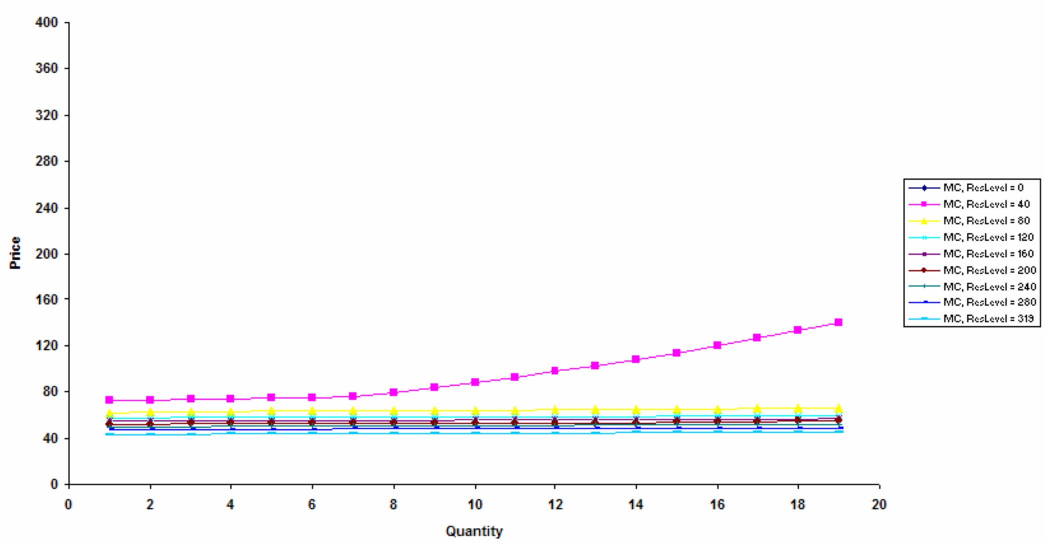

Figure 9.26 MC Curves for Period 1 on Branch 2

The MC curves for the higher-priced branch are clearly higher, reflecting the greater opportunity costs of generation discussed above.

9.3 Incorporating Ramp Rates

In this section, we consider the addition of ramp rate restrictions to the unit rules VC model described in Section 9.2. As stated in that section, the state space for the DP is three dimensional. The dimensions are

1. Previous market outcome
2. Reservoir level
3. Unit state

In order to consider ramp rates, we need to add an extra dimension:

4. Previous dispatch level

In other words, the generation levels that are able to be achieved in the current period are dependant on the generation level at which the unit was dispatched in the previous period. Under the DP with a three-dimensional state space, there was a different value curve for each combination of previous market outcome and unit state, defined over reservoir level. This new DP requires a four-dimensional state space, and therefore a different value *surface* is needed for each combination of previous market outcome, unit state, defined over the dimensions of reservoir level and previous dispatch level.

In this section we will show how the four-dimensioned DP algorithm can be made to work efficiently, such that an extra dimension does not influence the computational time. The examples and analysis that we present here ignore the unit state dimension for simplicity, but are equally applicable to models with this dimension included.

9.3.1 Constructing the MC Curve for a Given State with no Ramp Rates

Recall from Section 6.6, that if we do not consider ramp rates then we construct p different Value and MV curves for the following period, $t+1$ (Section 1 of Figure 9.27).

For a given previously observed RD curve (period *t-1*), we then created a single weighted MV curve for the following period (*t+1*), based on the Markov probabilities of each RD occurring in the current period, *t* (Section 2 of Figure 9.27). Then, depending on the current reservoir level, the MOC curve is produced directly from this weighted MV curve (Section 3 of Figure 9.27).

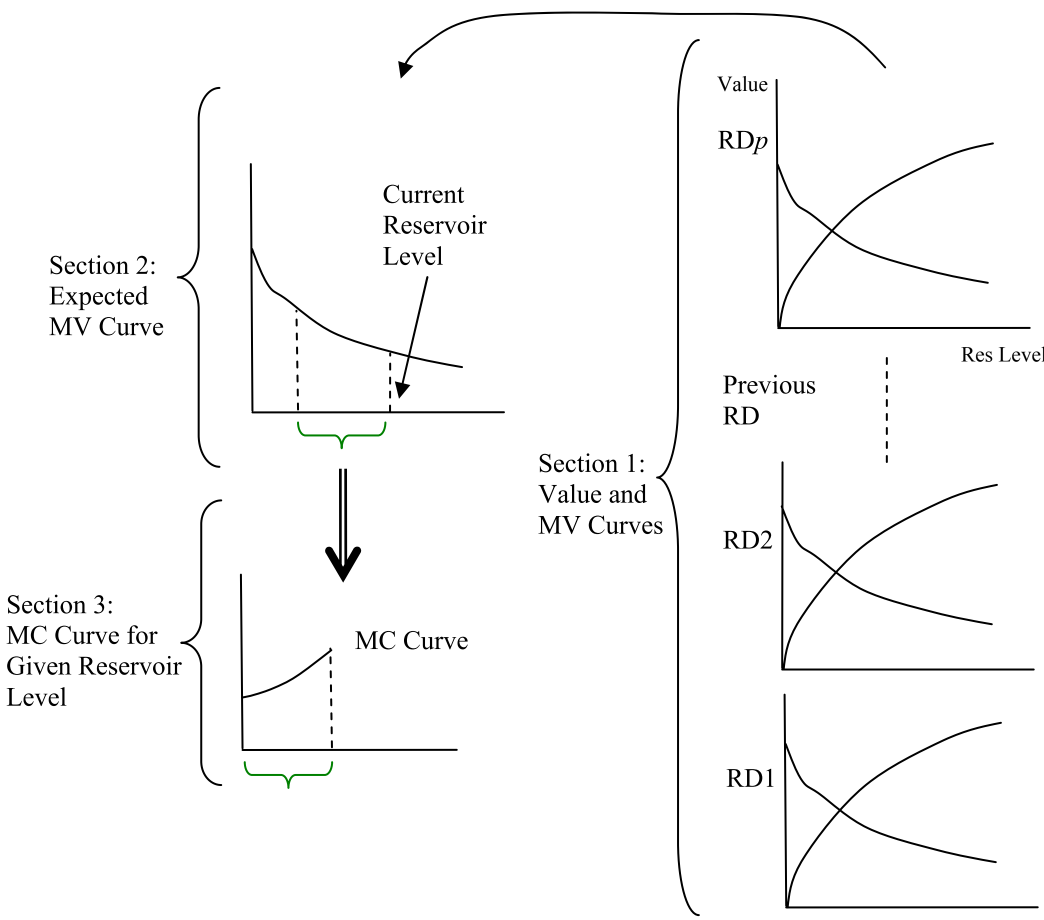

Figure 9.27 Construction of MC Curve with No Ramp Rates

9.3.2 Constructing the MC Curve for a Given State with Ramp Rates

As explained in Section 9.3, when ramp rates are considered, a set of three-dimensional value surfaces defined over reservoir level and previous dispatch level are required in each period, one for each possible previous RD curve (Section 1 of Figure 9.28). For each of these previous RD curves, we must then create a weighted value surface based on the Markov chain probabilities of each of the RD curves occurring in the current period (Section 2 of Figure 9.28).

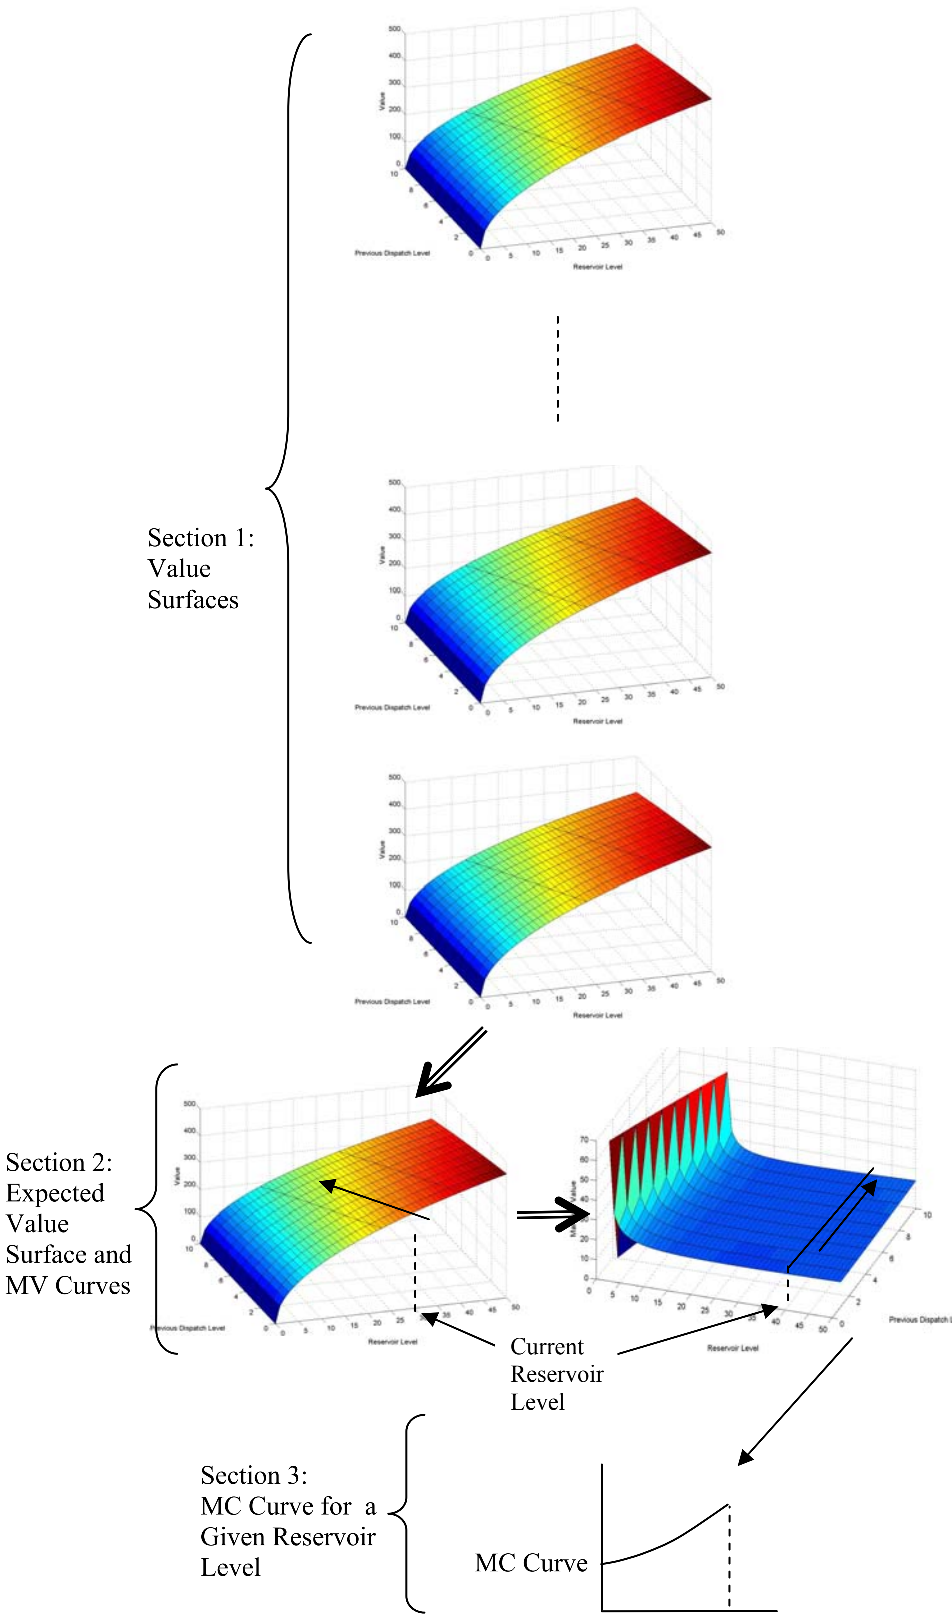

Figure 9.28 Construction of MC Curve with Ramp Rates

In this section we explain how the calculation of the MV curves (and hence the MOC curves for the previous period) is performed from these weighted value surfaces.

Start by considering a generator that is heading towards an off-peak period, under a given previous RD curve, with a ramp rate of ± 2 MW/period. The value curve will, of course, be higher as the reservoir level increases (when you hold previous dispatch constant), and (generally speaking) higher for lower previous dispatch levels (when you hold the reservoir level constant, as it is desirable to be at a lower dispatch level when trying to ramp down). This situation is represented in a single three-dimensional Value curve, in Figure 9.29. Figure 9.30 shows a view of Figure 9.29 from above, where the line indicates all possible (reservoir level, previous dispatch level) states that could be reached in that period from the (23,0) state in the previous period (assuming no inflows).

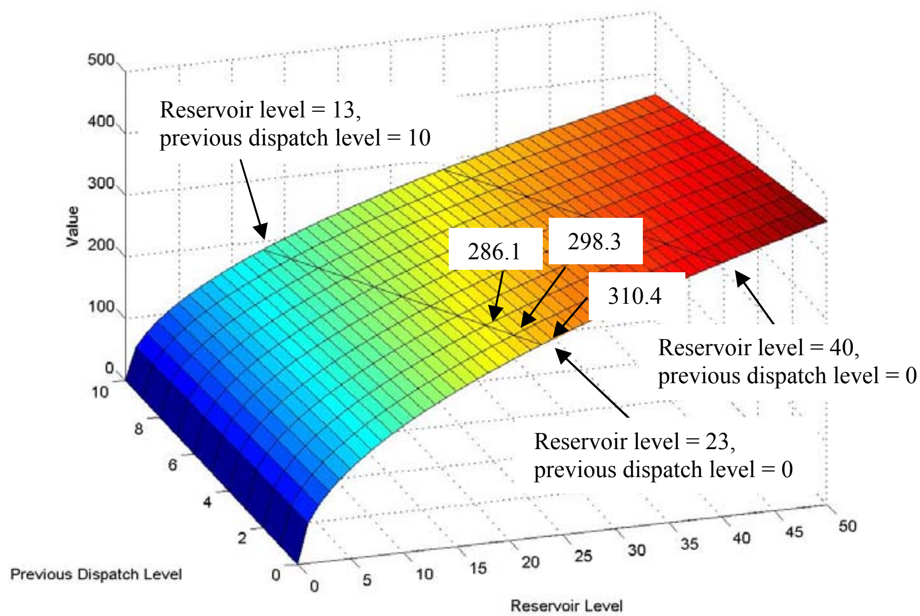

Figure 9.29 3D Value Surface for Period (t+1) as Off-Peak is Approached

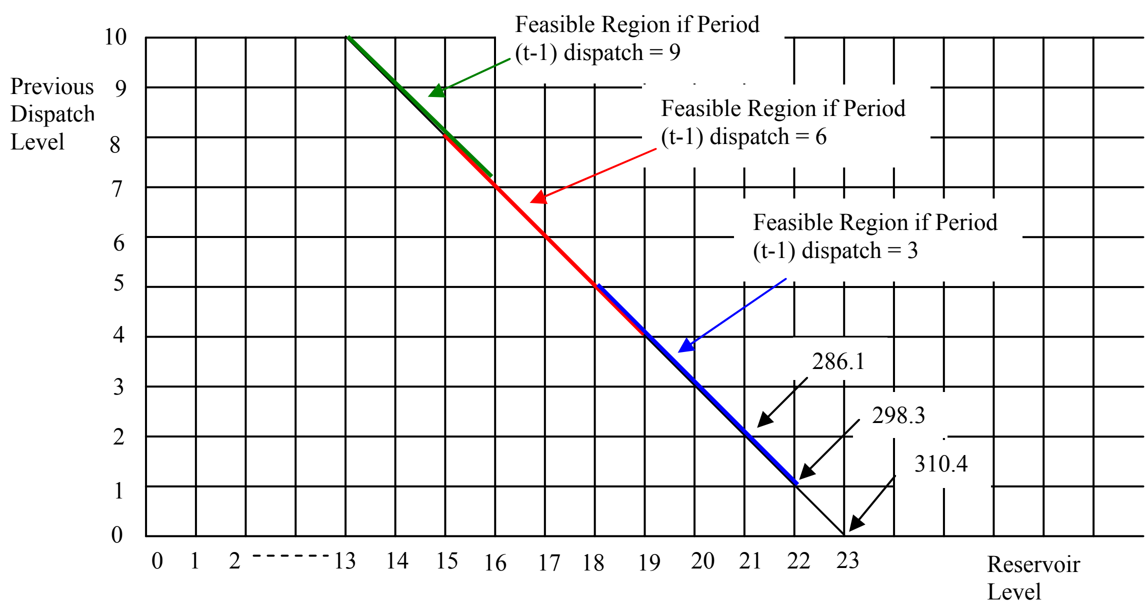

Figure 9.30 Value Surface for Period (t+1) Viewed from Above, with Ramp Rate of ± 2 MW/Period

Recall that in the case without ramp rates, if we are at reservoir level *r* in period *t*, then the MOC of the first unit of generation comes from the difference between the expected value at reservoir level *r* and level (*r*-1) in period (*t*+1), and so on. In the ramp rate model, the previous dispatch level must also be considered. In other words, if the reservoir is at level *r* in period *t*:

MOC of the first unit of generation: comes from the difference between the expected value at reservoir level *r* under a previous dispatch level of "0" and the expected value at reservoir level (*r*-1) under a previous dispatch level of "1" in period (*t*+1)

MOC of the second unit of generation: comes from the difference between the expected value at reservoir level (*r*-1) under a previous dispatch level of "1" and the expected value at reservoir level (*r*-2) under a previous dispatch level of "2" in period (*t*+1).

and so on.

In other words, the MOC curve for a given (reservoir level, previous market outcome) state in period t is calculated in a 'diagonal' fashion, from the expected value *surface* for period (t+1) (Section 3 of Figure 9.28). For example, consider the case provided in Figure 9.29 and Figure 9.30, where the reservoir level is 23 in period t. If we were to generate one unit of electricity in this period, the future value component would fall from 310.4 to 298.3. Thus, the difference between these two values (12.1) gives our MOC of generating that first unit. Likewise, the MOC of generating the second unit is 298.3 minus 286.1 (12.2). This process can be extended to provide the MOC for each feasible level of generation, for the given reservoir level (23) in period t.

If there are no (or highly unrestrictive) ramp rates restrictions, then the value surface takes the same value regardless of the previous dispatch level. In this case, the MOC curve that we would produce from a given reservoir level collapses down to that which we would get if we only considered a two-dimensional value curve.

Of course, when there are ramp rate restrictions, it is only a certain range of this MOC curve that is valid (or feasible). For example, in the case above with a maximum ramp rate of ± 2 MW/period, if the previous level of generation was 6MW, then only the section of the MOC curve corresponding to generation levels from 4 – 8MW are of any interest in the current period (as highlighted in red in Figure 9.30). The implication is that for a given (previous RD, reservoir level) state, if the previous dispatch level (t-1) moves from r to (r+1), then the MC of dispatching the (r+2)th unit in the current period (t) would not change (assuming that the generation level is feasible in both cases in terms of ramping); it still comes from the same point on the expected value surface (for period (t+1)). This is important as it means computational complexity is not increased as much as one would normally expect as a result of adding a further dimension in the DP algorithm. This is discussed further in Section 9.3.5. Note also that the computational efficiency is greater, the tighter the ramp rates.

Note that virtually all of the values in the matrix that are used to represent the Value surface are potentially used in determining the MC of generation for some period t state or other. Therefore, all these values need to be calculated (or interpolated as discussed in Chapter 8).

Unfortunately, because the MOC curves are constructed diagonally, some of the efficiency related to sliding along a single overall MOC curve is lost. This means that now, rather than having MOC curves for consecutive reservoir levels parallel to one another, they may take any form, so long as the MOC curve for the higher reservoir level is below that for the lower reservoir level. This concept is demonstrated in Figure 9.31, where the figure on the left shows the set of MOC curves for a subset of reservoir levels ignoring ramp rates, and the figure on the right shows an equivalent set of MOC curves considering ramp rates.

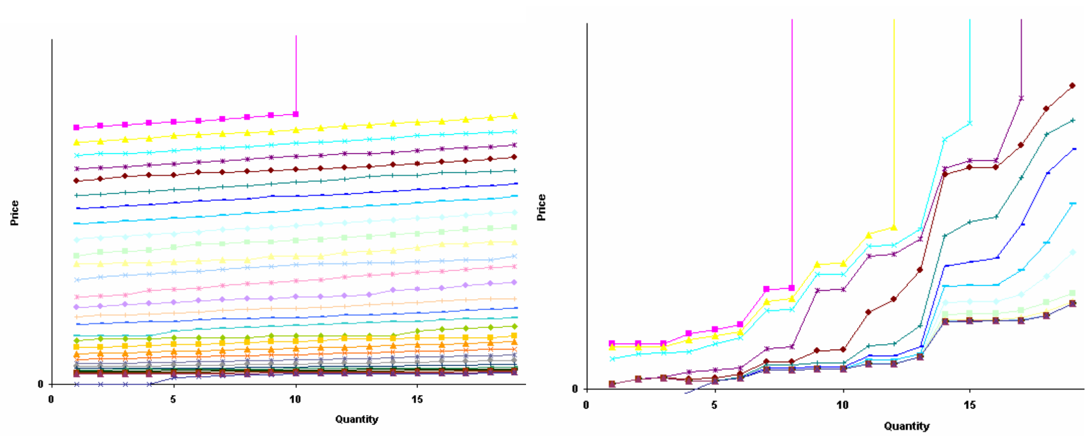

Figure 9.31 MOC Curve Comparison

9.3.3 Monotonicity in the MOC curve under the Ramp Rate Model

If we assume no ramp rate restrictions, then in order for the MOC curves to be monotone, the only condition was that the two-dimensional value curve for the following period had to be concave (i.e. the MV of storage was falling). Now that the MOC curves are being produced from the diagonal of a value surface, this condition becomes slightly more complex.

Ramping Down:

If the generator is heading towards an off-peak period, and trying to ramp down as fast as possible, then the value surface for period (t+1) will be decreasing in the previous dispatch level and increasing in the reservoir level dimensions. In particular, let us assume that the value surface is *concave in both dimensions*. This is the case presented in Figure 9.29. Therefore, we know that as dispatch is increased (from any current reservoir level), future value is falling off at a faster rate in each of the two dimensions that it is defined over. Therefore, the MOC must be monotonically increasing as dispatch level increases.

Ramping Up:

If the generator is heading towards a peak period, and trying to ramp up as fast as possible, then the value surface for period (t+1) will be increasing in the previous dispatch level.

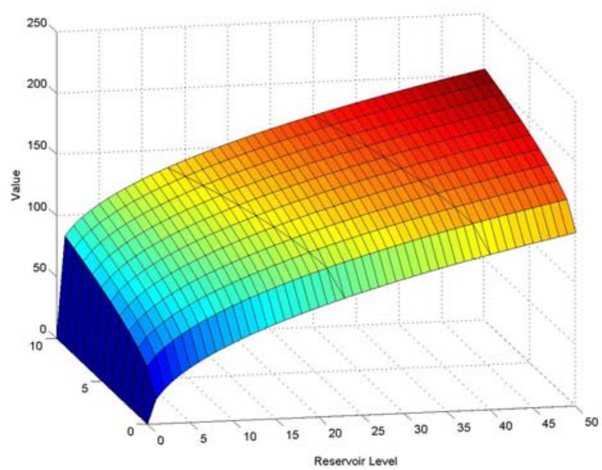

Figure 9.32 Value Curve for Period (t+1) for a Unit Ramping Up

Let us consider a specific example, where the value curve shown in Figure 9.32 is given by the equation:

$$\text{Value} = 20 * \sqrt{\text{Reservoir Level}} + 20 * \sqrt{\text{Previous Dispatch}}$$

In this case, the contribution to the future value from the reservoir level and the previous dispatch (period *t*) level can be treated independently, as we try to determine whether the MV/MOC curve will be monotonically increasing in the period *t* dispatch level. Therefore, we will also consider the contributions to the MOC to be independent of one another.

Consider now a different example, where the future value contribution from the reservoir level decreases at an increasing rate as we increase the current dispatch, and the contribution from the previous dispatch level is increasing at a decreasing rate (i.e. the value surface is concave in both dimensions, and meets the conditions of the scenario that we are considering). A possible set of values are presented in Table 9.2.

Dispatch (x)	Reservoir Level Value Contribution	Previous Dispatch Level Value Contribution	Sum	MV/MOC = Sum(x-1) – Sum(x)
0	500	20	520	
1	480	120	600	-80
2	450	200	650	-50
3	410	260	670	-20
4	350	310	660	10
5	270	350	620	40
6	170	380	550	70
7	50	400	450	100

Table 9.2 An Example of MOC Curve Production

In this case, we are approaching a peak period, and thus it is desirable for the unit to be ramped up to a certain extent. This is reflected in the fact that the MOC of the first few units of generation are actually negative (in order to ensure that dispatch in period *t* is at least at this level). In addition, we can observe that the MOC is monotonically increasing.

We can prove that this is the case for all examples where the contributions from reservoir and previous dispatch levels are concave.

If we consider the components independently, we can show that the MOC contributions from both the reservoir level and from the previous dispatch level are increasing as the generation level increases.

Reservoir Level MOC Contribution: We know that the value curve is increasing and concave in the reservoir level variable. Therefore the MV must be falling as we increase the reservoir level or, alternatively, that the marginal value (MOC) must be monotonically non-decreasing as we increase the generation level in period t (thus reducing the reservoir level).

Previous Dispatch Level Contribution: Because we are ramping up, we know that we would prefer to be at a higher generation level as we approach period (t+1). However, as the surface is concave in this variable, we know that the marginal benefit of being at higher and higher generation levels is falling. Therefore, the MOC contribution must be negative over the entire range of generation levels, but with the values approaching zero as the previous dispatch level increases. In other words, although the contribution of this term is negative, it is monotonically non-decreasing as the previous dispatch level increases.

Therefore, as both terms that contribute to the MOC are monotonically non-decreasing, then the sum of the two terms must also be monotonically non-decreasing. This will be true as long as the value surface is concave in both dimensions, regardless of whether the value functions are increasing or decreasing. For example, if the most desirable previous dispatch level to be at is somewhere in the middle of the feasible dispatch range, then the value surface may form a concave parabola around that previous dispatch level, and this would still give a monotone MOC curve.

Therefore, this has shown that as long as the value surface is concave in both dimensions, then the offers produced will be monotone, as they are required to be.

Unfortunately, although we need the value surface to be concave in both dimensions in order to produce increasing MOC curves (and thus monotone offers), due to the discretisation of the various problem dimensions, the value surfaces that are produced will not always take this form[70]. In these cases, we must perform an adjustment to the surface, so that they become concave, and will produce the required form of offers.

9.3.4 Process for Producing Offers for all Previous Dispatch Levels with Ramp Rates

In this section, we outline the process for producing the optimal offers under all previous dispatch levels for a given market and reservoir level state. The steps in this process are:

1. Produce the MOC curve for a given (reservoir level, previous RD) state using the approach described in Section 9.3.2.
2. Produce an unrestricted offer for the given state, based on this MC curve, as shown in Figure 9.33. In other words, this is an offer that ignores any ramping restrictions and thus goes all the way from zero to the maximum generation level (*d*).
3. Work out the expected payoff from the dispatch at the intersection of each of the RD curves with the unrestricted offer (the red stars in Figure 9.33).
4. Work through the previous dispatch levels, producing the offer as the feasible section (with respect to ramp rates) of the unrestricted offer, along with vertical and horizontal segments such that dispatch can only fall within this feasible

[70] As was the case without ramp rates, demonstrated in Section 6.6.1.

range, regardless of the RD curve outcome (as shown in green in Figure 9.33)[71]. This will mean that the horizontal segment below (above) the lowest (highest) feasible generation level will be lower (higher) than any price that a RD curve could intersect within this range.

5. Calculate the value of each offer. This is very computationally efficient when using the following steps.
    a. For any RD curves that naturally intersect the unrestricted offer in the feasible generation range (the middle two red stars in Figure 9.33), use the expected payoffs found in step 2.
    b. For the other RD curves, the resulting dispatch must either be at the lowest or the highest feasible generation level, therefore Q is easy to find, and hence so is P (as we have a defined P for each Q level under each RD curve).

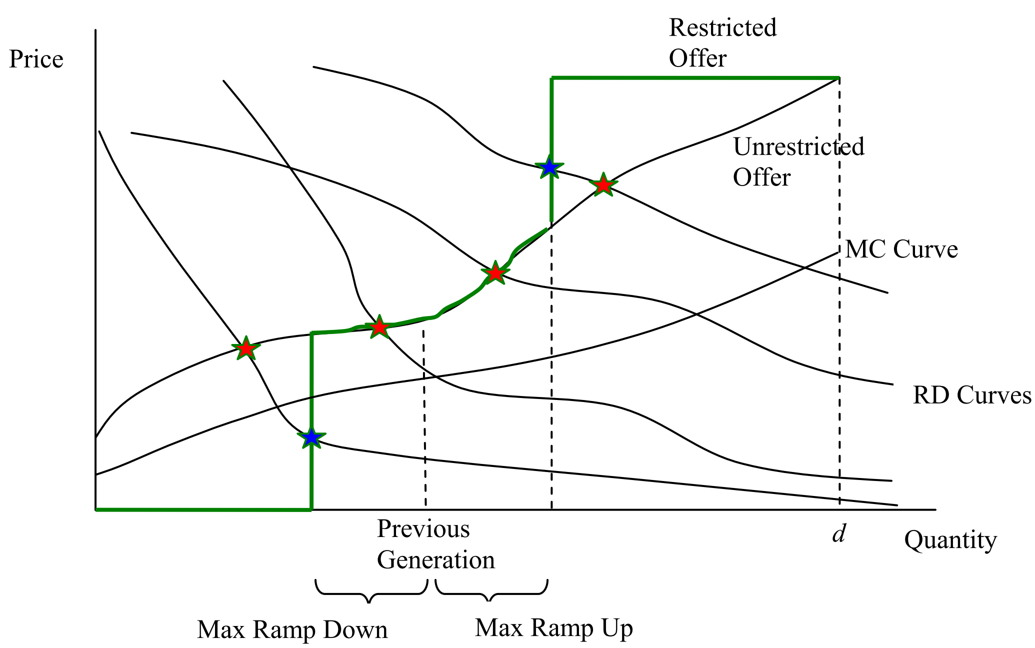

Figure 9.33 Unrestricted and Ramp Rate Restricted Offers

[71] Note that we know this must be the optimal offer for the given previous generation level as the method explicitly considers the future value of getting to a higher (or lower) generation level in the current period (t) when determining the MOC of each additional unit of generation. Hence, the method deals with the strategic value of pushing generation higher or lower before it might be immediately desirable to do so.

9.3.5 Calculation Complexity Implications

By considering the previous dispatch level, an extra dimension is added to the DP state space. Generally, adding dimensions in DP leads to the "curse of dimensionality", meaning that computational times would also increase by a full extra dimension. However, in this section we show that because of the various structures of the problem faced here, this is not the case.

The first point to note is that there will be no effect on the pre-processing phase of the algorithm. The same set of MC levels can still be used, and the UMSs have not changed.

In Section 9.2.3 we explained that there is no need to produce a completely separate offer for each previous dispatch level under a given reservoir level. A single offer needs to be produced, ignoring the ramping restrictions, and then the feasible ranges of this become the optimal offer for each possible previous dispatch level. In addition, Section 9.3.4 demonstrated that the calculation of the value of these partial offers is very computationally efficient.

Unfortunately, under the ramp-rate model, separate MV/MOC curves need to be created for each reservoir level, rather than just a single one defined for all possible reservoir level positions.

The overall computational complexity of the branched approach including ramp rates can be approximately summarised by Equation 9.1.

$$0.1tb\begin{pmatrix} 100 + 160P + 10b + 40Pd + 490Prd + 30P^2 + 10bP + 65PM + 190P^2 r \\ + 20Prd\text{infl} + 100Pr + 30Prb + 100P^2 rd + 28PMd \end{pmatrix}$$

Equation 9.1 Computational Complexity of the Branched RT Phase with Ramp Rates in Full

Simplifying this down to the terms of the largest orders of magnitude, we can see that the algorithm is of order of complexity as shown in Equation 9.2.

$$tpr(Pd + b) + tpMd$$

Equation 9.2 Order of Computational Complexity for Branched RT Phase with Ramp Rates

Recall the simplified computational complexity of the branched model without ramp rates, given in Equation 7.10 from Chapter 7:

$$tpr(P + d + b) + tpMd$$

Equation 7.10 Computational Complexity for Branched RT Phase without Ramp Rates

When the previous dispatch level dimension is considered, it might be expected that the computational time would increase by a factor given by the size of this new DP dimension (i.e. the number of different possible dispatch levels, or *d*). However, the only difference between the complexity orders is that the terms $tprP + tprd$ have changed to a single term, $tprPd$. While this is still a significant increase, it is certainly less than that of a whole dimension[72].

9.3.6 Further Experimental Evidence of Offer Construction Theories

In this section, we will show the effect of implementing the ramp rate conditions on the offers that are made to the market, under three example scenarios. The first example has one very high-price period in the middle of the off-peak section of the planning horizon, the second example has two relatively low-price periods in the middle of the peak section of the planning horizon, while the third example has half the horizon with very low prices (or RD curves), and then suddenly jumps up to high prices for the remainder of the horizon. Note that, as with the cases presented in Section 9.2.3, all scenarios

[72] In fact, in preliminary testing we have found that the percentage increase in computational time may be anywhere from 10 – 25% of its potential expected increase, as given by the term *d*.

presented are relatively extreme cases, enabling us to demonstrate the effect of the ramp rate constraints clearly.

Example 9.3.6.1: Anomalous High-Price Period in the Off-Peak Range

The first example that we use is the same as Example 9.2.3.1, except that we ignore the unit rules, and consider ramp rates only. Figure 9.34 demonstrates the optimal offer to provide to the market at a reservoir level of 90 in period 11 assuming that there are no ramp rate restrictions, while Figure 9.35 demonstrates the equivalent optimal offer assuming ramp rates must be considered.

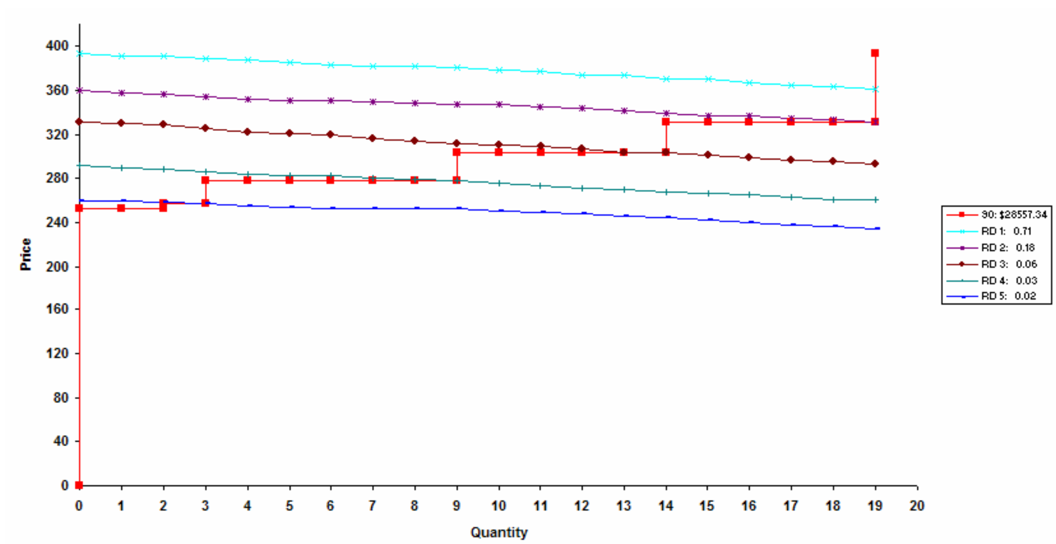

Figure 9.34 Offer for Reservoir Level 90 in Period 11 under No Ramp Rates

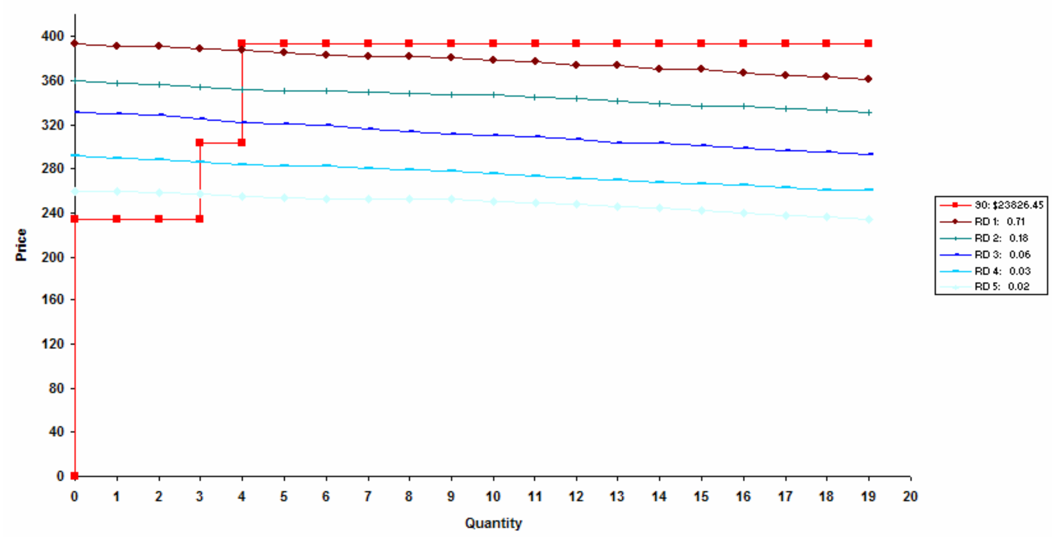

Figure 9.35 Offer for Reservoir Level 90 in Period 11 under Ramp Rates

It is clear that if ramp rates do not need to be considered, then the optimal dispatch depends on the RD curve and can range anywhere from 3 to 19 MW (the maximum capacity of the unit considered). However, if ramp rates do exist, then the optimal offer is much more restrictive (with dispatch no higher than 4 MW), recognising that a low-price period is to follow, and that it will be costly to be ramped too high at this point. Again consider the MC curves that lead to these two offers, as shown in Figure 9.36. In the ramp rate case, the model recognises the cost in being ramped up to a high dispatch level at the end of the period, and thus the MC curve is very high across most of the feasible generation range. This leads to the much more restrictive offer demonstrated in Figure 9.35.

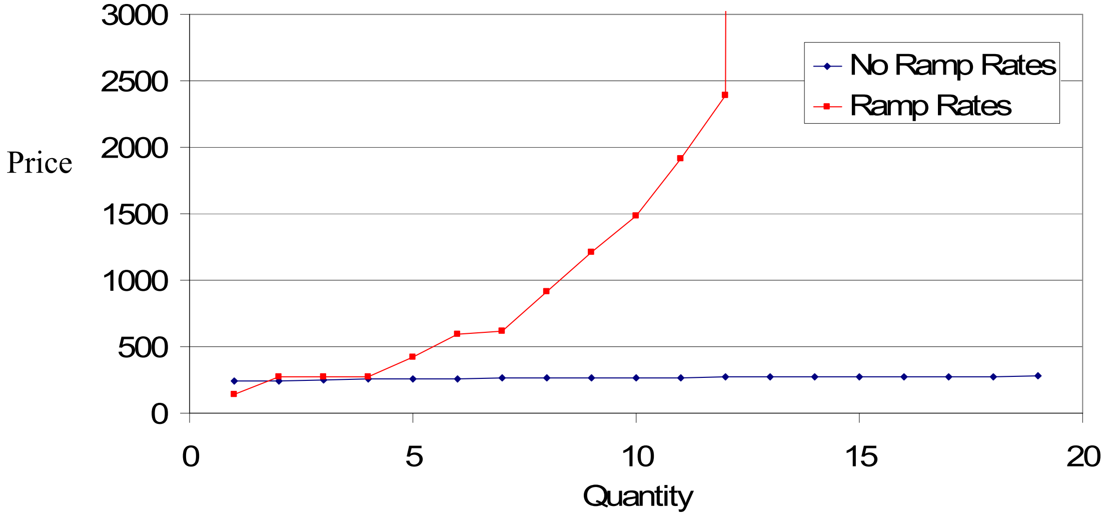

Figure 9.36 MC Curve for Reservoir Level 90 in Period 11 under No Ramp Rates

Again it is interesting to observe how the expected payoff over the horizon is affected by the restrictiveness of the ramping conditions. Table 9.3 shows that, as would be expected from standard optimisation theory, the expected payoff gets progressively lower as we decrease the generation range flexibility.

Ramp Up/Down Limit	Expected Payoff
No Restrictions	35540.95
20	35540.95
10	34614.22
5	33051.04
2	29190.62
1	25428.88

Table 9.3 Expected Payoff vs Ramp Rate Restrictions

Finally, Figure 9.37 shows an entire value surface for the end of period 11, considering ramp rates, for the above example. Note how the value surface is significantly greater at lower previous dispatch levels, indicating the benefit of being ramped down to a low dispatch level at this point. This is the reason for the very high MC curve in this case.

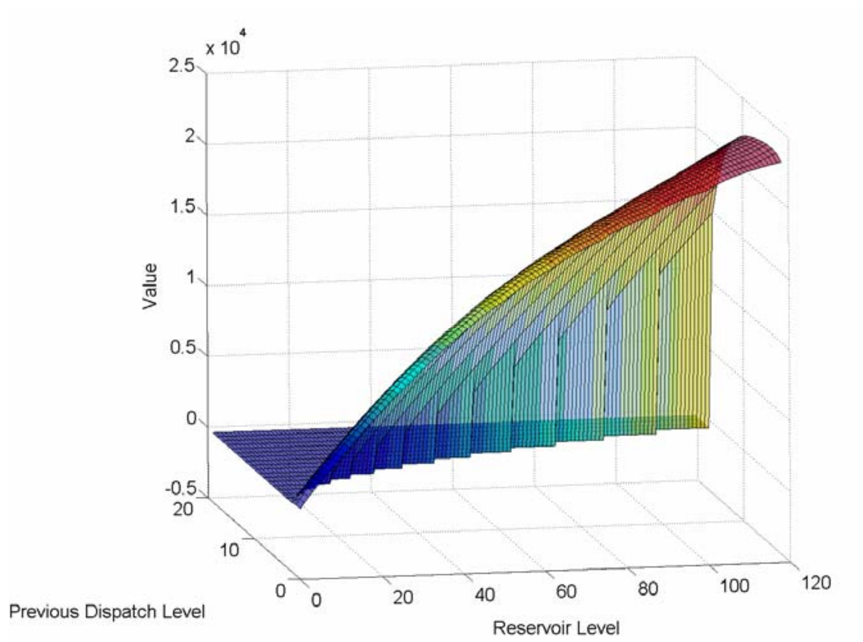

Figure 9.37 Value Surface with Ramp Rates just before an Off-Peak Period

This example has demonstrated that the form of the offers is affected when ramp rates are considered, as the impact of current generation levels on the future must be considered. In addition, it has demonstrated that by progressively tightening the ramp rate restrictions, the negative impact on expected payoff escalates very quickly. The same general results, under different scenarios, are also produced in the following two examples.

Example 9.3.6.2: Anomalous Low-Price Periods in the Peak Range

In this second scenario, the general height of the RD curves, or the "peakiness" of each period is shown in Figure 9.38. Clearly, there are two anomalous low-price periods in the middle of the peak section of the planning horizon.

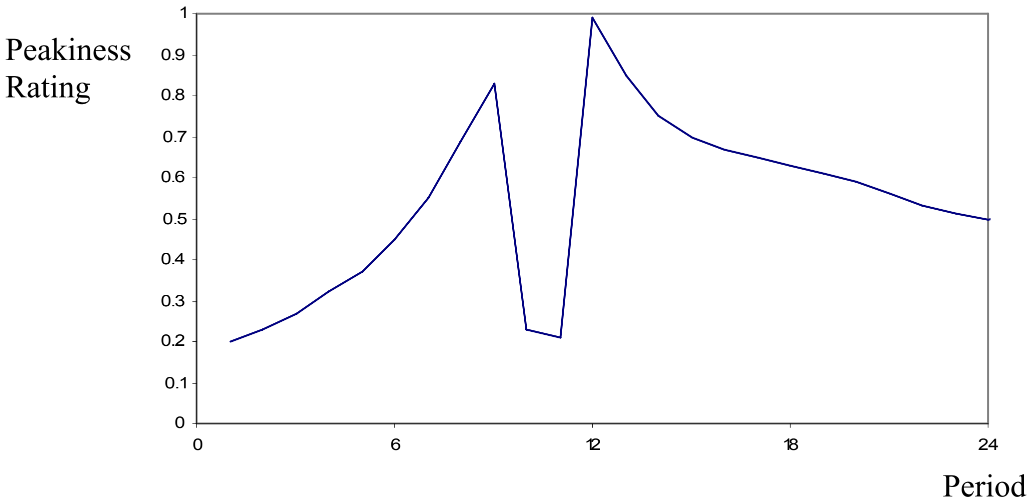

Figure 9.38 Peakiness Ratings for Scenario 1

Figure 9.39 demonstrates the optimal offer to provide to the market at a reservoir level of 90 in period 11 assuming that there are no ramp rate restrictions, while Figure 9.40 demonstrates the equivalent optimal offer assuming ramp rates must be considered.

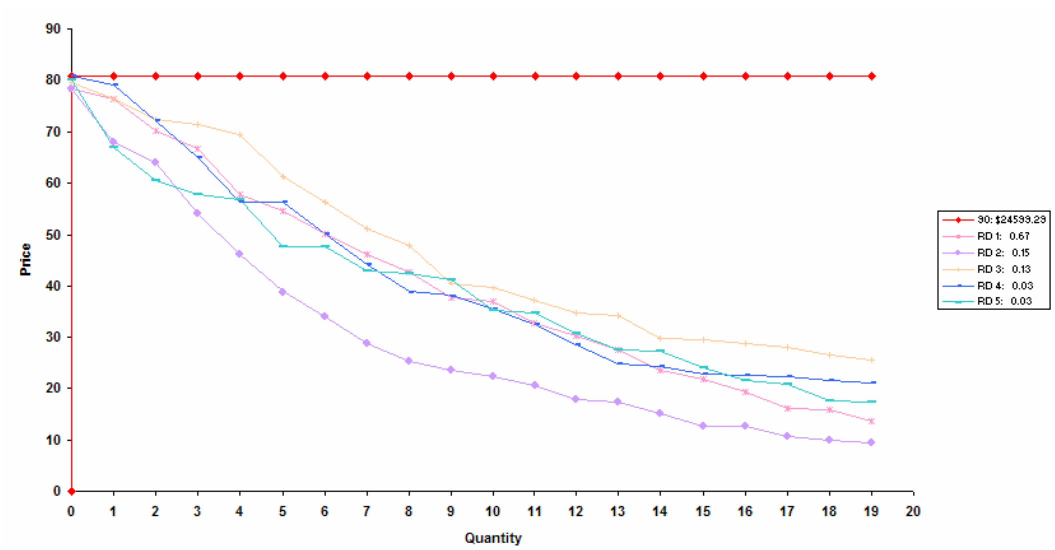

Figure 9.39 Offer for Reservoir Level 90 in Period 11 under No Ramp Rates

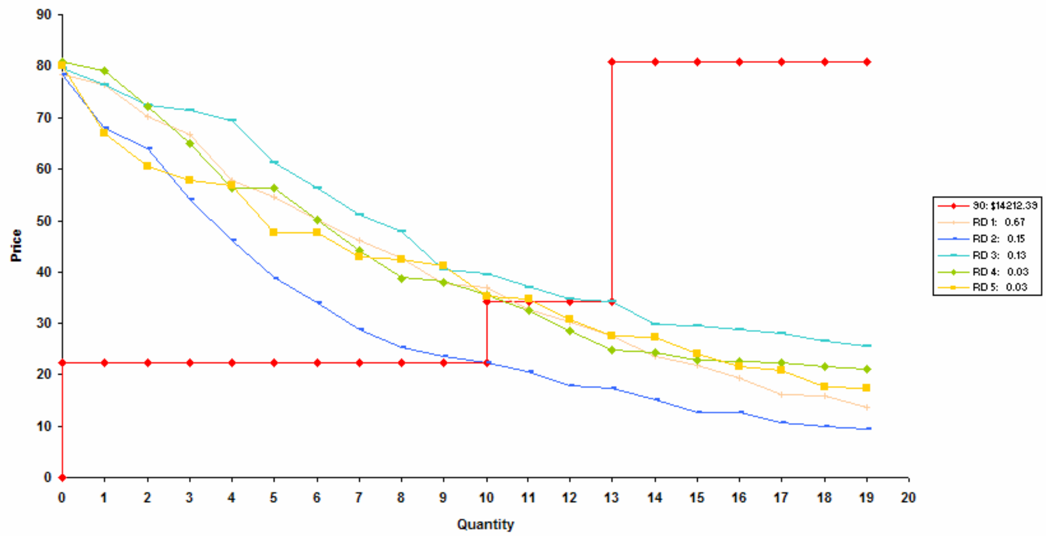

Figure 9.40 Offer for Reservoir Level 90 in Period 11 under Ramp Rates

It is clear that if ramp rates do not need to be considered, then it is not optimal to dispatch any electricity (release any fuel) in this low-price period. However, if ramp rates do exist, then there will be a positive offer, leading to dispatch between 10 and 13 MW, depending on the RD curve that occurs. The reason for this is that the following period (12) is a very high price period, under which a high dispatch level is likely to be desired. Therefore, under ramp rate restrictions, the dispatch level in the current period must be relatively high, to enable the generator to reach these high desired dispatch levels in the following period. Another way of thinking about this is to consider that the offers provided in Figure 9.39 and Figure 9.40 are based on the MC curves shown in Figure 9.41. In the ramp rate case, the model recognises the value in being ramped up to a relatively high dispatch level at the end of the period, and thus the MC curve is very low across most of the feasible generation range. This leads to the much more generous offer demonstrated in Figure 9.40.

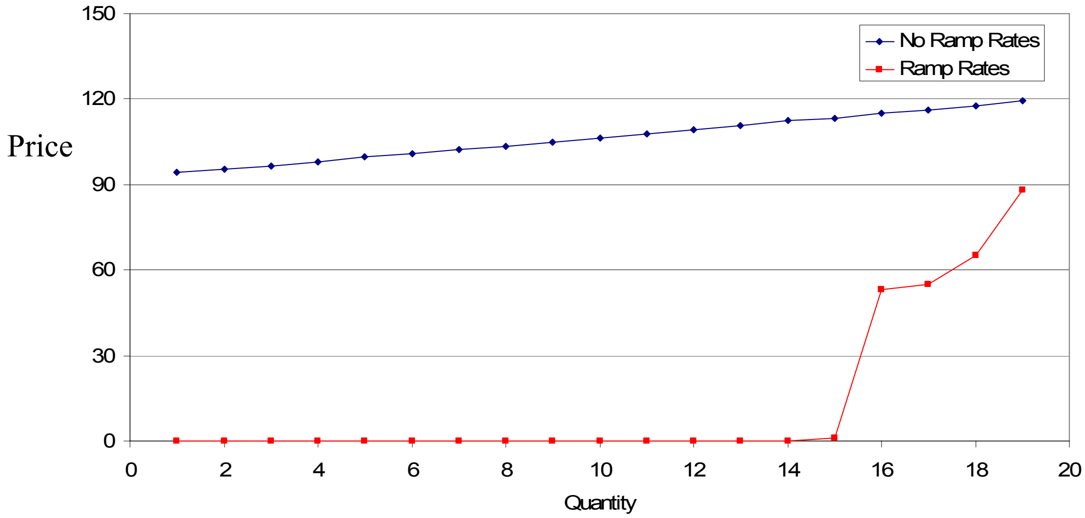

Figure 9.41 MC Curve for Reservoir Level 90 in Period 11 under No Ramp Rates

It is also interesting to observe how the expected payoff over the horizon is affected by the restrictiveness of the ramping conditions. Table 9.4 shows that in this case, as we decrease the generation range flexibility, the expected payoff gets progressively lower.

Ramp Up/Down Limit (MW/period)	Expected Payoff
No Restrictions	30575.06
20	30575.06
10	29471.61
5	25639.81
2	19765.59
1	16943.66

Table 9.4 Expected Payoff vs Ramp Rate Restrictions

Finally, Figure 9.42 shows an entire value surface for the end of period 11, considering ramp rates, for the above example. The ‘zero’ values indicate infeasible regions due to ramp rate restrictions combined with limited fuel. Note how, particularly at high reservoir levels, the value surface is significantly greater at higher previous dispatch levels, indicating the benefit of being ramped up to a high dispatch level at this point. This is the reason for the very low MC curve in this case.

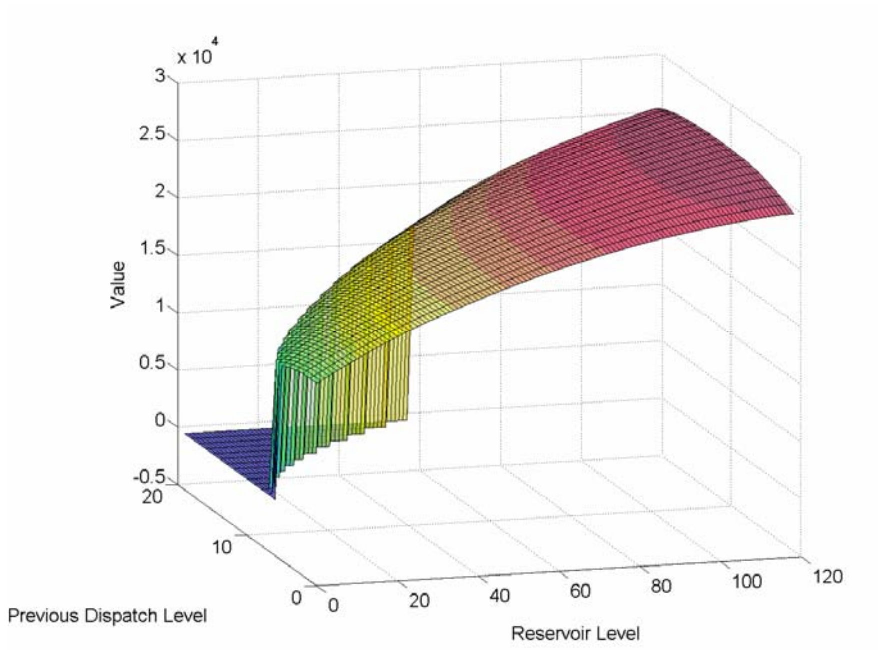

Figure 9.42 Value Surface with Ramp Rates just before a Peak Period

Example 9.3.6.3: Low-Price Periods Changing Suddenly to High-Price Periods

The third example that we consider is the same as Example 9.2.3.3, except that we ignore the unit rules, and consider ramp rates only. Figure 9.43 demonstrates the optimal offer to provide to the market at a reservoir level of 90 in period 13 assuming that there are no ramp rate restrictions, while Figure 9.44 demonstrates the equivalent optimal offer assuming ramp rates must be considered.

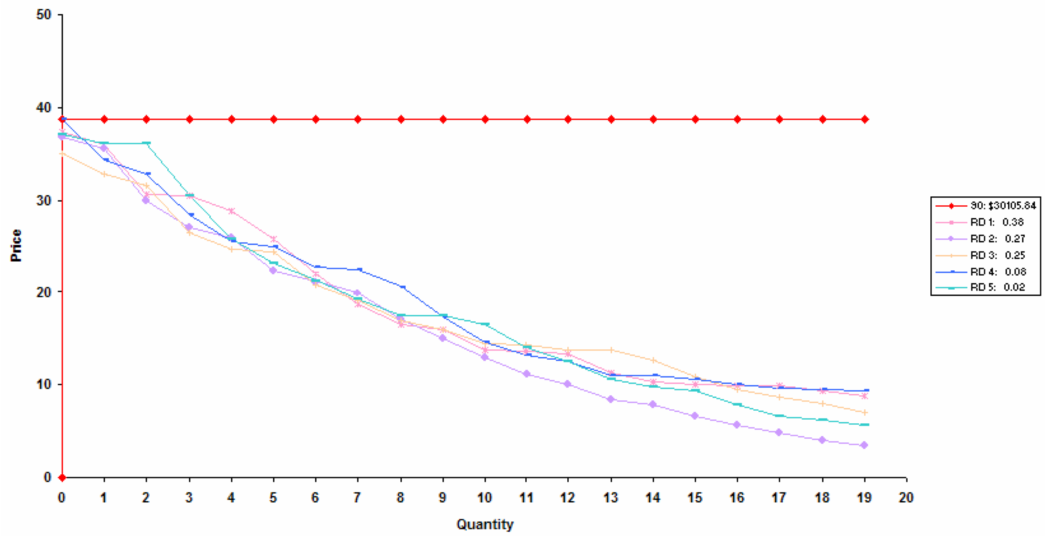

Figure 9.43 Offer for Reservoir Level 90 in Period 13 under No Ramp Rates

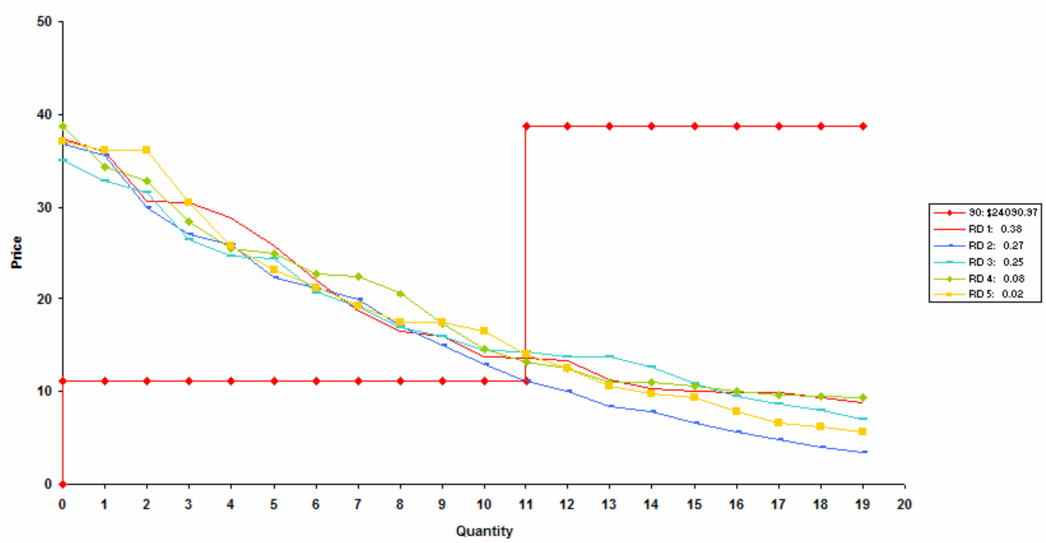

Figure 9.44 Offer for Reservoir Level 90 in Period 13 under Ramp Rates

If ramp rates do not need to be considered, then it is optimal to not dispatch any electricity (or release any fuel) in this final low-price period. However, if ramp rates do exist, then there will be a positive offer, leading to dispatch off 11 MW under all RD curves. This enables a relatively high level of dispatch in the first high price period, which otherwise would not have been possible. Again, look at the MC curves that lead to this offers, as shown in Figure 9.45. We can see that up to a quantity of around 11 or 12 MW, the MC curve for ramp rates is lower, while beyond this point it is higher.

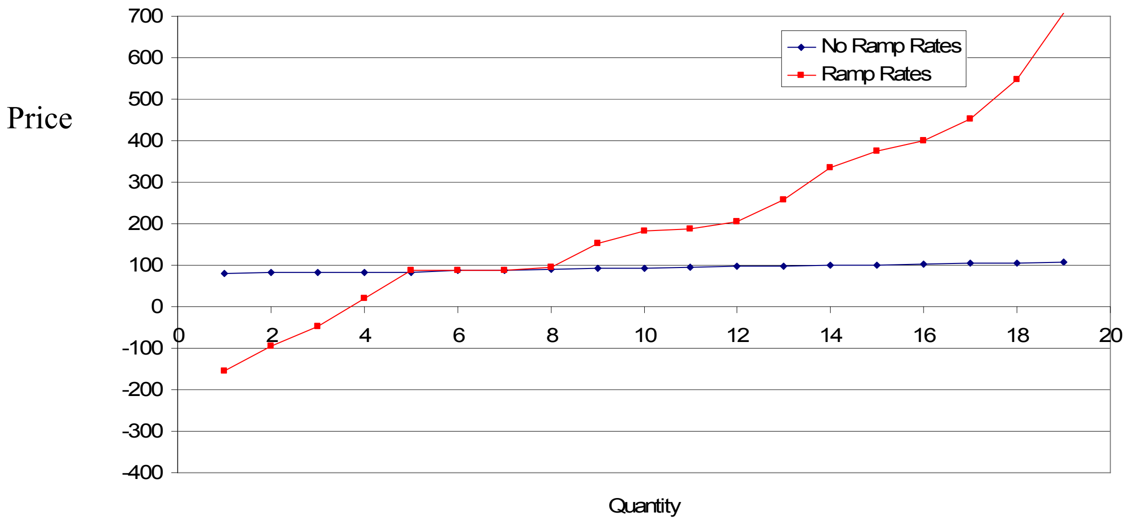

Figure 9.45 MC Curve for Reservoir Level 90 in Period 13 under No Ramp Rates

The reason for the crossover at this point is demonstrated in Figure 9.46. We can see that, with ramp rates, the optimal dispatch level in period 14 is between 14 and 15 MW. Given that this particular example was solved assuming a ramp limit of 3 MW/period, the optimal dispatch level of 11 MW in period 13 makes sense. The relative MC curves are clearly reflected in the offers shown above.

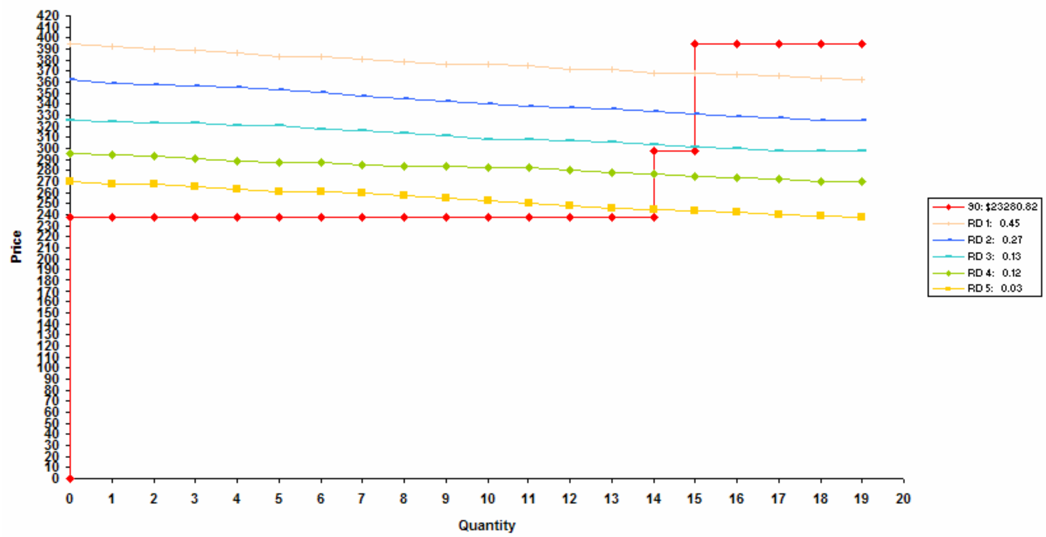

Figure 9.46 Offer for Reservoir Level 90 in Period 14 under No Ramp Rates

9.4 Summary and Conclusions

In this chapter we have demonstrated how the Value Curve approach presented in the previous chapter can be modified to account of unit rules, including:

1. Minimum feasible generation level
2. Start-up and shut-down costs
3. Minimum up and down times
4. Fixed start-up process
5. Ramp rate restrictions

We have also demonstrated that although adding these additional complexities adds two dimensions to the state space of the DP (unit state and previous dispatch level), the computational complexity increases at a much slower rate than that which would be expected from the "curse of dimensionality". In the case of the unit state dimension, this is because much of the computational complexity is in the production of the offers and their values, but for many of the unit states the offers are fixed and thus their value is easy to determine. In the case of the previous dispatch level dimension under ramp rate restrictions, we have shown that for a given reservoir level, a single overall offer can be constructed and the feasible segment of this overall offer used as the optimal offer. In other words, there is no need to construct a completely separate offer from each previous dispatch level.

We have demonstrated many of the offering behaviours hypothesised in Chapter 3 are as expected under stochastic scenarios, in addition to the deterministic scenarios presented there. With respect to the unit rules, we have shown how a unit may move through its start-up process during low-price periods (unprofitable in the short-term), in order to capitalise on high-price periods when they arrive. Alternatively, we have shown that a unit may remain switched on over very low price periods when it is costly to shut-down and start-up again, if they expect to be generating again in the short-term. We have also shown that ramp rate restrictions can have similar effects, promoting a generator to ramp up sufficiently in time for a peak set of periods, or alternatively, to ramp down in sufficient time for a set of off-peak periods. A final interesting behaviour that we have demonstrated is that a generator may, counter intuitively, provide a more generous offer under low RD curves than under high RD curves, if the correlation between the heights of these RD curves is high.

In the following chapter we present extensions to the optimal offering models presented in this thesis, along with possible approaches to dealing with these extensions.

Chapter 10

EXTENSIONS AND FUTURE RESEARCH

10.1 Introduction

In this chapter, we begin by describing some additional applications of the simulation model presented in Section 8.3, which would enable a generation firm to analyse the potential performance of their offering plan in advance of a planning horizon, or to make after-the-fact analysis of their performance over the planning horizon. The remainder of the chapter is dedicated to proposing possible extensions to the model presented in this thesis, and suggesting methods of addressing some of these. These extensions include accounting for contracts for the future supply of electricity, dealing with the cases of multiple reservoirs in parallel or in series and using the model for gas generators who are able to buy more fuel at certain points in the planning horizon if they wish.

10.2 Additional Applications of the Simulation Model

In addition to the uses of the simulation described and applied already in this thesis, there are three further valuable applications of the simulation model, which a generating firm could employ either in advance to analyse the expected performance of a fixed offering plan that they are considering, or to perform after-the-fact analysis of their offering performance over a planning horizon.

10.2.1 Finding Expected Payoff of a Given Offer Set

The simulation model can be used to simulate the payoff that could be expected for a single instance over a planning horizon for a given set of offers (one offer for each period), as opposed to an offering strategy (one offer for each possible state in each period) determined by one of the models discussed in this thesis. This given set of offers could be the offers that a generating firm is considering providing to the market for the horizon, or could come from any other sort of simple decision rule.

Therefore, this feature could be used by a generator in advance of the horizon to estimate the difference in the expected payoff between what could be achieved using the offer optimisation model and what could be achieved using the current offers under consideration.

Additionally, this feature would help in developing the input data to the optimisation model in preparation for online application. We would always expect the optimisation to produce a better quality of solution than a handmade offer set (if its input data accurately reflects the real world). If the handmade offer set is better, then this could be used as a feedback mechanism to improve the performance of the model (by improving the quality of its dataset).

10.2.2 Finding Payoff for a Given Residual Demand and Inflow Set

The simulation model enables the user to provide a set of RD curve outcomes for the planning horizon and make a comparison between:

1. The optimal decisions and payoff under Perfect Information (i.e. if we knew these sequences in advance)
2. The optimal decisions and payoff under uncertainty that could have been made given the stochastic market and inflow information that could reasonably have been available at the time when the decisions were made
3. The actual decisions and payoffs that were made and achieved in these periods

These comparisons could be made in hindsight to compare how well the generator operator did with their offers (3) compared to how well they should have been able to do (2), had they used the offer strategy construction approach proposed in this thesis (given the stochastic information that they could reasonably have had at the time). As such, this would help a generator to evaluate the potential value-added of having this model over using their current strategy.

Note that this feature also enables comparisons between the actual payoffs and the absolute optimal that could have been achieved under Perfect Information (1).

10.2.3 Evaluating Market Power Exploitation

As described in Section 8.3, the simulation model can be used to produce the optimal offering strategy of a generator that behaves as a perfect competitor, always offering at marginal cost. As such, comparisons can be made between these perfectly competitive offers and those actually made over a given horizon by a particular generator. Such an analysis would be valuable from two perspectives.

1. A market regulator could use this type of analysis to demonstrate that a generator's offers are significantly raised above marginal cost and is thus exploiting market power, and should be further restricted in its offering behaviour, or
2. A generator could use this type of analysis to defend its position that it is offering at or close to marginal cost and thus not exploiting market power, and the regulator thus has no reason to further restrict its behaviour.

Clearly, the extent of market power that the model implies will depend on the assumptions made with respect to the input data (such as any generation costs other than opportunity costs, value placed on water, and expectations with respect to possible residual curves) used in the model, as well as the structure of the model (branched versus non-branched). The careful choice of such data and structure could produce either result as desired by the user, and so the choice of inputs would likely be the point of significant debate if this technique were used in such an inquiry. In addition, the generators would argue that it must be *long-term* marginal cost that is considered, as returns made must be sufficient to encourage necessary levels of investment in generation.

10.3 Accounting for Contracts

In the wholesale electricity market, generators may hold contracts for differences (CfDs) with various clients, effectively locking in a price for a certain proportion of their generation. If a generator faces known CfD levels in each period of the planning horizon, it is a very simple matter to adjust the revenue function used in the construct of the offers accordingly. Let us define:

R_t = Revenue in period t

Q_t = Quantity dispatch in period t

$P_t(Q_t)$ = Price level in period t as a result of dispatch level Q_t

K_t = CfD level in period t

Revenue in period t was determined by:

$$R_t = Q_t * P_t(Q_t)$$

Accounting for CfDs, it is now:

$$R_t = [Q_t - K_t] * P_t(Q_t)$$ [73]

In other words, if the dispatch level is greater than the CfD level, then the generator receives the market price for the excess generation. On the other hand, if the CfD level is greater than the dispatch level, the generator must purchase electricity from the market, at the market price, to meet its contractual supply commitments.

Consider a very simple example with just one possible RD curve of the form $P_t = 120 - 2Q_t$ and a contract to supply 20MW of electricity in period t, as shown in Figure 10.1.

[73] Note that the pre-determined revenue associated with the CfDs is not considered here because it is independent of any of the operating decisions that we make.

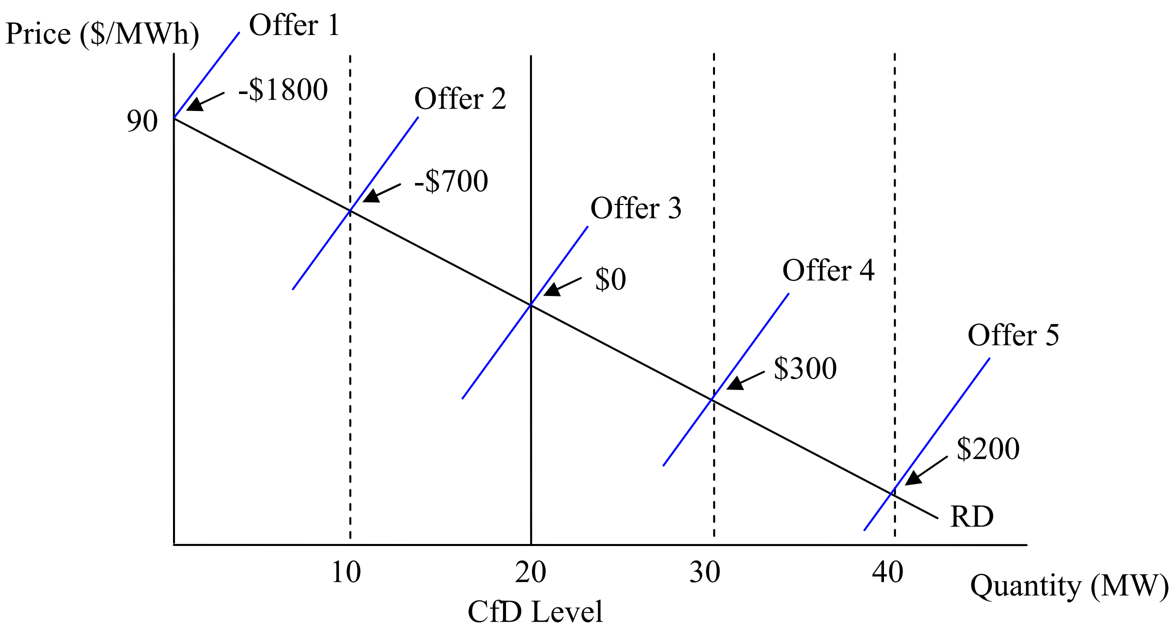

Figure 10.1 Revenue under CfDs

Consider five different possible offers (1-5), passing through the RD curve at quantities of 0, 10, 20, 30, and 40 respectively. Note that there is only a single possible RD curve in this case, this is equivalent to just making a deterministic dispatch decision.

Offer	Market Clearing Price ($)	Dispatch Level (MW)	Excess Supply (MW)	Revenue ($)
1	90	0	-20	-1800
2	70	10	-10	-700
3	50	20	0	0
4	30	30	10	300
5	10	40	20	200

Table 10.1 Net Outcome Possibilities

For this example, assuming zero marginal cost, the optimal decision out of these five options is clearly to provide Offer 4, leading to a dispatch level of 30MW, 10MW higher than the contracted level.

Note that the modification that would be required under the R&A approach discussed earlier in this thesis would not be as straightforward. The reason for this relates to the inaccuracy of that algorithm as explained in Section **6.7.1**. This inaccuracy means that it is possible to construct offers that do not register an intersection with all RD curves within the (price, quantity) space. Because many of the revenues at the points of intersection will be negative when considering contracts, as shown in the example above, the incentives for the algorithm to avoid an intersection with RD curves is greatly increased, and offers produced by this approach would be distorted as such. In order to mitigate this problem, a constant term would need to be added to all payoffs ($1800 in this case) in order to make all payoffs non-negative.

10.4 Physical Marginal Generation Cost Curve

In the preceding chapters, we have assumed that the marginal cost associated with generation from a fuel-constrained generation unit comes solely from the cost associated with foregoing future use of that fuel (i.e. the marginal opportunity cost). However, it is possible that there may be other marginal costs associated with generation that also need to be considered. These are the more direct, physical marginal costs of generation, incurred in the actual operation of the generation unit itself, and are much more likely to occur in thermal generation units than hydro.

Dealing with these physical marginal costs of generation requires a very simple adjustment to the RT phase algorithm presented in Section 6.6.2. Previously, the offers constructed were based directly on the MOC curve from the given state. When considering physical marginal generation costs, we must base offers on the total MC curve, which is the simple vertical sum of the MOC curve and the physical MC curve, as shown in Figure 10.2. We can easily determine that the total MC must be the vertical sum of the MOC and physical MC curves by considering the marginal cost to the generator of each successive unit of generation; there is a well defined marginal opportunity cost and marginal physical cost for each.

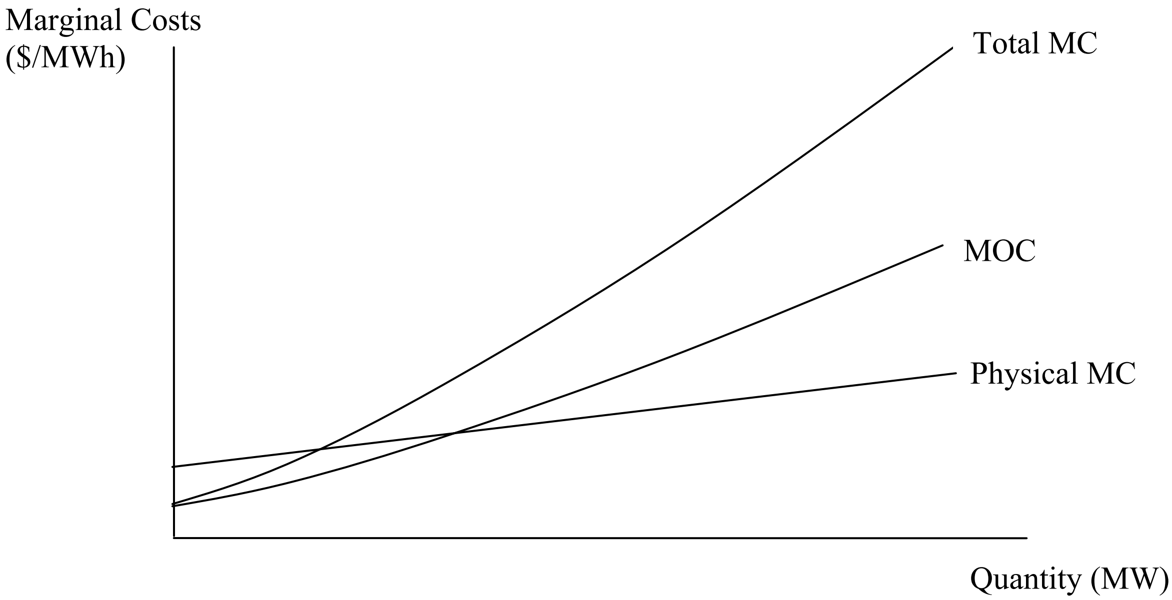

Figure 10.2 Total MC Curve

10.5 Multiple Reservoirs

Whether we have generation units attached to multiple reservoirs in parallel or in series, in order to apply the approach developed in this thesis, a single, collective offer to the market must be provided, rather than one for each unit. One of these offers must be provided for each state, as defined by the combination of the previous market outcome and the levels of all reservoirs ((n+1)-dimensional state space, where there are n separate reservoirs).

The reason that a single offer must be provided from each state is that the generation company is offering into the same uncertain market scenario from all reservoirs, assuming block dispatch[74]. If separate offers were provided from each unit, then this would mean the residual situation faced by each unit would be affected, when the offers from all other units are adjusted.

[74] Block dispatch allows the generator to provide its required dispatch from any of the units within a local block, regardless of which of these units provided the offer that led to the dispatch in the market clearing process.

In this section, we present the issues related to the construction of offers from multiple reservoirs that are either in parallel or in series, and suggest possible methods of solving these problems. For simplicity, the examples presented consider just two reservoirs.

10.5.1 Parallel Reservoirs

In order to extend the approach presented in this thesis to multiple parallel reservoirs, the value curve for each period would have to become multi-dimensional, defined over the storage levels for all reservoirs. The reason for this is that the total value of storage depends not only on the total level of storage amongst all reservoirs (as you might expect when all reservoirs are offering electricity to the same market node), but rather, the value depends on the split of storage between these reservoirs. To demonstrate this, consider the simplest case of two identical parallel reservoirs with maximum storage level *M* (total storage of *2M*), each with a single generation unit with maximum generation level $Q_{\max}$. Assume that the generator is coming into a peak period and it is desirable for it to generate as much electricity as possible from both reservoirs. If there was total available storage of *M*, then it would be more valuable to have each reservoir sitting at *½M*, than to have one sitting at *M* and the other empty and unable to produce any electricity. The construction of the MC curves from such value surfaces is far more complex than from a single value curve.

To begin to deal with this problem, we must consider the desirable balance between storage in the multiple parallel reservoirs. In a very simple case, such as that presented above, it is clear that it is optimal to distribute storage as evenly as possible. To this end, we construct a "Target Balance Line", as demonstrated in Figure 10.3 (the original theory behind this concept is discussed in Read (1989)).

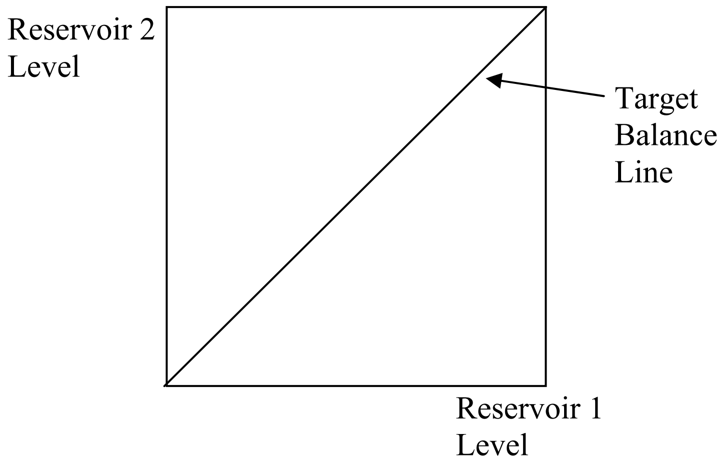

Figure 10.3 Simple Balance Line

At each period, when determining the dispatch from each unit, we will be aiming to get as near to this target balance line as possible, subject to expectations about inflows at the respective reservoirs within the next period. Assuming that there are no generation costs other than opportunity costs, then in this case, we can effectively just aggregate the two reservoirs into one, with minimal loss of accuracy[75].

However, as the situation becomes more complex, these target balance lines will take different shapes, leaving us no longer able to simply aggregate the reservoir levels, and so the construction of the MOC curves on which to base the offers becomes much more difficult. As a simple example, if the generation unit at one of the reservoirs had a greater capacity than the one at the other reservoir, then it would be desirable to have a greater storage level at the reservoir with the larger-capacity unit. This is demonstrated in Figure 10.4, where reservoir 1 has the generation unit with the larger capacity.

[75] As with the rest of this thesis, we assume here that the efficiency of the generation units is linear. Note, however, that the efficiencies of the different units do not need to be identical (i.e. generation unit 1 may require i cumecs of water flow to produce 1MW of electricity, while unit 2 may require $j \neq i$ cumecs). This can be dealt with by simply referring to reservoir levels in terms of possible units of electricity that can be produced, rather than volumes of water.

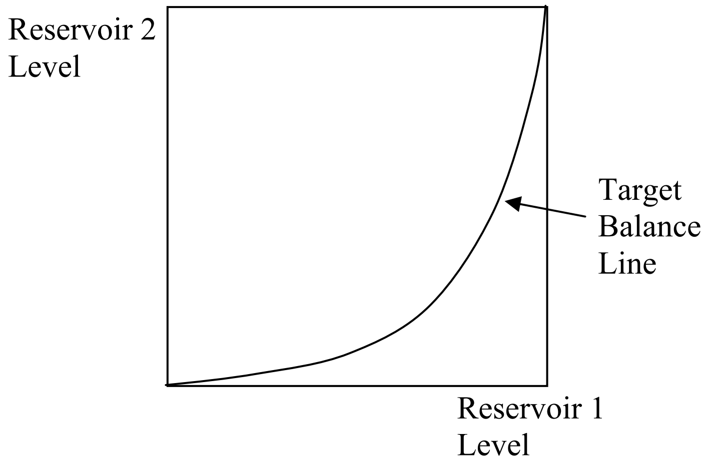

Figure 10.4 Balance Line under Non-Uniform Generation Units

This concept of target balance lines is similar to the PRISM model's "balance zone" approach (Read, Culy, Halliburton, & Winter (1988)) for balancing storage between the North and South Islands of New Zealand, but more complex, because in this thesis we are providing offers with uncertain release levels, rather than directly deciding release.

One possible alternative to explicitly modelling an *n* dimensional value surface is to have a two dimensional value surface, where the dimensions are defined as "total storage" (the sum over all reservoirs), and "relative balance", where the relative balance variable indicates how close to the target balance line you will likely end up, as a result of offering in the current period. This approach is considered beyond the scope of this thesis and is left to future research.

10.5.2 Reservoirs in Series

To consider reservoir in series (or reservoir chains), the situation is a lot more complex than the parallel reservoir case, as flows (and flow delays) between consecutive reservoirs must be modelled and optimised. This case is applicable only to hydro units, as thermal units will always be in parallel. The general concept of the approach that is described for parallel reservoirs is still valid; the value curve will need to be multi-dimensional and the target balance would (in most circumstances) favour the bulk of storage in the up-stream reservoir (as it could be used later to generate additional electricity at the downstream generator). However, the detailed application of this approach to the case of a reservoir chain is left to future research.

10.6 Gas Generators

In this section, we consider how the algorithm presented in this thesis might be applied to a gas generator which, rather than having a fixed fuel availability with uncertain additional inflows over the planning horizon, has a steady contracted flow of fuel (gas) into its "stockpile" over time, and at regular intervals can choose to either supplement this flow by purchasing gas from the gas spot market, or reduce the flow by selling gas on this spot market. We assume that flexibility in drawing gas from the pipeline effectively means that we have a small reservoir, with deterministic inflows each period equal to the net supply from the gas market. This means that gas does not need to be used immediately at the time at which it is injected at the start of the pipeline[76], and thus can be saved for use in the short- to medium-term (depending on the physical constraints of the system).

Depending on the market setup and the technology involved, those regular intervals might be once every eight hours, when the worker's shift changes and the valve for the pipeline leading to the generator could be physically changed. We define a "gas supply period" to be the set of periods between each opportunity to adjust the gas inflow (16 half-hour offering periods, in this case). This is demonstrated in Figure 10.5.

[76] In fact, due to gas pipeline transit times, it could not be used immediately anyway.

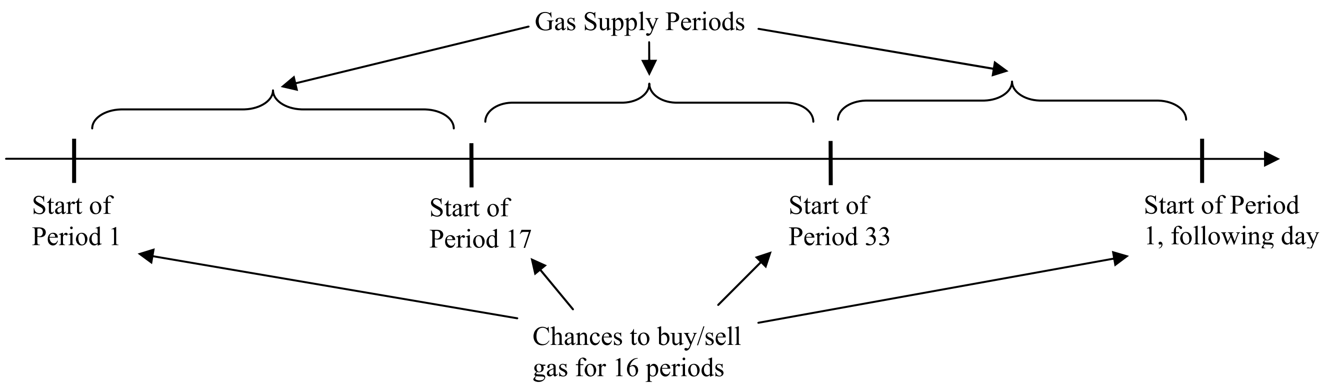

Figure 10.5 Gas Market Purchase Level Alteration Opportunities - Timeline

Before proceeding with a possible approach to this problem, we need to consider whether the gas generator has market power in the gas market (in addition to the electricity market). If they do not have market power, then in each period there will be fixed prices for purchasing gas from or selling gas to the gas market in any desired quantity, most likely with a spread between the prices. A deterministic case is shown in Figure 10.6.

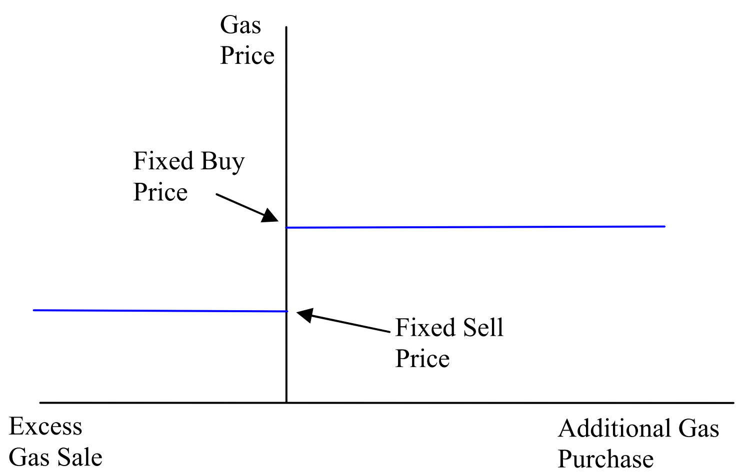

Figure 10.6 Fixed Gas Price Spread - No Market Power

If the generator does have market power, then there will be a residual supply curve that defines a changing buy or sell price, which is dependant on the quantity of gas that we trade with the market. Figure 10.7 shows how a deterministic residual supply curve is constructed by horizontally subtracting the rest-of-market demand curve from the supply curve (this is the exact opposite of the method for constructing the RD curves considered throughout this thesis).

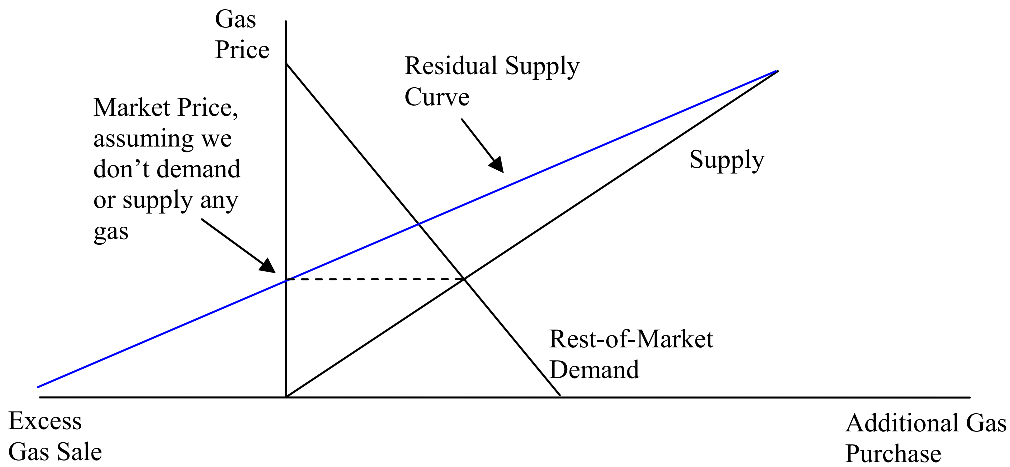

Figure 10.7 Residual Supply Curve for Gas - Market Power

In this section we discuss the approach to this problem under two assumptions. The first is that there are a discrete number of possible decisions that can be made at the beginning of each gas supply period that will determine the flow for the following eight hours. The approach that we suggest to determining the optimal decision is based on the decision branch method presented in Section 7.2. The second assumption that we consider is that a different level of flow is possible in each period of the horizon. Under this assumption, we suggest optimising the desired flow separately for each period, which is a more accurate, but consequently slower, method.

We recognise that the methods presented in this section deal with a very simplistic representation of the physical complexities of a gas generator and its supply pipeline, and leave more detailed analysis to further research.

10.6.1 Discrete Decisions – Branching Approach

The approach that we present here assumes that there are a discrete number of possible decisions at the start of each gas supply period, regarding the (positive or negative) change of flow for that period. Under a discretised situation such as this, we are able to almost directly apply the branched version of the Value Curve algorithm presented in Figure 7.24 in Chapter 7, where we have one branch associated with each possible decision. At the source of those branches (the beginning of the gas supply period), we must decide, for each reservoir level, which of those branches is most beneficial to traverse, given that the future scenario from the end of the gas supply period will be identical under all options (i.e. it is a non-transitive branching tree which converges between each gas supply period). An example of this with just three possible decisions (decrease flow by 10 units per period, no change, and increase flow by 10 units per period) is shown in Figure 10.8.

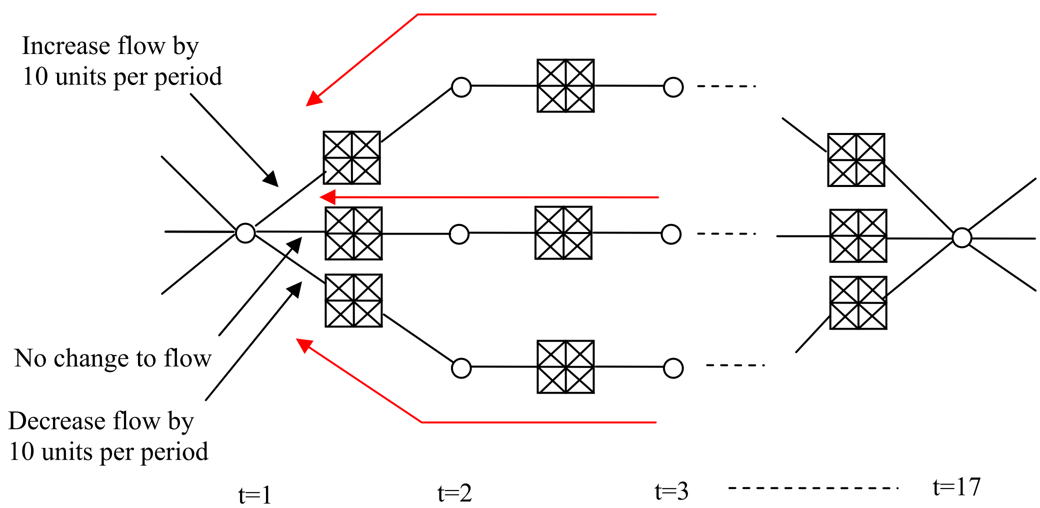

Figure 10.8 Non-Transitive Branching Tree for Discrete Gas Market Decisions

The only change that needs to be made to the original algorithm is in the construction of the value curves as the algorithm works backwards through the gas supply period. These values must be adjusted each period to account for the expected cost or revenue associated with any additional purchases or sales of gas within the gas market, from each state. For the deterministic cases presented in Figure 10.6 and Figure 10.7, it is very simple to observe the price that will result from any given buy/sell decision, and thus to establish the value curve adjustments required. Under stochastic gas market scenarios, we simply take the expected price level at our given buy/sell decision, and base the value curve adjustments on these levels.

Note that the computational requirements for the PP phase are not affected at all by considering interaction with the gas market, as the same uncertain electricity market scenarios on which the OCFs are constructed apply under all branches.

Note also that this branched non-transitive structure would make it very easy to deal with a condition that required a generator to lock in its flow change decision a certain number of offering periods out from the start of a gas supply period. The decision would be based on the branch that has the highest value at the start of the gas supply period at the reservoir level that we expect to be at, at that time.

10.6.2 Continuous Decisions – Optimisation

If a generator wishes to have complete control over the amount of gas that they purchase from or sell to the market, the decision tree approach cannot be applied, as there would be a continuum of possible decisions, and thus an infinite number of possible branches. As such, the decision would need to be optimised on a period by period basis, given the expectations for both the gas and electricity markets. Although the decision is optimised period-by-period in this manner, we recognise that the overall gas flow decision should be a constant over the entire gas supply period. If we ignore this requirement and assume that the flow can be adjusted on a period-by-period basis, every half hour, once the generator knows its electricity dispatch for that period, then we can use the following approach.

In order to construct the optimal offer from a given state, we must begin by constructing the MC curve on which to base this offer. In most of the models considered throughout this thesis to this point, this MC curve is fully defined by the monotonically non-decreasing MOC curve (which considers the future alternative uses of the fuel). In this case though, we must adjust this curve to consider possible interaction with the market, both buying and selling. To start with, let us consider the MC curve associated with the residual supply curve presented in Figure 10.7. We know that the MC curve should be twice as steep as the residual supply curve, with the same intercept on the vertical axis (in the same way that the MR curve is twice as steep as RD). We can show this by taking the derivative of the total cost function:

$$TotalCost = TC = Q * P(Q) = Q(a + bQ) = aQ + bQ^2$$

$$\therefore MC = \frac{\partial TC}{\partial Q} = a + 2bQ$$

This is demonstrated in Figure 10.9.

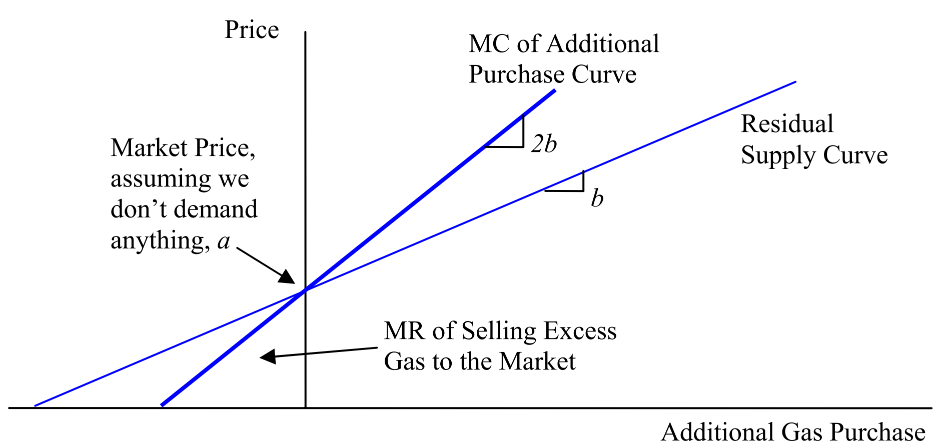

Figure 10.9 MC Curve for Gas Market Interaction

Therefore, for each additional unit of generation, we can decide whether to source the necessary gas from our reservoir (at the cost indicated by the MOC) or from the gas market (at the cost indicated by the MC of additional purchase curve). As such, the combined MC curve is the horizontal sum of the MOC and the MC of purchasing from the gas market. The values below a quantity of zero are the incremental value that could be gained by selling units of gas to the market.

Figure 10.10 demonstrates how these two MC curves can be combined, leading to an offer that results in a “buy” decision with respect to the gas market. Figure 10.11 shows a case where the optimal offer based on the combined MC curve results in a “sell” decision with respect to the gas market.

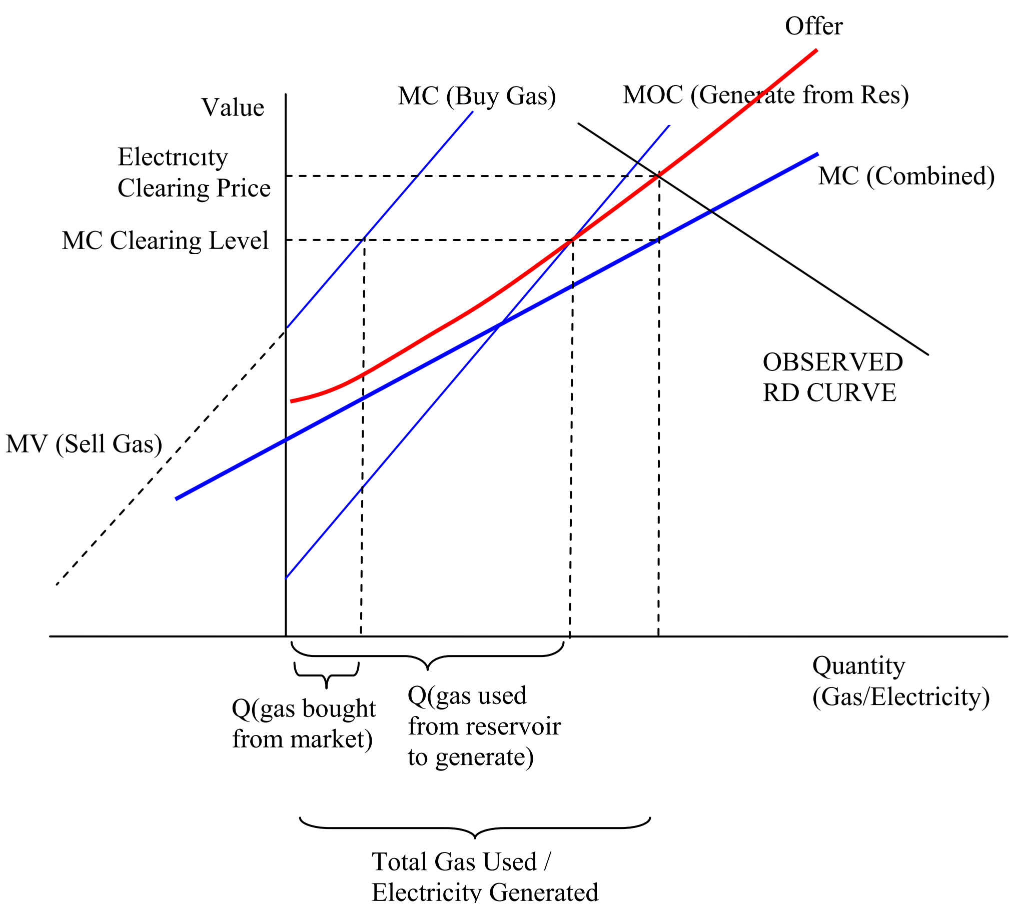

Figure 10.10 Example of Gas Market Interaction – Buy Gas from Market

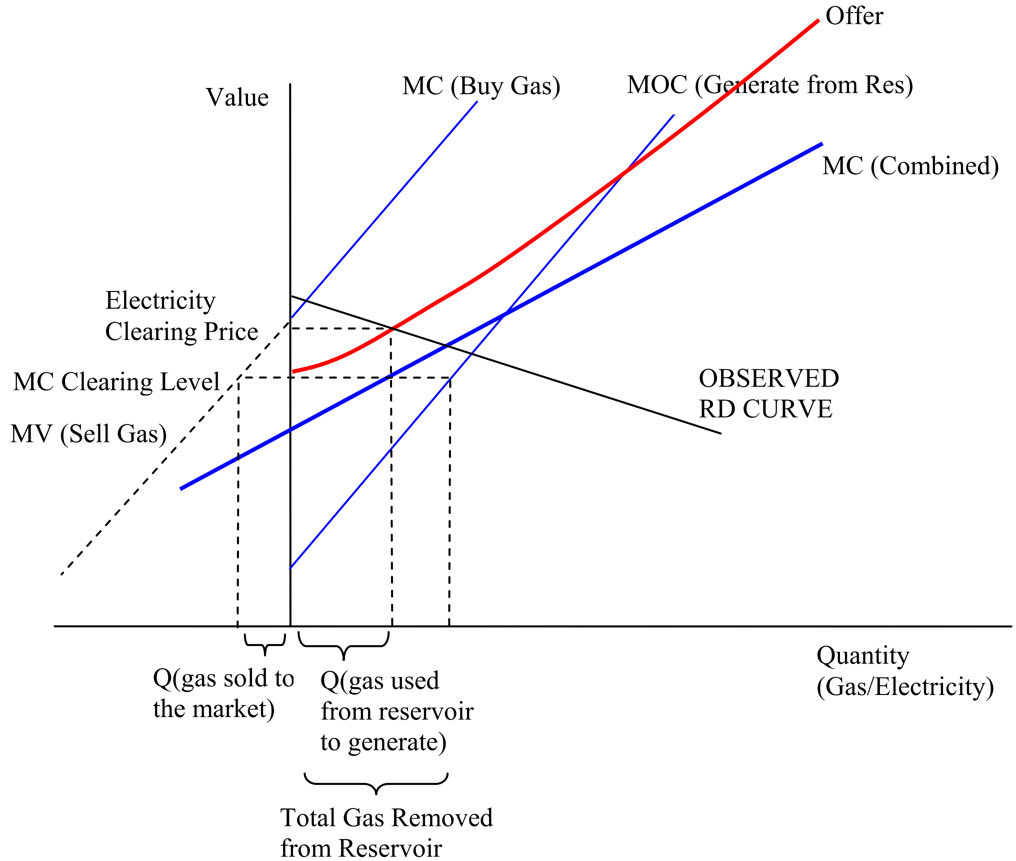

Figure 10.11 Example of Gas Market Interaction – Sell Gas to Market

Note that if we were to make different assumptions regarding market power or the deterministic/stochastic nature of the gas market to those demonstrated in this example, the only effect would be to change the construction of this combined MC curve. We could potentially even consider problems where there is a correlation between the electricity and gas markets just by changing the way that we construct this curve.

Finally, as with the branching approach to gas markets, the only difference between the approach suggested here and the original Value Curve algorithm, other than the MC curve construction, is that the value at each state will need to be adjusted to account for any revenues or costs associated with the optimal gas market interactions chosen at each state.

We leave the case of constant flows under continuous decisions to future research.

10.7 Other Extensions

Other possible extensions that we have not dealt with specifically in this thesis include:

Inflow Correlation

The type of DP structure used in this thesis is ideally suited to modelling the correlation between inflows, but to do so would add an additional “previous inflow” dimension to the DP. We have chosen to exclude this as inflows are relatively predictable over the short-term horizon that these models consider, and as such, the additional computational burden is not justified.

Offer Funnelling Rules

Offer funnelling rules are rules that limit a generator’s ability to adjust their offers as real-time approaches, without having to explain their actions to the market authorities. We propose that pre-dispatch market outcomes could be used to assist in the optimisation of offers subject to these rules by helping a generator to predict the approximate desired position of their offer in these future periods. However, we consider this extension to be outside the scope of this thesis.

Offer Finalisation Rules

An offer finalisation rule requires a generator to “lock in” their offer a certain number of periods ahead of the period in which the offer will apply. The difficulty in doing this is that the generator does not know what their reservoir level will be at that time (or even what “macro-state” the market will be in), and so is uncertain which of their future offers determined within the RT phase algorithm should be provided. One approach to this may be to combine, in some way, the offers that should be provided in that period, under the different states that could be reached at that time, from the current state. However, we consider detailed analysis of this issue to be beyond the scope of this thesis.

10.8 Summary and Conclusion

This chapter started by demonstrating that the simulation model presented in Section 8.3 can be applied in three further valuable ways to those which have already been discussed. Both of these applications could be used as intermediate steps towards an online implementation of the offer construction algorithm presented in this thesis. The first application would help a user to analyse possible offers in advance of a planning horizon, and thus help construct the required data sets, while the second application involved analysing offers after-the-fact, comparing how well a generator's offers performed, with how well they could have been expected to perform under both uncertainty and certainty.

We then discussed accounting for contracts for differences, and dealing with a physical marginal generation cost curve, and showed that they could be incorporated within the algorithm presented in this thesis, with only minor adjustments.

Finally, we have proposed more complex extensions, including multiple reservoirs and gas market interaction, which are significantly more complex in all but the most basic cases. We consider these extensions to be beyond the scope of this thesis.

Chapter 11

CONCLUSIONS

11.1 Introduction

Over the past 20 years, electricity markets around the world have shifted from government-coordinated central operation to deregulated, competitive markets where retailers and generators compete for the rights to buy and sell electricity. As such, there have been significant changes to the incentives and objectives of the players in the industry. Specifically, generating companies are now required to provide an offer into the market that describes the prices at which they are prepared to produce successive blocks of electricity. They provide these offers based on profit-driven operational strategies that are subject to various behavioural constraints and other considerations. In particular, these considerations include intertemporal constraints that the generator may face over time, such as fuel limitations and thus conservation, market uncertainty and correlation, uncertain fuel inflows, and unit operational rules.

It is the optimisation of the offering process subject to these intertemporal constraints that has been addressed in this thesis. A review of the literature in Chapter 3 identified the need for further research into this area, in order to optimise the behaviour of the generators considering all these intertemporal issues simultaneously, while simplifying the computational complexity of the method over those that have previously addressed similar issues.

This thesis has presented a multi-dimensioned dynamic program for constructing optimal offers over a planning horizon, which uses a two-stage approach whose computational time increases at a substantially slower rate than previously published algorithms as the problem size increases, and thus is able to solve more realistically sized problems in real time. Another significant contribution of this thesis is the combination of dynamic programming with a decision analysis or branching structure, where the branches represent overall *macro-states* that either the market or the generator can be in at any given time. This has been developed as a general solution approach, and as such its application is not limited to the offer optimisation problem explored in this thesis.

Section 11.2 outlines the main results and insights presented in this thesis, ignoring unit operation rules. Section 11.3 outlines the insights gained in this thesis, considering unit operation rules. Section 11.4 summarises the key research contributions. Section 11.5 presents suggestions for future research. Section 11.6 provides some concluding reflections.

11.2 Multi-Period Offer Strategy Optimisation

In this thesis, we have presented a two-phase approach to dynamic programming to optimise the behaviour of generators over a short-term planning horizon. The first phase involves pre-processing a large number of offers, which can be combined together in the second, real-time phase, to give the optimal offers for any given system state. The purpose of the phase separation was to bring the computationally intensive offer formation process out of real time, enabling this component of the algorithm to be performed in advance. The consequence of this separation is that problem scenarios of greater size and detail can now potentially be solved on-line, in a ‘reasonable’ amount of time, from the perspective of a generator providing offers into the market every half hour. The algorithm developed in Chapter 6 and Chapter 7 considers the construction of a set of sculpted offers for a generator with market power for the planning horizon, which are monotonically non-decreasing in both the price and quantity dimensions. These offers are optimised while considering the intertemporal complexities of limited fuel, residual demand curve uncertainty within a period, residual demand curve correlation between periods, and the existence of a branch structure in the real-world that defines the overall macro-state of the market at any given point in time.

In Chapter 8, comparisons on the two performance measures of computational time and expected payoff were made between our algorithm and that presented in Rajaraman & Alvarado (2003), which was the only other algorithm in the literature that considered this specific problem. We showed that over a large representative set of scenarios, our model produced expected payoffs that were slightly closer to the optimal behaviour under perfect information, than those achieved by the R&A algorithm, and that it produced significantly better payoffs under the type of residual demand curves that we believe to most accurately represent those that a generator would face in a real market. In terms of computational times, we found that substantial gains could be made by converting the two-level nested dynamic program method of the R&A algorithm into the two-phase approach that was the basis of our algorithm. Specifically:

- If the real-world inherently reflected a branched structure (where you could identify the market as being in an overall *macro-state*), but this structure was ignored, then there was an average real-time (Phase 2) time saving of around 99.88%.
- If the real-world inherently reflected a branched structure and this structure was considered, there was an average real-time (Phase 2) time saving of around 99.91%.
- The main time saving by considering an underlying branched structure in the real-world was in Phase 1, the Pre-Processing phase, which is performed just once in advance of each planning horizon.
- Ignoring the branched structure (or where the branched structure does not exist), the average Pre-Processing time required for our algorithm (4578 seconds, 76 minutes) was on average around 5.7% of the time requirement for the R&A algorithm in each period (80324 seconds, 22.3 hours).
- Considering a branched structure, the average Pre-Processing time (235 seconds) fell to just 0.3% of the time requirement for the R&A algorithm in each period.

Note that for both phases, the largest percentage time savings were under the larger, more realistic scenarios.

We then compared our algorithm to various simplified versions, under the assumption that the real-world contains a branched structure, in order to identify the value in considering each of the complexities that our algorithm deals with. The simplified models and the conclusions were:

Price-Based Offers:

We found that if we restricted the set of feasible offers, requiring a generator to offer all their available capacity at a single price level, there would be only a small loss in expected payoff, but the real-time computational time required would increase, as a two-phase algorithm would not produce single-price offers and thus was not appropriate.

Quantity-Based Offers:

By requiring a generator to provide a fixed quantity to the market, regardless of the resulting price, there was a significant fall in the expected payoffs observed, but no significant change to the computational time required.

Naïve Correlation:

This simplification assumed that the decision maker ignored the impact of the current market state on the future market states which may occur. This had the effect of significantly reducing the computational time required, but at the expense of significantly worse expected payoffs.

Naïve Water:

By ignoring water limitations, the marginal opportunity cost of water falls to zero, thus leading to much more generous offers than those that would be provided were the generator to base their offers on the true marginal opportunity cost. As the offers are based on a single constant marginal cost, the computational time requirements in real-time are negligible, and thus significantly less than those required by the algorithm that considers water limitations. However, the expected payoff is significantly and substantially lower under this simplification.

Perfect Competitor:

If the generator were to assume that they had no market power, then their optimal behaviour would be to offer their capacity into the market at marginal cost. The result of this was a significant and substantially lower expected payoff, but with no significant difference in computational times.

Value Curve Interpolation:

Rather than determining the optimal offer and value of that offer for all possible reservoir levels in all periods of the horizon, this simplification required that offers were determined for just a subset of possible reservoir levels, and the value curve interpolated between this subset of points. As expected, we found that as the subset of points got smaller and smaller, the computational time fell, but at the expense of reduced expected payoff.

We then produced an efficient frontier over the two performance dimensions of expected payoff and computational time for all these algorithms and concluded that there were only six algorithms worth considering, and that the selection between them is dependent on the specific scenario that is being faced. Depending on the time available and the size of the instance being solved, the best choice of algorithm according to this efficient frontier (from slowest with best mean percentage error to fastest with worst mean percentage error) would be one of:

1. Value Curve approach with branching and with no interpolation,
2. Value Curve approach with branching, solving only every 2nd reservoir level and interpolating the value curve between these points, and
3. Value Curve approach with branching, solving only every 5th reservoir level and interpolating the value curve between these points, and
4. Value Curve approach with branching, solving only every 10th reservoir level and interpolating the value curve between these points, and
5. Naïve Correlation Approach.
6. Naïve Water Approach

11.3 Multi-Period Offer Strategy Optimisation considering Unit Rules

In Chapter 9, we showed how the algorithm presented could be extended to consider unit operating rules. These included:

1. Minimum feasible generation level
2. Start-up and shut-down costs
3. Minimum up and down times
4. Fixed start-up process
5. Ramp rate restrictions

We demonstrated that, although adding these additional complexities adds two dimensions to the state space of the DP (unit state and previous dispatch level), the computational complexity increases at a much slower rate than that which would be expected from the curse of dimensionality. In the case of the unit state dimension, this is because much of the computational complexity is in the production of the offers and their values, but for many of the unit states the offers are fixed and thus their value is easy to determine. In the case of ramp rate restrictions, we have shown that for a given reservoir level, a single overall offer can be constructed and the feasible segment of this overall offer used as the optimal offer. In other words, there is no need to construct a completely separate offer for each previous dispatch level.

As expected, our results have illustrated that behaviour which is suboptimal in the short-term may be the optimal behaviour when considering these additional intertemporal complexities. For example, with respect to the unit rules, we showed how a unit may move through its start-up process during low-price periods (unprofitable in the short-term), in order to capitalise on high-price periods when they arrive. Alternatively, we showed that a unit may remain switched on over very low price periods when it is costly to shut-down and start-up again, if they expect to be generating again soon. We also illustrated how ramp rate restrictions can have similar effects, promoting a generator to ramp up sufficiently in time for a peak set of periods, or alternatively, to ramp down in sufficient time for a set of off-peak periods.

11.4 Research Contributions

To summarise the key contributions made in this thesis, we have:

1. Demonstrated that when taking intertemporal constraints into account in their offering process, a generator may, counter intuitively, provide a more generous offer under low RD curves than under high RD curves, if the correlation between the heights of these RD curves is high. This impact of correlations was previously hypothesised in Drayton-Bright (1997), but has not been demonstrated in the literature. The effect of this is that it is not possible to provide an offer to the market that is valid and optimal over multiple periods (as was required in the old England/Wales market).
2. Developed a marginal cost based offer patching theory, which can significantly reduce the computational complexity of the offer strategy development problem. Based on this theory, we have developed and implemented a two-phase dynamic programming algorithm, which separates the construction of offers into separate Pre-Processing and Real-Time phases, in a manner not previously suggested in the literature.
3. Determined theoretical extensions to this algorithm to deal with a combination of intertemporal constraints and considerations, including market correlation and uncertainty, fuel limitations, and unit rules (including ramp rates) that have not previously been considered together, while still maintaining feasible computational times, with respect to realistic marketplace requirements. These extensions have also been implemented.
4. Developed an additional new technique that combines the concepts of dynamic programming and decision analysis by building in a branching structure to the overall state of the generator and the market at any point in time. As well as

being able to better represent a scenario that has this type of structure, this branching significantly reduces the computational complexity of the algorithm.

11.5 Future Work

The algorithms developed in this thesis have some limitations. Here we summarise potential research avenues that could improve electricity market generator offer strategy development. Some of these avenues were discussed in more detail in Chapter 10.

- The techniques developed in this thesis apply to a generating unit or set of units attached to a single fuel reservoir. This is an accurate representation of a fuel-constrained thermal station, but a hydro generator will generally possess multiple reservoirs, either along a river chain, or attached to separate rivers. We believe that extending the algorithms to consider multiple reservoirs simultaneously would be an important extension.
- In Chapter 7, we developed a general description of a method that combines the techniques of dynamic programming and decision analysis and then applied this to the offer optimisation problem. This technique has much wider potential applications, and future research into these applications would be beneficial.
- The marginalistic dynamic programs presented in this thesis have, out of necessity, involved a discretisation of the primal state-space (reservoir levels, and market states). A valuable course of future work would be to explore the possibility of generalising these approaches to a continuous state-space, as in the method of dual dynamic programming.
- In this thesis, we have assumed that generators have a risk-neutral perspective. A valid future extension of this work would be to consider how offering strategies would change if risk aversion were to be considered.
- Finally, in this thesis we have ignored two offering constraints, the first of which occurs commonly in real-world markets, and the second of which could be implemented in future markets or market modifications. Offer finalisation rules require a final offer to be provided a certain number of periods out from real-

time (for example, four periods or two hours ahead in New Zealand), which can then only be changed in exceptional circumstances. Offer funnelling rules require that as real-time approaches, the offers made by generators cannot change too dramatically, thus encouraging correct signalling by generators in advance of finalised offers. It would be valuable to extend the methods developed in this thesis to consider these two constraints.

11.6 Concluding Reflections

In this thesis, we have seen that pre-computation is a valid and worthwhile approach to dealing with the offer strategy problem faced by generators in the deregulated electricity markets that currently exist around the world. We have also seen that there is value in applying a structured analytical approach, incorporating elements of both decision analysis and dynamic programming, in order to both create a better representation of the problem faced and to reduce the computational burden of solving that problem. Combined, these two developments mean that it is now practical to optimise an offer strategy in real-time, taking into account correlations and intertemporal issues, and considering a realistically detailed representation of the real-world and its complexities, at least for a single reservoir or fuel stockpile.

While further developments will be required before these approaches can be applied to the real problems faced by many electricity generators, the results have already provided insights about the way in which optimal offering strategies should be structured, and evolve, when inter-temporal constraints are important. And it also seems worthwhile noting that the basic analytical strategies applied here seem equally applicable to a wide variety of problems outside the electricity sector.

BIBLIOGRAPHY

Alvey, T., Goodwin, D., Ma, X., Streiffert, D., & Sun, D. (1998). A Security-Constrained Bid-Clearing System for the New Zealand Wholesale Electricity Market. *IEEE Transactions on Power Systems, 13*(2).

Anderson, E. J., & Philpott, A. B. (2001). *Estimation of Electricity Market Distribution Functions*: Australian Graduate School of Management.

Anderson, E. J., & Philpott, A. B. (2002a). Optimal Offer Construction in Electricity Markets. *Mathematics of Operations Research, 27*(1), 82-100.

Anderson, E. J., & Philpott, A. B. (2002b). Using Supply Functions for Offering Market Generation into an Electricity Market. *Operations Research, 50*(3), 477-489.

Baillo, A., Ventosa, M., Ramos, A., & Rivier, M. (2001). Strategic Unit Commitment for Generation in Deregulated Electricity Markets. In B. F. Hobbs, M. H. Rothkopf, R. P. O'Neill & H. P. Chao (Eds.), *The Next Generation of Electric Power Unit Commitment Models*. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Baillo, A., Ventosa, M., Rivier, M., & Ramos, A. (2004). Optimal Offering Strategies for Generation Companies Operating in Electricity Spot Markets. *IEEE Transactions on Power Systems, 19*(2), 745-753.

Baron, D. P. (1972). Incentive Contracts and Competitive Bidding. *American Economic Review, 62*, 384-394.

Barroso, L. A., Fampa, M. H. C., Kelman, R., Pereira, M. V. F., & Lino, P. (2002). Market Power Issues in Bid-Based Hydrothermal Dispatch. *Annals of Operations Research, 117*, 247-270.

Barton, B. (1998). From Public Service to Market Commodity: Electricity and Gas Law in New Zealand. *Journal of Energy and Natural Resources Law, 16*(4), 351-388.

Barton, B. (1999). Risk and Promise in Energy Market Liberalization: Consumer Choice in Buying Electricity. *Applied Energy, 64*, 275-288.

Barton, B. (2000a). Governance in the electricity industry. *New Zealand Law Journal*, 300-303.

Barton, B. (2000b). Responsive Regulation in the Electricity Industry. *New Zealand Law Journal*(2000b), 300-347.

Bellman, R. E., & Dreyfus, S. E. (1962). *Applied Dynamic Programming*. New Jersey: Princeton University Press.

Borenstein, S., Bushnell, J., & Stoft, S. (2000). The competitive effects of transmission capacity in a deregulated electricity industry. *Rand Journal of Economics, 31*(2), 294-325.

Bower, J., & Bunn, D. (1999). *A Model-Based Comparison of Pool and Bilateral Market Mechanisms for Electricity Trading*. London: London Business School.

Bunn, D. W. (1984). *Applied Decision Analysis*: McGraw-Hill.

Counsell, K., & Evans, L. (2003). Day Ahead Markets: Is There a Place for a Day-Ahead Market in the NZEM? http://www.iscr.org.nz/navigation/research.html Retrieved May 30, 2006

Culy, J., & Read, E. G. (1994). *Short Run Gaming in the Victorian Wholesale Electricity Market*. Wellington: NZ Institute of Economic Research.

David, A. K. (1993). Competitive bidding in electricity supply. *IEE Proceedings 1993, 140*(5), 421-426.

David, A. K., & Wen, F. (2000). Strategic Bidding in Competitive Electricity Markets: a Literature Survey. *Proceedings of IEEE PES 2000 Summer Meeting, Seattle, USA, 4*(July 2000), 2168-2173.

Day, C. J., Hobbs, B. F., & Pang, J. S. (2002). Oligopolistic Competition in Power Networks: A Conjectured Supply Function Approach. *IEEE Transactions on Power Systems, 17*(3), 597-607.

Denton, M. J., Rassenti, S. J., & Smith, V. L. (2001). Spot Market Mechanism Design and Competitivity Issues in Electric Power. *Journal of Economic Behaviour and Organisation, 44*(4), 435-453.

Doty, D. S. (2002). *Genetic Algorithm-Based Simulation of Electric Power Markets.* Masters dissertation. Iowa State University, Ames, Iowa.

Drayton-Bright, G. R. (1997). *Coordinating Energy and Reserves in a Wholesale Electricity Market.* University of Canterbury, Christchurch.

Erev, I., & Roth, A. E. (1995). *On the need for low rationality, cognitive game theory: Reinforcement learning in experimental games with unique, mixed strategy equilibria*: University of Pittsburgh.

Fatai, K., Oxley, L., & Scrimgeour, F. G. (2003). Modeling and Forecasting the Demand for Electricity in New Zealand: A Comparison of Alternative Approaches. *The Energy Journal, 24*(1), 75-102.

Ferrero, R. W., Rivera, J. F., & Shahidehpour, S. M. (1998). Application of Games with Incomplete Information for Pricing Electricity in Deregulated Power Pools. *IEEE Transactions on Power Systems, 13*(1), 184-189.

Ferrero, R. W., Shahidehpour, S. M., & Ramesh, V. C. (1997b). Transaction Analysis in Deregulated Power Systems Using Game Theory. *IEEE Transactions on Power Systems, 12*(3), 1340-1347.

Finlay, D. J. (1995). *Optimal Bidding Strategies in Competitive Power Pools.* Master of Science dissertation. University of Illinois, Urbana.

Friedman, L. (1956). A Competitive-Bidding Strategy. *Operations Research, 4*(1), 104-112.

Gajjar, G. R., Khaparde, S. A., Nagaraju, P., & Soman, S. A. (2003). Application of Actor-Critic Learning Algorithm for Optimal Bidding Problem of a Genco. *IEEE Transactions on Power Systems, 18*(1), 11-18.

Garcia, A., Campos-Nanez, E., & Reitzes, J. (2005). Dynamic Pricing and Learning in Electricity Markets. *Operations Research, 53*(2), 231-241.

Green, R. J. (1996). Increasing Competition in the British Electricity Spot Market. *The Journal of Industrial Economics, 44*(2), 205-216.

Green, R. J., & Newberry, D. M. (1992). Competition in the British Electricity Spot Market. *Journal of Political Economy, 100*(5), 929-953.

Gross, G., & Finlay, D. J. (2000). Generation Supply Bidding in Perfectly Competitive Electricity Markets. *Computational and Mathematical Organization Theory, 6*, 83-98.

Gross, G., Finlay, D. J., & Deltas, G. (1998). Strategic Bidding in Electricity Generation Supply Markets. *Proceedings of IEEE PES 1998 Winter Meeting*, 309-315.

Hammons, T. J., Rudnick, H., & Barroso, L. A. (2002). Latin America: Deregulation in a Hydro-Dominated Market. *HRW, 10*(4), 20-27.

Hao, S. (1999). A Study of Basic Bidding Strategy in Clearing Pricing Auctions. *Proceedings of IEEE PES Power Industry Computer Applications Conference (PICA '99)*, 55-60.

Harraway, J. (1997). *Introductory Statistical Methods for Biological, Health and Social Sciences*: University of Otago Press.

Hernaez, E. C., Gil, J. B., Leon, J. I. D. L. F., Roque, A. M. S., Rodriguez, M. J. V., Gonzalez, J. G., et al. (2004). Competitors' Response Representation for Market Simulation in the Spanish Daily Market. In D. W. Bunn (Ed.), *Modelling Prices in Competitive Electricity Markets* (pp. 22-68): John Wiley & Sons.

Hobbs, B. F. (2001). *The Next Generation of Electric Power Unit Commitment Models*. Dordrecht: Kluwer Academic Publishers.

Hobbs, B. F., & Kelly, K. A. (1992). Using game theory to analyze electric transmission pricing policies in the United States. *European Journal of Operational Research, 56*, 154-171.

Hobbs, B. F., Metzler, C. B., & Pang, J. S. (2000). Strategic Gaming Analysis for Electric Power Systems: An MPEC Approach. *IEEE Transactions on Power Systems, 15*(2), 638-645.

Hogan, W. W. (1997). A Market Power Model with Strategic Interaction in Electricity Networks. *The Energy Journal, 18*(4), 107-141.

Kelman, R., Barroso, L. A., & Pereira, M. V. F. (2001). *Market Power Assessment and Mitigation in Hydrothermal Systems*: Power System Research Inc.

Kelman, R., & Pereira, M. V. F. (1998). *Strategic Pricing in Hydrothermal Systems*: Power System Research Inc.

Klemperer, P. D., & Meyer, M. A. (1989). Supply Function Equilibria in Oligopoly Under Uncertainty. *Econometrica, 57*(6), 1243-1277.

Krishna, V., & Ramesh, V. C. (1997a, 30/4/03). *Intelligent Agents for Negotiations in Market Games, Part 1: Model.* Paper presented at the Proceedings of PICA 1997.

Krishna, V., & Ramesh, V. C. (1997b, 30/4/03). *Intelligent Agents for Negotiations in Market Games, Part 2: Application.* Paper presented at the Proceedings of PICA 1997.

Ladurantaye, D. D., Gendreau, M., & Potvin, J. V. (2006). *Strategic Bidding for Price-Taker Hydroelectricity Producers*: Centre for Research on Transportation.

Lamont, J. W., & Rajan, S. (1997). Strategic Bidding in an Energy Brokerage. *IEEE Transactions on Power Systems, 12*(4), 1729-1733.

Li, C., Hsu, E., Svoboda, A. J., Tseng, C. L., & Johnson, R. B. (1997). Hydro Unit Commitment in Hydro-Thermal Optimization. *IEEE Transactions on Power Systems, 12*(2), 764-769.

Li, C., Svoboda, A. J., Guan, X., & Singh, H. (1999). Revenue Adequate Bidding Strategies in Competitive Electricity Markets. *IEEE Transactions on Power Systems, 14*(2), 492-497.

Mco. (2003). NZEM Pricing: How do New Zealand Electricity Prices Work. http://www.nzelectricity.co.nz/downloads/NZEMPricing.pdf Retrieved February 20, 2003

MED. (2000). Ministerial Inquiry into the Electricity Industry. http://www.electricityinquiry.govt.nz/reports/final/final-03.html Retrieved May 30, 2006

Mielczarski, W., Michalik, G., & Widjaja, M. (1999a). Bidding Strategies in Electricity Markets. *Proceedings of 21st IEEE International Conference, 1999*, 71-76.

Murray, K., & Stevenson, T. (2005). State of Competition and Investment and Entry Barriers to New Zealand's Wholesale and Retail Electricity Markets. http://www.electricitycommission.govt.nz/pdfs/submissions/pdfsretail/pdsfcompetition/Final-recommend.pdf Retrieved May 30, 2006

Natarajan, H. (2003). *Modeling Electricity Spot Price Behaviour: Seasonality, Mean Reversion, Heteroscedasticity and Jumps*. Bangalore: Indian Institute of Management Bangalore.

Neame, P. J., Philpott, A. B., & Pritchard, G. (2001). *Long-Term Hydro Scheduling in a Pool Market.* Paper presented at the 36th Annual Conference of the Operational Research Society of New Zealand, Christchurch, NZ.

Neame, P. J., Philpott, A. B., & Pritchard, G. (2003). Offer Stack Optimisation in Electricity Pool Markets. *Operations Research, 51*(3), 397-408.

Neame, P. J., Philpott, A. B., & Pritchard, G. (2005). Hydroelectric Reservoir Optimization in a Pool Market. *Mathematical Programming, 103*, 445-461.

Newberry, D. M. (1991). *Capacity-constrained Supply Function Equilibria: Competition and Entry in the Electricity Spot Market*. Cambridge: Cambridge Center for Economic Policy Research, Working Paper 557.

Norris, J. R. (1997). *Markov Chains*. Cambridge: Press Syndicate of the University of Cambridge.

O'Neill, R. P., Sotkiewicz, P. M., Hobbs, B. F., Rothkopf, M. H., & Stewart, R. J. J. (2005). Efficient Market-Clearing Prices in Markets with Non-Convexities. *EJORS, 164*, 269-285.

Oren, S. S., & Rothkopf, M. H. (1975). Optimal Bidding in Sequential Auctions. *Operations Research, 23*(6).

Pereira, M. V. F., Campodonico, N., & Kelman, R. (1998). *Long-term Hydro Scheduling based on Stochastic Models.* Paper presented at the EPSOM, Zurich.

Pereira, M. V. F., Campodonico, N., & Kelman, R. (1999). *Application of Stochastic Dual DP and Extensions to Hydrothermal Scheduling*: Power System Research Inc.

Pereira, M. V. F., Granville, S., Dix, R., & Barroso, L. A. (2004a). *Strategic Bidding Under Uncertainty: A mixed Integer Linear Programming Approach*: Power System Research Inc.

Pereira, M. V. F., Granville, S., Dix, R., & Barroso, L. A. (2004b). *Nash Equilibrium in Strategic Bidding: A Binary Expansion Approach*: Power System Research Inc.

Pereira, M. V. F., Granville, S., Dix, R., & Barroso, L. A. (2004c). *Strategic Bidding Under Uncertainty: A Comparison Between The Binary Expansion Approach and Non Linear Optimization Methods*: Power System Research Inc.

Pereira, M. V. F., & Pinto, L. M. V. G. (1991). Multi-Stage Stochastic Optimization Applied to Energy Planning. *Mathematical Programming, 52*, 359-375.

Petrov, V. (2002). *Exploring Computational Power Markets with Evolutionary Algorithms.* Masters dissertation. Iowa State University, Ames, Iowa.

Philpott, A. B. (2004). *Single Unit-Commitment Problems in Electricity Pool Markets.* Paper presented at the 39th Annual Conference of the Operational Research Society of New Zealand, Auckland, NZ.

Philpott, A. B., Craddock, M., & Waterer, H. (2000). Hydro-electric unit commitment subject to uncertain demand. *European Journal of Operational Research, 125*, 410-424.

Philpott, A. B., Pritchard, G., Neame, P., & Zakeri, G. (2002). *Market Distribution Functions in the Electric Power Industry*. Auckland: Electric Power Optimisation Centre, University of Auckland.

Philpott, A. B., & Schultz, R. (2004). *Unit Commitment in Electricity Pool Markets*. Auckland: Electric Power Optimisation Centre, University of Auckland.

Philpott, A. B., & Schultz, R. (2006). *Unit Commitment in Electricity Pool Markets*. Auckland: Electric Power Optimisation Centre, University of Auckland.

Pokharel, B. K., Shrestha, G. B., Lie, T. T., & Fleten, S. E. (2004, 4/7/05). *Profit Based Unit Commitment in Competitive Markets.* Paper presented at the 2004 International Conference on Power System Technology, Singapore.

Pokharel, B. K., Shrestha, G. B., Lie, T. T., & Fleten, S. E. (2005). *Price Based Unit Commitment for Gencos in Deregulated Markets.* Paper presented at the IAEE Power Engineering Society General Meeting, USA.

Pritchard, G., Zakeri, G., & Philpott, A. B. (2002). *Estimation of Market Distribution Functions in Electricity Pool Markets.* Paper presented at the Proceedings of the 37th Annual Conference of the Operational Research Society of New Zealand, Auckland, 2002.

Pritchard, G., Zakeri, G., & Philpott, A. B. (2004). *HERO (Hydro-Electric reservoir optimization).* Paper presented at the 39th Annual Conference of the Operational Research Society of New Zealand, Auckland, NZ.

Pritchard, G., Zakeri, G., & Philpott, A. B. (2005). *Non-Parametric Estimation of Market Distribution Functions in Electricity Pool Markets*. Auckland: EPOC, University of Auckland.

Rajaraman, R., & Alvarado, F. (2003). *Optimal Bidding Strategy in Electricity Markets Under Uncertain Energy and Reserve Prices*: Power Systems Engineering Research Centre.

Rardin, R. L., & Uzsoy, R. (2001). Experimental Evaluation of Heuristic Optimisation Algorithms: A Tutorial. *Journal of Heuristics, 7*, 261-304.

Read, E. G. (1979). *Optimal Operation of Power Systems.* University of Canterbury, Christchurch.

Read, E. G. (1985). *A New Variant of Stochastic DP For Multi-Reservoir Release Scheduling.* Paper presented at the Operational Research Society of New Zealand Annual Conference.

Read, E. G. (1989). A Dual Approach to Stochastic Dynamic Programming for Reservoir Release Scheduling. In A. O. Esogbue (Ed.), *Dynamic Programming for Optimal Water Resources System Management* (pp. 361-372). New York: Prentice-Hall.

Read, E. G. (2001). Update on Offer Curve Development. In H. T. Warren Clark (Ed.).

Read, E. G., Culy, J. G., Halliburton, T. S., & Winter, N. L. (1988). A Simulation Model for Electricity Planning in New Zealand. *Operational Research '87, 16*(2), 493-507.

Read, E. G., & George, J. A. (1986). Dual Dynamic Programming For Linear Production/Inventory Systems. *New Zealand Operational Research, 14*(2), 133-136.

Read, E. G., George, J. A., Lamar, B. W., & Rosenthal, R. E. (1994). *Hydro Unit Commitment/Dispatch Scheduling*. Christchurch: Energy Modelling Research Group, University of Canterbury.

Read, E. G., George, J. A., & Macgregor, A. D. (1994, 10/6/05, 25/1/2005). *Stochastic Dual Dynamic Programming with Lagged Variables.* Paper presented at the 30th Annual Conference of the Operational Research Society of New Zealand, Wellington, N.Z.

Richter, C. W. J., & Sheble, G. B. (1998). Genetic Algorithm Evolution of Utility Bidding Strategies for the Competitive Marketplace. *IEEE Transactions on Power Systems, 13*(1), 256-262.

Richter, C. W. J., Sheble, G. B., & Ashlock, D. (1999). Comprehensive Bidding Strategies with Genetic Programming/Finite State Automata. *IEEE Transactions on Power Systems, 14*(4), 1207-1212.

Roth, A. E., & Erev, I. (1995). Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behaviour, 8*(Special Issue: Nobel Symposium), 164-212.

Rudkevich, A., Duckworth, M., & Rosen, R. (1998). Modeling Electricity Pricing in a Deregulated Generation Industry: The Potential for Oligopoly Pricing in a Poolco. *The Energy Journal, 19*(3).

Sakk, E., Thomas, R. J., & Zimmerman, R. (1997). Power System Bidding Tournaments for a Deregulated Environment. *Proceedings of the Thirtieth Hawaiian International Conference on System Sciences, 5*, 681-686.

Scott, T. (1998). *Hydro Reservoir Management for an Electricity Market with Long-Term Contracts.* PhD dissertation. University of Canterbury, Christchurch.

Seeley, K., Lawarree, J., & Liu, C. C. (2000). Analysis of Electricity Market Rules and Their Effects on Strategic Behaviour in a Non-Congestive Grid. *IEEE Transactions on Power Systems, 15*(1), 157-162.

Sheble, G. B. (1996a). Priced Based Operation in an Auction Market Structure. *IEEE Transactions on Power Systems, 11*(4), 1770-1777.

Sheble, G. B. (1996b). Agent Based Economics. In M. Ilac, F. Galiana & L. Fink (Eds.), *Power System Restructuring: Engineering and Economics* (pp. 187-242): Kluwer Academic Publishers.

Singh, H. (1999a). Introduction to Game Theory and its Application in Electric Power Markets. *IEEE Computer Applications in Power*, 18-22.

Singh, H. (1999b). The California Electricity Market: An Assessment of the First Year of Operations. *Power Engineering Society Summer Meeting 1999, 2*, 753-755.

Singh, H., Hao, S., & Papalexopoulos, A. (1998). Transmission Congestion in Competitive Electricity Markets. *IEEE Transactions on Power Systems, 13*(2), 672-680.

Song, H., Liu, C.C., Lawarree, J, Dahlgren, R.W. (2000a). Optimal Electricity Supply Bidding by Markov Decision Process. *IEEE Transactions on Power Systems, 15*(2), 618-624.

Stothert, A., & MacLeod, I. (2000). Competitive Bidding as a Control Problem. *IEEE Transactions on Power Systems, 15*(1), 88-94.

Sugianto, L. F., & Widjaja, M. (2001). Optimizing Bidding Strategy in the Australian National Electricity Market. *International Journal of Fuzzy Systems, 3*(4), 532-540.

Tipping, J. P., Read, E. G., & McNickle, D. C. (2004). *The Incorporation of Hydro Storage into a Spot Price Model for the New Zealand Electricity Market.* Paper presented at the 6th IAEE European Conference, Zurich, Switzerland.

Villar, J., & Rudnick, H. (2003). Hydrothermal Market Simulator Using Game Theory: Assessment of Market Power. *IEEE Transactions on Power Systems, 18*(1), 91-98.

Visudhipan, P., & Ilic, M. D. (1999). Dynamic Games-Based Modeling of Electricity Markets. *Proceedings of IEEE PES 1999 Summer Meeting, 1*, 274-281.

von der Fehr, N. M., & Harbord, D. (1993). Spot Market Competition in the UK Electricity Industry. *Economics Journal, 103*, 531-546.

Wen, F., & David, A. K. (2002). Coordination of Bidding Strategies in Day-Ahead Energy and Spinning Reserve Markets. *Electrical Power and Energy Systems, 24*, 251-261.

Wilson, R. (1997, 21/3/03). *Activity Rules for a Power Exchange.* Paper presented at the University of California Energy Institute Power Conference, Berkeley.

Wolak, F. A. (2000). *An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behaviour in a Competitive Market*. Stanford, California: Department of Economics, Stanford University.

Yakowitz, S. (1982). Dynamic Programming Applications in Water Resources. *Water Resources Research, 18*(4), 673-696.

Yang, M. (1995). *Dual Dynamic Programming for Reservoir Management with Correlated Inflows.* University of Canterbury, Christchurch.

Yang, M., Wen, F., Wu, F. F., & Ni, Y. (2002). Development of Optimal Bidding Strategies for Competitive Generation Companies Using Possibility Theory.

Yeh, W. W. G. (1985). Reservoir Management and Operations Models: A State-of-the-Art Review. *Water Resources Research, 21*(12), 1797-1818.

Zhang, D., Wang, Y., & Luh, P. B. (2000). Optimization Based Bidding Strategies in the Deregulated Market. *IEEE Transactions on Power Systems, 15*(3), 981-986.

APPENDICES

A. CONTRIBUTORY LINES

The purpose of this appendix is to determine the form of the contributory line associated with a certain pair of marginal cost curve and residual demand curve slope.

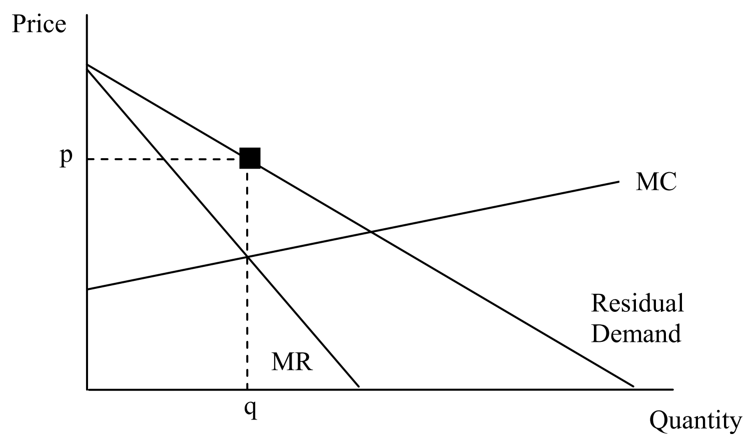

Figure A.1 MR and MC Curves

Consider the cost and revenue curves shown in Figure A.1. The form of the marginal cost curve is:

$$MC = MC_{\mathrm{int}} + MC_{slp} q$$

The form of the residual demand curve is:

$$RD = a(sc,q) - b(sc,q)q$$

and hence the form of the marginal revenue curve (the derivative of the equation, quantity multiplied by price, as defined by the residual demand curve), is:

$$MR = a(sc,q) - 2b(sc,q)q$$

where, under each residual demand curve scenario, the *a* and *b* functions are just constant values.

The optimal dispatch quantity is given by the point where marginal revenue is equal to marginal cost. Therefore:

$$a(sc,q) - 2b(sc,q)q = MC_{\mathrm{int}} + MC_{slp} q$$

$$q\left(MC_{slp} + 2b(sc,q)\right) = a(sc,q) - MC_{\mathrm{int}}$$

$$q = \frac{a(sc,q) - MC_{\mathrm{int}}}{MC_{slp} + 2b(sc,q)}$$

The optimal price point to offer is then found at this quantity level on the residual demand curve:

$$RD = a(sc,q) - b(sc,q)q$$

$$RD = p = a(sc,q) - b(sc,q)\left[\frac{a(sc,q) - MC_{\text{int}}}{MC_{slp} + 2b(sc,q)}\right]$$

Now that we know the analytic form of the optimal dispatch point under any residual demand curve position, we can establish the form of the contributory line by finding its intercept and slope.

To find the intercept of the CL, we simply need to determine the price level to offer where q = 0. Therefore, for q = 0, the optimal quantity equation tells us that:

$$q = \frac{a(sc,q) - MC_{\text{int}}}{MC_{slp} + 2b(sc,q)} = 0$$

$$a(sc,q) = MC_{\text{int}}$$

We now need to substitute this into the optimal price equation:

$$p = a(sc,q) - b(sc,q)\left[\frac{a(sc,q) - MC_{\text{int}}}{MC_{slp} + 2b(sc,q)}\right] = MC_{\text{int}} - b(sc,q)\left[\frac{MC_{\text{int}} - MC_{\text{int}}}{MC_{slp} + 2b(sc,q)}\right]$$

$$p = MC_{\text{int}} - b(sc,q)\left[\frac{0}{MC_{slp} + 2b(sc,q)}\right] = MC_{\text{int}}$$

Therefore, the intercept of the CL is $(0, MC_{\text{int}})$.

Now, we need to know the slope of the CL. We find this by taking the derivative of the price and quantity each with respect to the parameter that is changing, $a(sc,q)$, and then compare the derivatives.

$$\frac{\partial q}{\partial a(sc,q)} = \frac{1}{2b + MC_{slp}}$$

$$\frac{\partial p}{\partial a(sc,q)} = 1 - \frac{b}{2b + MC_{slp}}$$

$$\text{Slope of CL} = \frac{rise}{run} = \frac{\dfrac{\partial q}{\partial a(sc,q)}}{\dfrac{\partial p}{\partial a(sc,q)}} = \left[1 - \frac{b}{2b - MC_{sl}}\right]\left[2b + MC_{sl}\right] = b + MC_{sl}$$

Therefore, the slope of the CL is $b + MC_{sl}$, giving the equation for the form of the CL to be:

$$p = MC_{\text{int}} + \left[b + MC_{sl}\right]q$$

B. DIRECT MV CURVE APPROACHES UNDER INFLOW UNCERTAINTY

This appendix reports the method used to construct MV curves for a given period directly from the MV curves for the following period, under the simple cases of no uncertainty and inflow uncertainty and correlation. Here, the decision is a fixed level of release for each period, rather than a more complex offer, as considered in Chapter 6.

The following notation is used throughout this appendix:

MVR_t	MV of releasing an additional unit of storage in period t
MVS_t	MV of storing an additional unit of storage in period t for use in a later period
$EMVS_t$	Expected MV or storing an additional unit of storage in period t for use in a later period
R_t	Storage level in period t
Q_t	Release quantity in period t

TR_{t+1}	Target reservoir level at period *t+1*
f_i	Level of inflow *i*
i	Inflow index
I	Number of possible inflow levels per period
π_i	Probability of inflow *i*

1. Deterministic Case

With no uncertainty, the optimality condition for determining release over a planning horizon is that the MV in all periods must align (ignoring reservoir limits). Therefore, in each period the condition is:

$$MVR_t = MVS_t$$

2. Inflow Uncertainty

When inflow uncertainty is considered, the release level is chosen such that:

$$MVR_t = EMVS_t$$

$EMVS_t$ is calculated for all target storage levels as the weighted average of the MVS_t that could be achieved as a result of inflow uncertainty if we targeted this storage level.

$$EMVS_t(TR_t) = \sum_{i=1}^{I} \pi_i MVS_t(TR_{t+1} + f_i)$$

where

$$TR_{t+1} = R_t - Q_t$$

This is demonstrated with a numerical example in Figure 1442HB.1.

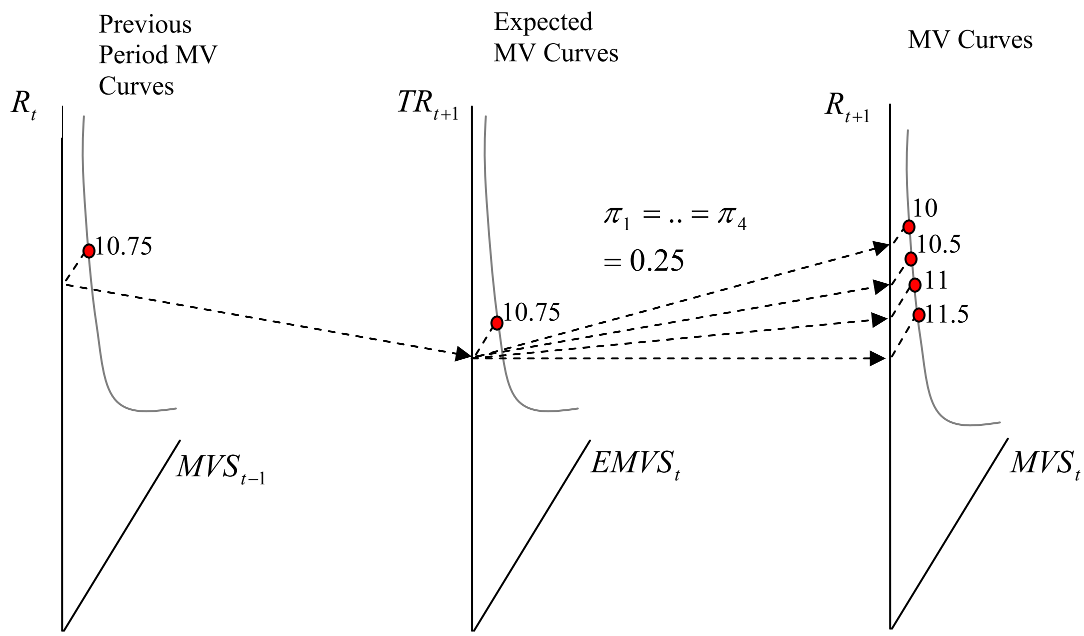

Figure B.1 Equating MV's under Inflow Uncertainty - Forwards

Therefore, from a given reservoir level in a single period, we work forwards, equating $EMVS_t$ and MVR_t, where MVR_t is determined from either:

1. The cost of the other units that we don't have to run if we generate with this hydro unit (if it is a cost minimisation problem), or
2. The marginal revenue (MR) associated with a single linear RD curve (if it is a revenue maximisation problem)

$EMVS_t$ will be upward sloping in release, as greater release leads to a lower EOP reservoir level, while MVR_t will be downward sloping because as the release is increased, generation begins to be substituted away from the more efficient alternative hydro generation units (or the MR falls if it is a revenue maximisation problem). This is shown in Figure B.2.

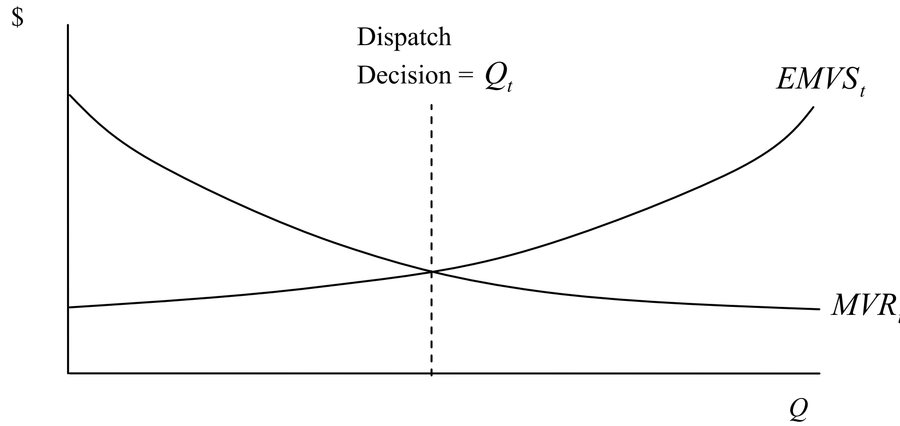

Figure B.2 Equating $EMVS_t$ and MVR_t in a Single Period

3. Inflow Uncertainty and Correlation

There are now *I* different *MVS* curves in every period, indexed by the previously observed inflow, *i*. Figure B.3 provides an example.

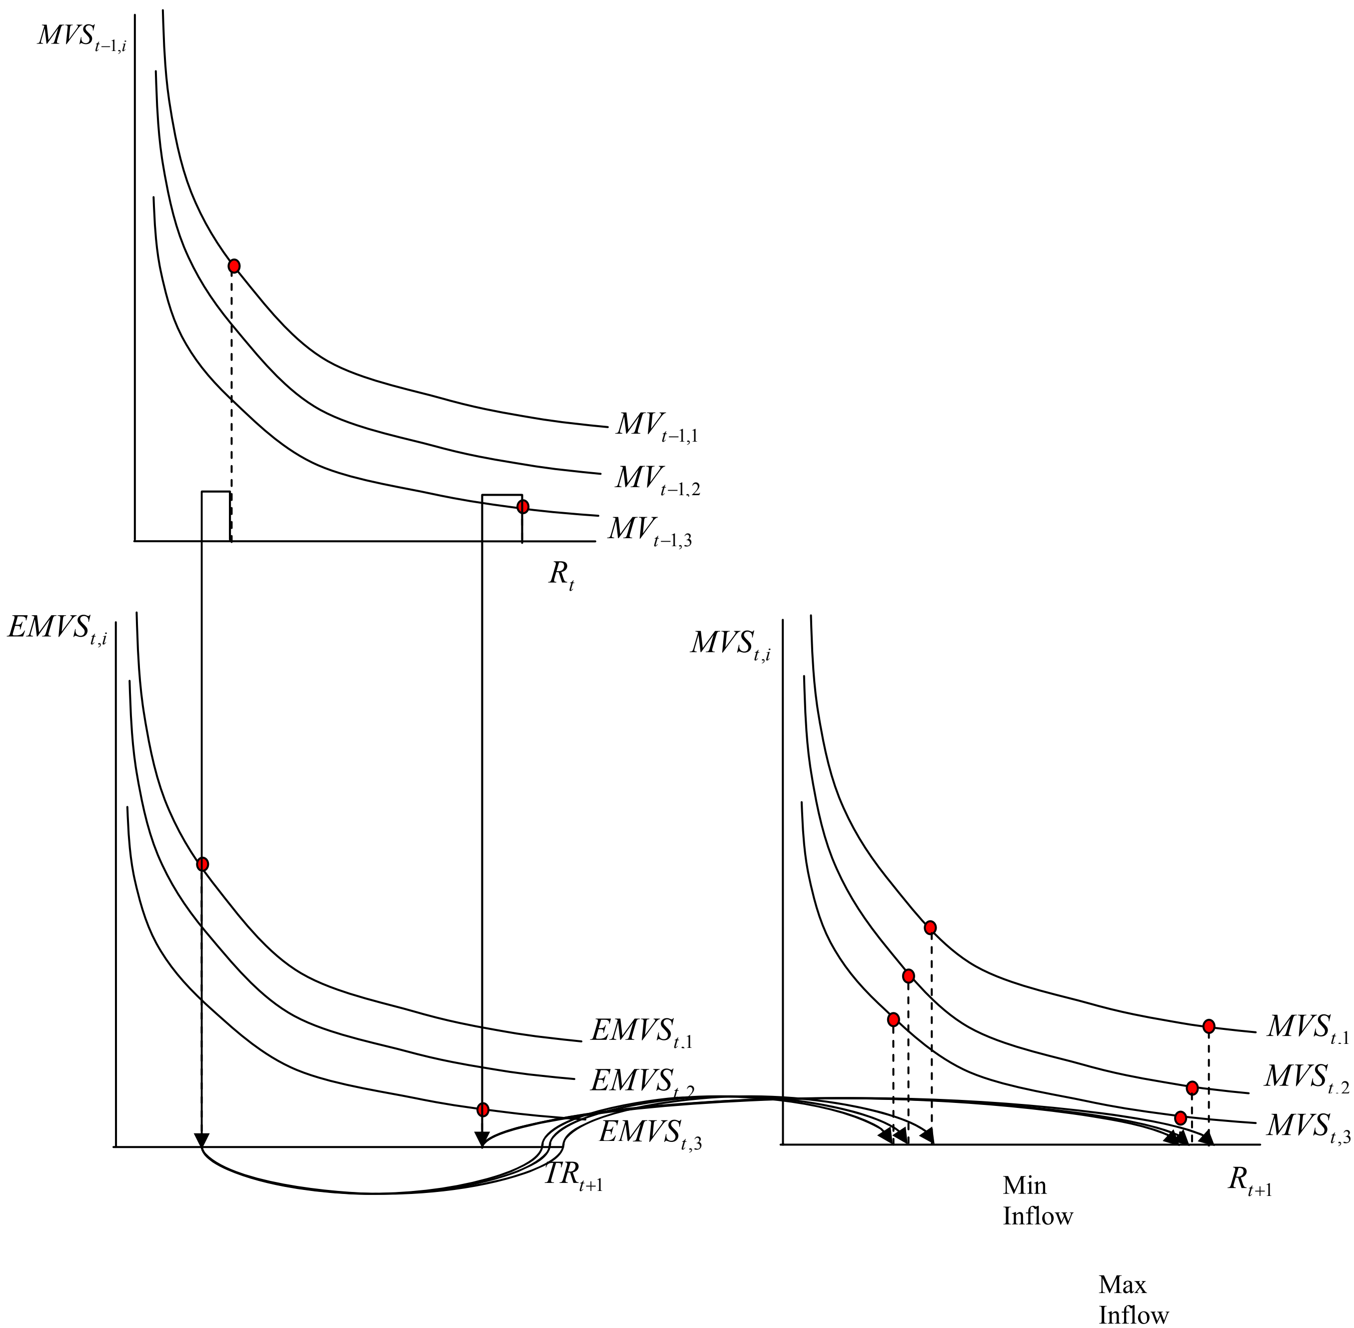

Figure B.3 Equating MV's under Inflow Uncertainty and Correlation

Each previous inflow level has its own $EMVS_t$, which is calculated as a weighted sum of the MVS_t curves (with probabilities depending on the previous inflow level). Release decisions are then made in the same way as the case without correlation.

The approach to constructing MV curves under market uncertainty is very similar to that under inflow uncertainty, as it does not really matter where the distribution of end-of-period reservoir level comes from.

The difficulty in extending this approach to market uncertainty, though, is that the decision becomes a market offer (rather than a fixed release) and as such, release depends on the market outcome. Therefore, under this structure, the concept of "target storage" no longer applies.

C. Construction of the Beginning-of-Period MOC Curve

If there is no correlation between market outcomes in consecutive periods, then in any given period there is a single value curve for the following period, based on expectations of what the market will do from then on and how we will optimally be able to respond to them. This gives a single MV curve, and when limited fuel is considered, this is what provides us with our Marginal Opportunity Cost (MOC) curve for generation from a given reservoir level in the current period.

However, when we have correlation between market uncertainty, a different value curve is needed for the following period for each RD curve (market state) that could possibly occur in the current period (number of possible RD curves denoted by p). This means that we have p MV curves for the following period and hence p MOC curves for the current period and state. In other words, for each RD curve that could occur in this period, there is a different associated MOC curve.

This implies that to be theoretically correct, we cannot simply base our offer for this period on a single MOC curve[77]. In this appendix, we will consider:

1. How the offer should be constructed, given that the MOC curve is dependent on the RD curve that occurs, and is this approach practical and able to be solved in a reasonable time frame?
2. What alternative methods could give similar results in a shorter time frame?

1. How the Optimal Offer should be Constructed

As mentioned above, the MOC curve faced by a generator in any given state is different, depending on which RD curve occurs in this state. As such, we would ideally like to construct an offer such that dispatch is set to its optimal position for each RD curve, based on each RD curve's associated MOC curve.

Figure 1450HC.1 demonstrates this process under a scenario with three possible RD curves in each period, with the following steps:

i. For a reservoir level of x at the beginning of period t, with previous RD curve r, identify the three possible RD curves with associated probabilities of $pr_{r,1}$, $pr_{r,2}$ and $pr_{r,3}$.
ii. Each of these RD curves has its own associated MV curve for period $t+1$, which can be used to construct the associated MOC curves. Determine these MOC curves.
iii. Determine the optimal dispatch point under each RD curve separately in their own plane, based on their own individual MOC curves[78].
iv. Bring these dispatch points into a single plane, to determine the overall optimal offer.

[77] Note that the R&A approach does not face this problem because it is a completely primal approach, and as such decisions are not based on a MOC curve
[78] This quantity can be determined by the intersection between the marginal revenue (MR) curve and the MOC curve, while the price is given by the RD curve at this quantity.

v. Calculate the expected value of this offer and the associated future value and insert into the value curve for period t for the market state index r and reservoir level x.

If the set of points found in step (iii) is monotone, then the overall optimal offer determined in step (iv) is simply this set along with arbitrary monotone connections, and is completely consistent with the possible MV_{t+1} curves that could occur. However, as shown in Figure C.1, it is quite likely that they will not form a monotone set, and hence some way of determining the best trade-off within the regions of uncertainty (the shaded square) is needed.

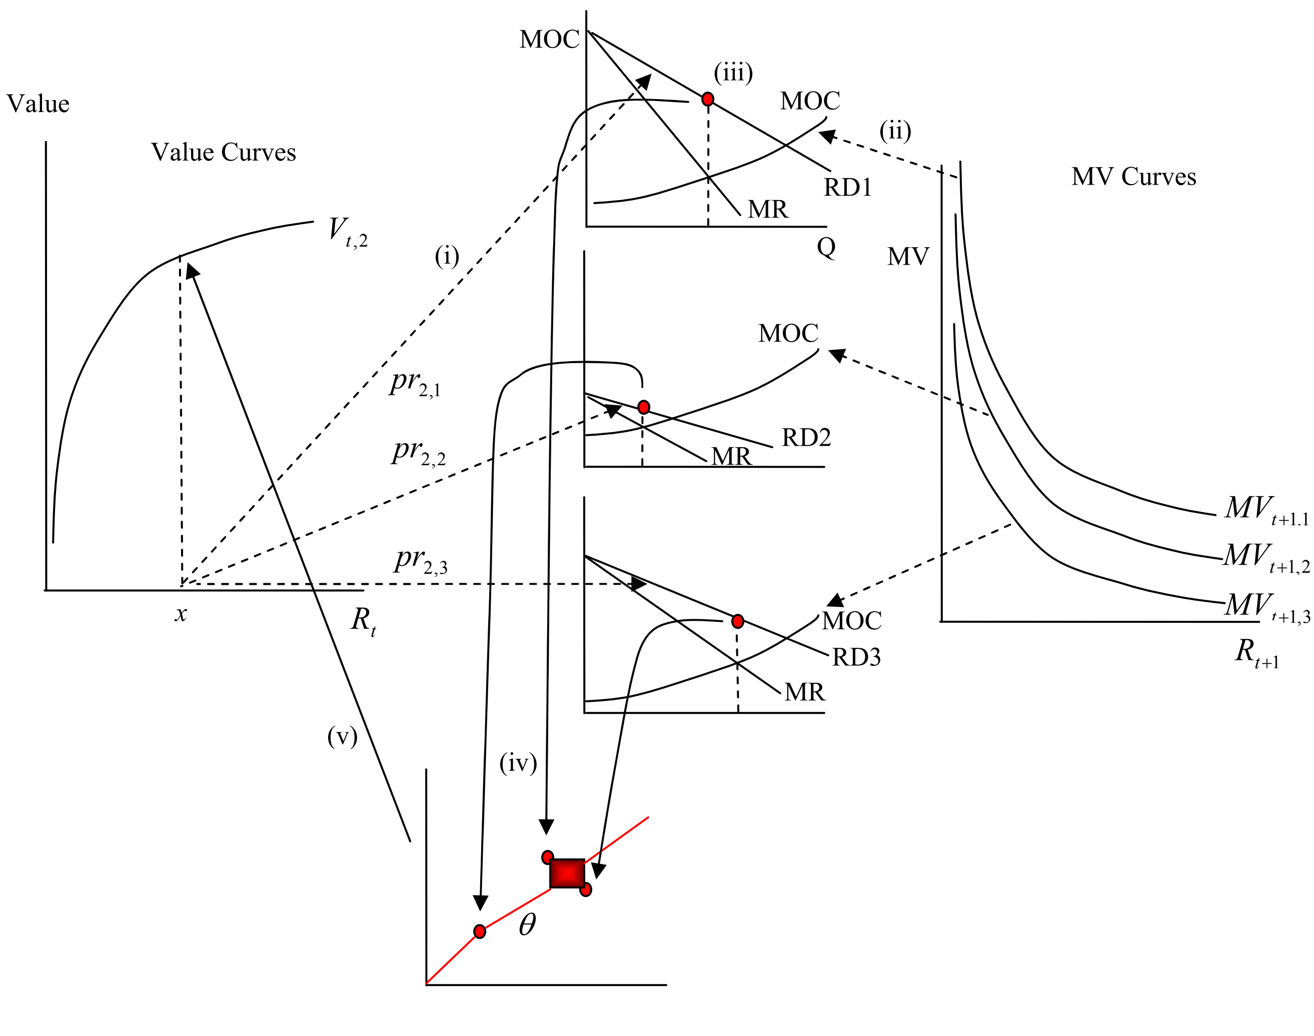

Figure C.1 Constructing an Offer that is Consistent with the MV_{t+1} Curves

If we assume that the payoff along each of the RD curves is unimodal (which is true for a linear RD curve), then it is desirable to have the dispatch as close to its optimal point as possible, on each RD curve. As such, the approach required to determine the optimal monotone offer is very similar to that shown in Section 4.5, where optimality conditions were established to determine the best location of these horizontal and vertical segments over regions of uncertainty, under kinked RD curves. In that case it was the kink and the different RD slopes that led to the non-monotonicity, whereas in this case it is the different MOC curves associated with each RD curve that lead to the non-monotonicity. The optimal monotonised offer will therefore contain vertical and/or horizontal steps through the region of uncertainty.

Unfortunately, as the number of RD curves considered increases, the number and size of regions of uncertainty will increase, and hence so does the complexity of the optimality conditions and the difficulty of performing this monotonisation, causing the approach to be very slow and thus impractical[79]. Therefore, we need an approach that will produce similar results but in significantly less computational time.

2. Alternative Methods

An obvious alternative to the approach described above is to take a vertically weighted average of the MV_{t+1} curves, and base the offer purely on the single weighted MOC that results. This is demonstrated in Figure C.2, with the following steps:

i. At state *x* now in period *t*, with previous RD curve 2, there are three possible RD curves here, with associated probabilities of $pr_{2,1}$, $pr_{2,2}$ and $pr_{2,3}$.

ii. Construct the $EMV_{t,r}$ curve for this state based on these probabilities.

[79] Alternatively, a DP could be performed over each area of uncertainty to find the optimal offer path. This is effectively what the R&A approach does when performing a DP over (price, quantity) space to find the optimal offer, and would mean that the benefits of a Two-Phase approach would be substantially lost.

iii. Produce the optimal offer over all possible RD curves using the DP approach described in this thesis, given the weighted MOC associated with the $EMV_{t,r}$ curve.

iv. Calculate the combined expected current and future value of this offer and insert into the value curve for period t for the market state index $r = 2$ and reservoir level x.

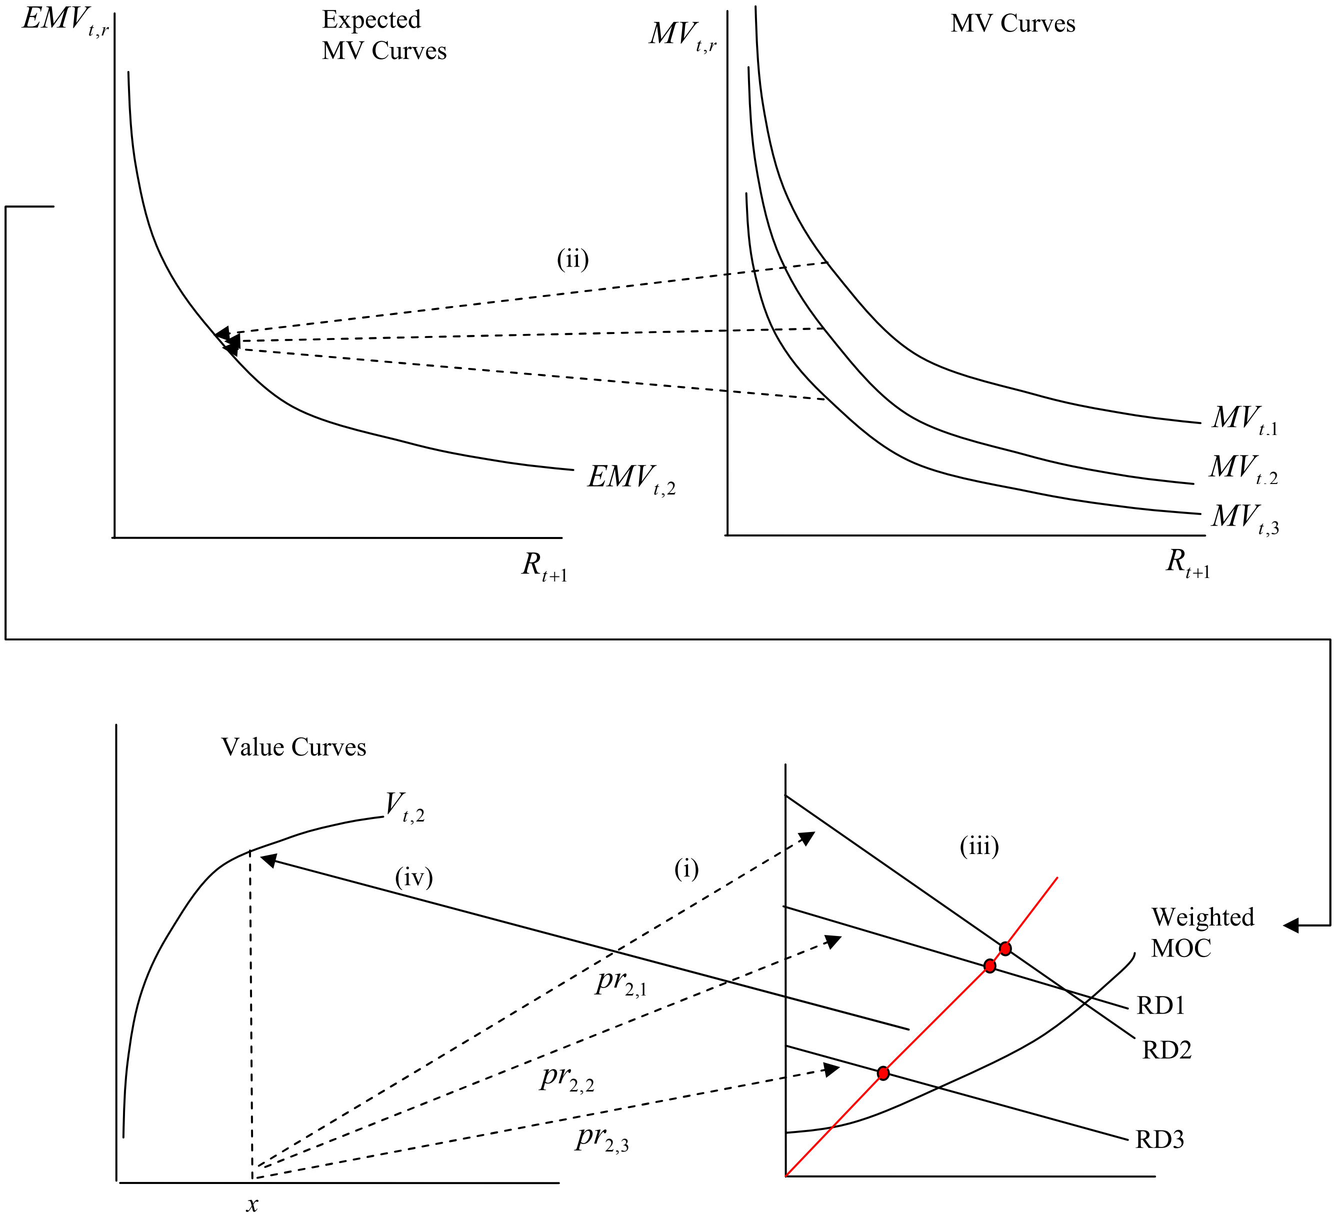

Figure C.2 Constructing an Offer based on a Weighted MV_{t+1} Curve

This approach is clearly far simpler than explicitly considering the separate MOC curves associated with each of the RD curves. In addition, this approach allows the algorithm to be split into a Two-Phase approach, as presented in this thesis. Before considering this approach to be acceptable, we need to know how much the quality of the solutions is going to be affected by this simplification, and then decide whether this affect will be acceptable.

Figure C.3 presents a simple example where there are two equally likely possible linear RD curves, and constant MOC levels of $30 and $10 associated with them, respectively.

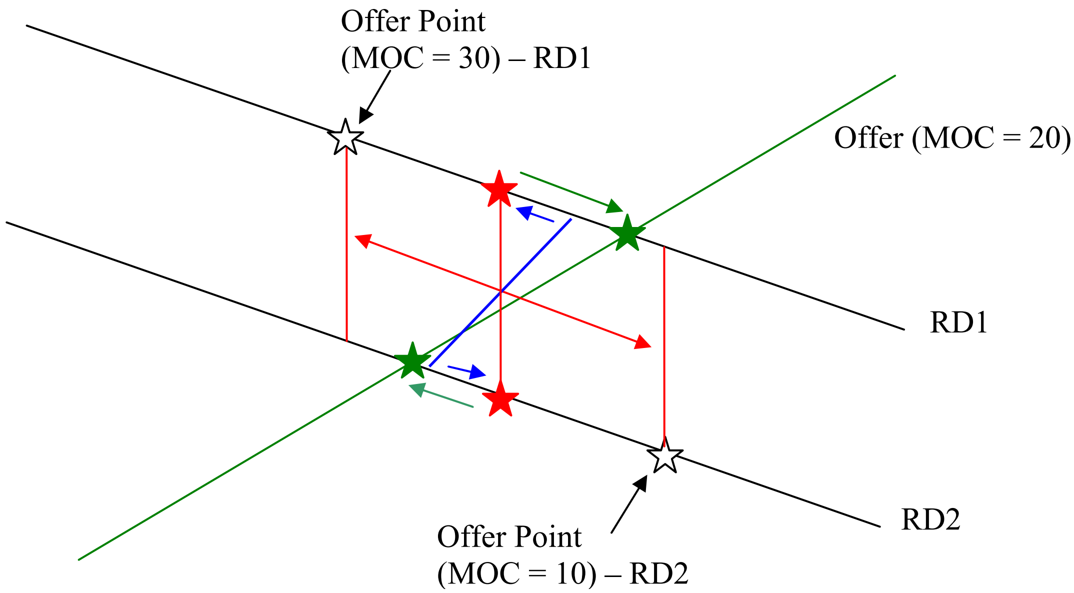

Figure C.3 Comparing the two Approaches

We can see that the two optimal points under the individual MOC curves (white stars) do not form a monotone set, and so the offer cannot pass through both these points, and must be adjusted accordingly. Given that the payoff along the RD curves is unimodel, we can show that the optimal adjusted offer will contain a vertical step[80] (red lines indicate possible examples) between these two RD curves, somewhere in the range indicated by the red arrow. The specific position of this vertical step will depend on the relative probabilities and values of being dispatched on either of the two RD curves. Recall from Section 4.5, that we know the offer must pass vertically between these two RD curves, because if we take any forward sloping segment between the two curves (as shown in blue), then we can get the points of intercept with the RD curves closer to their

[80] Note that there could also be times when a horizontal segment would be needed.

optimal positions (red stars) by making it vertical (moving in direction indicated by blue arrows).

On the other hand, the optimal overall offer to provide, based on a weighted MOC of 20 is shown in green, along with the associated dispatch points on the two RD curves (green stars). We can clearly see that each of the red stars are in slightly better positions than the green stars, closer to the optimal dispatch points (white stars) under the given RD curves. This is as expected, but the question is how much worse off is the generator by providing the green offer, compared to the red?

To answer this question, consider what would happen if the probability of RD1 was significantly higher than that of RD2. The optimal position of the vertical step (red line) would move to the left, towards the white star on RD1 with the now higher expected payoff (and further from the white star on RD2). Additionally, the weighted MOC would increase to closer to 30, moving the offer based on this weighted MOC upwards, and hence also closer to the white star on RD1 (and further away from the white star on RD2). In other words, the general directions of movement of the offers in response to changes in probability are the same for the two approaches.

Conclusion

In this appendix, we have shown two ways under which offers could be formed when there is correlation between RD curves in consecutive periods, both based on MOC curves.

The first method will work and produce a completely consistent outcome, but would not be particularly practical unless the optimal points under each of the market scenarios formed a monotone set.

It is a relatively common simplification made in the literature to use the MV_{t+1} curve as the MV_t curve because they are likely to be pretty similar. This assumption is fairly reasonable in practice, especially as time discretisation gets finer, yet not strictly accurate. This is effectively what the second method proposes, although it uses a *weighted average* of the EOP MV curves as the BOP MV curve. This simplification essentially implies that release is independent of the resulting market state.

Therefore, if we perform this simplification and use an overall weighted MOC curve, we can expect to get a very similar result to the completely consistent approach, but for significantly less computational effort. As such, for we conclude that this simplification is not unreasonable, and that it will allow application of the resulting algorithm to much larger-scale scenarios than could be considered without the simplification.

D. Pairwise Comparison Tests for Computational Time on Full Models Non-Branched

Paired T-Test and CI: NumPeriods = 20 VCN, NumPeriods = 40 VCN

```
Paired T for NumPeriods = 20 VCN - NumPeriods = 40 VCN_1

                   N      Mean     StDev  SE Mean
NumPeriods = 20   160    59.806    55.505    4.388
NumPeriods = 40   160   137.944   150.161   11.871
Difference        160  -78.1375  106.7098   8.4362


95% CI for mean difference: (-94.7989, -61.4761)
T-Test of mean difference = 0 (vs not = 0): T-Value = -9.26  P-Value = 0.000
```

Paired T-Test and CI: NumPeriods = 40 R&A, NumPeriods = 20 R&A

```
Paired T for NumPeriods = 40 R&A - NumPeriods = 20 R&A

                   N     Mean     StDev  SE Mean
NumPeriods = 40   160   106939    180030    14233
NumPeriods = 20   160    53709     77103     6095
Difference        160  53230.1  112035.6   8857.2


95% CI for mean difference: (35737.2, 70723.0)
T-Test of mean difference = 0 (vs not = 0): T-Value = 6.01  P-Value = 0.000
```

Paired T-Test and CI: NumPeriods = 20 R&A, NumPeriods = 20 VCN

```
Paired T for NumPeriods = 20 R&A - NumPeriods = 20 VCN

                   N     Mean    StDev  SE Mean
NumPeriods = 20  160  53708.7  77102.5   6095.5
NumPeriods = 20  160     59.8     55.5      4.4
Difference       160  53648.9  77054.4   6091.7


95% CI for mean difference: (41617.8, 65679.9)
T-Test of mean difference = 0 (vs not = 0): T-Value = 8.81  P-Value = 0.000
```

Paired T-Test and CI: NumPeriods = 40 R&A, NumPeriods = 40 VCN

```
Paired T for NumPeriods = 40 R&A - NumPeriods = 40 VCN

                   N    Mean   StDev  SE Mean
NumPeriods = 40  160  106939  180030    14233
NumPeriods = 40  160     138     150       12
Difference       160  106801  179918    14224


95% CI for mean difference: (78709, 134893)
T-Test of mean difference = 0 (vs not = 0): T-Value = 7.51  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 VCN, NumResLevels = 100 VCN

```
Paired T for NumResLevels = 300 VCN - NumResLevels = 100  VCN

                   N     Mean     StDev  SE Mean
NumResLevels = 3  319  138.571   149.494    8.370
NumResLevels = 1  319   59.730    56.444    3.160
Difference        319  78.8401  107.6241   6.0258


95% CI for mean difference: (66.9847, 90.6956)
T-Test of mean difference = 0 (vs not = 0): T-Value = 13.08  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 R&A, NumResLevels = 100 R&A

```
Paired T for NumResLevels = 300 R&A - NumResLevels = 100 R&A

                   N     Mean     StDev  SE Mean
NumResLevels = 3  319   114996    181423    10158
NumResLevels = 1  319    46147     66708     3735
Difference        319  68848.3  125061.0   7002.1


95% CI for mean difference: (55072.1, 82624.6)
T-Test of mean difference = 0 (vs not = 0): T-Value = 9.83  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 100 R&A, NumResLevels = 100 VCN

```
Paired T for NumResLevels = 100 R&A - NumResLevels = 100  VCN

                   N     Mean    StDev  SE Mean
NumResLevels = 1  319  46147.3  66708.0   3734.9
```

```
NumResLevels = 1  319     59.7     56.4      3.2
Difference        319  46087.6  66662.1   3732.4


95% CI for mean difference: (38744.3, 53430.8)
T-Test of mean difference = 0 (vs not = 0): T-Value = 12.35  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 R&A, NumResLevels = 300 VCN

```
Paired T for NumResLevels = 300 R&A - NumResLevels = 300 VCN

                    N    Mean   StDev  SE Mean
NumResLevels = 3  319  114996  181423    10158
NumResLevels = 3  319     139     149        8
Difference        319  114857  181312    10152


95% CI for mean difference: (94884, 134830)
T-Test of mean difference = 0 (vs not = 0): T-Value = 11.31  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 40 VCN, NumDispLevels = 20 VCN

```
Paired T for NumDispLevels = 40 VCN - NumDispLevels = 20  VCN

                    N     Mean     StDev  SE Mean
NumDispLevels =   478  134.515   148.134    6.776
NumDispLevels =   478   63.964    64.801    2.964
Difference        478  70.5502  103.2597   4.7230


95% CI for mean difference: (61.2698, 79.8306)
T-Test of mean difference = 0 (vs not = 0): T-Value = 14.94  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 40 R&A, NumDispLevels = 20 R&A

```
Paired T for NumDispLevels = 40 R&A - NumDispLevels = 20 R&A

                    N     Mean     StDev  SE Mean
NumDispLevels =   478   121992    181482     8301
NumDispLevels =   478    39310     58041     2655
Difference        478  82681.1  137417.3   6285.3


95% CI for mean difference: (70330.8, 95031.4)
T-Test of mean difference = 0 (vs not = 0): T-Value = 13.15  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 20 R&A, NumDispLevels = 20 VCN

```
Paired T for NumDispLevels = 20 R&A - NumDispLevels = 20  VCN

                    N     Mean    StDev  SE Mean
NumDispLevels =   478  39310.4  58040.7   2654.7
NumDispLevels =   478     64.0     64.8      3.0
Difference        478  39246.4  57994.7   2652.6


95% CI for mean difference: (34034.2, 44458.7)
```

T-Test of mean difference = 0 (vs not = 0): T-Value = 14.80 P-Value = 0.000

Paired T-Test and CI: NumDispLevels = 40 R&A, NumDispLevels = 40 VCN

Paired T for NumDispLevels = 40 R&A - NumDispLevels = 40 VCN

	N	Mean	StDev	SE Mean
NumDispLevels =	478	121992	181482	8301
NumDispLevels =	478	135	148	7
Difference	478	121857	181370	8296

95% CI for mean difference: (105556, 138158)
T-Test of mean difference = 0 (vs not = 0): T-Value = 14.69 P-Value = 0.000

Paired T-Test and CI: NumRDs = 100 VCN, NumRDs = 25 VCN

Paired T for NumRDs = 100 VCN - NumRDs = 25 VCN

	N	Mean	StDev	SE Mean
NumRDs = 100 VCN	637	143.141	146.014	5.785
NumRDs = 25 VCN	637	55.433	58.785	2.329
Difference	637	87.7080	112.5218	4.4583

95% CI for mean difference: (78.9533, 96.4627)
T-Test of mean difference = 0 (vs not = 0): T-Value = 19.67 P-Value = 0.000

Paired T-Test and CI: NumRDs = 100 R&A, NumRDs = 25 R&A

Paired T for NumRDs = 100 R&A - NumRDs = 25 R&A

	N	Mean	StDev	SE Mean
NumRDs = 100 R&A	637	129527	179509	7112
NumRDs = 25 R&A	637	31856	52180	2067
Difference	637	97670.8	146326.4	5797.7

95% CI for mean difference: (86286.0, 109055.7)
T-Test of mean difference = 0 (vs not = 0): T-Value = 16.85 P-Value = 0.000

Paired T-Test and CI: NumRDs = 25 R&A, NumRDs = 25 VCN

Paired T for NumRDs = 25 R&A - NumRDs = 25 VCN

	N	Mean	StDev	SE Mean
NumRDs = 25 R&A	637	31855.7	52180.3	2067.5
NumRDs = 25 VCN	637	55.4	58.8	2.3
Difference	637	31800.3	52137.5	2065.8

95% CI for mean difference: (27743.7, 35856.8)
T-Test of mean difference = 0 (vs not = 0): T-Value = 15.39 P-Value = 0.000

Paired T-Test and CI: NumRDs = 100 R&A, NumRDs = 100 VCN

Paired T for NumRDs = 100 R&A - NumRDs = 100 VCN

	N	Mean	StDev	SE Mean
NumRDs = 100 R&A	637	129527	179509	7112
NumRDs = 100 VCN	637	143	146	6
Difference	637	129383	179401	7108

95% CI for mean difference: (115425, 143342)
T-Test of mean difference = 0 (vs not = 0): T-Value = 18.20 P-Value = 0.000

E. MINITAB OUTPUT FOR SOLUTION QUALITY FOR FULL MODELS NON-BRANCHED

General Linear Model: Percentage Error versus RD Type, Approach, ...

Factor	Type	Levels	Values
RD Type	fixed	5	Branch Different, Naturally Monotone, Single Different, Single Similar, Stepped
Approach	fixed	2	R&A Approach, Value Curve Approach (not branched)
NumPeriods	fixed	2	20, 40
NumResLevels	fixed	2	100, 300
NumDispLevels	fixed	2	20, 40
NumRDs	fixed	2	25, 100
NumInflows	fixed	2	1, 5
TransitionCorrelationType	fixed	2	Strong, Weak

Analysis of Variance for Percentage Error, using Adjusted SS for Tests

Source	DF	Seq SS	Adj SS	Adj MS	F
RD Type	4	6665.53	6665.53	1666.38	432.93
Approach	1	1744.47	1744.47	1744.47	453.22
RD Type*Approach	4	3255.90	3255.90	813.97	211.47
NumPeriods	1	10.03	10.03	10.03	2.61
Approach*NumPeriods	1	0.18	0.18	0.18	0.05
NumResLevels	1	120.60	120.60	120.60	31.33
Approach*NumResLevels	1	17.08	17.08	17.08	4.44
NumDispLevels	1	375.50	375.50	375.50	97.55
Approach*NumDispLevels	1	194.35	194.35	194.35	50.49
NumRDs	1	29.72	29.72	29.72	7.72
Approach*NumRDs	1	22.52	22.52	22.52	5.85

```
NumInflows                           1     17.11    17.11    17.11    4.45
Approach*NumInflows                  1      0.03     0.03     0.03    0.01
TransitionCorrelationType            1    184.15   184.15   184.15   47.84
Approach*TransitionCorrelationType   1    120.27   120.27   120.27   31.25
Error                              618   2378.73  2378.73     3.85
Total                              639  15136.17

Source                                 P
RD Type                            0.000
Approach                           0.000
RD Type*Approach                   0.000
NumPeriods                         0.107
Approach*NumPeriods                0.827
NumResLevels                       0.000
Approach*NumResLevels              0.036
NumDispLevels                      0.000
Approach*NumDispLevels             0.000
NumRDs                             0.006
Approach*NumRDs                    0.016
NumInflows                         0.035
Approach*NumInflows                0.931
TransitionCorrelationType          0.000
Approach*TransitionCorrelationType 0.000
Error
Total


S = 1.96191   R-Sq = 84.28%   R-Sq(adj) = 83.75%


Least Squares Means for Percentage Error

RD Type*Approach
Branch Different R&A Approach                             5.3760   0.2452
Branch Different Value Curve Approach (not branched)      2.0745   0.2452
Naturally Monotone R&A Approach                           1.2484   0.2452
Naturally Monotone Value Curve Approach (not branched)    1.1570   0.2452
Single Different R&A Approach                             1.7344   0.2452
Single Different Value Curve Approach (not branched)      0.9386   0.2452
Single Similar R&A Approach                               1.0772   0.2452
Single Similar Value Curve Approach (not branched)        0.7825   0.2452
Stepped      R&A Approach                                15.4798   0.2452
Stepped      Value Curve Approach (not branched)          3.4535   0.2452

Approach*NumDispLevel
R&A Approach 20                                           6.3002   0.1551
R&A Approach 40                                           3.6661   0.1551
Value Curve Approach (not branched) 20                    1.8961   0.1551
Value Curve Approach (not branched) 40                    1.4663   0.1551

Approach*TransitionCo
R&A Approach Strong                                       4.0133   0.1551
R&A Approach Weak                                         5.9531   0.1551
Value Curve Approach (not branched) Strong                1.5783   0.1551
Value Curve Approach (not branched) Weak                  1.7841   0.1551

Approach*NumResLevels
R&A Approach 100                                          5.5806   0.1551
R&A Approach 300                                          4.3857   0.1551
Value Curve Approach (not branched) 100                   1.9519   0.1551
Value Curve Approach (not branched) 300                   1.4105   0.1551

Bonferroni Simultaneous Tests
```

Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*NumDispLevels
Approach = R&A Approach
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference
R&A Approach	40	-2.634	0.2193
Value Curve Approach (not branched)	20	-4.404	0.2193
Value Curve Approach (not branched)	40	-4.834	0.2193

Approach	NumDispLevels	T-Value	Adjusted P-Value
R&A Approach	40	-12.01	0.0000
Value Curve Approach (not branched)	20	-20.08	0.0000
Value Curve Approach (not branched)	40	-22.04	0.0000

Approach = R&A Approach
NumDispLevels = 40 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference
Value Curve Approach (not branched)	20	-1.770	0.2193
Value Curve Approach (not branched)	40	-2.200	0.2193

Approach	NumDispLevels	T-Value	Adjusted P-Value
Value Curve Approach (not branched)	20	-8.07	0.0000
Value Curve Approach (not branched)	40	-10.03	0.0000

Approach = Value Curve Approach (not branched)
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference
Value Curve Approach (not branched)	40	-0.4298	0.2193

Approach	NumDispLevels	T-Value	Adjusted P-Value
Value Curve Approach (not branched)	40	-1.959	0.3030

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of RD Type*Approach
RD Type = Branch Different
Approach = R&A Approach subtracted from:

RD Type	Approach	Difference of Means	SE of Difference
Branch Different	Value Curve Approach (not branched)	-3.302	0.3468
Naturally Monotone	R&A Approach	-4.128	0.3468
Naturally Monotone	Value Curve Approach (not branched)	-4.219	0.3468
Single Different	R&A Approach	-3.642	0.3468
Single Different	Value Curve Approach (not branched)	-4.437	0.3468
Single Similar	R&A Approach	-4.299	0.3468
Single Similar	Value Curve Approach (not branched)	-4.594	0.3468
Stepped	R&A Approach	10.104	0.3468
Stepped	Value Curve Approach (not branched)	-1.923	0.3468

Adjusted

RD Type	Approach	T-Value	P-Value
Branch Different	Value Curve Approach (not branched)	-9.52	0.0000
Naturally Monotone	R&A Approach	-11.90	0.0000
Naturally Monotone	Value Curve Approach (not branched)	-12.16	0.0000
Single Different	R&A Approach	-10.50	0.0000
Single Different	Value Curve Approach (not branched)	-12.79	0.0000
Single Similar	R&A Approach	-12.40	0.0000
Single Similar	Value Curve Approach (not branched)	-13.24	0.0000
Stepped	R&A Approach	29.13	0.0000
Stepped	Value Curve Approach (not branched)	-5.54	0.0000

RD Type = Branch Different
Approach = Value Curve Approach (not branched) subtracted from:

RD Type	Approach	Difference of Means	SE of Difference
Naturally Monotone	R&A Approach	-0.826	0.3468
Naturally Monotone	Value Curve Approach (not branched)	-0.917	0.3468
Single Different	R&A Approach	-0.340	0.3468
Single Different	Value Curve Approach (not branched)	-1.136	0.3468
Single Similar	R&A Approach	-0.997	0.3468
Single Similar	Value Curve Approach (not branched)	-1.292	0.3468
Stepped	R&A Approach	13.405	0.3468
Stepped	Value Curve Approach (not branched)	1.379	0.3468

RD Type	Approach	T-Value	Adjusted P-Value
Naturally Monotone	R&A Approach	-2.382	0.7888
Naturally Monotone	Value Curve Approach (not branched)	-2.645	0.3766
Single Different	R&A Approach	-0.981	1.0000
Single Different	Value Curve Approach (not branched)	-3.275	0.0502
Single Similar	R&A Approach	-2.876	0.1878
Single Similar	Value Curve Approach (not branched)	-3.725	0.0096
Stepped	R&A Approach	38.652	0.0000
Stepped	Value Curve Approach (not branched)	3.976	0.0035

RD Type = Naturally Monotone
Approach = R&A Approach subtracted from:

RD Type	Approach	Difference of Means	SE of Difference
Naturally Monotone	Value Curve Approach (not branched)	-0.0914	0.3468
Single Different	R&A Approach	0.4860	0.3468
Single Different	Value Curve Approach (not branched)	-0.3098	0.3468
Single Similar	R&A Approach	-0.1712	0.3468
Single Similar	Value Curve Approach (not branched)	-0.4659	0.3468
Stepped	R&A Approach	14.2314	0.3468
Stepped	Value Curve Approach (not branched)	2.2050	0.3468

RD Type	Approach	T-Value	Adjusted P-Value
Naturally Monotone	Value Curve Approach (not branched)	-0.264	1.0000
Single Different	R&A Approach	1.401	1.0000
Single Different	Value Curve Approach (not branched)	-0.893	1.0000
Single Similar	R&A Approach	-0.494	1.0000
Single Similar	Value Curve Approach (not branched)	-1.343	1.0000
Stepped	R&A Approach	41.034	0.0000
Stepped	Value Curve Approach (not branched)	6.358	0.0000

RD Type = Naturally Monotone

Approach = Value Curve Approach (not branched) subtracted from:

RD Type	Approach	Difference of Means	SE of Difference
Single Different	R&A Approach	0.5774	0.3468
Single Different	Value Curve Approach (not branched)	-0.2184	0.3468
Single Similar	R&A Approach	-0.0798	0.3468
Single Similar	Value Curve Approach (not branched)	-0.3745	0.3468
Stepped	R&A Approach	14.3228	0.3468
Stepped	Value Curve Approach (not branched)	2.2964	0.3468

RD Type	Approach	T-Value	Adjusted P-Value
Single Different	R&A Approach	1.665	1.0000
Single Different	Value Curve Approach (not branched)	-0.630	1.0000
Single Similar	R&A Approach	-0.230	1.0000
Single Similar	Value Curve Approach (not branched)	-1.080	1.0000
Stepped	R&A Approach	41.297	0.0000
Stepped	Value Curve Approach (not branched)	6.621	0.0000

RD Type = Single Different
Approach = R&A Approach subtracted from:

RD Type	Approach	Difference of Means	SE of Difference
Single Different	Value Curve Approach (not branched)	-0.7958	0.3468
Single Similar	R&A Approach	-0.6572	0.3468
Single Similar	Value Curve Approach (not branched)	-0.9519	0.3468
Stepped	R&A Approach	13.7454	0.3468
Stepped	Value Curve Approach (not branched)	1.7191	0.3468

RD Type	Approach	T-Value	Adjusted P-Value
Single Different	Value Curve Approach (not branched)	-2.295	0.9941
Single Similar	R&A Approach	-1.895	1.0000
Single Similar	Value Curve Approach (not branched)	-2.745	0.2805
Stepped	R&A Approach	39.633	0.0000
Stepped	Value Curve Approach (not branched)	4.957	0.0000

RD Type = Single Different
Approach = Value Curve Approach (not branched) subtracted from:

RD Type	Approach	Difference of Means	SE of Difference
Single Similar	R&A Approach	0.1386	0.3468
Single Similar	Value Curve Approach (not branched)	-0.1561	0.3468
Stepped	R&A Approach	14.5412	0.3468
Stepped	Value Curve Approach (not branched)	2.5149	0.3468

RD Type	Approach	T-Value	Adjusted P-Value
Single Similar	R&A Approach	0.3997	1.0000
Single Similar	Value Curve Approach (not branched)	-0.4500	1.0000
Stepped	R&A Approach	41.9273	0.0000
Stepped	Value Curve Approach (not branched)	7.2512	0.0000

RD Type = Single Similar
Approach = R&A Approach subtracted from:

Difference SE of

```
RD Type         Approach                                of Means  Difference
Single Similar  Value Curve Approach (not branched)      -0.2947      0.3468
Stepped         R&A Approach                             14.4026      0.3468
Stepped         Value Curve Approach (not branched)       2.3763      0.3468

                                                                Adjusted
RD Type         Approach                             T-Value    P-Value
Single Similar  Value Curve Approach (not branched)  -0.8497     1.0000
Stepped         R&A Approach                         41.5276     0.0000
Stepped         Value Curve Approach (not branched)   6.8516     0.0000


RD Type = Single Similar
Approach = Value Curve Approach (not branched)  subtracted from:

                                              Difference       SE of
RD Type  Approach                               of Means  Difference  T-Value
Stepped  R&A Approach                             14.697      0.3468   42.377
Stepped  Value Curve Approach (not branched)       2.671      0.3468    7.701

                                              Adjusted
RD Type  Approach                              P-Value
Stepped  R&A Approach                           0.0000
Stepped  Value Curve Approach (not branched)    0.0000


RD Type = Stepped
Approach = R&A Approach  subtracted from:

                                              Difference       SE of
RD Type  Approach                               of Means  Difference  T-Value
Stepped  Value Curve Approach (not branched)      -12.03      0.3468   -34.68

                                              Adjusted
RD Type  Approach                              P-Value
Stepped  Value Curve Approach (not branched)    0.0000

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*TransitionCorrelationType
Approach = R&A Approach
TransitionCorrelationType = Strong  subtracted from:

                                                                  Difference
Approach                             TransitionCorrelationType      of Means
R&A Approach                         Weak                              1.940
Value Curve Approach (not branched)  Strong                           -2.435
Value Curve Approach (not branched)  Weak                             -2.229

                                                                       SE of
Approach                             TransitionCorrelationType    Difference
R&A Approach                         Weak                             0.2193
Value Curve Approach (not branched)  Strong                           0.2193
Value Curve Approach (not branched)  Weak                             0.2193


Approach                             TransitionCorrelationType  T-Value
R&A Approach                         Weak                          8.84
Value Curve Approach (not branched)  Strong                      -11.10
Value Curve Approach (not branched)  Weak                        -10.16

                                                                Adjusted
Approach                             TransitionCorrelationType   P-Value
```

R&A Approach	Weak	0.0000
Value Curve Approach (not branched)	Strong	0.0000
Value Curve Approach (not branched)	Weak	0.0000

Approach = R&A Approach
TransitionCorrelationType = Weak subtracted from:

Approach	TransitionCorrelationType	Difference of Means
Value Curve Approach (not branched)	Strong	-4.375
Value Curve Approach (not branched)	Weak	-4.169

Approach	TransitionCorrelationType	SE of Difference
Value Curve Approach (not branched)	Strong	0.2193
Value Curve Approach (not branched)	Weak	0.2193

Approach	TransitionCorrelationType	T-Value
Value Curve Approach (not branched)	Strong	-19.94
Value Curve Approach (not branched)	Weak	-19.01

Approach	TransitionCorrelationType	Adjusted P-Value
Value Curve Approach (not branched)	Strong	0.0000
Value Curve Approach (not branched)	Weak	0.0000

Approach = Value Curve Approach (not branched)
TransitionCorrelationType = Strong subtracted from:

Approach	TransitionCorrelationType	Difference of Means
Value Curve Approach (not branched)	Weak	0.2058

Approach	TransitionCorrelationType	SE of Difference
Value Curve Approach (not branched)	Weak	0.2193

Approach	TransitionCorrelationType	T-Value
Value Curve Approach (not branched)	Weak	0.9383

Approach	TransitionCorrelationType	Adjusted P-Value
Value Curve Approach (not branched)	Weak	1.000

F. PAIRWISE COMPARISON TESTS FOR COMPUTATIONAL TIME ON FULL MODELS

Paired T-Test and CI: Time - MVC, Time - VCN

```
Paired T for Time - MVC - Time - VCN

               N     Mean     StDev  SE Mean
Time - MVC   320  999.438  1296.340   72.468
Time - VCN   320   98.875   119.606    6.686
Difference   320  900.563  1208.984   67.584


95% CI for mean difference: (767.595, 1033.530)
T-Test of mean difference = 0 (vs not = 0): T-Value = 13.33  P-Value = 0.000
```

Paired T-Test and CI: Time - MVC, Time - VC

```
Paired T for Time - MVC - Time - VC

               N     Mean     StDev  SE Mean
Time - MVC   320  999.438  1296.340   72.468
Time - VC    320   74.166    73.562    4.112
Difference   320  925.272  1241.824   69.420


95% CI for mean difference: (788.693, 1061.851)
T-Test of mean difference = 0 (vs not = 0): T-Value = 13.33  P-Value = 0.000
```

Paired T-Test and CI: Time - VCN, Time - VC

```
Paired T for Time - VCN - Time - VC

               N     Mean     StDev  SE Mean
Time - VCN   320  98.8750  119.6056   6.6862
Time - VC    320  74.1656   73.5616   4.1122
Difference   320  24.7094   61.8994   3.4603


95% CI for mean difference: (17.9015, 31.5172)
T-Test of mean difference = 0 (vs not = 0): T-Value = 7.14  P-Value = 0.000
```

Paired T-Test and CI: RD = Stepped - MVC, RD = Single Sim - MVC

```
Paired T for RD = Stepped - MVC - RD = Single Sim - MVC

                   N     Mean    StDev  SE Mean
RD = Stepped - M  64  1570.98  1833.23   229.15
RD = Single Sim   64   971.23  1257.71   157.21
Difference        64  599.750  765.597   95.700


95% CI for mean difference: (408.510, 790.990)
T-Test of mean difference = 0 (vs not = 0): T-Value = 6.27  P-Value = 0.000
```

Paired T-Test and CI: RD = Stepped - MVC, RD = Single Diff - MVC

```
Paired T for RD = Stepped - MVC - RD = Single Diff - MVC

                   N     Mean     StDev  SE Mean
RD = Stepped - M  64  1570.98   1833.23   229.15
RD = Single Diff  64   740.48    965.97   120.75
Difference        64  830.500  1071.154  133.894


95% CI for mean difference: (562.934, 1098.066)
T-Test of mean difference = 0 (vs not = 0): T-Value = 6.20  P-Value = 0.000
```

Paired T-Test and CI: RD = Stepped - MVC, RD = Branch Diff - MVC

```
Paired T for RD = Stepped - MVC - RD = Branch Diff - MVC

                   N     Mean    StDev  SE Mean
RD = Stepped - M  64  1570.98  1833.23   229.15
RD = Branch Diff  64   971.17  1124.36   140.55
Difference        64  599.813  851.073  106.384


95% CI for mean difference: (387.221, 812.404)
T-Test of mean difference = 0 (vs not = 0): T-Value = 5.64  P-Value = 0.000
```

Paired T-Test and CI: RD = Stepped - MVC, RD = Nat Monotone - MVC

```
Paired T for RD = Stepped - MVC - RD = Nat Monotone - MVC

                   N     Mean    StDev  SE Mean
RD = Stepped - M  64  1570.98  1833.23   229.15
```

RD = Nat Monoton	64	743.31	949.26	118.66
Difference	64	827.672	974.979	121.872

95% CI for mean difference: (584.129, 1071.214)
T-Test of mean difference = 0 (vs not = 0): T-Value = 6.79 P-Value = 0.000

Paired T-Test and CI: RD = Single Sim - MVC, RD = Single Diff - MVC

Paired T for RD = Single Sim - MVC - RD = Single Diff - MVC

	N	Mean	StDev	SE Mean
RD = Single Sim	64	971.234	1257.712	157.214
RD = Single Diff	64	740.484	965.971	120.746
Difference	64	230.750	721.365	90.171

95% CI for mean difference: (50.558, 410.942)
T-Test of mean difference = 0 (vs not = 0): T-Value = 2.56 P-Value = 0.013

Paired T-Test and CI: RD = Single Sim - MVC, RD = Branch Diff - MVC

Paired T for RD = Single Sim - MVC - RD = Branch Diff - MVC

	N	Mean	StDev	SE Mean
RD = Single Sim	64	971.234	1257.712	157.214
RD = Branch Diff	64	971.172	1124.360	140.545
Difference	64	0.062500	561.576801	70.197100

95% CI for mean difference: (-140.215211, 140.340211)
T-Test of mean difference = 0 (vs not = 0): T-Value = 0.00 P-Value = 0.999

Paired T-Test and CI: RD = Single Sim - MVC, RD = Nat Monotone - MVC

Paired T for RD = Single Sim - MVC - RD = Nat Monotone - MVC

	N	Mean	StDev	SE Mean
RD = Single Sim	64	971.234	1257.712	157.214
RD = Nat Monoton	64	743.313	949.264	118.658
Difference	64	227.922	671.589	83.949

95% CI for mean difference: (60.164, 395.680)
T-Test of mean difference = 0 (vs not = 0): T-Value = 2.72 P-Value = 0.009

Paired T-Test and CI: RD = Branch Diff - MVC, RD = Single Diff - MVC

Paired T for RD = Branch Diff - MVC - RD = Single Diff - MVC

	N	Mean	StDev	SE Mean
RD = Branch Diff	64	971.172	1124.360	140.545
RD = Single Diff	64	740.484	965.971	120.746
Difference	64	230.688	531.927	66.491

95% CI for mean difference: (97.816, 363.559)

T-Test of mean difference = 0 (vs not = 0): T-Value = 3.47 P-Value = 0.001

Paired T-Test and CI: RD = Branch Diff - MVC, RD = Nat Monotone - MVC

Paired T for RD = Branch Diff - MVC - RD = Nat Monotone - MVC

	N	Mean	StDev	SE Mean
RD = Branch Diff	64	971.172	1124.360	140.545
RD = Nat Monoton	64	743.313	949.264	118.658
Difference	64	227.859	418.298	52.287

95% CI for mean difference: (123.372, 332.347)
T-Test of mean difference = 0 (vs not = 0): T-Value = 4.36 P-Value = 0.000

Paired T-Test and CI: RD = Nat Monotone - MVC, RD = Single Diff - MVC

Paired T for RD = Nat Monotone - MVC - RD = Single Diff - MVC

	N	Mean	StDev	SE Mean
RD = Nat Monoton	64	743.313	949.264	118.658
RD = Single Diff	64	740.484	965.971	120.746
Difference	64	2.82812	572.63644	71.57956

95% CI for mean difference: (-140.21220, 145.86845)
T-Test of mean difference = 0 (vs not = 0): T-Value = 0.04 P-Value = 0.969

Paired T-Test and CI: NumPeriods = 40 MVC, NumPeriods = 20 MVC

Paired T for NumPeriods = 40 MVC - NumPeriods = 20 MVC

	N	Mean	StDev	SE Mean
NumPeriods = 40	160	1335.48	1586.48	125.42
NumPeriods = 20	160	663.40	792.08	62.62
Difference	160	672.075	919.727	72.711

95% CI for mean difference: (528.471, 815.679)
T-Test of mean difference = 0 (vs not = 0): T-Value = 9.24 P-Value = 0.000

Paired T-Test and CI: NumResLevels = 300 - MVC, NumResLevels = 100 - MVC

Paired T for NumResLevels = 300 - MVC - NumResLevels = 100 - MVC

	N	Mean	StDev	SE Mean
NumResLevels = 3	160	1550.24	1594.14	126.03
NumResLevels = 1	160	448.63	468.72	37.06
Difference	160	1101.61	1189.44	94.03

95% CI for mean difference: (915.90, 1287.33)
T-Test of mean difference = 0 (vs not = 0): T-Value = 11.72 P-Value = 0.000

Paired T-Test and CI: NumDispLevels = 40 - MVC, NumDispLevels = 20 - MVC

```
Paired T for NumDispLevels = 40 - MVC - NumDispLevels = 20 - MVC

                    N      Mean    StDev  SE Mean
NumDispLevels =   160    867.41  1122.09    88.71
NumDispLevels =   160   1131.47  1441.32   113.95
Difference        160  -264.063  591.733   46.781


95% CI for mean difference: (-356.454, -171.671)
T-Test of mean difference = 0 (vs not = 0): T-Value = -5.64  P-Value = 0.000
```

Paired T-Test and CI: NumRDs = 100 - MVC, NumRDs = 25 - MVC

```
Paired T for NumRDs = 100 - MVC - NumRDs = 25 - MVC

                    N     Mean    StDev  SE Mean
NumRDs = 100 - M  160  1764.07  1465.21   115.84
NumRDs = 25 - MV  160   234.81   219.18    17.33
Difference        160  1529.26  1291.99   102.14


95% CI for mean difference: (1327.54, 1730.99)
T-Test of mean difference = 0 (vs not = 0): T-Value = 14.97  P-Value = 0.000
```

Paired T-Test and CI: NumPeriods = 40 VC, NumPeriods = 20 VC

```
Paired T for NumPeriods = 40 VC - NumPeriods = 20 VC

                    N     Mean    StDev  SE Mean
NumPeriods = 40   160  98.8062  88.8617   7.0251
NumPeriods = 20   160  49.5250  41.6926   3.2961
Difference        160  49.2813  50.2805   3.9750


95% CI for mean difference: (41.4306, 57.1319)
T-Test of mean difference = 0 (vs not = 0): T-Value = 12.40  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 - VC, NumResLevels = 100 - VC

```
Paired T for NumResLevels = 300 - VC - NumResLevels = 100 - VC

                    N     Mean    StDev  SE Mean
NumResLevels = 3  160  105.938   87.075    6.884
NumResLevels = 1  160   42.394   35.257    2.787
Difference        160  63.5438  53.9706   4.2668


95% CI for mean difference: (55.1169, 71.9706)
T-Test of mean difference = 0 (vs not = 0): T-Value = 14.89  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 40 - VC, NumDispLevels = 20 - VC

```
Paired T for NumDispLevels = 40 - VC - NumDispLevels = 20 - VC

                    N     Mean    StDev  SE Mean
NumDispLevels =   160  93.9625  87.7461   6.9369
NumDispLevels =   160  54.3688  48.6672   3.8475
```

```
Difference        160  39.5938  42.1814   3.3347


95% CI for mean difference: (33.0077, 46.1798)
T-Test of mean difference = 0 (vs not = 0): T-Value = 11.87  P-Value = 0.000
```

Paired T-Test and CI: NumRDs = 100 - VC, NumRDs = 25 - VC

```
Paired T for NumRDs = 100 - VC - NumRDs = 25 - VC

                    N     Mean    StDev  SE Mean
NumRDs = 100 - V  160  119.706   79.740    6.304
NumRDs = 25 - VC  160   28.625   18.004    1.423
Difference        160  91.0812  62.9260   4.9747


95% CI for mean difference: (81.2562, 100.9063)
T-Test of mean difference = 0 (vs not = 0): T-Value = 18.31  P-Value = 0.000
```

G. MINITAB OUTPUT FOR SOLUTION QUALITY FOR FULL MODELS

General Linear Model: Percentage Error versus Approach, NumPeriods, ...

```
Factor                      Type   Levels  Values
Approach                    fixed       3  Direct MV Curve Approach, Value Curve
                                           Approach (branched), Value Curve
                                           Approach (not branched)
NumPeriods                  fixed       2  20, 40
NumResLevels                fixed       2  100, 300
NumDispLevels               fixed       2  20, 40
NumRDs                      fixed       2  25, 100
NumInflows                  fixed       2  1, 5
TransitionCorrelationType   fixed       2  Strong, Weak
RD Type                     fixed       5  Branch Different, Naturally Monotone,
                                           Single Different, Single Similar,
                                           Stepped


Analysis of Variance for Percentage Error, using Adjusted SS for Tests

Source                              DF    Seq SS   Adj SS   Adj MS       F
Approach                             2    24.870   24.870   12.435   19.86
NumPeriods                           1     1.948    1.948    1.948    3.11
Approach*NumPeriods                  2     1.918    1.918    0.959    1.53
NumResLevels                         1   105.947  105.947  105.947  169.19
Approach*NumResLevels                2     2.612    2.612    1.306    2.09
NumDispLevels                        1     6.922    6.922    6.922   11.05
Approach*NumDispLevels               2     8.325    8.325    4.162    6.65
NumRDs                               1     1.400    1.400    1.400    2.24
Approach*NumRDs                      2     0.063    0.063    0.031    0.05
NumInflows                           1    18.270   18.270   18.270   29.18
```

Approach*NumInflows	2	0.781	0.781	0.391	0.62
TransitionCorrelationType	1	4.681	4.681	4.681	7.48
Approach*TransitionCorrelationType	2	0.667	0.667	0.334	0.53
RD Type	4	601.130	601.130	150.282	239.99
Approach*RD Type	8	23.400	23.400	2.925	4.67
Error	927	580.484	580.484	0.626	
Total	959	1383.417			

Source	P
Approach	0.000
NumPeriods	0.078
Approach*NumPeriods	0.217
NumResLevels	0.000
Approach*NumResLevels	0.125
NumDispLevels	0.001
Approach*NumDispLevels	0.001
NumRDs	0.135
Approach*NumRDs	0.951
NumInflows	0.000
Approach*NumInflows	0.536
TransitionCorrelationType	0.006
Approach*TransitionCorrelationType	0.587
RD Type	0.000
Approach*RD Type	0.000
Error	
Total	

S = 0.791325   R-Sq = 58.04%   R-Sq(adj) = 56.59%

Least Squares Means for Percentage Error

Approach	Mean	SE Mean
Direct MV Curve Approach	1.4037	0.04424
Value Curve Approach (branched)	1.2999	0.04424
Value Curve Approach (not branched)	1.6812	0.04424
NumPeriods		
20	1.5066	0.03612
40	1.4165	0.03612
Approach*NumPeriods		
Direct MV Curve Approach 20	1.4175	0.06256
Direct MV Curve Approach 40	1.3898	0.06256
Value Curve Approach (branched) 20	1.3130	0.06256
Value Curve Approach (branched) 40	1.2869	0.06256
Value Curve Approach (not branched) 20	1.7895	0.06256
Value Curve Approach (not branched) 40	1.5730	0.06256
NumResLevels		
100	1.7938	0.03612
300	1.1294	0.03612
Approach*NumResLevels		
Direct MV Curve Approach 100	1.8019	0.06256
Direct MV Curve Approach 300	1.0054	0.06256
Value Curve Approach (branched) 100	1.6276	0.06256
Value Curve Approach (branched) 300	0.9723	0.06256
Value Curve Approach (not branched) 100	1.9519	0.06256
Value Curve Approach (not branched) 300	1.4105	0.06256

NumDispLevel

```
20                                                     1.5465  0.03612
40                                                     1.3767  0.03612

Approach*NumDispLevel
Direct MV Curve Approach 20                            1.4053  0.06256
Direct MV Curve Approach 40                            1.4020  0.06256
Value Curve Approach (branched) 20                     1.3382  0.06256
Value Curve Approach (branched) 40                     1.2617  0.06256
Value Curve Approach (not branched) 20                 1.8961  0.06256
Value Curve Approach (not branched) 40                 1.4663  0.06256

NumRDs
 25                                                    1.4234  0.03612
100                                                    1.4998  0.03612

Approach*NumRDs
Direct MV Curve Approach  25                           1.3647  0.06256
Direct MV Curve Approach 100                           1.4426  0.06256
Value Curve Approach (branched)  25                    1.2523  0.06256
Value Curve Approach (branched) 100                    1.3476  0.06256
Value Curve Approach (not branched)  25                1.6533  0.06256
Value Curve Approach (not branched) 100                1.7091  0.06256

NumInflows
1                                                      1.3236  0.03612
5                                                      1.5996  0.03612

Approach*NumInflows
Direct MV Curve Approach 1                             1.3060  0.06256
Direct MV Curve Approach 5                             1.5013  0.06256
Value Curve Approach (branched) 1                      1.1405  0.06256
Value Curve Approach (branched) 5                      1.4594  0.06256
Value Curve Approach (not branched) 1                  1.5244  0.06256
Value Curve Approach (not branched) 5                  1.8380  0.06256

TransitionCo
Strong                                                 1.3918  0.03612
Weak                                                   1.5314  0.03612

Approach*TransitionCo
Direct MV Curve Approach Strong                        1.3653  0.06256
Direct MV Curve Approach Weak                          1.4420  0.06256
Value Curve Approach (branched) Strong                 1.2317  0.06256
Value Curve Approach (branched) Weak                   1.3681  0.06256
Value Curve Approach (not branched) Strong             1.5783  0.06256
Value Curve Approach (not branched) Weak               1.7841  0.06256

RD Type
Branch Different                                       1.7042  0.05711
Naturally Monotone                                     1.1851  0.05711
Single Different                                       0.8643  0.05711
Single Similar                                         0.6727  0.05711
Stepped                                                2.8816  0.05711

Approach*RD Type
Direct MV Curve Approach Branch Different              1.5689  0.09892
Direct MV Curve Approach Naturally Monotone            1.2449  0.09892
Direct MV Curve Approach Single Different              0.9121  0.09892
Direct MV Curve Approach Single Similar                0.6710  0.09892
Direct MV Curve Approach Stepped                       2.6212  0.09892
Value Curve Approach (branched) Branch Different       1.4692  0.09892
Value Curve Approach (branched) Naturally Monotone     1.1535  0.09892
Value Curve Approach (branched) Single Different       0.7422  0.09892
```

Value Curve Approach (branched)	Single Similar	0.5646	0.09892
Value Curve Approach (branched)	Stepped	2.5702	0.09892
Value Curve Approach (not branched)	Branch Different	2.0745	0.09892
Value Curve Approach (not branched)	Naturally Monotone	1.1570	0.09892
Value Curve Approach (not branched)	Single Different	0.9386	0.09892
Value Curve Approach (not branched)	Single Similar	0.7825	0.09892
Value Curve Approach (not branched)	Stepped	3.4535	0.09892

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*RD Type
Approach = Direct MV Curve Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Direct MV Curve Approach	Naturally Monotone	-0.324	0.1399
Direct MV Curve Approach	Single Different	-0.657	0.1399
Direct MV Curve Approach	Single Similar	-0.898	0.1399
Direct MV Curve Approach	Stepped	1.052	0.1399
Value Curve Approach (branched)	Branch Different	-0.100	0.1399
Value Curve Approach (not branched)	Branch Different	0.506	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Direct MV Curve Approach	Naturally Monotone	-2.316	1.0000
Direct MV Curve Approach	Single Different	-4.695	0.0003
Direct MV Curve Approach	Single Similar	-6.419	0.0000
Direct MV Curve Approach	Stepped	7.522	0.0000
Value Curve Approach (branched)	Branch Different	-0.713	1.0000
Value Curve Approach (not branched)	Branch Different	3.614	0.0334

Approach = Direct MV Curve Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Direct MV Curve Approach	Single Different	-0.3328	0.1399
Direct MV Curve Approach	Single Similar	-0.5739	0.1399
Direct MV Curve Approach	Stepped	1.3763	0.1399
Value Curve Approach (branched)	Naturally Monotone	-0.0915	0.1399
Value Curve Approach (not branched)	Naturally Monotone	-0.0879	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Direct MV Curve Approach	Single Different	-2.379	1.0000
Direct MV Curve Approach	Single Similar	-4.103	0.0047
Direct MV Curve Approach	Stepped	9.839	0.0000
Value Curve Approach (branched)	Naturally Monotone	-0.654	1.0000
Value Curve Approach (not branched)	Naturally Monotone	-0.628	1.0000

Approach = Direct MV Curve Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Direct MV Curve Approach	Single Similar	-0.2411	0.1399
Direct MV Curve Approach	Stepped	1.7091	0.1399
Value Curve Approach (branched)	Single Different	-0.1700	0.1399
Value Curve Approach (not branched)	Single Different	0.0264	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Direct MV Curve Approach	Single Similar	-1.724	1.0000
Direct MV Curve Approach	Stepped	12.217	0.0000
Value Curve Approach (branched)	Single Different	-1.215	1.0000
Value Curve Approach (not branched)	Single Different	0.189	1.0000

Approach = Direct MV Curve Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Direct MV Curve Approach	Stepped	1.9502	0.1399
Value Curve Approach (branched)	Single Similar	-0.1064	0.1399
Value Curve Approach (not branched)	Single Similar	0.1115	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Direct MV Curve Approach	Stepped	13.9411	0.0000
Value Curve Approach (branched)	Single Similar	-0.7607	1.0000
Value Curve Approach (not branched)	Single Similar	0.7968	1.0000

Approach = Direct MV Curve Approach
RD Type = Stepped subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (branched)	Stepped	-0.051	0.1399
Value Curve Approach (not branched)	Stepped	0.832	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (branched)	Stepped	-0.36	1.0000
Value Curve Approach (not branched)	Stepped	5.95	0.0000

Approach = Value Curve Approach (branched)
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (branched)	Naturally Monotone	-0.3158	0.1399
Value Curve Approach (branched)	Single Different	-0.7270	0.1399
Value Curve Approach (branched)	Single Similar	-0.9046	0.1399
Value Curve Approach (branched)	Stepped	1.1010	0.1399
Value Curve Approach (not branched)	Branch Different	0.6052	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (branched)	Naturally Monotone	-2.257	1.0000
Value Curve Approach (branched)	Single Different	-5.197	0.0000
Value Curve Approach (branched)	Single Similar	-6.467	0.0000
Value Curve Approach (branched)	Stepped	7.871	0.0000
Value Curve Approach (not branched)	Branch Different	4.327	0.0018

Approach = Value Curve Approach (branched)
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (branched)	Single Different	-0.4113	0.1399

Value Curve Approach (branched)	Single Similar	-0.5888	0.1399
Value Curve Approach (branched)	Stepped	1.4168	0.1399
Value Curve Approach (not branched)	Naturally Monotone	0.0036	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (branched)	Single Different	-2.940	0.3531
Value Curve Approach (branched)	Single Similar	-4.209	0.0030
Value Curve Approach (branched)	Stepped	10.128	0.0000
Value Curve Approach (not branched)	Naturally Monotone	0.025	1.0000

Approach = Value Curve Approach (branched)
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (branched)	Single Similar	-0.1775	0.1399
Value Curve Approach (branched)	Stepped	1.8280	0.1399
Value Curve Approach (not branched)	Single Different	0.1964	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (branched)	Single Similar	-1.269	1.0000
Value Curve Approach (branched)	Stepped	13.068	0.0000
Value Curve Approach (not branched)	Single Different	1.404	1.0000

Approach = Value Curve Approach (branched)
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (branched)	Stepped	2.0056	0.1399
Value Curve Approach (not branched)	Single Similar	0.2179	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (branched)	Stepped	14.337	0.0000
Value Curve Approach (not branched)	Single Similar	1.557	1.0000

Approach = Value Curve Approach (branched)
RD Type = Stepped subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (not branched)	Stepped	0.883	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (not branched)	Stepped	6.31	0.0000

Approach = Value Curve Approach (not branched)
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (not branched)	Naturally Monotone	-0.917	0.1399
Value Curve Approach (not branched)	Single Different	-1.136	0.1399
Value Curve Approach (not branched)	Single Similar	-1.292	0.1399
Value Curve Approach (not branched)	Stepped	1.379	0.1399

Adjusted

Approach	RD Type	T-Value	P-Value
Value Curve Approach (not branched)	Naturally Monotone	-6.558	0.0000
Value Curve Approach (not branched)	Single Different	-8.120	0.0000
Value Curve Approach (not branched)	Single Similar	-9.236	0.0000
Value Curve Approach (not branched)	Stepped	9.858	0.0000

Approach = Value Curve Approach (not branched)
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (not branched)	Single Different	-0.2184	0.1399
Value Curve Approach (not branched)	Single Similar	-0.3745	0.1399
Value Curve Approach (not branched)	Stepped	2.2964	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (not branched)	Single Different	-1.562	1.0000
Value Curve Approach (not branched)	Single Similar	-2.677	0.7932
Value Curve Approach (not branched)	Stepped	16.416	0.0000

Approach = Value Curve Approach (not branched)
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach (not branched)	Single Similar	-0.1561	0.1399
Value Curve Approach (not branched)	Stepped	2.5149	0.1399

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach (not branched)	Single Similar	-1.116	1.0000
Value Curve Approach (not branched)	Stepped	17.978	0.0000

Approach = Value Curve Approach (not branched)
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value
Value Curve Approach (not branched)	Stepped	2.671	0.1399	19.09

Approach	RD Type	Adjusted P-Value
Value Curve Approach (not branched)	Stepped	0.0000

H. Pairwise Comparison Tests for Computational Time on Simplified Models

Paired T-Test and CI: Time - VC, Time - Quantity-Based Offers

```
Paired T for Time - VC - Time - Quantity-Based Offers

                    N       Mean      StDev   SE Mean
Time - VC         1600    52.2431    47.5270    1.1882
Time - Quantity-  1600    52.2950    51.8753    1.2969
Difference        1600  -0.051875  16.797523  0.419938


95% CI for mean difference: (-0.875562, 0.771812)
T-Test of mean difference = 0 (vs not = 0): T-Value = -0.12  P-Value = 0.902
```

Paired T-Test and CI: Time - VC, Time - Price-Based Offers

```
Paired T for Time - VC - Time - Price-Based Offers

                    N      Mean    StDev  SE Mean
Time - VC         1600    52.243   47.527    1.188
Time - Price-Bas  1600   409.603  573.801   14.345
Difference        1600  -357.359  534.177   13.354


95% CI for mean difference: (-383.553, -331.165)
T-Test of mean difference = 0 (vs not = 0): T-Value = -26.76  P-Value = 0.000
```

Paired T-Test and CI: Time - VC, Time - Naïve Correlation Approa

```
Paired T for Time - VC - Time - Naïve Correlation Approa

                     N      Mean    StDev  SE Mean
Time - VC         1600  52.2431  47.5270   1.1882
Time - Naïve Cor  1600   1.4319   1.1011   0.0275
Difference        1600  50.8113  46.8511   1.1713


95% CI for mean difference: (48.5138, 53.1087)
T-Test of mean difference = 0 (vs not = 0): T-Value = 43.38  P-Value = 0.000
```

Paired T-Test and CI: Time - VC, Time - Perfectly Competitive Ap

```
Paired T for Time - VC - Time - Perfectly Competitive Ap

                     N      Mean     StDev  SE Mean
Time - VC         1600  52.2431   47.5270   1.1882
Time - Perfectly  1600  50.2250   58.0830   1.4521
Difference        1600  2.01812  32.80256  0.82006


95% CI for mean difference: (0.40961, 3.62664)
T-Test of mean difference = 0 (vs not = 0): T-Value = 2.46  P-Value = 0.014
```

Paired T-Test and CI: NumPeriods = 40 QB, NumPeriods = 20 QB

```
Paired T for NumPeriods = 40 QB - NumPeriods = 20 QB

                    N     Mean    StDev  SE Mean
NumPeriods = 40   160  101.806   99.369    7.856
NumPeriods = 20   160   51.844   45.199    3.573
Difference        160  49.9625  70.5057   5.5740


95% CI for mean difference: (38.9539, 60.9711)
T-Test of mean difference = 0 (vs not = 0): T-Value = 8.96  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 - QB, NumResLevels = 100 - QB

```
Paired T for NumResLevels = 300 - QB - NumResLevels = 100 - QB

                    N     Mean    StDev  SE Mean
NumResLevels = 3  160  109.038   95.240    7.529
NumResLevels = 1  160   44.613   44.877    3.548
Difference        160  64.4250  68.0197   5.3774


95% CI for mean difference: (53.8046, 75.0454)
T-Test of mean difference = 0 (vs not = 0): T-Value = 11.98  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 40 - QB, NumDispLevels = 20 - QB

```
Paired T for NumDispLevels = 40 - QB - NumDispLevels = 20 - QB

                    N     Mean    StDev  SE Mean
NumDispLevels =   160  96.3063  97.1735   7.6822
```

```
NumDispLevels =   160  57.3438  54.4660   4.3059
Difference        160  38.9625  63.8207   5.0455


95% CI for mean difference: (28.9977, 48.9273)
T-Test of mean difference = 0 (vs not = 0): T-Value = 7.72  P-Value = 0.000
```

Paired T-Test and CI: NumRDs = 100 - QB, NumRDs = 25 - QB

```
Paired T for NumRDs = 100 - QB - NumRDs = 25 - QB

                    N     Mean    StDev  SE Mean
NumRDs = 100 - Q  160  124.125   90.662    7.167
NumRDs = 25 - QB  160   29.525   21.232    1.679
Difference        160  94.6000  75.8213   5.9942


95% CI for mean difference: (82.7615, 106.4385)
T-Test of mean difference = 0 (vs not = 0): T-Value = 15.78  P-Value = 0.000
```

Paired T-Test and CI: NumPeriods = 40 PB, NumPeriods = 20 PB

```
Paired T for NumPeriods = 40 PB - NumPeriods = 20 PB

                    N     Mean    StDev  SE Mean
NumPeriods = 40   160  1049.62  1277.09   100.96
NumPeriods = 20   160   610.56   728.26    57.57
Difference        160  439.063  952.969   75.339


95% CI for mean difference: (290.269, 587.856)
T-Test of mean difference = 0 (vs not = 0): T-Value = 5.83  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 - PB, NumResLevels = 100 - PB

```
Paired T for NumResLevels = 300 - PB - NumResLevels = 100 - PB

                    N     Mean     StDev  SE Mean
NumResLevels = 3  160  1275.66   1313.15   103.81
NumResLevels = 1  160   384.52    366.59    28.98
Difference        160  891.137  1058.636   83.693


95% CI for mean difference: (725.845, 1056.430)
T-Test of mean difference = 0 (vs not = 0): T-Value = 10.65  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 40 - PB, NumDispLevels = 20 - PB

```
Paired T for NumDispLevels = 40 - PB - NumDispLevels = 20 - PB

                    N     Mean     StDev  SE Mean
NumDispLevels =   160  937.625  1211.408   95.770
NumDispLevels =   160  722.550   876.086   69.261
Difference        160  215.075   855.950   67.669
```

```
95% CI for mean difference: (81.429, 348.721)
T-Test of mean difference = 0 (vs not = 0): T-Value = 3.18  P-Value = 0.002
```

Paired T-Test and CI: NumRDs = 100 - PB, NumRDs = 25 - PB

```
Paired T for NumRDs = 100 - PB - NumRDs = 25 - PB

                    N     Mean    StDev  SE Mean
NumRDs = 100 - P  160  1459.68  1196.73    94.61
NumRDs = 25 - PB  160   200.49   168.47    13.32
Difference        160  1259.19  1100.95    87.04


95% CI for mean difference: (1087.29, 1431.09)
T-Test of mean difference = 0 (vs not = 0): T-Value = 14.47  P-Value = 0.000
```

Paired T-Test and CI: NumPeriods = 40 NC, NumPeriods = 20 NC

```
Paired T for NumPeriods = 40 NC - NumPeriods = 20 NC

                    N      Mean     StDev   SE Mean
NumPeriods = 40   160   1.94375   1.33764   0.10575
NumPeriods = 20   160   1.00625   0.82031   0.06485
Difference        160  0.937500  1.257368  0.099404


95% CI for mean difference: (0.741178, 1.133822)
T-Test of mean difference = 0 (vs not = 0): T-Value = 9.43  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 - NC, NumResLevels = 100 - NC

```
Paired T for NumResLevels = 300 - NC - NumResLevels = 100 - NC

                    N     Mean    StDev  SE Mean
NumResLevels = 3  160  2.08125  1.23356  0.09752
NumResLevels = 1  160  0.86875  0.80190  0.06340
Difference        160  1.21250  1.16223  0.09188


95% CI for mean difference: (1.03103, 1.39397)
T-Test of mean difference = 0 (vs not = 0): T-Value = 13.20  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 40 - NC, NumDispLevels = 20 - NC

```
Paired T for NumDispLevels = 40 - NC - NumDispLevels = 20 - NC

                    N      Mean     StDev   SE Mean
NumDispLevels =   160   1.79375   1.42351   0.11254
NumDispLevels =   160   1.15625   0.82070   0.06488
Difference        160  0.637500  1.266340  0.100113


95% CI for mean difference: (0.439777, 0.835223)
T-Test of mean difference = 0 (vs not = 0): T-Value = 6.37  P-Value = 0.000
```

Paired T-Test and CI: NumRDs = 100 - NC, NumRDs = 25 - NC

```
Paired T for NumRDs = 100 - NC - NumRDs = 25 - NC

                    N      Mean     StDev   SE Mean
NumRDs = 100 - N  160   1.68750   1.29919   0.10271
NumRDs = 25 - NC  160   1.26250   1.06096   0.08388
Difference        160  0.425000  1.168507  0.092379


95% CI for mean difference: (0.242553, 0.607447)
T-Test of mean difference = 0 (vs not = 0): T-Value = 4.60  P-Value = 0.000
```

Paired T-Test and CI: NumPeriods = 40 PC, NumPeriods = 20 PC

```
Paired T for NumPeriods = 40 PC - NumPeriods = 20 PC

                    N     Mean     StDev  SE Mean
NumPeriods = 40   160  106.213   123.291    9.747
NumPeriods = 20   160   50.831    50.547    3.996
Difference        160  55.3813  103.6858   8.1971


95% CI for mean difference: (39.1920, 71.5705)
T-Test of mean difference = 0 (vs not = 0): T-Value = 6.76  P-Value = 0.000
```

Paired T-Test and CI: NumResLevels = 300 - PC, NumResLevels = 100 - PC

```
Paired T for NumResLevels = 300 - PC - NumResLevels = 100 - PC

                    N     Mean     StDev  SE Mean
NumResLevels = 3  160  112.456   119.902    9.479
NumResLevels = 1  160   44.588    51.037    4.035
Difference        160  67.8688  103.8586   8.2107


95% CI for mean difference: (51.6526, 84.0849)
T-Test of mean difference = 0 (vs not = 0): T-Value = 8.27  P-Value = 0.000
```

Paired T-Test and CI: NumDispLevels = 40 - PC, NumDispLevels = 20 - PC

```
Paired T for NumDispLevels = 40 - PC - NumDispLevels = 20 - PC

                    N     Mean     StDev  SE Mean
NumDispLevels =   160  92.0875   91.8964   7.2650
NumDispLevels =   160  64.9563  102.3890   8.0946
Difference        160  27.1313   92.6026   7.3209


95% CI for mean difference: (12.6725, 41.5900)
T-Test of mean difference = 0 (vs not = 0): T-Value = 3.71  P-Value = 0.000
```

Paired T-Test and CI: NumRDs = 100 - PC, NumRDs = 25 - PC

```
Paired T for NumRDs = 100 - PC - NumRDs = 25 - PC

                    N     Mean    StDev  SE Mean
```

```
NumRDs = 100 - P  160  129.906  116.451    9.206
NumRDs = 25 - PC  160   27.138   20.594    1.628
Difference        160  102.769  107.396    8.490


95% CI for mean difference: (86.000, 119.537)
T-Test of mean difference = 0 (vs not = 0): T-Value = 12.10  P-Value = 0.000
```

I. MINITAB OUTPUT ON SOLUTION QUALITY FOR SIMPLIFIED MODELS

General Linear Model: Percentage Error versus Approach, NumPeriods, ...

Factor	Type	Levels	Values
Approach	fixed	6	Naive Correlation Approach, Naive Water Approach, Perfectly Competitive Approach, Price-Based Offer Approach, Quantity-Based Offer Approach, Value Curve Approach
NumPeriods	fixed	2	20, 40
NumResLevels	fixed	2	100, 300
NumDispLevels	fixed	2	20, 40
NumRDs	fixed	2	25, 100
NumInflows	fixed	2	1, 5
TransitionCorrelationType	fixed	2	Strong, Weak
RD Type	fixed	5	Branch Different, Naturally Monotone, Single Different, Single Similar, Stepped

Analysis of Variance for Percentage Error, using Adjusted SS for Tests

Source	DF	Seq SS	Adj SS	Adj MS	F
Approach	5	1429626	1429626	285925	11486.08
NumPeriods	1	1979	1979	1979	79.51
Approach*NumPeriods	5	3539	3539	708	28.43
NumResLevels	1	7724	7724	7724	310.28
Approach*NumResLevels	5	6836	6836	1367	54.92
NumDispLevels	1	3293	3293	3293	132.30
Approach*NumDispLevels	5	33092	33092	6618	265.87
NumRDs	1	1	1	1	0.03

Approach*NumRDs	5	48	48	10	0.39
NumInflows	1	899	899	899	36.10
Approach*NumInflows	5	61	61	12	0.49
TransitionCorrelationType	1	259	259	259	10.42
Approach*TransitionCorrelationType	5	3325	3325	665	26.72
RD Type	4	24861	24861	6215	249.68
Approach*RD Type	20	543371	543371	27169	1091.40
Error	9534	237332	237332	25	
Total	9599	2296245			

Source	P
Approach	0.000
NumPeriods	0.000
Approach*NumPeriods	0.000
NumResLevels	0.000
Approach*NumResLevels	0.000
NumDispLevels	0.000
Approach*NumDispLevels	0.000
NumRDs	0.860
Approach*NumRDs	0.858
NumInflows	0.000
Approach*NumInflows	0.787
TransitionCorrelationType	0.001
Approach*TransitionCorrelationType	0.000
RD Type	0.000
Approach*RD Type	0.000
Error	
Total	

S = 4.98931   R-Sq = 89.66%   R-Sq(adj) = 89.59%

Least Squares Means for Percentage Error

Approach	Mean	SE Mean
Naive Correlation Approach	3.0799	0.12473
Naive Water Approach	36.0085	0.12473
Perfectly Competitive Approach	9.3935	0.12473
Price-Based Offer Approach	2.3040	0.12473
Quantity-Based Offer Approach	4.1324	0.12473
Value Curve Approach	1.1858	0.12473
NumPeriods		
20	8.8966	0.07201
40	9.8048	0.07201
Approach*NumPeriods		
Naive Correlation Approach 20	2.6400	0.17640
Naive Correlation Approach 40	3.5199	0.17640
Naive Water Approach 20	36.1739	0.17640
Naive Water Approach 40	35.8431	0.17640
Perfectly Competitive Approach 20	7.7169	0.17640
Perfectly Competitive Approach 40	11.0701	0.17640
Price-Based Offer Approach 20	2.2456	0.17640
Price-Based Offer Approach 40	2.3623	0.17640
Quantity-Based Offer Approach 20	3.4939	0.17640
Quantity-Based Offer Approach 40	4.7710	0.17640
Value Curve Approach 20	1.1094	0.17640
Value Curve Approach 40	1.2621	0.17640
NumResLevels		
100	10.2477	0.07201

300	8.4537	0.07201
Approach*NumResLevels		
Naive Correlation Approach 100	4.0319	0.17640
Naive Correlation Approach 300	2.1279	0.17640
Naive Water Approach 100	35.6363	0.17640
Naive Water Approach 300	36.3807	0.17640
Perfectly Competitive Approach 100	11.8138	0.17640
Perfectly Competitive Approach 300	6.9732	0.17640
Price-Based Offer Approach 100	3.0989	0.17640
Price-Based Offer Approach 300	1.5091	0.17640
Quantity-Based Offer Approach 100	5.3171	0.17640
Quantity-Based Offer Approach 300	2.9478	0.17640
Value Curve Approach 100	1.5880	0.17640
Value Curve Approach 300	0.7835	0.17640
NumDispLevel		
20	8.7650	0.07201
40	9.9364	0.07201
Approach*NumDispLevel		
Naive Correlation Approach 20	3.2244	0.17640
Naive Correlation Approach 40	2.9355	0.17640
Naive Water Approach 20	31.3793	0.17640
Naive Water Approach 40	40.6378	0.17640
Perfectly Competitive Approach 20	10.4836	0.17640
Perfectly Competitive Approach 40	8.3034	0.17640
Price-Based Offer Approach 20	2.0530	0.17640
Price-Based Offer Approach 40	2.5549	0.17640
Quantity-Based Offer Approach 20	4.3205	0.17640
Quantity-Based Offer Approach 40	3.9444	0.17640
Value Curve Approach 20	1.1291	0.17640
Value Curve Approach 40	1.2424	0.17640
NumInflows		
1	9.0447	0.07201
5	9.6567	0.07201
TransitionCo		
Strong	9.1864	0.07201
Weak	9.5150	0.07201
Approach*TransitionCo		
Naive Correlation Approach Strong	3.2499	0.17640
Naive Correlation Approach Weak	2.9099	0.17640
Naive Water Approach Strong	35.5163	0.17640
Naive Water Approach Weak	36.5008	0.17640
Perfectly Competitive Approach Strong	9.8821	0.17640
Perfectly Competitive Approach Weak	8.9049	0.17640
Price-Based Offer Approach Strong	2.3938	0.17640
Price-Based Offer Approach Weak	2.2141	0.17640
Quantity-Based Offer Approach Strong	2.8218	0.17640
Quantity-Based Offer Approach Weak	5.4430	0.17640
Value Curve Approach Strong	1.2542	0.17640
Value Curve Approach Weak	1.1174	0.17640
RD Type		
Branch Different	11.4421	0.11386
Naturally Monotone	9.8792	0.11386
Single Different	9.4473	0.11386
Single Similar	9.5084	0.11386
Stepped	6.4765	0.11386

Approach*RD Type

Naive Correlation Approach Branch Different	3.1456	0.27891
Naive Correlation Approach Naturally Monotone	1.7683	0.27891
Naive Correlation Approach Single Different	1.7827	0.27891
Naive Correlation Approach Single Similar	1.0896	0.27891
Naive Correlation Approach Stepped	7.6135	0.27891
Naive Water Approach Branch Different	51.3956	0.27891
Naive Water Approach Naturally Monotone	31.8425	0.27891
Naive Water Approach Single Different	45.5510	0.27891
Naive Water Approach Single Similar	47.6272	0.27891
Naive Water Approach Stepped	3.6263	0.27891
Perfectly Competitive Approach Branch Different	7.5427	0.27891
Perfectly Competitive Approach Naturally Monotone	11.3875	0.27891
Perfectly Competitive Approach Single Different	2.4614	0.27891
Perfectly Competitive Approach Single Similar	5.8485	0.27891
Perfectly Competitive Approach Stepped	19.7274	0.27891
Price-Based Offer Approach Branch Different	2.3133	0.27891
Price-Based Offer Approach Naturally Monotone	4.8011	0.27891
Price-Based Offer Approach Single Different	2.6113	0.27891
Price-Based Offer Approach Single Similar	0.9725	0.27891
Price-Based Offer Approach Stepped	0.8217	0.27891
Quantity-Based Offer Approach Branch Different	2.4309	0.27891
Quantity-Based Offer Approach Naturally Monotone	7.7895	0.27891
Quantity-Based Offer Approach Single Different	3.2839	0.27891
Quantity-Based Offer Approach Single Similar	0.8252	0.27891
Quantity-Based Offer Approach Stepped	6.3327	0.27891
Value Curve Approach Branch Different	1.8245	0.27891
Value Curve Approach Naturally Monotone	1.6861	0.27891
Value Curve Approach Single Different	0.9934	0.27891
Value Curve Approach Single Similar	0.6875	0.27891
Value Curve Approach Stepped	0.7374	0.27891

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*NumPeriods
Approach = Naive Correlation Approach
NumPeriods = 20 subtracted from:

Approach	NumPeriods	Difference of Means	SE of Difference	T-Value
Naive Correlation Approach	40	0.880	0.2495	3.527

Approach	NumPeriods	Adjusted P-Value
Naive Correlation Approach	40	0.0279

Approach = Naive Water Approach
NumPeriods = 20 subtracted from:

Approach	NumPeriods	Difference of Means	SE of Difference	T-Value
Naive Water Approach	40	-0.33	0.2495	-1.3

Approach	NumPeriods	Adjusted P-Value
Naive Water Approach	40	1.0000

Approach = Naive Water Approach

NumPeriods = 40 subtracted from:

Approach = Perfectly Competitive Approach
NumPeriods = 20 subtracted from:

Approach	NumPeriods	Difference of Means	SE of Difference	T-Value
Perfectly Competitive Approach	40	3.353	0.2495	13.44

Approach	NumPeriods	Adjusted P-Value
Perfectly Competitive Approach	40	0.0000

Approach = Price-Based Offer Approach
NumPeriods = 20 subtracted from:

Approach	NumPeriods	Difference of Means	SE of Difference	T-Value
Price-Based Offer Approach	40	0.117	0.2495	0.468

Approach	NumPeriods	Adjusted P-Value
Price-Based Offer Approach	40	1.0000

Approach = Quantity-Based Offer Approach
NumPeriods = 20 subtracted from:

Approach	NumPeriods	Difference of Means	SE of Difference	T-Value
Quantity-Based Offer Approach	40	1.277	0.2495	5.120

Approach	NumPeriods	Adjusted P-Value
Quantity-Based Offer Approach	40	0.0000

Approach = Value Curve Approach
NumPeriods = 20 subtracted from:

Approach	NumPeriods	Difference of Means	SE of Difference	T-Value	Adjusted P-Value
Value Curve Approach	40	0.1527	0.2495	0.6121	1.000

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*NumResLevels
Approach = Naive Correlation Approach
NumResLevels = 100 subtracted from:

Approach	NumResLevels	Difference of Means	SE of Difference	T-Value
Naive Correlation Approach	300	-1.904	0.2495	-7.63

Approach	NumResLevels	Adjusted P-Value
Naive Correlation Approach	300	0.0000

Approach = Naive Water Approach
NumResLevels = 100 subtracted from:

Approach	NumResLevels	Difference of Means	SE of Difference	T-Value
Naive Water Approach	300	0.74	0.2495	3.0

Approach	NumResLevels	Adjusted P-Value
Naive Water Approach	300	0.1882

Approach = Perfectly Competitive Approach
NumResLevels = 100 subtracted from:

Approach	NumResLevels	Difference of Means	SE of Difference	T-Value
Perfectly Competitive Approach	300	-4.84	0.2495	-19.40

Approach	NumResLevels	Adjusted P-Value
Perfectly Competitive Approach	300	0.0000

Approach = Price-Based Offer Approach
NumResLevels = 100 subtracted from:

Approach	NumResLevels	Difference of Means	SE of Difference	T-Value
Price-Based Offer Approach	300	-1.590	0.2495	-6.373

Approach	NumResLevels	Adjusted P-Value
Price-Based Offer Approach	300	0.0000

Approach = Quantity-Based Offer Approach
NumResLevels = 100 subtracted from:

Approach	NumResLevels	Difference of Means	SE of Difference	T-Value
Quantity-Based Offer Approach	300	-2.369	0.2495	-9.50

Approach	NumResLevels	Adjusted P-Value
Quantity-Based Offer Approach	300	0.0000

Approach = Value Curve Approach
NumResLevels = 100 subtracted from:

Approach	NumResLevels	Difference of Means	SE of Difference	T-Value	Adjusted P-Value
Value Curve Approach	300	-0.8045	0.2495	-3.225	0.0834

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*TransitionCorrelationType

```
Approach = Naive Correlation Approach
TransitionCorrelationType = Strong  subtracted from:

                                                         Difference
Approach                        TransitionCorrelationType    of Means
Naive Correlation Approach      Weak                           -0.340

                                                             SE of
Approach                        TransitionCorrelationType  Difference  T-Value
Naive Correlation Approach      Weak                           0.2495   -1.363

                                                         Adjusted
Approach                        TransitionCorrelationType   P-Value
Naive Correlation Approach      Weak                         1.0000


Approach = Naive Water Approach
TransitionCorrelationType = Strong  subtracted from:

                                                         Difference
Approach                        TransitionCorrelationType    of Means
Naive Water Approach            Weak                             0.98

                                                             SE of
Approach                        TransitionCorrelationType  Difference  T-Value
Naive Water Approach            Weak                           0.2495      3.9

                                                         Adjusted
Approach                        TransitionCorrelationType   P-Value
Naive Water Approach            Weak                         0.0053


Approach = Perfectly Competitive Approach
TransitionCorrelationType = Strong  subtracted from:

                                                         Difference
Approach                        TransitionCorrelationType    of Means
Perfectly Competitive Approach  Weak                           -0.977

                                                             SE of
Approach                        TransitionCorrelationType  Difference  T-Value
Perfectly Competitive Approach  Weak                           0.2495    -3.92

                                                         Adjusted
Approach                        TransitionCorrelationType   P-Value
Perfectly Competitive Approach  Weak                         0.0060

Approach = Price-Based Offer Approach
TransitionCorrelationType = Strong  subtracted from:

                                                       Difference
Approach                      TransitionCorrelationType    of Means
Price-Based Offer Approach    Weak                           -0.180
                                                            SE of
Approach                      TransitionCorrelationType  Difference  T-Value
Price-Based Offer Approach    Weak                           0.2495   -0.720

                                                       Adjusted
Approach                      TransitionCorrelationType   P-Value
Price-Based Offer Approach    Weak                         1.0000


Approach = Quantity-Based Offer Approach
```

TransitionCorrelationType = Strong subtracted from:

Approach	TransitionCorrelationType	Difference of Means
Quantity-Based Offer Approach	Weak	2.621

Approach	TransitionCorrelationType	SE of Difference	T-Value
Quantity-Based Offer Approach	Weak	0.2495	10.507

Approach	TransitionCorrelationType	Adjusted P-Value
Quantity-Based Offer Approach	Weak	0.0000

Approach = Value Curve Approach
TransitionCorrelationType = Strong subtracted from:

Approach	TransitionCorrelationType	Difference of Means	SE of Difference
Value Curve Approach	Weak	-0.1367	0.2495

Approach	TransitionCorrelationType	T-Value	Adjusted P-Value
Value Curve Approach	Weak	-0.5482	1.000

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*NumDispLevels
Approach = Naive Correlation Approach
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference	T-Value
Naive Correlation Approach	40	-0.289	0.2495	-1.158

Approach	NumDispLevels	Adjusted P-Value
Naive Correlation Approach	40	1.0000

Approach = Naive Water Approach
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference	T-Value
Naive Water Approach	40	9.26	0.2495	37.1

Approach	NumDispLevels	Adjusted P-Value
Naive Water Approach	40	0.0000

Approach = Perfectly Competitive Approach
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference	T-Value
Perfectly Competitive Approach	40	-2.180	0.2495	-8.74

Approach	NumDispLevels	Adjusted P-Value
Perfectly Competitive Approach	40	0.0000

Approach = Price-Based Offer Approach
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference	T-Value
Price-Based Offer Approach	40	0.5019	0.2495	2.012

Approach	NumDispLevels	Adjusted P-Value
Price-Based Offer Approach	40	1.0000

Approach = Quantity-Based Offer Approach
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference	T-Value
Quantity-Based Offer Approach	40	-0.376	0.2495	-1.51

Approach	NumDispLevels	Adjusted P-Value
Quantity-Based Offer Approach	40	1.0000

Approach = Value Curve Approach
NumDispLevels = 20 subtracted from:

Approach	NumDispLevels	Difference of Means	SE of Difference	T-Value	Adjusted P-Value
Value Curve Approach	40	0.1133	0.2495	0.4542	1.000

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*RD Type
Approach = Naive Correlation Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Correlation Approach	Naturally Monotone	-1.377	0.3944
Naive Correlation Approach	Single Different	-1.363	0.3944
Naive Correlation Approach	Single Similar	-2.056	0.3944
Naive Correlation Approach	Stepped	4.468	0.3944
Value Curve Approach	Branch Different	-1.321	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Correlation Approach	Naturally Monotone	-3.492	0.2096
Naive Correlation Approach	Single Different	-3.455	0.2400
Naive Correlation Approach	Single Similar	-5.212	0.0001
Naive Correlation Approach	Stepped	11.327	0.0000
Value Curve Approach	Branch Different	-3.349	0.3538

Approach = Naive Correlation Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Correlation Approach	Single Different	0.014	0.3944

Naive Correlation Approach	Single Similar	-0.679	0.3944
Naive Correlation Approach	Stepped	5.845	0.3944
Value Curve Approach	Naturally Monotone	-0.082	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Correlation Approach	Single Different	0.036	1.0000
Naive Correlation Approach	Single Similar	-1.721	1.0000
Naive Correlation Approach	Stepped	14.819	0.0000
Value Curve Approach	Naturally Monotone	-0.208	1.0000

Approach = Naive Correlation Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Correlation Approach	Single Similar	-0.693	0.3944
Naive Correlation Approach	Stepped	5.831	0.3944
Value Curve Approach	Single Different	-0.789	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Correlation Approach	Single Similar	-1.757	1.0000
Naive Correlation Approach	Stepped	14.783	0.0000
Value Curve Approach	Single Different	-2.001	1.0000

Approach = Naive Correlation Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Correlation Approach	Stepped	6.5239	0.3944
Value Curve Approach	Single Similar	-0.4021	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Correlation Approach	Stepped	16.540	0.0000
Value Curve Approach	Single Similar	-1.019	1.0000

Approach = Naive Correlation Approach
RD Type = Stepped subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach	Stepped	-6.876	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Stepped	-17.43	0.0000

Approach = Naive Water Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Water Approach	Naturally Monotone	-19.55	0.3944
Naive Water Approach	Single Different	-5.84	0.3944
Naive Water Approach	Single Similar	-3.77	0.3944
Naive Water Approach	Stepped	-47.77	0.3944

Value Curve Approach	Branch Different	-49.57	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Water Approach	Naturally Monotone	-49.6	0.0000
Naive Water Approach	Single Different	-14.8	0.0000
Naive Water Approach	Single Similar	-9.6	0.0000
Naive Water Approach	Stepped	-121.1	0.0000
Value Curve Approach	Branch Different	-125.7	0.0000

Approach = Naive Water Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Water Approach	Single Different	13.71	0.3944
Naive Water Approach	Single Similar	15.78	0.3944
Naive Water Approach	Stepped	-28.22	0.3944
Value Curve Approach	Naturally Monotone	-30.16	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Water Approach	Single Different	34.75	0.0000
Naive Water Approach	Single Similar	40.02	0.0000
Naive Water Approach	Stepped	-71.53	0.0000
Value Curve Approach	Naturally Monotone	-76.45	0.0000

Approach = Naive Water Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Water Approach	Single Similar	2.08	0.3944
Naive Water Approach	Stepped	-41.92	0.3944
Value Curve Approach	Single Different	-44.56	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Water Approach	Single Similar	5.3	0.0001
Naive Water Approach	Stepped	-106.3	0.0000
Value Curve Approach	Single Different	-113.0	0.0000

Approach = Naive Water Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Naive Water Approach	Stepped	-44.00	0.3944
Value Curve Approach	Single Similar	-46.94	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Naive Water Approach	Stepped	-111.6	0.0000
Value Curve Approach	Single Similar	-119.0	0.0000

Approach = Naive Water Approach
RD Type = Stepped subtracted from:

		Difference	SE of

Approach	RD Type	of Means	Difference
Value Curve Approach	Stepped	-2.889	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Stepped	-7.324	0.0000

Approach = Perfectly Competitive Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Perfectly Competitive Approach	Naturally Monotone	3.845	0.3944
Perfectly Competitive Approach	Single Different	-5.081	0.3944
Perfectly Competitive Approach	Single Similar	-1.694	0.3944
Perfectly Competitive Approach	Stepped	12.185	0.3944
Value Curve Approach	Branch Different	-5.718	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Perfectly Competitive Approach	Naturally Monotone	9.75	0.0000
Perfectly Competitive Approach	Single Different	-12.88	0.0000
Perfectly Competitive Approach	Single Similar	-4.30	0.0077
Perfectly Competitive Approach	Stepped	30.89	0.0000
Value Curve Approach	Branch Different	-14.50	0.0000

Approach = Perfectly Competitive Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Perfectly Competitive Approach	Single Different	-8.93	0.3944
Perfectly Competitive Approach	Single Similar	-5.54	0.3944
Perfectly Competitive Approach	Stepped	8.34	0.3944
Value Curve Approach	Naturally Monotone	-9.70	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Perfectly Competitive Approach	Single Different	-22.63	0.0000
Perfectly Competitive Approach	Single Similar	-14.04	0.0000
Perfectly Competitive Approach	Stepped	21.14	0.0000
Value Curve Approach	Naturally Monotone	-24.60	0.0000

Approach = Perfectly Competitive Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Perfectly Competitive Approach	Single Similar	3.387	0.3944
Perfectly Competitive Approach	Stepped	17.266	0.3944
Value Curve Approach	Single Different	-1.468	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Perfectly Competitive Approach	Single Similar	8.587	0.0000
Perfectly Competitive Approach	Stepped	43.774	0.0000
Value Curve Approach	Single Different	-3.722	0.0865

Approach = Perfectly Competitive Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Perfectly Competitive Approach	Stepped	13.879	0.3944
Value Curve Approach	Single Similar	-5.161	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Perfectly Competitive Approach	Stepped	35.19	0.0000
Value Curve Approach	Single Similar	-13.08	0.0000

Approach = Perfectly Competitive Approach
RD Type = Stepped subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach	Stepped	-18.99	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Stepped	-48.14	0.0000

Approach = Price-Based Offer Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Price-Based Offer Approach	Naturally Monotone	2.488	0.3944
Price-Based Offer Approach	Single Different	0.298	0.3944
Price-Based Offer Approach	Single Similar	-1.341	0.3944
Price-Based Offer Approach	Stepped	-1.492	0.3944
Value Curve Approach	Branch Different	-0.489	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Price-Based Offer Approach	Naturally Monotone	6.307	0.0000
Price-Based Offer Approach	Single Different	0.756	1.0000
Price-Based Offer Approach	Single Similar	-3.399	0.2953
Price-Based Offer Approach	Stepped	-3.782	0.0682
Value Curve Approach	Branch Different	-1.239	1.0000

Approach = Price-Based Offer Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Price-Based Offer Approach	Single Different	-2.190	0.3944
Price-Based Offer Approach	Single Similar	-3.829	0.3944
Price-Based Offer Approach	Stepped	-3.979	0.3944
Value Curve Approach	Naturally Monotone	-3.115	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Price-Based Offer Approach	Single Different	-5.55	0.0000
Price-Based Offer Approach	Single Similar	-9.71	0.0000
Price-Based Offer Approach	Stepped	-10.09	0.0000
Value Curve Approach	Naturally Monotone	-7.90	0.0000

Approach = Price-Based Offer Approach

RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Price-Based Offer Approach	Single Similar	-1.639	0.3944
Price-Based Offer Approach	Stepped	-1.790	0.3944
Value Curve Approach	Single Different	-1.618	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Price-Based Offer Approach	Single Similar	-4.155	0.0143
Price-Based Offer Approach	Stepped	-4.537	0.0025
Value Curve Approach	Single Different	-4.102	0.0180

Approach = Price-Based Offer Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Price-Based Offer Approach	Stepped	-0.1508	0.3944
Value Curve Approach	Single Similar	-0.2850	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Price-Based Offer Approach	Stepped	-0.3824	1.0000
Value Curve Approach	Single Similar	-0.7225	1.0000

Approach = Price-Based Offer Approach
RD Type = Stepped subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach	Stepped	-0.0843	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Stepped	-0.2137	1.0000

Approach = Quantity-Based Offer Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Quantity-Based Offer Approach	Naturally Monotone	5.359	0.3944
Quantity-Based Offer Approach	Single Different	0.853	0.3944
Quantity-Based Offer Approach	Single Similar	-1.606	0.3944
Quantity-Based Offer Approach	Stepped	3.902	0.3944
Value Curve Approach	Branch Different	-0.606	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Quantity-Based Offer Approach	Naturally Monotone	13.586	0.0000
Quantity-Based Offer Approach	Single Different	2.163	1.0000
Quantity-Based Offer Approach	Single Similar	-4.071	0.0206
Quantity-Based Offer Approach	Stepped	9.892	0.0000
Value Curve Approach	Branch Different	-1.537	1.0000

Approach = Quantity-Based Offer Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Quantity-Based Offer Approach	Single Different	-4.506	0.3944
Quantity-Based Offer Approach	Single Similar	-6.964	0.3944
Quantity-Based Offer Approach	Stepped	-1.457	0.3944
Value Curve Approach	Naturally Monotone	-6.103	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Quantity-Based Offer Approach	Single Different	-11.42	0.0000
Quantity-Based Offer Approach	Single Similar	-17.66	0.0000
Quantity-Based Offer Approach	Stepped	-3.69	0.0967
Value Curve Approach	Naturally Monotone	-15.47	0.0000

Approach = Quantity-Based Offer Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Quantity-Based Offer Approach	Single Similar	-2.459	0.3944
Quantity-Based Offer Approach	Stepped	3.049	0.3944
Value Curve Approach	Single Different	-2.291	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Quantity-Based Offer Approach	Single Similar	-6.233	0.0000
Quantity-Based Offer Approach	Stepped	7.729	0.0000
Value Curve Approach	Single Different	-5.807	0.0000

Approach = Quantity-Based Offer Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Quantity-Based Offer Approach	Stepped	5.5075	0.3944
Value Curve Approach	Single Similar	-0.1377	0.3944

Approach	RD Type	T-Value	Adjusted P-Value
Quantity-Based Offer Approach	Stepped	13.9627	0.0000
Value Curve Approach	Single Similar	-0.3490	1.0000

Approach = Quantity-Based Offer Approach
RD Type = Stepped subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value
Value Curve Approach	Branch Different	-4.508	0.3944	-11.43
Value Curve Approach	Naturally Monotone	-4.647	0.3944	-11.78
Value Curve Approach	Single Different	-5.339	0.3944	-13.54
Value Curve Approach	Single Similar	-5.645	0.3944	-14.31
Value Curve Approach	Stepped	-5.595	0.3944	-14.19

Approach	RD Type	Adjusted P-Value
Value Curve Approach	Branch Different	0.0000
Value Curve Approach	Naturally Monotone	0.0000
Value Curve Approach	Single Different	0.0000
Value Curve Approach	Single Similar	0.0000
Value Curve Approach	Stepped	0.0000

Approach = Value Curve Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value
Value Curve Approach	Naturally Monotone	-0.138	0.3944	-0.351
Value Curve Approach	Single Different	-0.831	0.3944	-2.107
Value Curve Approach	Single Similar	-1.137	0.3944	-2.883
Value Curve Approach	Stepped	-1.087	0.3944	-2.756

Approach	RD Type	Adjusted P-Value
Value Curve Approach	Naturally Monotone	1.000
Value Curve Approach	Single Different	1.000
Value Curve Approach	Single Similar	1.000
Value Curve Approach	Stepped	1.000

Approach = Value Curve Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value
Value Curve Approach	Single Different	-0.6928	0.3944	-1.756
Value Curve Approach	Single Similar	-0.9986	0.3944	-2.532
Value Curve Approach	Stepped	-0.9488	0.3944	-2.405

Approach	RD Type	Adjusted P-Value
Value Curve Approach	Single Different	1.000
Value Curve Approach	Single Similar	1.000
Value Curve Approach	Stepped	1.000

Approach = Value Curve Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value	Adjusted P-Value
Value Curve Approach	Single Similar	-0.3058	0.3944	-0.7753	1.000
Value Curve Approach	Stepped	-0.2560	0.3944	-0.6490	1.000

Approach = Value Curve Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value	Adjusted P-Value
Value Curve Approach	Stepped	0.04984	0.3944	0.1264	1.000

J. Pairwise Comparison Tests for Computational Time on Interpolation Models

Paired T-Test and CI: Time - No Interp, Time - Interp 2

```
Paired T for Time - No Interp - Time - Interp 2

                    N     Mean    StDev  SE Mean
Time - No Interp  1600  52.2431  47.5270   1.1882
Time - Interp 2   1600  42.2406  64.4952   1.6124
Difference        1600  10.0025  44.4326   1.1108


95% CI for mean difference: (7.8237, 12.1813)
T-Test of mean difference = 0 (vs not = 0): T-Value = 9.00  P-Value = 0.000
```

Paired T-Test and CI: Time - Interp 2, Time - Interp 5

```
Paired T for Time - Interp 2 - Time - Interp 5

                   N     Mean    StDev  SE Mean
Time - Interp 2  1600  42.2406  64.4952   1.6124
Time - Interp 5  1600  25.4319  42.7665   1.0692
Difference       1600  16.8088  32.8591   0.8215


95% CI for mean difference: (15.1975, 18.4200)
T-Test of mean difference = 0 (vs not = 0): T-Value = 20.46  P-Value = 0.000
```

Paired T-Test and CI: Time - Interp 5, Time - Interp 10

```
Paired T for Time - Interp 5 - Time - Interp 10

                     N     Mean     StDev  SE Mean
Time - Interp 5   1600  25.4319   42.7665   1.0692
Time - Interp 10  1600  18.9200   33.8102   0.8453
Difference        1600  6.51187  20.22444  0.50561


95% CI for mean difference: (5.52014, 7.50361)
T-Test of mean difference = 0 (vs not = 0): T-Value = 12.88  P-Value = 0.000
```

K. MINITAB Output on Solution Quality for Interpolation Models

General Linear Model: Simulated Value versus Approach, NumPeriods, ...

Factor	Type	Levels	Values
Approach	fixed	4	Value Curve Approach, x Interp 2nd - Value Approach, y Interp 5th - Value Approach, z Interp 10th - Value Approach
NumPeriods	fixed	2	20, 40
NumResLevels	fixed	2	100, 300
NumDispLevels	fixed	2	20, 40
NumRDs	fixed	2	25, 100
NumInflows	fixed	2	1, 5
TransitionCorrelationType	fixed	2	Strong, Weak
RD Type	fixed	5	Branch Different, Naturally Monotone, Single Different, Single Similar, Stepped

Analysis of Variance for Simulated Value, using Adjusted SS for Tests

Source	DF	Seq SS	Adj SS	Adj MS	F
Approach	3	460.91	460.91	153.64	295.45
NumPeriods	1	25.24	25.24	25.24	48.55
Approach*NumPeriods	3	1.22	1.22	0.41	0.79
NumResLevels	1	1625.71	1625.71	1625.71	3126.30
Approach*NumResLevels	3	71.44	71.44	23.81	45.79
NumDispLevels	1	22.96	22.96	22.96	44.16
Approach*NumDispLevels	3	0.22	0.22	0.07	0.14
NumRDs	1	0.32	0.32	0.32	0.62
Approach*NumRDs	3	0.01	0.01	0.00	0.01

NumInflows	1	833.52	833.52	833.52	1602.89
Approach*NumInflows	3	19.42	19.42	6.47	12.45
TransitionCorrelationType	1	25.33	25.33	25.33	48.71
Approach*TransitionCorrelationType	3	0.23	0.23	0.08	0.15
RD Type	4	1475.90	1475.90	368.97	709.55
Approach*RD Type	12	121.17	121.17	10.10	19.42
Error	6356	3305.19	3305.19	0.52	
Total	6399	7988.79			

Source	P
Approach	0.000
NumPeriods	0.000
Approach*NumPeriods	0.502
NumResLevels	0.000
Approach*NumResLevels	0.000
NumDispLevels	0.000
Approach*NumDispLevels	0.936
NumRDs	0.430
Approach*NumRDs	0.999
NumInflows	0.000
Approach*NumInflows	0.000
TransitionCorrelationType	0.000
Approach*TransitionCorrelationType	0.932
RD Type	0.000
Approach*RD Type	0.000
Error	
Total	

S = 0.721118   R-Sq = 58.63%   R-Sq(adj) = 58.35%

Least Squares Means for Simulated Value

Approach	Mean	SE Mean
Value Curve Approach	1.1858	0.01803
x Interp 2nd - Value Approach	1.2424	0.01803
y Interp 5th - Value Approach	1.4871	0.01803
z Interp 10th - Value Approach	1.8670	0.01803
NumPeriods		
20	1.3828	0.01275
40	1.5084	0.01275
NumResLevels		
100	1.9496	0.01275
300	0.9416	0.01275
Approach*NumResLevels		
Value Curve Approach 100	1.5880	0.02550
Value Curve Approach 300	0.7835	0.02550
x Interp 2nd - Value Approach 100	1.6667	0.02550
x Interp 2nd - Value Approach 300	0.8181	0.02550
y Interp 5th - Value Approach 100	2.0060	0.02550
y Interp 5th - Value Approach 300	0.9681	0.02550
z Interp 10th - Value Approach 100	2.5375	0.02550
z Interp 10th - Value Approach 300	1.1965	0.02550
NumDispLevel		
20	1.3857	0.01275
40	1.5055	0.01275
NumInflows		

1	1.0847	0.01275
5	1.8065	0.01275

Approach*NumInflows

Value Curve Approach 1	0.8230	0.02550
Value Curve Approach 5	1.5486	0.02550
x Interp 2nd - Value Approach 1	0.9048	0.02550
x Interp 2nd - Value Approach 5	1.5800	0.02550
y Interp 5th - Value Approach 1	1.1909	0.02550
y Interp 5th - Value Approach 5	1.7833	0.02550
z Interp 10th - Value Approach 1	1.4200	0.02550
z Interp 10th - Value Approach 5	2.3139	0.02550

TransitionCo

Strong	1.5085	0.01275
Weak	1.3827	0.01275

RD Type

Branch Different	2.0598	0.02016
Naturally Monotone	1.8955	0.02016
Single Different	1.4035	0.02016
Single Similar	1.0913	0.02016
Stepped	0.7777	0.02016

Approach*RD Type

Value Curve Approach Branch Different	1.8245	0.04031
Value Curve Approach Naturally Monotone	1.6861	0.04031
Value Curve Approach Single Different	0.9934	0.04031
Value Curve Approach Single Similar	0.6875	0.04031
Value Curve Approach Stepped	0.7374	0.04031
x Interp 2nd - Value Approach Branch Different	1.8702	0.04031
x Interp 2nd - Value Approach Naturally Monotone	1.7260	0.04031
x Interp 2nd - Value Approach Single Different	1.0862	0.04031
x Interp 2nd - Value Approach Single Similar	0.7886	0.04031
x Interp 2nd - Value Approach Stepped	0.7411	0.04031
y Interp 5th - Value Approach Branch Different	2.1073	0.04031
y Interp 5th - Value Approach Naturally Monotone	1.9181	0.04031
y Interp 5th - Value Approach Single Different	1.4791	0.04031
y Interp 5th - Value Approach Single Similar	1.1526	0.04031
y Interp 5th - Value Approach Stepped	0.7783	0.04031
z Interp 10th - Value Approach Branch Different	2.4374	0.04031
z Interp 10th - Value Approach Naturally Monotone	2.2516	0.04031
z Interp 10th - Value Approach Single Different	2.0553	0.04031
z Interp 10th - Value Approach Single Similar	1.7365	0.04031
z Interp 10th - Value Approach Stepped	0.8541	0.04031

Bonferroni Simultaneous Tests
Response Variable Simulated Value
All Pairwise Comparisons among Levels of Approach*NumResLevels
Approach = Value Curve Approach
NumResLevels = 100 subtracted from:

Approach	NumResLevels	Difference of Means	SE of Difference	T-Value
Value Curve Approach	300	-0.8045	0.03606	-22.31

Approach	NumResLevels	Adjusted P-Value
Value Curve Approach	300	0.0000

Approach = x Interp 2nd - Value Approach

```
NumResLevels = 100  subtracted from:

                                             Difference       SE of
Approach                        NumResLevels    of Means  Difference  T-Value
x Interp 2nd - Value Approach   300              -0.8486     0.03606   -23.54
                                             Adjusted
Approach                        NumResLevels   P-Value
x Interp 2nd - Value Approach   300              0.0000

Approach = y Interp 5th - Value Approach
NumResLevels = 100  subtracted from:

                                             Difference       SE of
Approach                        NumResLevels    of Means  Difference  T-Value
y Interp 5th - Value Approach   300               -1.038     0.03606   -28.79

                                             Adjusted
Approach                        NumResLevels   P-Value
y Interp 5th - Value Approach   300              0.0000

Approach = z Interp 10th - Value Approach
NumResLevels = 100  subtracted from:

                                             Difference       SE of
Approach                        NumResLevels    of Means  Difference  T-Value
z Interp 10th - Value Approach  300               -1.341     0.03606   -37.19

                                             Adjusted
Approach                        NumResLevels   P-Value
z Interp 10th - Value Approach  300              0.0000


Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*NumInflows
Approach = Value Curve Approach
NumInflows = 1  subtracted from:

                                           Difference       SE of
Approach                        NumInflows    of Means  Difference  T-Value
Value Curve Approach            5              0.72555     0.03606   20.123

                                           Adjusted
Approach                        NumInflows   P-Value
Value Curve Approach            5              0.0000


Approach = x Interp 2nd - Value Approach
NumInflows = 1  subtracted from:

                                           Difference       SE of
Approach                        NumInflows    of Means  Difference  T-Value
x Interp 2nd - Value Approach   5               0.6752     0.03606   18.726

                                           Adjusted
Approach                        NumInflows   P-Value
x Interp 2nd - Value Approach   5              0.0000


Approach = y Interp 5th - Value Approach
NumInflows = 1  subtracted from:

                                           Difference       SE of
```

Approach	NumInflows	of Means	Difference	T-Value
y Interp 5th - Value Approach	5	0.5924	0.03606	16.430

Approach	NumInflows	Adjusted P-Value
y Interp 5th - Value Approach	5	0.0000

Approach = z Interp 10th - Value Approach
NumInflows = 1 subtracted from:

Approach	NumInflows	Difference of Means	SE of Difference	T-Value
z Interp 10th - Value Approach	5	0.8939	0.03606	24.79

Approach	NumInflows	Adjusted P-Value
z Interp 10th - Value Approach	5	0.0000

Bonferroni Simultaneous Tests
Response Variable Percentage Error
All Pairwise Comparisons among Levels of Approach*RD Type
Approach = Value Curve Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach	Naturally Monotone	-0.138	0.05701
Value Curve Approach	Single Different	-0.831	0.05701
Value Curve Approach	Single Similar	-1.137	0.05701
Value Curve Approach	Stepped	-1.087	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Naturally Monotone	-2.43	1.0000
Value Curve Approach	Single Different	-14.58	0.0000
Value Curve Approach	Single Similar	-19.94	0.0000
Value Curve Approach	Stepped	-19.07	0.0000

Approach = Value Curve Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach	Single Different	-0.6928	0.05701
Value Curve Approach	Single Similar	-0.9986	0.05701
Value Curve Approach	Stepped	-0.9488	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Single Different	-12.15	0.0000
Value Curve Approach	Single Similar	-17.52	0.0000
Value Curve Approach	Stepped	-16.64	0.0000

Approach = Value Curve Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach	Single Similar	-0.3058	0.05701

Value Curve Approach	Stepped	-0.2560	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Single Similar	-5.364	0.0000
Value Curve Approach	Stepped	-4.490	0.0014

Approach = Value Curve Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
Value Curve Approach	Stepped	0.04984	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
Value Curve Approach	Stepped	0.8743	1.0000

Approach = x Interp 2nd - Value Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
x Interp 2nd - Value Approach	Naturally Monotone	-0.144	0.05701
x Interp 2nd - Value Approach	Single Different	-0.784	0.05701
x Interp 2nd - Value Approach	Single Similar	-1.082	0.05701
x Interp 2nd - Value Approach	Stepped	-1.129	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
x Interp 2nd - Value Approach	Naturally Monotone	-2.53	1.0000
x Interp 2nd - Value Approach	Single Different	-13.75	0.0000
x Interp 2nd - Value Approach	Single Similar	-18.97	0.0000
x Interp 2nd - Value Approach	Stepped	-19.81	0.0000

Approach = x Interp 2nd - Value Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
x Interp 2nd - Value Approach	Single Different	-0.6399	0.05701
x Interp 2nd - Value Approach	Single Similar	-0.9374	0.05701
x Interp 2nd - Value Approach	Stepped	-0.9849	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
x Interp 2nd - Value Approach	Single Different	-11.22	0.0000
x Interp 2nd - Value Approach	Single Similar	-16.44	0.0000
x Interp 2nd - Value Approach	Stepped	-17.28	0.0000

Approach = x Interp 2nd - Value Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
x Interp 2nd - Value Approach	Single Similar	-0.2975	0.05701
x Interp 2nd - Value Approach	Stepped	-0.3451	0.05701

Adjusted

Approach	RD Type	T-Value	P-Value
x Interp 2nd - Value Approach	Single Similar	-5.219	0.0000
x Interp 2nd - Value Approach	Stepped	-6.053	0.0000

Approach = x Interp 2nd - Value Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
x Interp 2nd - Value Approach	Stepped	-0.04757	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
x Interp 2nd - Value Approach	Stepped	-0.8345	1.0000

Approach = y Interp 5th - Value Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
y Interp 5th - Value Approach	Naturally Monotone	-0.189	0.05701
y Interp 5th - Value Approach	Single Different	-0.628	0.05701
y Interp 5th - Value Approach	Single Similar	-0.955	0.05701
y Interp 5th - Value Approach	Stepped	-1.329	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
y Interp 5th - Value Approach	Naturally Monotone	-3.32	0.1732
y Interp 5th - Value Approach	Single Different	-11.02	0.0000
y Interp 5th - Value Approach	Single Similar	-16.74	0.0000
y Interp 5th - Value Approach	Stepped	-23.31	0.0000

Approach = y Interp 5th - Value Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
y Interp 5th - Value Approach	Single Different	-0.439	0.05701
y Interp 5th - Value Approach	Single Similar	-0.765	0.05701
y Interp 5th - Value Approach	Stepped	-1.140	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
y Interp 5th - Value Approach	Single Different	-7.70	0.0000
y Interp 5th - Value Approach	Single Similar	-13.43	0.0000
y Interp 5th - Value Approach	Stepped	-19.99	0.0000

Approach = y Interp 5th - Value Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
y Interp 5th - Value Approach	Single Similar	-0.3264	0.05701
y Interp 5th - Value Approach	Stepped	-0.7008	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
y Interp 5th - Value Approach	Single Similar	-5.73	0.0000

y Interp 5th - Value Approach	Stepped	-12.29	0.0000

Approach = y Interp 5th - Value Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
y Interp 5th - Value Approach	Stepped	-0.3744	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
y Interp 5th - Value Approach	Stepped	-6.567	0.0000

Approach = z Interp 10th - Value Approach
RD Type = Branch Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
z Interp 10th - Value Approach	Naturally Monotone	-0.186	0.05701
z Interp 10th - Value Approach	Single Different	-0.382	0.05701
z Interp 10th - Value Approach	Single Similar	-0.701	0.05701
z Interp 10th - Value Approach	Stepped	-1.583	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
z Interp 10th - Value Approach	Naturally Monotone	-3.26	0.2143
z Interp 10th - Value Approach	Single Different	-6.70	0.0000
z Interp 10th - Value Approach	Single Similar	-12.29	0.0000
z Interp 10th - Value Approach	Stepped	-27.77	0.0000

Approach = z Interp 10th - Value Approach
RD Type = Naturally Monotone subtracted from:

Approach	RD Type	Difference of Means	SE of Difference
z Interp 10th - Value Approach	Single Different	-0.196	0.05701
z Interp 10th - Value Approach	Single Similar	-0.515	0.05701
z Interp 10th - Value Approach	Stepped	-1.398	0.05701

Approach	RD Type	T-Value	Adjusted P-Value
z Interp 10th - Value Approach	Single Different	-3.44	0.1095
z Interp 10th - Value Approach	Single Similar	-9.04	0.0000
z Interp 10th - Value Approach	Stepped	-24.51	0.0000

Approach = z Interp 10th - Value Approach
RD Type = Single Different subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value
z Interp 10th - Value Approach	Single Similar	-0.319	0.05701	-5.59
z Interp 10th - Value Approach	Stepped	-1.201	0.05701	-21.07

Approach	RD Type	Adjusted P-Value
z Interp 10th - Value Approach	Single Similar	0.0000
z Interp 10th - Value Approach	Stepped	0.0000

Approach = z Interp 10th - Value Approach
RD Type = Single Similar subtracted from:

Approach	RD Type	Difference of Means	SE of Difference	T-Value
z Interp 10th - Value Approach	Stepped	-0.8825	0.05701	-15.48

Approach	RD Type	Adjusted P-Value
z Interp 10th - Value Approach	Stepped	0.0000