School of Mathematics and Statistics, University of Canterbury,
Christchurch, New Zealand

CAMS 690 MSc Thesis

submitted in partial fulfilment of the requirements for

The Degree of Master of Science in Computation and Applied Mathematics

Development of an Optimisation Framework for the Investigation of Multiple Water Stores

by

Fletcher Frater

Supervisor: Associate Professor Rua Murray

In Association with Aqualinc Research Limited

2022

Abstract

This thesis develops a framework to investigate optimal storage distributions of water. This is developed through an investigation of current soil and agricultural models to include accurate details of a soil-crop-atmosphere system combined with farming practices and infrastructure. Irrigation scheduling is combined with the use of on-farm and bulk water storages in a Mixed Integer Non-Linear Program. This is developed in the *Python* programming language with the use of the `pyomo` library. The program is solved with the Outer Approximation decomposition technique.

The framework showed good agreement with current rule based decision scheduling, and was able to produce more optimal soil moisture dynamics. These dynamics reduced drainage losses and were able to meet the evapotranspiration demand equally well with less water use in some cases. This showed that with optimal management, more desirable dynamics are achievable. With access to additional storage, the model was able to utilise additional water in a cost efficient manner to improve crop production under varying source water availabilities. There was a high sensitivity to rainfall shown in the ability for soil moisture to hold enough water to provide sufficient supply to the crop. This prompted the conclusion that the model must be used on a case-by-case basis to investigate different strategies with fixed parameter combinations.

With the structure of the framework presented, it is possible to scale to include explicit details of large farms in order to run simulations. In the current set up, simulations work on a per hectare basis but this can be scaled as required to represent larger farming systems. Future work would involve using the `pyomo` model to add extra detail into the framework to allow larger scales to be represented more accurately.

Acknowledgements

I would first like to thank Associate Professor Rua Murray for his support, guidance, and time as my supervisor. Although I approached him with an industry project, he was able to offer advice and insight that I believe pushed me to deliver a comprehensive thesis, thank you Rua.

I would also like to thank Andrew and John of Aqualinc. This project has expanded my skill base, and I thank you both for the guidance throughout. I look forward to what we can achieve together in the future.

I must also thank my loving family for their constant support over the last five years. It has been a fantastic journey, and I am forever grateful to have you all backing me every step of the way.

Finally, I would like to extend gratitude to my friends over the last five years. Without you, this would have been a very different journey and I thank you all. In particular, I want to mention Ashton, Annabella, Alex, and Rudeep. I also want to say thanks to my fellow Mathematics and Statistics Masters students, Kira and Victor.

Contents

1 Introduction **1**

1.1 Environment and Climate . . . 1

1.2 Irrigation Systems and Scheduling . . . 2

1.3 Farming System Constraints . . . 2

1.4 System Optimisation Framework . . . 3

1.5 Thesis Overview . . . 4

2 Soil Hydrology, Agriculture, and Water Allocation **6**

2.1 Soil Hydrology . . . 6

2.1.1 Soil Profile . . . 6

2.1.2 Soil Moisture . . . 6

2.1.3 Root Zone, Water Holding Capacity, and Plant Available Water . . . 7

2.2 Soil-Crop-Atmosphere System . . . 8

2.2.1 Evapotranspiration . . . 8

2.2.2 Crop Factors . . . 9

2.2.3 Irrigation . . . 9

2.2.4 Drainage . . . 10

2.3 Water Supply and Demand . . . 11

2.3.1 Demand . . . 12

2.3.2 Supply . . . 12

3 Agricultural Modelling **13**

3.1 Crop Water Production Functions . . . 13

3.2 Evapotranspiration Models . . . 15

3.2.1 FAO Penman-Monteith Evapotranspiration Model . . . 16

3.2.2 Accurate Evapotranspiration . . . 17

3.3 Models of Soil Moisture Dynamics . . . 18

3.3.1 Mass Balance Model Development . . . 19

3.3.2 Irricalc Irrigation System Model . . . 20

3.4 Irrigation Distributions . . . 21

3.4.1 Water Infiltration 21
3.4.2 Efficiency and Adequacy of Irrigation Area 22
3.5 Data 24

4 Optimisation Methods and Framework **28**
4.1 Previous Modelling Approaches 28
4.2 Framework Approach 29
4.2.1 Objective Function 29
4.2.2 Moisture Dynamics 30
4.2.3 Irrigation and the Return Period 31
4.2.4 Drainage and Efficiency 32
4.2.5 Inclusion of Water Storages 33
4.3 Optimisation Methods 36
4.3.1 MindtPy Decomposition Method 36
4.3.2 Convexity of Nonlinear Functions 38
4.3.3 Mathematical Constraints 40
4.3.4 Sub Problem Solvers 42

5 Verification of Model Design **44**
5.1 Initial Simulations 44
5.2 Comparison with Irricalc Modelling 50
5.3 Higher Cost Emphasis 52
5.4 Rainfall Events and Modelled Drainage 55
5.5 Summary 55

6 Extending the Framework for Water Storage **58**
6.1 Water Storage Rules 58
6.2 Simulation Results 59
6.2.1 Initial on-farm Storage Levels 61
6.2.2 Rainfall Dependency 62
6.3 Sensitivity to Parameters 64
6.4 Benefit of Storage Access 68

7 Discussion **73**
7.1 Sensitivity to Pricing 74
7.2 Strategies 74
7.3 Viability of Soil Moisture as a Storage 75
7.4 Limitations and Future Work 76

8 Conclusion **78**

Bibliography **81**

A Constraint Formalism **89**

List of Figures

2.1 Typical soil profile; A – surface soil, B – subsoil, C – substratum. Picture adapted from Tomáš Kebert under the Creative Commons Attribution-Share Alike 4.0 International license: creativecommons.org/licenses/bysa/4.0/legalcode 7

3.1 Relationship between application efficiency, application uniformity and application depth. As the adequately irrigated area increases, the overall efficiency decreases. This is due to more water being required for less net gain due to losses. 23

3.2 Example of the time series input for the reference evapotranspiration, as computed by the Penman-Monteith FAO model, and the climate rainfall. Both inputs are given in millimeters which can be directly used in model computations. 24

3.3 Crop factor time series for pasture. This allows the calculation of potential evapotranspiration for pasture from the reference ET calculated by the FAO Penman-Monteith equation. 25

3.4 Example of the potential evapotranspiration for pasture. The upper plot represents the scaled evapotranspiration over the 2005/06 season. The lower plot is a closer look over the high demand section of the season. Also plotted is the average PET of 4 mm per day over this higher demand period in blue. This is the demand that must be met. 26

3.5 Example of river flow data for the 2005/06 season. In the upper plot the raw flow rate data is shown. This is measured data from a reference location on the river. The lower plot shows this flow rate as an allowed take. Also plotted is the cut off levels for assumed reliabilities. When the flow rate drops below these cut offs there is a zero allowable take. Lower reliabilities are cut off first. 26

3.6 Example of how the reliabilities translate to supply availability. The highest reliability has the supply available more often, while a lower reliability suffers frequent restrictions. 27

4.1 Efficiency curves for an assumed uniformity of 80%. In blue is the fitted linear function, while the red represents the constrained efficiency calculation. The vertical section of the modified curve represents the limiting ratio relationship, $I_t \leq 2(F_c - S_t)$. 34

5.1 Soil moisture trajectories for the 1997/98 season using 60% and 80% soil moisture triggers, alongside the optimal based soil moisture output. Also shown is the consistent supply for the season. The bottom three plots show the deficit in the AET compared to the PET. 46

5.2 Irrigation events for the 1997/98 season, note the decline in events towards the end of the season in the optimal approach. 46

5.3 Drainage amounts for the 1997/98 season. The higher trigger of 80% experiences more drainage due to the higher overall soil moisture level. The optimal approach still experiences drainage where the soil moisture level is too high to take full advantage of the rainfall; this can be expected at the start of the season where spring rainfall combined with a high soil moisture level lead to drainage. 47

5.4 Soil moisture trajectories for the 2000/01 season. A higher initial soil moisture before restriction, allows slightly more of the ET load to be met. Ultimately the final result of an extremely low moisture level is unavoidable. The ET load that is unmet over the restriction period is very close to the assumed daily demand of five mm per day. This will have a severe impact on the crop. 48

5.5 Irrigation events for the 2000/01 season. Very similar across the three methods when supply is available, the optimal approach varies the amount of water per irrigation event, but overall maintains a similar moisture level as the higher trigger method, as seen in Figure 5.4. 48

5.6 Drainage amounts for the three methods applied to the 2000/01 season. The optimal approach is again able to maintain less drainage through the main portion of the season. Drainage at the beginning is very similar across methods. 49

5.7 Soil moisture levels for the 20005/06 season. This season experiences a number of restrictions and the soil moisture levels reflect that. The lower trigger level sees greater overall dips in soil moisture due to the poorer starting position when entering the restricted period. A high soil moisture level appears advantageous here as it prevents a compounding of water deficits. This is further seen in the ability to meet the ET load. The optimal approach and higher trigger level have similar deficits in meeting the demand, while the lower trigger experiences a greater inability to meet the demand. 50

5.8 Irrigation events for the three approaches of the 2005/06 season. The irrigation amounts are very similar between methods throughout the middle of the season. This was seen to translate to different soil moisture levels. A lower trigger level meant the soil moisture may have gone into a restriction with an already lower soil moisture level. This manifested as the low trigger having a much lower soil moisture average, despite having similar water use. 51

5.9 Drainage amounts for the 2005/06 season. The optimal approach has less drainage, while the higher trigger level again experiences the greatest amount of water loss. 51

5.10 Soil moisture dynamics with increased irrigation cost for the 2000/01 season. There is a more pronounced reduction of irrigation at the start of the season, before reaching a higher soil moisture level as the restriction begins. The ability to meet the ET demand remains poor as the restriction takes hold. 53

5.11 Irrigation events with the increased irrigation cost for the 2000/01 season. There is less initial irrigation, before ramping up towards the heavy restriction. Overall there is only a slight decrease in the total irrigation used, compared to the cheaper costing. 53

5.12 Irrigation events for the 2005/06 season at an increased irrigation cost of three dollars per mm per hectare. There is a slight increase in the irrigation earlier in the season. 54

5.13 2005/06 season soil moisture with increased irrigation cost. A higher average soil moisture is achieved through more irrigation events earlier in the season. This further shows through better timing of irrigation events, the effects of restrictions can be managed. 54

5.14 Comparing soil moisture levels and irrigation amounts for the 2000/01 season when different rainfall amounts are used as input series. The low rainfall is a quarter of the regular amount. There is little difference observed in the soil moisture levels between the two simulations, but the lower rainfall requires more irrigation to account for less moisture. . . . 56

5.15 Soil moisture and irrigation for different rainfall profiles on the 2005/06 season. Much lower soil moisture recovery is observed with less rainfall when supply is available. This is despite similar amount of irrigation. This shows there is some dependency on additional rainfall in managing the soil moisture levels. 56

6.1 Simulation results for the 2000/01 season with access to on-farm and bulk storages. Towards the end of the season the actual ET falls away from the maximum potential, as expected with less water availability. Finally the rainfall over the season is shown in the last plot. 60

6.2 Output from the 2000/01 season with an increased cost to irrigate from the supply and bulk storage. An almost almost identical results as the cheaper cost simulation is observed. Irrigation stops earlier and there is more unmet ET demand at the very tail of the season. The normalised objective function is also reduced as a result, down to 0.84 from 0.93. . . 61

6.3 Simulation results for the 2005/06 season. Overall the available water is kept high, allowing the ET demand to be met. The on-farm storage is used over two irrigation events during a longer restriction towards the tail of the season. For this season there was a number of high rainfall events, it is possible that without these events the ability to meet the ET demand would require more storage input. 62

6.4 Simulation of the 2005/06 season with no initially stored water in the on-farm storage. The results show a similar pattern to the previous 2005/06 season simulations in terms of soil moisture and irrigation events. The storage pond is only filled up to 5 mm to allow for a single irrigation event. This irrigation bridges the first main restriction period during the middle of the season. The ET demand is satisfied similarly to the no storage case, and the attained normalised objective function value is very similar. . . 63

6.5 Dependency between total rainfall and optimal irrigation over 30 years. More rainfall refilling the soil moisture requires less irrigation. 63

6.6 Simulation of the 2000/01 season with reduced rainfall. The rainfall was scaled to a quarter of the original amount. A greater number of irrigation events compensate the reduced rainfall, and there are more irrigation events during the restriction. These irrigations are a combination of the on-farm storage and the bulk storage. This is in an effort to maintain ET in a cost effective way. The reduced rainfall peaks at roughly 8 mm. . . . 64

6.7 Low rainfall for the 2005/06 season. The on-farm storage is empty at the start of the season and this is refilled and then irrigated from during most restrictions. The ET demand is able to be met as well as with the original rainfall. Once again the bulk storage is used to stretch the on-farm stored water during irrigation. This increases the number of irrigation events during one of the longer restricted periods. 65

6.8 2005/06 season simulated with low rainfall, an empty on-farm storage, and higher irrigation costs. The changes are a lack of bulk storage being used, and an earlier decrease in irrigation at the tail of the season. The decrease in soil moisture towards the end of the season also reduces the ability to meet all of the ET demand, and there is an increase in the deficit between the maximum potential and AET. 65

6.9 Increasing cost of bulk storage irrigation for the 2000/01 season. Irrigation is plotted as total millimeters applied during the season against the respective cost of bulk storage. Average soil moisture is computed as the seasonal average of daily soil moisture levels. The objective function is normalised with respect to the maximum possible profit. There is initially no water in the on-farm storage. When the cost of the bulk storage exceeds the cost to fill and irrigate from the on-farm store, the irrigation split between the three supplies reshuffles to accommodate the lowest cost combination. 66

6.10 Increasing irrigation from on-farm storage cost for 2000/01 season. Total irrigation is again very similar with the distribution between the three supplies managing the increasing cost. Objective function decreases as a response to increasing price. 67

6.11 Increasing cost of Bulk storage access for the 2005/06 season. As with the 2000/01 season the total irrigation remains reasonably constant throughout, as does the soil moisture. This perhaps implies the total irrigation for a season is optimal and the distribution between irrigation is one that minimises cost. 68

6.12 Increasing cost of source irrigation and storage for the 2000/01 season. The bulk storage is priced as 50 cents higher than the source, and on-farm storage irrigation is fixed at \$0.25/mm/Ha, with an increasing cost to store. Irrigation amounts hold constant until around nine dollars, at which the cost to irrigate to sufficient soil moisture levels out weigh the profit generated. 69
6.13 Source irrigation cost increase for the 2005/06 season. As with the 2000/01 season, the irrigation levels hold constant until a cut off around \$10/mm/Ha. Around four dollars there is an increase in the amount of on-farm storage irrigation. No bulk storage irrigation is used with the shorter restriction periods found in the 2005/06 season. 69
6.14 Comparison of the 2000/01 season with and without access to additional storage. Top plot represents the soil moisture trajectories of the two cases. The middle plots show the irrigation, and in the additional storage case, which supply the irrigation is coming from. Finally at the bottom is the unmet ET demand. With the additional irrigation made possible with the storage access, the ET demand is met as required. 70
6.15 Comparison for the 2005/06 season. The additional supply with storage does not present a significant improvement. Both cases are able to meet the ET demand well. 70
6.16 Contrast between storage access for the 2000/01 season with greatly reduced rainfall. Similar dynamics are observed, but with a greater amount of irrigation across both cases. 71
6.17 Contrast between 2005/06 season with and without storage as well as reduced rainfall. Here the observed dynamics greatly differ. Without access to additional storage, the soil moisture is unable to maintain or recover as well as before. This leads to severe drop off in the ability to meet the demand. By contrast, with additional storage the demand is able to be met as well as with the original amount of rainfall. 72

This page intentionally left blank.

Chapter 1

Introduction

Water abstraction and the associated economic demand surrounding water use are fundamental ideas that have long been a subject of research within the agricultural sector. This research has looked into many areas of the agricultural cycle, and in particular, the use and management of irrigation. Irrigation plays a vital role in farming production, and this inherently relies on the ability to access water as and when needed. Irrigation provides the economic backbone many farms require to maintain productivity. The availability and supply of water for this irrigation is at the forefront of decision making: how can a farmer make informed decisions if their water supply is unreliable, and how can they meet the demand for their crops? It is within this supply-demand relationship that farmers must optimise their application and distribution of water.

1.1 Environment and Climate

In an ideal world, a farmer would have access to enough water to meet the demands of all crops at reasonable cost as and when needed. This is not the case for many farmers, where water supply and availability is subject to various climate factors and imposed limits. These climate factors affect the flow rate of rivers and the wider availability of water abstraction, while limits can restrict the amount of water available at any one time. Farmers must adhere to these limits as they authorise the take and use of water. These limits can be a seasonal total or related to flow rate. With these limitations in place, a farmer must make decisions on how they should manage their water in order to maximise crop production. To improve crop production, fertiliser can be applied during growth stages. Fertiliser typically contains nitrogen which is an essential part of the plant growth process. The application of fertiliser to crops is associated with environmental drawbacks. Excessive amounts of applied fertiliser is considered an environmental pollutant due to the corresponding nitrate leaching [1, 2]. This is linked to over-irrigation, where excess water is applied to the soil and increased drainage occurs [3]. Nitrate then leaches through the soil to groundwater and into waterways [4]. How water is utilised in irrigation offers the greatest potential for a reduction in nitrate leaching. More precisely, the water supply should be better matched to the exact demands of the crop to minimise excess irrigation [5].

1.2 Irrigation Systems and Scheduling

Irrigation scheduling is subject to multiple constraints, in particular the availability of water supply and the return period of the irrigator. The irrigator return period is the physical restriction on the time taken to reset the irrigator. A new irrigation event is unable to be scheduled until the return period has passed. To be able to irrigate there must also be an available supply of water. During dry periods in the growing season, this supply may become restricted. This prevents irrigation until the supply is available again, and may occur multiple times during a season for any length of time. The uncertainty surrounding restrictions can promote less than optimal irrigation scheduling. This scheduling can be influenced by the perceived availability of water, leading to preemptive decisions to irrigate. This can lead to excessive water loss and drainage. With a better approach to water supply management and irrigation, nitrate leaching can be reduced and crop production can be optimised. This can reduce water use while maintaining optimal crop growth [6]. Irrigation can be supplemented with water from different storage structures. This includes on-farm storage ponds and off-farm bulk storage structures. These storage options have different costs associated with water access, but are able to improve water supply reliability. On-farm ponds are limited in size and must give up potentially profitable land to be constructed. They require constant refilling from the water supply but can provide a small amount of additional water during a drought. This can help stabilise crop production under varying supply reliability [7]. Off-farm bulk storage (often a dam or lake) is able to supply multiple farms with highly reliable water. This reliability comes at a cost and bulk storage is often expensive to access for an irrigation season.

1.3 Farming System Constraints

To construct an accurate representation of a farming system, various constraints must be used. These constraints are able to be used in a model to help drive daily variables forward. They also restrict model outputs to be physical. Drainage is able to be incorporated into a farming model through the development of a statistical model that represents non-uniform application of water onto the soil. This allows the modelling to compute drainage, and as an extension, the nitrate leaching that must be avoided. This is developed in full in Section 3.4. The daily soil moisture is able to be represented by a simple vertical averaging model. This representation is a constraint that enforces the soil moisture level, S, for the next day's iteration in the model. This can be represented as,

$$\mathrm{S}_{t+1} = \mathrm{S}_t + R_t + I_t - D_t - \mathrm{AET}_t \tag{1.1}$$

and includes terms to account for rainfall, irrigation, drainage, and evapotranspiration (AET) for crop production. To be able to irrigate there must be a water supply available. This supply can come from the water source or additional storage. On any day t the irrigation amount applied must follow,

$$I_t \leq \text{available supply}_t \times \text{irrigation decision}_t \tag{1.2}$$

where the irrigation decision is a binary decision variable controlling the return period of the irrigation. The return period is able to be modelled as the sum of binary irrigation decisions,

$$\sum_{i}^{\text{return period}} \text{irrigation decision}_i \leq 1. \tag{1.3}$$

Additional constraints include physical limits on the amount of water in the soil, field capacity, and the amount of water lost to the crop and atmosphere each day. This is given as,

$$\text{AET}_t = \text{potential loss}_t \times \text{crop type} \times \text{loss factor}_t. \tag{1.4}$$

These constraints contain various nonlinear terms which are needed to model the water loss and drainage. These terms also allow an accurate representation of the water taken up by crops and the atmosphere.

1.4 System Optimisation Framework

With a model that represents the farming system, an optimisation layer is able to be incorporated to investigate different pricing parameter combinations and storage strategies. While a developed model could simulate different parameter combinations, there would exist a lack of mathematical direction. How can a simulated scenario and resulting decisions be considered optimal? With the inclusion of an optimisation strategy, different scenarios can be tested with the backing of an optimisation method to ensure decisions within the season are made optimally. To manage this, an objective function is developed to balance crop profit with water availability, drainage/leaching, and irrigation costs. In its simplest form, crop yield, Y_a, can be combined with drainage and irrigation to form a basic objective function,

$$\begin{aligned} &\max \pi, \\ &\text{s.t.} \quad \pi = c_1 Y_a - c_2 D - c_3 I. \end{aligned} \tag{1.5}$$

This is subject to the constraints outlined in equations 1.1, 1.2, 1.4 and is expanded in full detail in Section 4.2.1. Posing this to maximise economic benefits is in line with current modelling literature, where optimisation is often used from an economic perspective [8]. The framework is ultimately constructed from simple models that have been shown to be effective in similar modelling scenarios [9, 10].

1.5 Thesis Overview

This thesis develops an optimisation framework based on models that allow an accurate representation of a farming system. This includes return periods, irrigator application uniformity, and water loss associated with irrigation events. Additionally, relationships between soil moisture and the ability to support crop growth are incorporated alongside crop-specific parameters. This differs from the current literature as the models implemented capture the dynamics and restrictions of the soil-crop-atmosphere system, combined with physical constraints related to the operation of farming practices. This is built up from optimal irrigation scheduling, with constraints outlined at equations 1.1 and 1.3, where the optimisation goal is to maximise an equation of the form of equation 1.5. This scheduling is part of a larger framework that is developed to investigate water allocation between different water stores. The goal of this framework is to be able to investigate various farming scenarios associated with different water storages, pricing structures, and irrigation strategies. This will provide a more advanced simulation tool that is able to investigate storage options for water and potential investments in drought mitigation.

The optimisation meets specific objective goals related to crop growth, drainage reduction, and limitation of the total amount of irrigation. Specifically, the relationship between the soil-crop-atmosphere system, water supply availability, and different water storage options is utilised. The soil-crop-atmosphere is the interrelated dynamics of how the soil, crop, and atmosphere influence each other and is considered here as a key component in the supply-demand relationship between water source, soil moisture, and crop in the modelling. Further changes from the literature include a modification of the supply-demand relationship. Usually the soil moisture is taken to be the driving demand, here this is moved to the supply side instead. The demand is the replaced by the crop itself. River flow is used to determine the supply source reliability, which is the availability of a water to meet demand. This supply has water take rules applied to it and is used as the primary water supply for irrigation and storage.

Through optimisation, this framework was able to show that relying on the soil moisture as the only store of water is not an appropriate decision. The optimisation is able to show the benefits that storage access brings, and the situations this has the most economic benefit under.

Chapter 2 begins by explaining the relevant soil hydrology and agricultural background related to the project. Key terms used in the modelling are defined, and the background necessary to establish a model with physical significance detailed. The approach to the supply-demand relationship is established with reference to water supplies, storages, and the requirement that demand be satisfied.

In Chapter 3, different modelling equations that have been developed to describe processes in the soil-crop-atmosphere system are introduced. These models are used to describe the necessary interactions between different dynamics to represent real world processes in the modelling set up. In this chapter the data and parameters necessary to model the dynamics are given, and brief examples presented. Here, equations 1.1 and 1.4 are developed.

Chapter 4 presents the modelling framework and its development. Constraints such as 1.3 are established, and the objective function of equation 1.5 is presented in full detail. Relevant agricultural models are modified to conform to modelling decisions, and these are integrated into the framework. The optimisation approach is presented, as well as mathematical modifications to the framework to meet the conditions on the solver methods. This chapter also presents the chosen optimisation solver and the mathematical background on which solutions are found on.

In Chapter 5, initial optimisation results are detailed. These are compared to a state-of-the-art rule based model, irricalc, which determines irrigation scheduling based on soil moisture levels. For these initial simulations there are no additional storages and water supply is determined by the source availability. Shortcomings in the optimisation design are discussed and the effects of changes to the pricing structure on drainage and irrigation are explored. This chapter provides a sanity check that the models are working as expected and are able to represent physical situations.

Full simulation results with appropriate storage options are explored in Chapter 6 for two interesting water supply profiles. Modifications to different pricing structures are also analysed for their effect on a variety of parameters. A comparison between optimal approaches with and without storage access is presented. Results are described in detailed output plots that show water distribution, irrigation decision, and supply availability throughout the season.

In Chapter 7 simulation results are discussed with reference to real world implications. Irrigation strategies with reference to costings, crop demand, and rainfall are summarised. Conclusions are drawn about the viability of strictly using soil moisture as a supply of water and how this is influenced by rainfall. Limitations of the framework and future work are also discussed.

Chapter 2

Soil Hydrology, Agriculture, and Water Allocation

2.1 Soil Hydrology

Hydrology is considered one of the key aspects of managing a water resource [11, 12] and is the study of the occurrence and movement of water. Through modelling, the field of hydrology is able to advise on important issues and practices in the agricultural industry [13, 14]. This information allows informed decisions surrounding water management, irrigation practices, and crop production processes [15, 16, 17, 18].

2.1.1 Soil Profile

The soil profile is the layered soil structure below the ground surface (Figure 2.1). The profile can have large amounts of variability in the composition of the soil, both vertically and spatially. Various factors impact the movement of water through the profile. These factors can include the location or ground slope, the current moisture content, and different soil types. Different soil types have the greatest impact on the movement of water. This is related to porosity and pore size distribution. The size and distribution of pores (space between soil particles), changes the behaviour of water through the soil. The pores allow the movement of water between soil particles and directly influences the hydraulic conductivity of the soil [19]. Soil hydraulic conductivity represents the ability of moisture to move through the pores. A denser soil, such as clay, will have a low hydraulic conductivity, while a sandy soil will allow moisture to move more easily about the profile with a higher conductivity. Different soils present in the profile can change the movement of moisture between the surface layer, subsoil, and substratum.

2.1.2 Soil Moisture

Soil moisture and the infiltration of moisture through the soil profile are key concepts in hydrology [20]. Soil moisture is referred to as the amount of water held in the soil [21]. The amount of moisture present in the soil directly influences crop production [22]. This

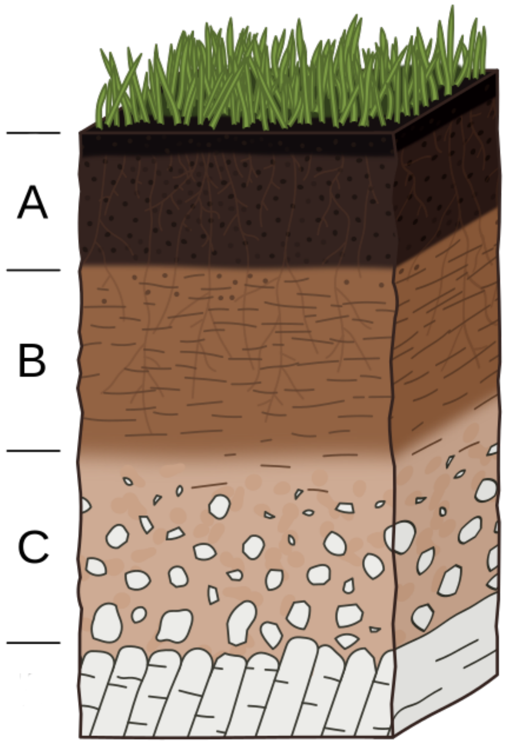

Figure 2.1: Typical soil profile; A – surface soil, B – subsoil, C – substratum.
Picture adapted from Tomáš Kebert under the Creative Commons Attribution-Share Alike 4.0 International license: creativecommons.org/licenses/bysa/4.0/legalcode

places a reliance on soil moisture for crop production and there exists a need to maintain sufficient moisture levels. A higher soil moisture reduces dehydration stress on the crop, and ensures reliable growth [23]. With such importance on understanding the soil moisture levels, various hydrological models have been developed and investigated to better understand these dynamics [24, 25, 26, 27]. These models are used to inform decisions surrounding the use and management of water on farms, as well as irrigation practices. Soil moisture levels are given as a measurement that is an amount of water depth per surface area: cubic meters of water per square meter of land surface. Often this is rescaled and referenced as a one dimensional measurement, millimeters of soil moisture.

2.1.3 Root Zone, Water Holding Capacity, and Plant Available Water

The root zone is the region of the soil profile which the crop roots occupy (Zone A and B in figure 2.1). The root zone is typically assumed to be within one meter of the soil surface, and most hydrological models assume the roots are contained in this upper region of the soil profile [28]. For example, pasture is taken to have a root depth of 600 millimeters. This region is defined as the area for which a crops root system is able to uptake water [29].

The maximum amount of water, that is able to be held in the soil profile, is determined by the **Field capacity** (Fc) of the soil. This is different for different soil types and soil profiles [30, 31]. The amount of water held per millimeter of soil, is defined by the **Water Holding Capacity** (WHC) of the soil. This is a ratio of inferred water depth to soil depth; millimeters of water per millimeter of soil. WHC of different soils is able to be referenced from any standard soil profile table, such as *Soil texture and water* by Irrigation New Zealand [32]. The WHC multiplied by the root depth, gives the **Plant Available Water** (PAW). PAW is the water that is available for plant growth and evaporation loss. More precisely this is the amount of water between the field capacity and permanent wilting point of the crop [33]. The **Permanent Wilting Point** (PWP) of a crop is defined as the soil moisture level at which the plant is unable to maintain transpiration. Soil moisture at or below the PWP induces maximum stress on the plant, and may result in irreversible loss of production. At the permanent wilting point there still exists moisture in the soil profile, but it is tightly bound to the soil structure. This makes it unavailable for transpiration and hence the soil is referred to having zero PAW.

2.2 Soil-Crop-Atmosphere System

Relevant Soil-Crop-Atmosphere science for this thesis is related to the transpiration of water for crop growth and the loss of water from soil through environmental factors. These processes are related to the development of leaf area and plant height. These two factors impact how water is captured and lost from the soil. This development for different crops is accounted for in models with the use of a crop factor. Crop development changes over a season and this is reflected in the crop factor. This section will also consider evaporation to the atmosphere, irrigation processes and the loss of excess water through drainage.

2.2.1 Evapotranspiration

A better understanding of how moisture is removed form the soil can lead to more effective management of irrigation scheduling. Crop transpiration is a critical component for the growth of plants, and has been subject of focus for many researchers [34]. The process of water evaporation to the atmosphere is a well understood process [35]. Water uptake by a crop through transpiration has also been extensively modelled [36]. There exists developed models that combine both these processes in a model of **Evapotranspiration** (ET). The **Food and Agriculture Organization of the United Nations** (FAO) publishes the standard model for these calculations. The currently accepted model and standard reference calculation method is published in the Irrigation and Drainage Paper No. 56 [37]. Further research on the calibration of this model has been undertaken [38], as well as its evaluation under limited data [39]. The model accounts for the current climate conditions and calculates a maximum **Potential Evapotranspiration** (PET). This is calculated from the radiation, heat flux, air temperature, wind speed, and vapour pressure at the location of interest.

Evapotranspiration has been investigated for effects on soil moisture and crop production [33]. In field situations the **Actual Evapotranspiration** (AET) is less than the reference PET calculated by the model. AET is a function of the currently available soil moisture level [34]. As soil moisture decreases, it becomes harder to extract water from the soil structure. Various developments have been made to estimate the AET from reference calculations of the FAO model. These are usually based on a linear dependence to the soil moisture. In 1974 Minhas, Parikh, and Srinivasan [22] developed a model to estimate the AET with an appropriate scaling parameter that was based on the exponential decay of soil moisture. Minhas et al. proposed a model based on the idea that water extraction should be proportional to the amount of water present in the soil profile. The benefit of this model is its ability to be calibrated to different soil profiles with the inclusion of a scaling parameter. Furthermore, this scaling parameter was shown to allow the model to match the results of previous studies that showed a relationship between ET and soil moisture. This research allows a better understanding of how AET is influenced by soil moisture.

2.2.2 Crop Factors

Different crops and pasture often have different growth conditions and growing seasons. An empirically derived coefficient called the crop factor accounts for these differences. The crop factor scales the reference ET calculated by the FAO method to the appropriate crop. This allows the PET to be defined for the crop of interest, relative to the reference ET. The crop factor can be dynamic due to changes in the plants structure, such as height and leaf area, as the crop grows throughout the season [37]. This is observed as an increasing crop factor from sowing to maximum growth. The crop factor is then usually flat until harvest time, where it will begin to decline. By contrast, pasture has a crop factor that is relatively flat across the year due to having a maintained growth height. Importantly there is the assumption that the crop is not overly stressed, so that the average growth is represented by the crop factor. These factors can be derived from Lysimeter data or reference tables. A Lysimeter is a device used to measure soil moisture levels. The amount of water lost from the soil due to ET can then be calculated from measured levels. These can be used on farms to provide a more accurate representation of water lost for a particular soil type or crop. It is possible to use a time series of crop factors through the season to get an estimation of the PET. Such an estimation allows informed decisions surrounding crop management and yield estimations.

2.2.3 Irrigation

An irrigation event is used to add water the soil profile. There are a number of irrigation methods including travelling/sprayline, micro drip, and centre pivots. Each system has its own benefits and drawbacks but the end goal of increasing the soil moisture is the same. These systems also have an associated frequency of irrigation. This can be due to requiring manual input to move locations in a paddock, or the time taken for a complete rotation of the field. This frequency of irrigation is related to the return period of the system. The return period is a physical constraint that defines the maximum frequency of irrigation events. It restricts the availability to irrigate and can require additional planning to optimally use. Another parameter that can be modified is the trigger point of an irrigation event. This can determine an irrigation event based on a particular value of the soil moisture level. The use of a trigger point is the method current rule based modelling systems utilise to schedule irrigation events.

One of the most common irrigators is the centre pivot [40]. Centre pivots have a large initial cost but require very little maintenance or farmer input, and are able to provide a very uniform application of water. The applied water can be varied and one centre pivot can irrigate multiple paddocks. The return period of a centre pivot is defined by the time taken to return to the start of the rotation. This can be between three and five days depending on the speed that is set. Due to the widespread use of centre pivots these are the irrigators of focus for the modelling framework. This defines the appropriate return periods and application amounts, however a modification to these parameters allow different irrigation systems to be incorporated in the model if needed.

To improve irrigation under limited supplied adopted practice is deficit irrigation. Deficit irrigation keeps the available water to the crop sub maximal. This allows for rainfall to be taken up by the soil without inducing excess drainage. This also uses less water, and is an appropriate irrigation method for restricted water supplies. Irrigation is applied such that it reduces moisture deficits that would otherwise limit crop productivity [41]. This practice requires careful management to ensure crop production is optimal [42].

2.2.4 Drainage

Drainage through the soil profile is a key component of the soil-crop-atmosphere system that must be managed and closely monitored. The drainage of interest is the loss of water through the bottom of the soil profile. Other forms of drainage include runoff, where water is not taken up by the soil, and specific drainage systems that are designed to remove water from crop zones. If the soil is at field capacity, any extra water is unable to be freely taken up. Instead water is lost from the soil profile and replaced by the new application. Drainage has both a negative impact on the environment and a cost to profits. Excessive water loss due to drainage causes nitrogen/nitrate leeching which causes negative impacts on the surrounding environment [13]. Furthermore, excessive drainage is a loss of applied water which has potentially been bought at a cost to the farm. These factors give a negative impact to drainage which needs to be mitigated. This has prompted modelling analysis of how leaching can be monitored and reduced [43, 4, 44]. Research has considered the effects of different crops and soil types on the amount of leaching that occurs [45, 46, 47].

In 1993 D. Scholefield et al. [1] investigated the relationship between grassland nitrate leaching and the drainage that occurs. Scholefield et al. utilised drainage testing to measure the amount of nitrate loss, subject to differences in crop, soil type, and fertilizer application. The amount of leaching was observed to be insensitive to the rainfall intensity. An implication of this is rainfall and the likes of irrigation events themselves do not directly contribute to the nitrate losses at the time of application. The primary concerns instead lies with the total amount of drainage the soil sustains. The study concludes by stating the nitrate loss into waterways observed was in great excess of the environmental drinking water standard and the exact amount of leaching was found to be influenced by a number of factors. This implied obtaining a representable model of nitrate loss could prove difficult. In 1994 Gaines and Gaines [45] considered the effects of different soil types on the amount of nitrate leaching. They showed that less dense soils, like that of a sandy composition, had greater leaching. This was directly related to the ability of the soil to drain water. Again, a higher drainage amount was correlated with greater amounts of leaching. It is widely accepted that leaching is of great concern, and although differences in soil type and crop change the exact amount of leaching, there is general consensus that leaching must be reduced. Farmers must be encouraged to adopt better irrigation practices, either through the type of irrigation used, or the general practices surrounding the use of fertiliser and irrigation itself [48, 49]. It is these scenarios that the framework will be able to investigate, and offer optimal strategy advice on.

2.3 Water Supply and Demand

There are various water supplies available to farms. These have associated costs and benefits that must be considered when irrigating or investigating the viability of a particular irrigation method. Water schemes are able to provide pressurised water either from lakes, dams or rivers. Water supplies usually have an associated reliability related to the available flow rate that can be taken. With less reliable supplies irrigation practices tend to use excess water to compensate for potential supply restrictions. This often promotes a "just in case" mentality, applying more water than is required and further contributing to drainage and the resultant leaching. A more optimal use of the water would instead use sufficient irrigation to meet the demand before a restriction [5]. However this can be difficult to get correct, along with the inability to irrigate every day, and a high level of expertise and experience is required. Crop quality can be another driving factor used make irrigation decisions. Different crops can require access to water at different times to produce desirable traits in the final product [50]. With high value crops, excess irrigation may also be applied to reduce the chance of a reduction in yield, or the death of the crop. Again, this can increase nitrate leaching and drainage as well as wasting water.

These water supplies are able to provide more reliable water from large bulk stores off farm, usually at an increased cost. Farmers can also have their own on-farm storage of water in large ponds. These two stores of water can be used to help decrease the severity of water restrictions [7, 51]. They can allow additional irrigation when the main supply of water is restricted. The bulk off farm storage and the on-farm storage ponds are considered to be two available storages of water. A third storage is considered to be the soil moisture itself. This represents a philosophical shift in the traditional supply-demand relationship. The soil moisture is usually considered as the demand of this relationship, where the supply is provided by water storage and the supply scheme. This demand traditionally comes from the need to manage the soil moisture level in order maintain crop production.

2.3.1 Demand

With the decision to move the soil moisture to the water supply side, the demand is replaced by the AET of an associated crop. This puts emphasis on meeting the ET demand of a particular crop with a supply of water from the soil moisture. It removes the need to maximise the soil moisture in the traditional sense, and allows for better flexibility with irrigation management. The daily water demand is the AET and this is able to be calculated from the PET, crop factor, and current soil moisture levels.

2.3.2 Supply

With the shift of soil moisture to the supply side, the available water supplies become the soil moisture, on-farm storage ponds, bulk storage ponds (off farm storage), and source water provided by a water scheme and individual supplies. Source water is the main supply of water to farms. This provides pressurised water from a water take at some cost. There are different cost structures to accessing this water. Commonly a farmer will pay a fixed rate to have access to a flow rate delivered to the farm. This flow rate is the water volume per second the farm has available for storage and irrigation, usually in cubic meters per second. Furthermore the contracted flow rate does not guarantee the reliability of the supply. Sources will often be heavily restricted during drought, either reducing the flow available or shutting off the supply completely. Because of this there exists uncertainty around restrictions during an irrigation season. This can promote excessive irrigation, regardless of soil moisture levels, under the presumption water could be restricted any day.

An external storage of water a farmer can access to mitigate water restrictions, is large off farm bulk water storages. These bulk storages often require a farmer to buy into the storage to have the option of using the water. The benefit to these bulk systems is that the water is more reliable than water provided by a source. This reliability is usually at a significantly higher cost than source water. The benefit to the bulk storage is that it has the best chance of mitigating an extended drought, and ensuring crop production remains consistent.

Farmers are able to build on-farm storage ponds to hold water for irrigation when it is available. These storages can have a high investment cost to set up as well as on-going costs to maintain. Farms that rely on source water for irrigation are limited in the freedom to irrigate as needed due to the associated restrictions. On-farm water storage is able to mitigate some of these restrictions. This requires careful planning of refilling and irrigation, to maximise benefits.

The soil moisture is a final store of water. The soil moisture is the direct influence in meeting the demand, and this can be refilled through irrigation and rainfall. The soil moisture is limited in its storage volume and needs to be monitored in order to ensure adequate water levels are maintained. In recent years this monitoring ability has improved significantly. Farms are able to access real time measurements through the use of techniques like neutron scattering, time domain reflectometry, and gamma attenuation [21]. This allows irrigation strategies to be modified as required, and the demand met.

Chapter 3

Agricultural Modelling

Modelling has long been used in the Agricultural industry, from determining soil moisture levels [52] to showing a relationship between irrigation and crop yield [53]. In recent years these models have been able to be critically analysed through computer simulation and the validity numerically assessed. Models that have been tested to represent the soil hydrology and soil-crop-atmosphere system are investigated, and used to develop the modelling framework.

Models that relate crop water demands and yield are of particular interest as they can inform on potential production loss due to insufficient water through crop water productivity models [54]. Evapotranspiration is also an area of constant refinement, as this can be a direct influence of irrigation practices and water storage strategies [12]. Through understanding ET there exist a number of models that aim to predict soil moisture levels [20]. The soil moisture balance can be important for modelling changes in strategies of both crop management and also irrigation tactics. With the latest developments in soil moisture monitoring techniques, like real time telemetry measurements, these models are able to have accurate up to date inputs helping their success. Soil moisture models also allow insight into drainage losses and are important in assessing both possible nitrate leeching and irrigation water loss [3]. The techniques and models discussed here represent the accepted and widely used equations in the agricultural field, backed up by real world testing and simulation at Aqualinc.

3.1 Crop Water Production Functions

Understanding the potential yield of a crop is valuable information for a farmer, especially if crop response to irrigation is hard to measure. The growth of a crop is related to the amount of transpiration it is able to sustain. If there is little moisture available in the soil, evapotranspiration can be greatly reduced, impacting the production yield for the season. To better understand how crop production and ET are related, **Crop Water Production Functions** (CWPF) have been developed. The key feature of these models is the relation between evapotranspiration and yield. Different CWPF models that have been developed incorporate different methods of relating these terms by considering the deficit from maximum potential or using various empirically derived coefficients. These models typically fall into two sub categories for determining how reduced evapotranspiration affects the production yield. This is either as a additive model or a model with multiplicative consequences.

In 1982 G. Tsakiris [55] developed a method for establishing and relating crop sensitivity to irrigation scheduling. Tsakiris noted that there exits two forms the CWPFs can take, either multiplicative or additive. Tsakiris establishes the functional form of these two types, and this is an idea that is built on in subsequent research. These two types of models relate to the compounding of water consumption between different growth stages. In particular, the additive type models accept that a lack of evapotranspiration in growth stages will potentially have severe impacts on the resulting yield. This is in contrast to multiplicative type models, where there will be no yield if in any stage there is no evapotranspiration. This study also shows how Water Production Functions can be used as a simple representation of complicated interdependence between growth stages. Tsakiris notes that the investigated sensitivity parameters would be suitable for use in irrigation systems. This was related to the idea that a flexible irrigation schedule can be optimised for water use when the available irrigation supply is restricted.

Building on the two types of Production Functions Tsakiris developed, Henry et al. [56] and Hayrettin Kuşçu [57] explore different CWPFs that make use of either the multiplicative or the additive response. Both papers have a focus on a particular crop. Kuşçu assesses the model performance on predicting the yield of a tomato crop, while Henry et al. considers maize. Both studies consider different irrigation schedules throughout the growth stages of the crop. Concluding remarks for both studies mention the results of the models were different. Care is needed to appropriately use the sensitivity coefficients in different locations through the use of calibration. These papers also recorded different rankings for model metrics and highlighted that a CWPF should be chosen based on the particular modelling scenario. Henry et al. furthers this statement by saying sensitivity coefficients need to be accurately applied to the appropriate growth stages. The exact growth stages of a crop cannot always be determined simply from a time of planting, due to the complex biological processes, as original noted by Tsakiris. The results of research by Tsakiris, Henry et al. and Kuşçu show that unless there is explicit reason to favour a particular CWPF over another, a simple model choice will work best for the modelling developed in this thesis. A commonly used CWPF is that of Jensen [58] which relates the water stress to yield in a multiplicative manner. This is given as,

$$\frac{Y_a}{Y_m} = \prod_{i=1}^{n} \left(\frac{\mathrm{ET}}{\mathrm{PET}}\right)_i^{\lambda_i}. \tag{3.1}$$

Y_a represents the actual yield with Y_m the maximum potential yield when soil moisture is optimal. ET and PET represent the actual and maximum potential evapotranspiration respectively. For this particular CWPF, i represents the current growth stage of the crop, λ_i the appropriate sensitivity of the crop to water stress in the particular growth stage, and n the number of growth stages under consideration. The inclusion of the maximum yield and potential evapotranspiration allow climate and soil differences to be accounted for. λ_i is strictly dependent on the particular crop, and it must be determined empirically. This was the sensitivity parameter of focus for Tsakiris in 1982. In 2016 Manijeh Varzi [53] reviews two different CWPFs. The two models are the 1968 model by Jensen and another by Doorenbos and Kassam. The model of Doorenbos and Kassam was originally presented in the FAO Irrigation and Drainage paper number 33 [59], and more recently revisited in the FAO number 66 publication [60]. The model by Doorenbos and Kassam is referenced as the FAO66 model and is a linear function given as,

$$1 - \frac{Y_a}{Y_m} = K_y \left(1 - \frac{\mathrm{ET}}{\mathrm{PET}}\right). \tag{3.2}$$

Here K_y represents the sensitivity index. Varzi shows this model can also be transformed to represent different growth stage sensitivities as an additive model. This takes the form of,

$$1 - \frac{Y_a}{Y_m} = \sum_{i=1}^{n} K_{y_i} \left(1 - \frac{\mathrm{ET}_i}{\mathrm{PET}_i}\right). \tag{3.3}$$

Varzi highlights that equation 3.3 overestimates the yield production due to the additive nature. The review hypothesise this is due to the models potential to still estimate an ET amount in later growth stages, even if the crop has died early. In contrast, a review by Zwart and Bastiaanssen [54] observed an underestimation in the yield using the same model. The more recent review by Varzi goes on to mention additional research that has observed correct yield predictions by the FAO66 CWPF (equation 3.2). A point of interest from the 2016 review is that there exists strong experimental evidence that evapotranspiration and yield are linearly related over the growing season. This further strengthens the idea that simple CWPF are beneficial to use, as these relate the ET and yield in a scaled linear manner. Varzi states in the review this is especially true for forage or pasture type crops, and is a result shown in the publications that observed correct yield estimates. These findings suggest that the CWPF as purposed by the FAO number 66 paper is a suitable choice for modelling the crop yield in the framework.

3.2 Evapotranspiration Models

Evapotranspiration plays the largest role in the regular removal of moisture from the soil profile. One of the first attempts of this was carried out by Howard Penman [61] in 1948. Penman combined an energy balance equation with a mass transfer method to derive an equation that computed the evaporation of an open water surface. This combined equation used standard climatological records of sunshine, temperature, humidity, and wind speed. The SI units form of the original Penman equation is given as,

$$E = \frac{\Delta R_n + 6.43\gamma(1 + 0.536 U_2)(e_s - e_a)}{\lambda_v(\Delta + \gamma)}, \tag{3.4}$$

E	evaporation rate [mm day^{-1}],
R_n	net irradiance [MJ m^{-2} day^{-1}],
e_s	saturation vapour pressure [kPa],
e_a	actual vapour pressure [kPa],
Δ	slope vapour pressure curve [kPa °C^{-1}],
γ	psychrometric constant [kPa °C^{-1}],
λ_v	latent heat of vaporisation [MJ kg^{-1}].

This formula was derived from the results of experiments that measured how water was evaporated from large bodies of water. This work was later expanded by John Monteith who applied this idea to vegetation. In a 1965 paper titled *Evaporation and Environment* [62], Monteith modified the original equation by Penman. This modification included terms that account for a crop covered soil, as well as water loss through crop transpiration. This is now known as the Penman-Monteith equation and is given as,

$$\mathrm{ET} = \frac{\Delta(R_n - G) + \rho_a\, c_p \frac{e_s - e_a}{r_a}}{\left(\Delta + \gamma\left(1 + \frac{r_s}{r_a}\right)\right) L_v}. \tag{3.5}$$

The main change to equation 3.4 is the inclusion of the term G which represents the soil heat flux. There is also the addition of the specific heat of the air, c_p as well as the inclusion of aerodynamic resistance terms to account for the crop's protrusion from the soil surface. The aerodynamic terms represent the (bulk) surface resistance, r_s and the crops aerodynamic resistance, r_a. Other terms are defined the same as the original Penman equation with the exception that ET is now the combined evapotranspiration. ET is transformed by the volumetric latent heat of vaporisation, L_v to be in terms of mm of water per second.

3.2.1 FAO Penman-Monteith Evapotranspiration Model

The need for a standardised method of computing the evapotranspiration was recognised by the FAO. Equation 3.5 has be modified to be valid ET method for a described reference crop. With different crop types having varying morphology that can contribute to a significantly different evapotranspiration rate, a hypothetical reference crop was chosen by the FAO. This was taken to have a height of 0.12m, a fixed surface resistance of 70sm^{-1}, and a heat reflection coefficient (albedo) of 0.23. Such parameters resemble that of a green grass of uniform height that is both actively growing and adequately watered. In this sense, adequately watered is to indicate the crop is able to grow without water induced stress under optimum soil conditions. This definition gives a reference evapotranspiration rate, $\mathrm{ET}_{\mathrm{ref}}$ and allows the calculation of ET of a desired crop by use of crop factors and climate data. This method is referred to as the FAO Penman-Monteith equation and is presented in the FAO Irrigation and Drainage paper number 56 [37],

$$\mathrm{ET}_{\mathrm{ref}} = \frac{0.408\Delta(R_n - G) + \gamma \frac{900}{T+273} u_2 (e_s - e_a)}{\Delta + \gamma(1 + 0.34 u_2)}, \tag{3.6}$$

ET_{ref} reference evapotranspiration [mm day^{-1}],
R_n net radiation at crop surface [MJ m^{-2} day^{-1}],
G soil heat flux density [MJ m^{-2} day^{-1}],
T mean daily air temperature at 2m height [$^\circ$C],
u_2 wind speed at 2m height [m s^{-1}],
e_s saturation vapour pressure [kPa],
e_a actual vapour pressure [kPa],
Δ slope vapour pressure curve [kPa $^\circ$C^{-1}],
γ psychrometric constant [kPa $^\circ$C^{-1}].

Despite the age of the model, there has been various verifications of the method [63] and it remains the standard for calculating evapotranspiration. The model presented in equation 3.6 is the evapotranspiration model that will be employed in the modelling framework. This decision is also due to the availability of climate data for the locations of interest. Reference evapotranspiration values are available as a daily time series as calculated by the FAO Penman-Monteith method. This greatly simplifies the computation required and can be used as a direct data input to the model.

3.2.2 Accurate Evapotranspiration

With a representation of the maximum PET, it is important to consider how this represents the actual water loss from ET. At reduced soil moisture levels, water becomes harder to extract. Therefore the AET that is lost from the soil profile is represented by a function dependent on the current soil moisture level. As explored above at equation 3.2, there is a relationship between crop yield and evapotranspiration that can be estimated through the use of linear models. This is further seen in the simplest model that aims to relate the actual and potential ET. This model stipulates that the amount of AET is directly proportional to the ratio of soil moisture level to field capacity. This is represented as,

$$\mathrm{AET} = \mathrm{PET}\frac{\text{Soil Moisture}}{\text{Field Capacity}}. \tag{3.7}$$

In 2013 Zhao et al. [34] consider various relationships for determining an actual evapotranspiration amount. Zhao et al. show there is actually a non linear relationship between actual and potential evapotranspiration, removing the direct proportionality. It is further shown that different soils have different nonlinear relationships. This implies soil moisture and actual evapotranspiration should be described by an equation that can be scaled to different soil types. In 1974 Minhas et al. [22] purposed such a model, along with similar nonlinear relationships. Minhas et al. also describe other purposed methods of evapotranspiration reduction and is able to relate these other methods with a single equation. The model for the reduction in potential evapotranspiration is given as,

$$f(\mathrm{S}, r) = \frac{1 - \exp(-r\,\mathrm{S})}{1 - 2\exp(-r\,F_c) + \exp(-r\,\mathrm{S})}, \tag{3.8}$$

S current soil moisture level [mm],
F_c Field capacity [mm],
r scalable soil parameter.

The scalable parameter r, can be based off empirical knowledge of the soil type or can be fitted to previous soil moisture data as Minhas et al. achieves. With reference evapotranspiration defined in equation 3.6 equation 3.8 can be used, along with an appropriate crop factor, to calculate the actual evapotranspiration as,

$$\begin{aligned} \text{AET} &= K_c\, f(\text{S}, r)\, \text{ET}_{\text{ref}} && (3.9) \\ &= f(\text{S}, r)\, \text{PET}. && (3.10) \end{aligned}$$

3.3 Models of Soil Moisture Dynamics

Optimal irrigation strategies and an understanding of the evapotranspiration cannot be used effectively if the level of the soil moisture is not known. As described in Chapter 2 there exist techniques to asses the real time soil moisture levels. This allows modification of existing irrigation strategies, but to simulate long term soil moisture dynamics, models are required. Models of water movement in the soil are commonly referred to as Mass Balance models or Soil Moisture Deficit models. They model the movement of water through the soil though crop demand, water application, and losses. These models need to be able to work from the real time data collected from on-site monitoring sites, as well take a set of initial conditions to describe the dynamics of the soil moisture at some future time. This allows different irrigation strategies to be tested, as well as any needed changes to existing infrastructure. For this purpose various hydrological models have been developed and investigated to better understand the soil moisture dynamics. Soil moisture models are now briefly detailed. These are to be used as the main constraint to drive daily updates in the framework. This requires a compatible model that fits the goals of the framework.

In 2002 Guswa [20] considered two particular models of soil moisture dynamics. These models were a mass balance approach, also referred to as a bucket model, and a model based on Richards equation. For the Richards model the soil column is taken as a vertical cross section of the soil profile with the moisture dynamics described by a modification of Richards equation. This modification added two terms to represent the evaporation and transpiration and is given as,

$$\frac{\partial \theta}{\partial t} = \frac{\partial}{\partial z}\left[K(\theta)\left(\frac{\partial h}{\partial z} + 1\right)\right] - e' - \sigma', \tag{3.11}$$

θ	soil moisture content,
$\frac{\partial \theta}{\partial t}$	rate of change of soil moisture,
z	soil depth,
K	unsaturated hydraulic conductivity,
e'	evaporation,
σ'	transpiration.

In this model, irrigation and rainfall is interpreted as a boundary condition at the surface of the soil column. Additionally, the water infill rate, either from irrigation or rainfall, is accounted by a boundary flux term. Depending on infill rates, some water may run off and drain if the rate exceeds the soil uptake capacity. Equation 3.11 allows for a better representation of a varying soil profile. Conversely, the mass balance model investigated by Guswa is a vertically averaged model. This bucket model considers a volume in volume out approach over the root zone. This can be slightly modified with the definition of root depth and water holding capacity to represent soil moisture dynamics in terms of mm of water. This is given as,

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = I(\theta, t) - L(\theta) - T(\theta) - E(\theta), \tag{3.12}$$

θ	soil moisture content,
I	water infiltration,
L	water loss,
E	evaporation,
T	transpiration.

For exact details the models presented by equations 3.11 and 3.12 the original paper *Models of soil moisture dynamics in ecohydrology: A comparative study* [20] by Guswa is recommended. This bucket approach gives a simple representation of the addition and loss of water from the soil profile. Perhaps surprisingly, Guswa concludes with the remarks that the more complicated model does not significantly out perform the bucket approach. Such a conclusion was also noted in 1983 by Calder et al [52] where more detailed calculations related to the movement of water in a soil profile, were not necessarily found to improve predictions. Guswa details the differences in the models to represent how the infiltration, runoff, and uptake of evapotranspiration are accounted for in each model. Of these differences, the variability of soil type in the profile can be considered to make up the largest difference. Therefore it can be expected the more detailed model will perform better when the soil profile is more complicated. However, this of course can only work where the soil profile is explicitly known. Soil heterogeneity makes this complicated and so it assumed there is a simple soil profile in the framework simulations, this is in line with the modelling approach by Peri et al. [64] who looked at optimal water infiltration depths. If less variability is assumed with the soil type, then the simple model is able to produce valid estimates of the dynamics. Furthermore, applying the previously determined equations for the evapotranspiration and splitting the infiltration term into rainfall and irrigation events, a model can be attained for the soil moisture dynamics that fits with the goals of this thesis.

3.3.1 Mass Balance Model Development

In 1974 Minhas et al. [22] presents the loss of moisture from the soil as $\frac{\mathrm{d}\theta}{\mathrm{d}t} = -A(t)f(\theta)$, where $A(t)$ represents the potential evapotranspiration and $f(\theta)$ is the reduction function defined in equation 3.8. From this equation Minhas et al. comments on how irrigation and rainfall can be incorporated, and presents a process to represent the potential evapotranspiration as a constant, by subdividing the continuous interval. Taking the idea presented by Minhas et al. and subdividing the time interval, a similar expression to that of equation 3.12 can be achieved, but with a discrete time interpretation on a daily time scale.

This model is the bases for the rule based modelling undertaken by Aqualinc. This is the **irricalc** model and can be used for estimating soil moisture dynamics along with simulating and modelling irrigation scheduling, drainage, and evapotranspiration dynamics. The $t+1$ day's **soil moisture** S can then be expressed as,

$$\mathrm{S}_{t+1} = \mathrm{S}_t + R_t + I_t - D_t - \mathrm{AET}_t, \tag{3.13}$$

S_t	soil moisture content [mm],
I_t	irrigation applied [mm],
D_t	drainage [mm],
AET_t	actual evapotranspiration [mm],
R_t	rainfall [mm].

This model is of the vertically averaged type and operates on the assumption any rainfall, R_t or irrigation, I_t increases the soil moisture content by the amount applied. The drainage component, D_t is non zero if the irrigation and rainfall cause the soil moisture to exceed the field capacity. The soil is taken to be freely draining and the drainage to occur in the same time period as the irrigation and rainfall events. The actual evapotranspiration term is calculated from equation 3.10 using a reference evapotranspiration from the FAO Penman-Monteith model. The rainfall component of the model is the rainfall received over the time period from the previous soil moisture calculation $t-1$ to the current calculation t. This is given as a mm amount, and it is assumed that all rainfall enters the soil profile. Rainfall, like reference ET data, is given by local climate data and is used as time series inputs.

3.3.2 Irricalc Irrigation System Model

The irrigation term in the soil moisture model 3.13 represents the amount of water applied to the soil profile. This is different to the total amount of water exiting the irrigator. Water is lost through inefficiencies in the irrigation process, both in the transportation of water to the irrigator, and fluctuations in the soil moisture levels. These fluctuations in the soil moisture can cause water to run off outside the irrigating area, in effect wasting water. To account for this there is various irrigation efficiency calculations that consider the water change in the soil profile and the total amount of water applied.

A rule based irrigation model is one of the methods used by irricalc to model irrigation water use with soil moisture dynamics. This uses rules that rely on current soil moisture levels and can be modified to suit particular situations. These modifications include the application depth of the irrigation (amount of water applied), as well as the frequency of irrigation events. The application depth is the average water depth applied to the soil profile to raise the soil moisture level. This is an important detail as the water applied is not necessarily the amount of increase in the soil moisture level. This value is usually given as a millimeter measurement. The frequency of irrigation events can be determined by the current soil moisture level, or specified on a rule basis, eg every third day. There is a physical limit to this and the frequency of irrigation is given by the return period of the system.

3.4 Irrigation Distributions

The distribution of water leaving an irrigator is not necessarily a uniform application to the soil surface. As the soil is also not homogeneous, there needs to be a way of accounting for different infills of water at different locations of irrigation. This is achieved through the use of a statistical model that accounts for these discrepancies through the use of uniformity coefficients and a random distribution of soil moisture deficits. Previous research has used a normal distribution for modelling the infiltration of water to the soil profile, allowing the use Christiansen's uniformity coefficient to describe the irrigation application [64]. This allows an efficiency of application to be defined which can calculate how much water has been lost in the irrigation application.

3.4.1 Water Infiltration

To determine the amount of water retained, and thus the amount of irrigation water that drains away, a statistical distribution of water application is employed. This is a function of the soil water deficit at the time of irrigation, the average application depth applied, and the uniformity of the irrigation application. The uniformity of the application is specified by Christiansen's uniformity coefficient [65]. Christiansen's uniformity coefficient is defined as,

$$\mathrm{CU} = 100\left(1 - \frac{\sum\limits_{i=1}^{n}|X_i - \mu|}{\sum\limits_{i=1}^{n} X_i}\right). \tag{3.14}$$

Where n is the number of depth measurements taken, X_i the measured application depth and μ the average application depth. This type of equation is used to determine the application uniformity in a controlled test. Once determined, the value can be used for actual irrigation events. This provides an idea on the distribution of irrigated water applied to the surface of the soil. There is an assumption that the spatial distribution of the applied water follows a normal distribution when the uniformity is greater that 70% [66]. For usual irrigation practices the uniformity of application is usually $\sim\ 80\%$. From this assumption a model to determine the retained moisture after irrigation can be constructed using a normal distribution, and the ideas behind the measurement process of Christiansen's uniformity coefficient.

In 1979 Gideon Peri et al. [64] investigates the optimal depth of irrigation through the idea of a system optimal depth. Peri et al. aim to approach this with a cost perspective, where they note there is an average depth of water applied that will bring about maximum net income. They also note that this will require accountability for the excess water in irrigation. They detail an approach for this with a focus the operating conditions of irrigation events. A method for calculating this accountability by means of an efficiency of application is incorporated here. This is more in line with the framework goals and is described in the research thesis by John Bright [67]. A reconstruction of this method is presented and detailed below.

3.4.2 Efficiency and Adequacy of Irrigation Area

A probability density function (PDF) of the infiltrated water is taken to be normally distributed about the irrigated water amount, the average application. Using a similar idea as equation 3.14 an amount of water at point i is considered. This represents the amount of water that has infiltrated into the soil profile and is retained, u_i. Assuming the infiltrated water will be normally distributed about the average application,

$$f(u_i) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2}\frac{(u_i - \bar{u})^2}{\sigma^2}\right). \tag{3.15}$$

$f(u_i)$ represents the PDF for an amount of water infiltrated, u_i, at the ith point. In this formulation $\bar{u}$ is the average application of water which represents the set amount that is desired to be put on. Also to be defined is $u_r = F_c - \mathrm{S}$ which represents the deficit in soil moisture level. The standard deviation, σ, is able to be related to the uniformity coefficient as was shown by Peri et al. and is given as,

$$\frac{\bar{u}}{\sigma\sqrt{2}} = \frac{1}{1 - \mathrm{CU}\sqrt{\pi}}. \tag{3.16}$$

Possible water infiltration patterns fall into two cases. If the infiltrated water is not sufficient in meeting the deficit at point i, the deficit is greater than the applied water, and all the water should be retained since $u_i < u_r$. The change in soil moisture from this can then be given by the expectation value of infiltrated water. Restricting this to physical values for the infiltration, between zero and the deficit u_r, this is given as,

$$Q_{u_r} = \int_0^{u_r} u_i\, f(u_i)\mathrm{d}u_i. \tag{3.17}$$

The other case represents the deficit being over filled at the particular point i. This is equivalent to $u_i \geq u_r$. This has the deficit met by the infiltrated amount of water, as well as some excess water loss. The excess water that exceeds the deficit is taken to have drained away from the soil profile. Considering the probability that the applied water will exceed the deficit, a formula for the amount of water in excess to the deficit can be constructed. This can be achieved through the cumulative distribution function of $f(u_i)$. This is represented as,

$$F(u_i) = \int_{-\infty}^{u_i} f(x)\mathrm{d}x. \tag{3.18}$$

More specifically there is interest in the probability the applied irrigation exceeds the deficit, $P(u_i > u_r)$ which will be given by $(1 - F(u_r))$. Putting these two cases together allows the change in soil moisture, ΔQ ,to be written when accounting for water amounts. This equation can be used to find efficiencies of different irrigation events.

$$\Delta Q = Q_{u_r} + (1 - F(u_r))\, u_r. \tag{3.19}$$

The measure of efficiency is the change in soil moisture compared to the actual applied amount of water. The average application amount is taken to be the amount of applied

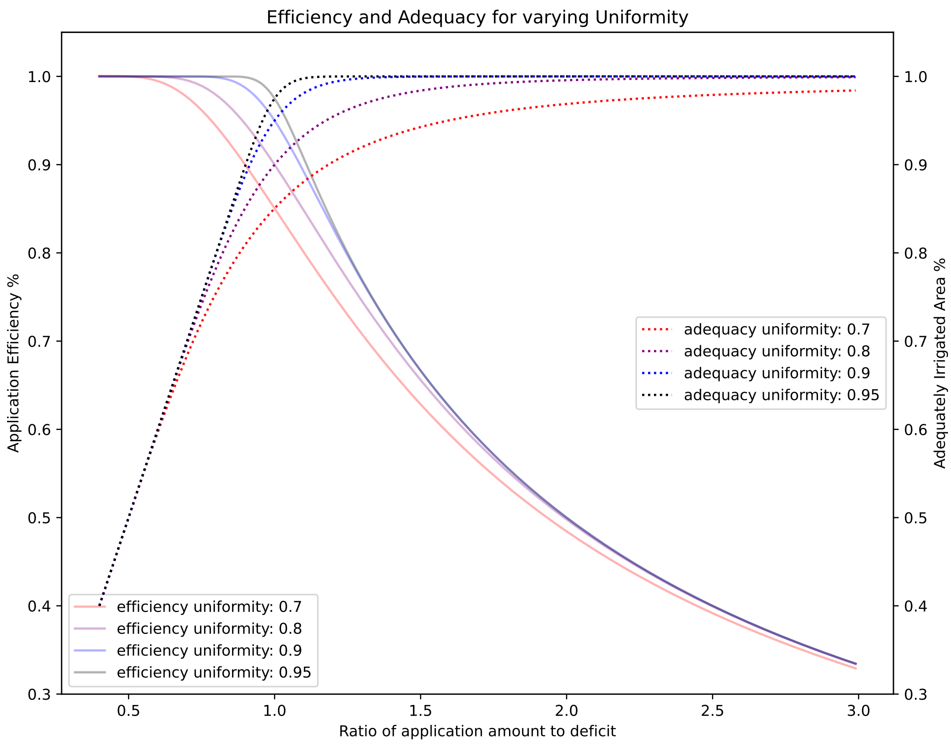

Figure 3.1: Relationship between application efficiency, application uniformity and application depth. As the adequately irrigated area increases, the overall efficiency decreases. This is due to more water being required for less net gain due to losses.

water. This is then compared to the change in soil moisture levels as given by equation 3.19. This gives $e_p = \frac{\Delta Q}{\bar{u}}$, the irrigation efficiency. This efficiency is tied to the uniformity of application and can help define excess water loss. The inclusion of the efficiency allows for a more accurate tracking of the applied water.

Through a similar method, the adequacy of irrigation can be calculated. This is a measure of how much of the water deficit is satisfied during an irrigation event. This can be calculated as $a_i = \frac{\Delta Q}{u_r}$. This can inform on the role the efficiency plays on the change in soil moisture level across the irrigated area. While it can be important to meet an adequately irrigated area, there becomes a trade off with the overall efficiency of irrigation. Intuitively this makes sense, as more water is required to reach a totally irrigated area, based on the decay of the efficiency. This trade off is shown at Figure 3.1. From the graph it can be seen that to achieve a total irrigated area at a uniformity of 0.8, twice the deficit would need to be applied. This would result in an efficiency of 50%, or a water loss equal to half the applied amount. The adequacy of irrigation is chosen to not be explicitly included in the modelling framework, since deficit irrigation strategies rarely meet the total deficit under optimal situations. The efficiency of application is explicitly included, as this directly influences the amount of water loss.

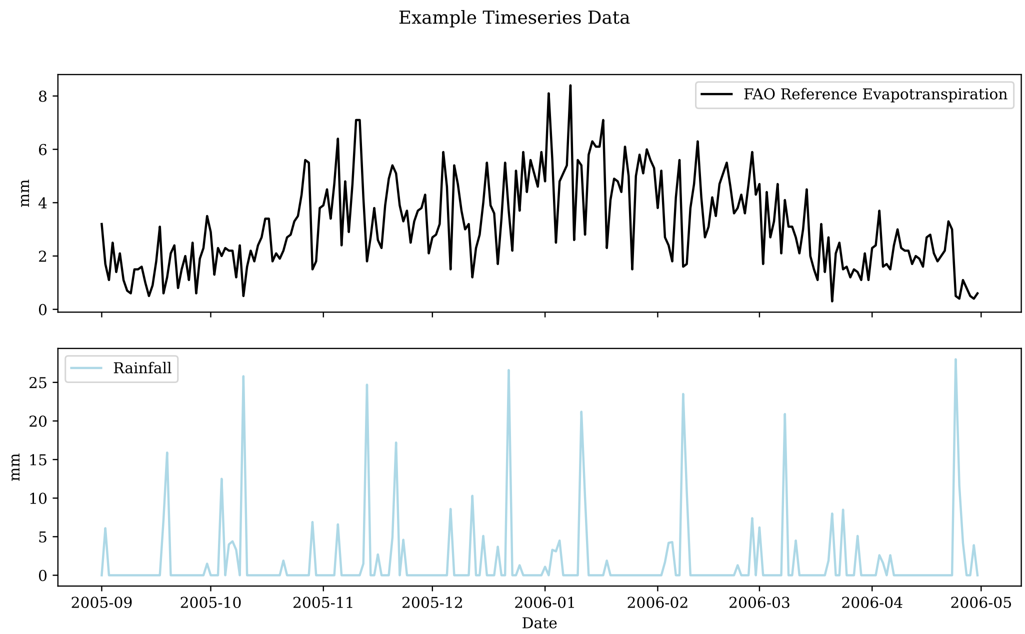

Figure 3.2: Example of the time series input for the reference evapotranspiration, as computed by the Penman-Monteith FAO model, and the climate rainfall. Both inputs are given in millimeters which can be directly used in model computations.

3.5 Data

There are various datasets used in determining the irrigation events in the rule based model outlined above. The data consists of local climate data such as rainfall and predetermined reference evapotranspiration values as time series. Water availability data is also of importance as this gives an allowed take from a water source on a particular day, and is how restrictions on water are incorporated into the modelling framework. For the rule based approach, this availability time series determines if irrigation is possible, however modelling can also be carried out with an assumed 100% availability. The climate and reference ET data is provided by Aqualinc and is composed of historic time series for a farm in Canterbury, New Zealand. An example of the 2005/06 season is shown at Figure 3.2

There is a time series needed to account for the crop factor. As mentioned in Chapter 2 this data can come from different sources, either as a standardised crop profile across the growing season, or derived from Lysimeter data. As there is interest in modelling pasture, a time series crop factor that represents the growing cycles of a pasture is used. This is shown at Figure 3.3 and is provided by Aqualinc.

Unlike other crops pasture is maintained year round and has a relatively flat crop profile across the year. The demand that must be met is calculated from this crop factor and the reference evapotranspiration data. This potential evapotranspiration is the upper limit to the loss of water from the soil. Figure 3.4 shows this for the 2005/06 season. Of particular interest is the load over the summer period, around the start of December to the end of February. The evapotranspiration load during this time is highest and it is also the period where water restrictions become an issue. There must be sufficient water applied to meet

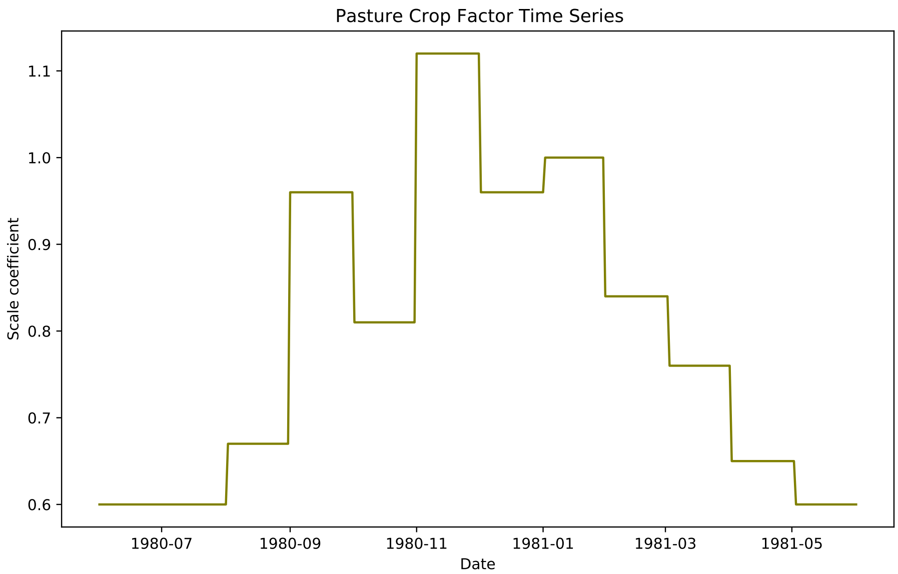

Figure 3.3: Crop factor time series for pasture. This allows the calculation of potential evapotranspiration for pasture from the reference ET calculated by the FAO Penman-Monteith equation.

the load, or the crop will suffer. In the lower part of Figure 3.4 this higher load is shown, as well as the average load across this period.

River flow time series data is used to determine the availability of water supply. These time series represent the flow rate of a river at a particular measurement point. It is this flow that determines if restrictions should be put in place on the water supply. In essence, this represents the allowed take of water from the water supply, shown in Figure 3.5. Also shown in the same figure is the scaled take allowed by a resource consent. Different reliabilities for the supply are shown in the bottom plot. River flow data is again provided by Aqualinc with flow data recorded from the Waimakariri river located in Canterbury, New Zealand. These plots show the cut off values for when the water supply will become restricted. These can be converted into a flow availability by normalising by the largest available flow, which is also the consented limit. This gives an of allowed water take for a day in the season. This is done on the assumption that a farmer will have a constrained physical capacity to take water from a supply. The percent of allowed flow multiplied by the system capacity gives the available water for irrigation and storage. These flow rates can be further manipulated by assuming a particular reliability. If the flow is assumed to be available with 95% reliability the upper plot of Figure 3.6 is produced for the 2005/6 season again. The lower plot at Figure 3.6 assumes a lower reliability of 75%. This reliability takes effect on the flow rate of the river, ie 95% allows water to still be taken when at a lower river flow rate.

Figure 3.6 shows the importance of understanding flow rate availability. The 95% plot has water almost readily available for use on demand, while reducing that rate to 75% produces large discrepancies in the available water, requiring careful planning.

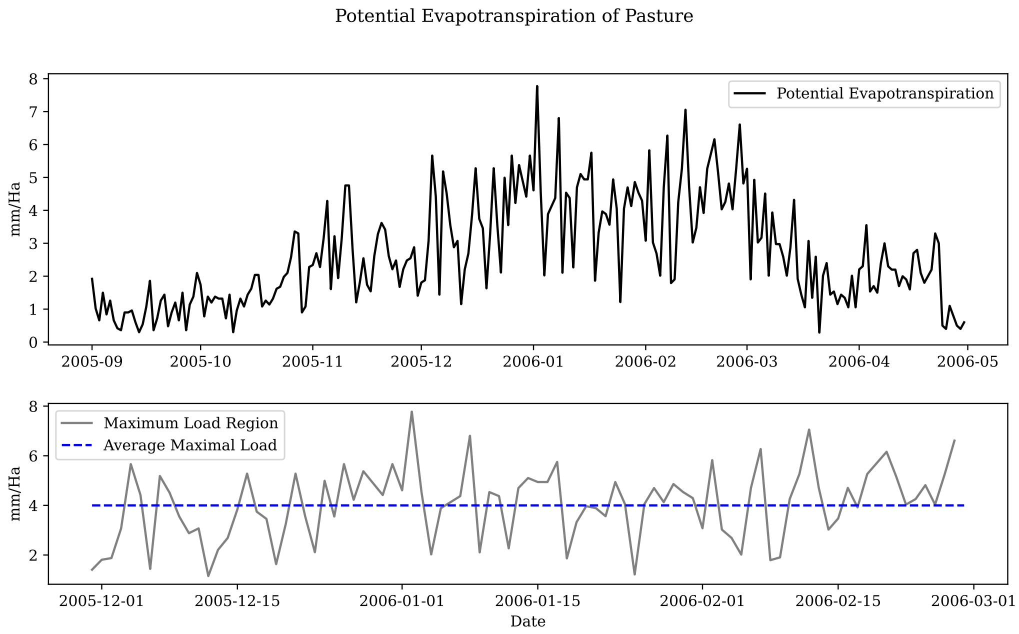

Figure 3.4: Example of the potential evapotranspiration for pasture. The upper plot represents the scaled evapotranspiration over the 2005/06 season. The lower plot is a closer look over the high demand section of the season. Also plotted is the average PET of 4 mm per day over this higher demand period in blue. This is the demand that must be met.

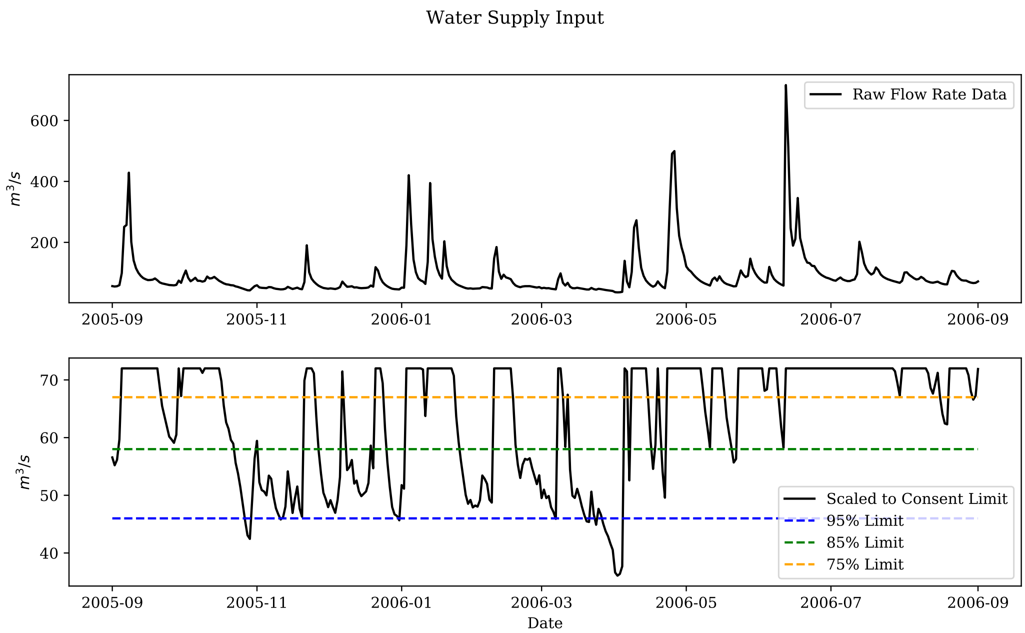

Figure 3.5: Example of river flow data for the 2005/06 season. In the upper plot the raw flow rate data is shown. This is measured data from a reference location on the river. The lower plot shows this flow rate as an allowed take. Also plotted is the cut off levels for assumed reliabilities. When the flow rate drops below these cut offs there is a zero allowable take. Lower reliabilities are cut off first.

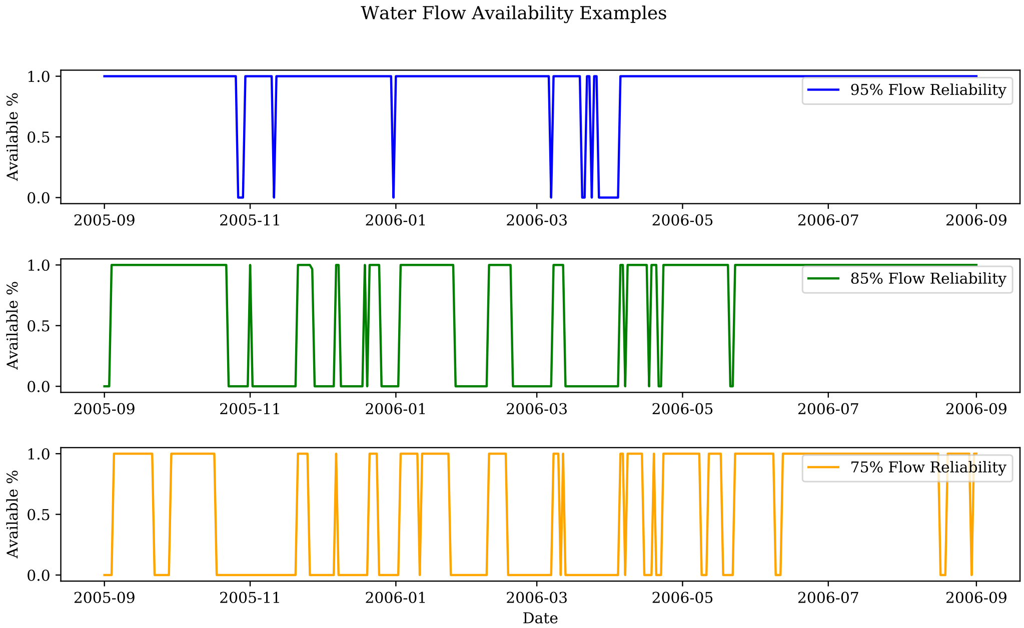

Figure 3.6: Example of how the reliabilities translate to supply availability. The highest reliability has the supply available more often, while a lower reliability suffers frequent restrictions.

There are various other parameters that need to be considered for use in the models outlined in this chapter. These parameters ultimately tie the mathematical functions to real world outputs. These parameters are provided by Aqualinc and are a representation of previous work or a range of values that have been observed through data monitoring and testing. These values are summarised in Table 3.1.

Table 3.1: Modelling parameter values and suggested ranges provided by Aqualinc.

Plant Available Water	80mm - 120mm
Return Period	3 days
Irrigation System Limit	15mm/Ha
Pump Limit	20-50mm/Ha
Uniformity Coefficient	0.8
Potential Yield	16000Kg/Ha
Crop Sensitivity	1.0
Soil r Value	10
Yield Profit (c_1)	0.3 - 0.5 $/Kg
Water Price (c_2, c_4, c_5)	1-8 $/mm/Ha
Drainage Cost (c_3)	1-10 $/mm/Ha

Chapter 4

Optimisation Methods and Framework

With the relevant modelling background required to represent a farming system established, this is now incorporated into the optimisation modelling framework. These models are used to explore how to distribute water between possible stores, while maintaining either economic (profit) or environmental (drainage reduction) goals. This is detailed to define the optimisation program, the constraints that must be accounted for, and an approach towards solutions with a particular choice of optimisation solver.

4.1 Previous Modelling Approaches

Models have been developed that look at the allocation of water resources [68, 69, 70]. These models aim to determine the best distribution of water resources from supplies that maximise potential profits between all demands. They are utilised in the operations of large reservoirs that provide the supply water. At the other end of the supply-demand chain, there has been extensive research into the optimal scheduling of irrigation for crops. In 1981 Bras and Cordova [71] developed a model to determined the optimal intraseasonal water allocation. This was applied to deficit irrigation to minimise the total water use. The presented methodology utilised stochastic processes and dynamic programming to unite multiple rules concerning the availability of water, crop sensitivity, and rainfall. This model was one of the first examples to incorporate multiple factors/equations in the soil-crop interaction. The results of the study presented a methodology to determine the optimal irrigation schedule, given an amount of water allocation, subject to crop demand. In 1986 research by Bright [67] applies similar ideas to Bras and Cordova to determine the optimal irrigation of a crop, under a varied water supply.

More recently, the optimisation technique of Non-Linear Programming has been utilised further to determine irrigation scheduling under a variety of conditions. Garg and Dadhich [8] utilise Non-Linear Programming to optimise deficit irrigation. The model results show it is possible to maximise profit of crop production, with a limited resources, via deficit irrigation. In a similar manner two years later, Linker et al. [72] use Non-Linear

Programming to determine the maximal profit, given a particular water allocation. In 2017 a more advanced revisit of a stochastic model was attempted by Anvari et al [9]. A variety of modelling dynamics were considered in this model, including crop details, evapotranspiration of the crop, and water scarcity. This research draws focus back to maximising crop growth under varying restrictions of supply. In 2021 a similar suite of dynamic equations to Anvari et al. was implemented by Galioto and Adriano [10] to combine economic modelling with binary irrigation scheduling. Galioto and Adriano applied simple soil moisture models to determine an irrigation scheduling that maximised overall profit. Galioto and Adriano incorporated different factors that could determine an economically beneficial irrigation decision, such as the labour cost involved in a irrigation event. These factors influenced whether an irrigation event was scheduled, and the soil moisture dynamics were simulated as a result of the decision. This paper in particular showed that using simple soil moisture models to economically maximise irrigation decisions is a viable modelling approach.

4.2 Framework Approach

The focus was to develop a framework that could investigate an optimal distribution of water between different water stores. Of interest was the soil moisture, actual ET and inclusion of on-farm and bulk storage supplies. An approach was used that was able to take time series data and initial parameters, to return a time series of irrigation events and storage use strategies. More specifically, this model output shows how the soil moisture dynamics develop over time, and how additionally available water can be taken advantage of. A framework that only operates on source supply was initially constructed. This allowed a comparison to the current state of the art rule based modelling, irricalc. This was to provide a sanity check that the optimisation was working as expected. This could then be expanded to simulate additional storage access.

The optimisation model is written using an existing optimisation library in *Python*. This library has the ability to define modelling objective functions, constraints, and decision variables. This is the *Python Optimization Modelling Objects* (`pyomo`) library [73]. `Pyomo` allows the creation of an optimisation program that can be passed to a variety of optimisation solvers. It is this flexibility to pick and choose solvers, alongside the ability to incorporate dynamic variables, that makes `pyomo` a compelling choice. An objective function was created to maximise the sum of particular decision variables. The problem was then constrained with a variety of physical equations and practical restrictions. The program along with input data defined a model object that could be passed to solvers.

4.2.1 Objective Function

An object function for the optimisation was chosen to incorporate profit for crop yield and costs associated with drainage and irrigation. Through the use of parameters this object function was able to be formed in units of dollars per hectare. This placed a tangible relation on model outputs, objective function results, and real world practices.

The profit component of the objective function is given by the CWPF as discussed in Chapter 3. This is modified to represent the total ET, instead of total growing stages. It was assumed that the sensitivity of the crop to water stress remains unchanged across the season. Daily variable values are summed over the season of interest. This gives the definition that $\mathrm{AET} = \sum_t^T \mathrm{AET}_t = \sum_t^T f(\mathrm{S}_t, r)\,\mathrm{PET}_t$, where f is the ET reduction function defined in equation 3.8. The total potential ET for the season is given by PET, and so the actual yield of the season is represented in the framework as,

$$Y_a = Y_0 \left(1 - K_s \left(1 - \frac{\mathrm{AET}}{\mathrm{PET}}\right)\right). \tag{4.1}$$

Y_0 is the potential yield and K_s represents the sensitivity of the crop to water stress. This results in actual yield being in units of kilograms. The actual yield is used in the objective function by associating a profit parameter with this value. Defining c_1 to be the profit of yield in $\$/Kg$ allows the AET to be a direct component of the objective function. Cost terms included in the objective function are for the total irrigation and drainage. Each of these has an associated cost in $/mm and are summed over all days in a respective season. Cost parameters associated with the drainage and irrigation were c_2 and c_3 respectively. The basic objective function is therefore represented as,

$$\pi = c_1 Y_a - c_2 D - c_3 I. \tag{4.2}$$

Where $D = \sum_{t=1}^T D_t$ for drainage, and $I = \sum_{t=1}^T I_t$ for the amount of irrigation. The optimisation maximises the value of π. This follows the transition of shifting the soil moisture to the supply side, as the focus in the objective function is placed on the ET, the demand of the crop. This particular focus was to allow the optimisation to better follow the water demands of the crop. For example towards the end of the season when the PET drops off, high soil moisture levels are not as essential, and irrigation could be reduced as needed.

4.2.2 Moisture Dynamics

A large component of the modelling is the ability to understand how the decision to irrigate impacts the soil moisture level, and is related to soil moisture dynamics such as ET. This was incorporated in the framework through the use of variables to represent the soil moisture. This allowed for the use of mass balance rules to calculate soil moisture levels for the next day, and enforce the related dynamics. This was achieved by defining a function that updated the next day's soil moisture level based on the current soil moisture and the net change in moisture for the day. If $t + 1 < T$ then S_{t+1} is given as,

$$\mathrm{S}_{t+1} = \mathrm{S}_t + R_t + I_t - D_t - \mathrm{AET}_t. \tag{4.3}$$

If $t = T$ the season is finished and the update function terminated. This was implemented in `pyomo` as a constraint, where the soil moisture levels are represented by the mass balance equation. This defines the main update constraint in the program and propagates irrigation decisions to subsequent days. This function defines the soil moisture on any given day, and allows the ET to be calculated. Similarly, ET was be represented in the model as a variable. This was defined by equation 3.10, and tied into both the soil moisture level and soil moisture update function. The AET represented by this variable could be calculated with the reference ET data and the reduction function,

$$\begin{aligned} \mathrm{AET}_t &= K_c\, f(\mathrm{S}_t, r)\ (\mathrm{ET}_{\mathrm{ref}})_t && (4.4)\\ &= f(\mathrm{S}_t, r)\, \mathrm{PET}_t. && (4.5) \end{aligned}$$

4.2.3 Irrigation and the Return Period

The role of irrigation is to refill the soil moisture level, and is therefore one of the important decision variables used in the model. The optimisation is able to determine the amount of water that should be applied as well, as the scheduling of irrigation events. The associated cost parameter, c_3 acts to regulate the total amount of irrigation applied at any time, through the objective function. Furthermore, there are additional constraints on the irrigation amount as restricted by both the available supply and max flow rate of the system. For example, a return period of three days implies the maximum amount of water that can be applied is 15 mm. This is included as an upper bound constraint of the irrigation amount.

The inability to irrigate everyday, as defined by the return period, is an important structure that must be accurately included. A decision to irrigate affects the next possible irrigation event, and must be propagated through the model. This is achieved by the use of a binary decision variable to allow irrigation. This binary variable allows a constraint that the sum of binary decision variables over the number of days in the return period, must be equal to or less than one. This restricts the model to a single irrigation event in the duration of the return period. This represents a constraint and is given as,

$$\sum_{i}^{\text{return period}} \text{irrigation decision}_i \leq 1. \tag{4.6}$$

The decision variable supersedes the amount of irrigation to apply through the use of constraints. When a decision to not irrigate is made, the irrigation amount is required to be zero. The final piece of the irrigation component is the supply availability and irrigation system capacity. With the available water supply given as an availability time series, this can be further scaled to define the allowed water take for irrigation. This scaling is achieved though the use of the irrigation system capacity parameter, and modifies the available supply to be an available amount of water for the irrigation system. This is implemented as a constraint, where the optimal irrigation amount is constrained by the water available to the system, and the decision to irrigate.

$$I_t \leq \text{available supply}_t \times \text{irrigation decision}_t. \tag{4.7}$$

4.2.4 Drainage and Efficiency

The constraints defined above give a good representation on the dynamics in the optimisation framework. The final piece of the soil moisture balance is the drainage component. There is interest in understanding the drainage as this affects the potential for nitrate leching. Ultimately the amount of drainage should be minimised, and so the amount of drainage is incorporated as a variable in the optimisation model. Instead of constraining the drainage explicitly, it is left as an implicit product of the other constraints in conjunction with the field capacity. The field capacity is defined as a constraint in the model where the daily total soil moisture must not exceed this value. Since the other variables in the soil moisture dynamic are explicitly constrained, the drainage is therefore the regulator of the field capacity constraint. This is in essence how the rule based modelling operates. The rule approach considers the total soil moisture level, and sets the drainage equal to the amount it exceeds the field capacity by. The benefit to this implicit definition is the removal of explicitly checking if the soil moisture balance exceeds the field capacity. The regulatory nature of this implicit definition is possible due to the cost associated with drainage in the objective function. The optimisation solution will look to reduce the amount of drainage to maximise the associated profit.

To account for the inefficiencies associated with irrigating near field capacity, a new variable, W, is introduced. Chapter 3 showed there can be an amount of water wasted in an irrigation event. This is dependent on the soil moisture level and the irrigation amount. Wasted water is considered to be drainage as it is lost from the soil when applying a particular irrigation amount. The efficiency is represented by defining a new variable, eff. The change to the soil moisture levels can be calculated by incorporating this efficiency measure on the irrigated amount in the soil moisture dynamics. This defines the water loss variable, which can then be included in the soil moisture balance. This maintains the cost of irrigation in the objective function, as the full amount of irrigation water is still bought at a cost to the farm. Formally the water loss is included as a drainage cost in the object function and is defined as,

$$W = \sum_{t}^{T} (1 - \text{eff}_t)\, I_t. \tag{4.8}$$

The eff on some day t, is given by a function dependent on S_t and I_t.

One of the drawbacks of the `pyomo` framework is the inability to define all *Python* objects in the optimisation program. This is due to the optimisation program being complied to a static file that is passed to solvers. The integral calculations for efficiency could not be directly implemented. To achieve a similar efficiency curve it was assumed that irrigation events will have a fixed uniformity coefficient of 0.8. An approximation to the curve was then found through curve fitting methods. Different approximations to this curve will impact the soil moisture dynamics. However, the efficiency has only a significant effect when the application amount is close to, or exceeds the deficit. With this in mind a simple representation was achieved through a linear fit. The important dynamics that were able to be captured with this function were close to the application amount being equal to the deficit, namely $I_t \sim F_c - \mathrm{S}_t$. At much higher application to deficit ratios, there is significant water loss and this should be discouraged, both in the modelling and in real world practices. To maintain this idea, and avoid potential issues with the efficiency approximation tending to zero faster than the actual efficiency curve, the irrigation applied was restricted to at most be twice the deficit. Mathematical this is of the form $I_t \leq 2(F_c - \mathrm{S}_t)$.

This linear function was specifically fit to the efficiency curve between an application to deficit ratio of 0.9 and 1.5. This produced a fit that was very similar to the actual efficiency curve when the application amount is close to the deficit. This fitted function is used as a constraint in the framework to represent the efficiency of an irrigation event. This then defines the amount of water loss. The fitted linear function for efficiency can be described as,

$$\mathrm{eff}_t = a - b\frac{I_t}{F_c - \mathrm{S}_t}, \tag{4.9}$$

where a and b represent positive fitted coefficients. As the applied amount increases, the approximated efficiency is poorer than the true efficiency. Since most irrigation events occur at a far lower deficit ratio, this was acceptable. At these lower ratio amounts it was possible to enforce a 100% efficiency by constraining the efficiency to the modified approximation shown in Figure 4.1. Also shown is the original efficiency curve, and the fitted linear approximation. The vertical line section represents the limited irrigation ratio discussed above.

4.2.5 Inclusion of Water Storages

The model was extended to include the use of on-farm storage ponds and the bulk off-farm storage. This modified the available water for irrigation events and allowed for irrigation when the main source was unavailable. Further changes included modifying the objective function to include associated costs with the additional water supplies. To investigate how these storages are best used, additional irrigation variables were introduced that were related to specific water supplies. Constraints that manage appropriate restrictions relating to return periods and system capacities were also designed.

The bulk storage is treated as a water supply with 100% reliability. This was regulated in the objective function by the much higher costing per mm of water used. As a consequence, the bulk storage should only be utilised when there is a net benefit. Such a case may be an extended drought, since the on-farm storage has a limited capacity that must be refilled. The addition of these extra irrigation methods required additional restrictions to adhere to the system capacity. This was enforced by a fixed total irrigation per day, irrespective of source.

The daily water level of the on-farm storage can be modeled the same as the soil moisture balance. **On-farm storage irrigation**, $I^{(\mathrm{of})}$, will decrease the amount of water available in the storage. The available on-farm storage volume is defined as,

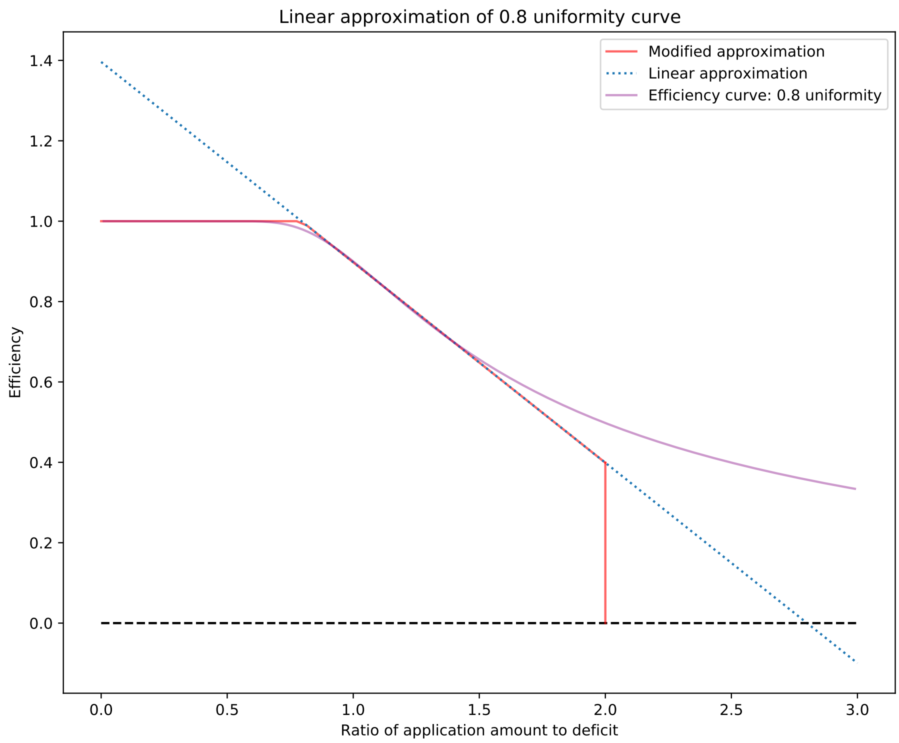

Figure 4.1: Efficiency curves for an assumed uniformity of 80%. In blue is the fitted linear function, while the red represents the constrained efficiency calculation. The vertical section of the modified curve represents the limiting ratio relationship, $I_t \leq 2(F_c - S_t)$.

$$V_{t+1} = V_t + \text{store}_t + \text{store}_t^{(\text{b})} - I_t^{(\text{of})}, \tag{4.10}$$

V_t	on-farm storage available water [mm],
store_t	water stored from source supply [mm],
$\text{store}_t^{(\text{b})}$	water stored from bulk storage supply [mm],
$I_t^{(\text{of})}$	on-farm storage irrigation [mm].

Multiple costs are associated with the storing and irrigating of water from the on-farm storage. These were again added to the objective function to regulate the amount of water used. Both of these storages were incorporated into the objective function with their respective cost variables. The **bulk storage irrigation**, $I_t^{(\text{b})}$ cost is included as,

$$I^{(\text{b})} = \sum_t^T I_t^{(\text{b})}. \tag{4.11}$$

The inclusion of the on-farm storage is twofold, both the storing and the irrigation incur costs. Furthermore, water can be stored from either the bulk storage or source water - although the practical use here is to strictly irrigate from the bulk storage, and not use it to refill the on-farm storage. Costs to store the water are assumed to be the same as directly irrigating, while an additional cost must be incurred to further irrigate from the storage. This additional cost can be attributed to the cost of redistributing the water once stored in an unpressurised storage, and as a cost associated with maintaining the on-farm storage. This is represented as,

$$\text{on-farm storage costs} = c_3 \sum_t^T \text{store}_t \; + \; c_4 \sum_t^T \text{store}_t^{(\text{b})} \; + \; c_5 \sum_t^T I_t^{(\text{of})}. \tag{4.12}$$

With the inclusion of the additional water stores, the objective function is completed as,

$$\begin{aligned} \pi = c_1 Y_0 &\left(1 - K_s \left(1 - \frac{\text{AET}}{\text{PET}}\right)\right) \\ &- c_2 \left(D + W\right) - c_3 \left(I + \text{Store}\right) \\ &- c4 \left(I^{(\text{b})} + \text{Store}^{(\text{b})}\right) \; - \; c_5 I^{(\text{of})}, \end{aligned} \tag{4.13}$$

c_1, c_2, c_3, c_4	cost/profit parameters,
Y_0	maximum potential yield [Kg/Ha],
K_s	crop sensitivity,
AET	total actual ET [mm],
PET	total maximum potential ET [mm],
D	total drainage [mm],
W	total water loss [mm],
I	total source supply irrigation [mm],
$I^{(\text{b})}$	total bulk storage irrigation [mm],
$I^{(\text{of})}$	total on-farm storage irrigation [mm],
Store	total water stored from source supply [mm],
$\text{Store}^{(\text{b})}$	total water stored from bulk storage supply [mm].

4.3 Optimisation Methods

With the models used in the optimisation program set out above at equations 4.3, 4.5, 4.6, 4.7, 4.8, 4.9, 4.10, 4.13, the particular type of optimisation program being constructed must be classified. This influenced the choice of solver, as simple linear programs can be solved without advanced decomposition methods. The inclusion of the binary irrigation decision creates a mixed integer program. There is also a mixture of linear and non-linear constraints. These modelling choices defined the program to be a **Mixed-Integer Nonlinear Program** (MINLP). Optimisation techniques are able to exploit certain qualities of programs. One such advantage is if the program represents a minimisation of a convex objective function and associated convex constraints. The model outlined here is the negation of a minimisation. If convexity with respect to the negation can be shown, this can be taken advantage of with the choice of solver in order to obtain a solution. With the `pyomo` framework there is the ability to choose from a number of solvers. Since this problem is a MINLP, this choice becomes limited. The decision was to use the open source **Mixed-Integer Nonlinear Decomposition Toolbox for** `Pyomo` (MindtPy) [74]. This solver is built on the same framework as `pyomo` and was a natural choice given the restricted subset of available solvers. MindtPy is a highly customisable MINLP solver that allows different **Nonlinear Programming** (NLP) and **Mixed Integer Linear Programming** (MILP) solvers to be used together, in order to find a solution. Solver techniques available to be used in MindtPy include an implementation of **Outer Approximation** (OA) [75] and **Extended Cutting Plane** (ECP) [76] method.

4.3.1 MindtPy Decomposition Method

The Outer-Approximation algorithm was chosen to be used in MindtPy to find solutions to the optimisation program. OA in MindtPy is the implementation of the algorithm by Duran and Grossmann [75]. This choice was due to fast convergence shown by OA in numerical testing by Bernal et al. [74]. OA has also been successfully used in the modelling of water related processes, such as the placement of control valves in a water network [77] and groundwater management [78, 79]. The OA decomposition method takes advantage of an assumed convexity of problems. Although MINLP are inherently non convex due to integer variables, they can be classified by a relaxation on the integer variable restriction. If the relaxation is convex, the problem is said to be convex, where all feasible combinations of the integer variables gives rise to a convex problem.

With the additional complexity MINLP have, solutions require a mixture of techniques. The OA decomposition method iterates through two alternating sub problems, in order to reach optimality. Each sub problem provides an update on either the upper or lower bound of the objective function. Outer Approximation relies on some assumptions about the form of the MINLP. These assumptions demand non linear functions are convex and continuously differentiable for both the objective function and constraints, and linear constraints must define a nonempty compact set. For all feasible integer combinations there must exit continuous variables for which the problem is feasible, and a constraint qualification holds [80]. In the original 1987 paper, Duran and Grossmann use Slaters condition as the constraint qualification to prove convergence of the OA method.

An overview of the Outer Approximation decomposition in Mindtpy is now discussed. For a more complete reference with additional details such as convergence and optimality, as well as implemented methods to deal with infeasibility of the OA method, the original paper by Duran and Grossmann - *An Outer-Approximation Algorithm For A Class Of Mixed-Integer Nonlinear Programs* [75] is recommended. The OA method is detailed for an appropriate MINLP involving an objective function to be minimised. It is this minimisation that requires the convexity of nonlinear functions. For a maximisation problem it is possible to reformulated it as the negative of the equivalent minimisation.

Define a simple MINLP as,

$$\begin{aligned}
&\min_{\mathbf{x},\,\mathbf{y}} \; f(\mathbf{x},\mathbf{y}),\\
&\text{s.t.} \quad g_j(\mathbf{x},\mathbf{y}) \leq 0 \quad \forall j = 1,\ldots l,\\
&\qquad \mathbf{Ax} + \mathbf{By} \leq \mathbf{b},\\
&\qquad \mathbf{x} \in \mathbb{R}^n,\ \mathbf{y} \in \mathbb{Z}^m,\\
&\qquad f,\, g_1,\ldots,g_l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R},
\end{aligned} \tag{4.14}$$

where f and g_j are convex and continuously differentiable. There are different ways Mindtpy initialises the OA method. One such method is obtaining trial solutions to the integer variables by a relaxation on the integer variable restriction. Alternatively an initial trial solution can be directly passed into the solver. The trial solution obtained through the relaxation provides a valid lower bound on the convex minimisation problem 4.14 as well as an initial linearisation point. With this initial trial solution the problem can be linearised into a Mixed Integer Linear Program (MILP). This linearisation is possible due to both the assumed convexity and the required differentiability. Through a first order linearisation, the program 4.14 can be approximated with an appropriate linear objective function as

$$\begin{aligned}
&\min_{\mathbf{x},\,\mathbf{y},\mu} \; \mu,\\
&\text{s.t.} \quad f(\mathbf{x}^i,\mathrm{y}^i) + \nabla f(\mathbf{x}^i,\mathbf{y}^i)^{\mathrm{T}} \begin{bmatrix}\mathbf{x}-\mathbf{x}^i\\ \mathbf{y}-\mathbf{y}^i\end{bmatrix} \leq \mu \quad \forall i = 1,\ldots,k,\\
&\qquad g_j(\mathbf{x}^i,\mathbf{y}^i) + \nabla g_j(\mathbf{x}^i,\mathbf{y}^i)^{\mathrm{T}} \begin{bmatrix}\mathbf{x}-\mathbf{x}^i\\ \mathbf{y}-\mathbf{y}^i\end{bmatrix} \leq 0 \quad \forall j = 1,\ldots l, \forall i = 1,\ldots,k,\\
&\qquad \mathbf{Ax} + \mathbf{By} \leq \mathbf{b},\\
&\qquad \mathbf{x} \in \mathbb{R}^n,\ \mathbf{y} \in \mathbb{Z}^m,\ \mu \in \mathbb{R}.
\end{aligned} \tag{4.15}$$

∇ represents the Jacobian matrix evaluated at given $\mathbf{x}^i \in \mathbb{R}^n$ for the respective functions. This represents the linearisation of 4.14 at an iteration k of the OA method. This linearisation can be solved with Mixed Integer methods such as branch and bound enumeration [81]. Such solutions will determine a new set of trial solutions $\mathbf{y}^{k+1}$ for the next iteration of the method. Due to the assumed convexity, a solution to this problem will under approximate the true objective function. This is due to the linearisation at any point lying below the graph of the original function. This can provide an updated lower bound to the original problem formulation. With this new trial solution, the integer values are able to be fixed and the resulting Non Linear Problem solved. The NLP to be solved is of the form:

$$\begin{aligned}
\min_{\mathbf{x}} \; & f(\mathbf{x}, \mathbf{y}^{k+1}), \\
\text{s.t.} \quad & g_j(\mathbf{x}, \mathbf{y}^{k+1}) \leq 0 \quad \forall j = 1, \dots l, \\
& \mathbf{Ax} + \mathbf{By}^{k+1} \leq \mathbf{b}, \\
& \mathbf{x} \in \mathbb{R}^n.
\end{aligned} \tag{4.16}$$

This can be solved with any standard NLP technique. The solution then provides a new combination of $\mathbf{x}^{k+1}$, as well as an upper bound on the objective function. If the newly found upper bound and the previous lower bound are within a solution tolerance, the problem is solved to optimality. If they are outside the tolerance, the process is repeated from solving the MILP with the new combination of points, $\{\mathbf{x}^{k+1}, \mathbf{y}^{k+1}\}$ used for the linearisation.

4.3.2 Convexity of Nonlinear Functions

With some simplification, the MINLP developed in this thesis can be represented as,

$$\begin{aligned}
\max_{\mathbf{x}} \; & \pi(\mathbf{x}), \\
\text{s.t.} \quad & g_j(\mathbf{x}) \leq 0 \quad \forall j = 1, \dots l, \\
& \mathbf{Ax} + \mathbf{By} \leq \mathbf{b}, \\
& \mathbf{Cx} + \mathbf{Dy} = \mathbf{o}, \\
& \mathbf{x} \in \mathbb{R}^n, \; \mathbf{y} \in \{0, 1\}^n,
\end{aligned} \tag{4.17}$$

where $-\pi, -g_1, \dots, -g_l : \mathbb{R}^n \to \mathbb{R}$ are required to be convex differentiable functions to agree with the assumptions of the OA method. To make claim that computed solutions are the optimal ones, Slaters condition is also required to hold [75, 80]. The formal definition of a single variable convex function can be given as,

$$f\left(tx_1 + (1-t)x_2\right) \leq tf\left(x_1\right) + (1-t)f\left(x_2\right), \tag{4.18}$$

$$\forall\, 0 \leq t \leq 1, \quad \forall\, x_1, x_2 \in X, \tag{4.19}$$

where X is a convex set, and $f : X \to \mathbb{R}$. This represents a straight line connecting any two points on f lying entirely above or just meeting the graph of f. To determine the convexity of a single valued function it is sufficient to consider the second derivative. A non negative second order derivative for all values in an interval, implies the function is convex over the same range. In the case of a multivariate function (a vector valued function as is the case with the models used in the framework) the convexity can be classified from the Hessian matrix of second order partial derivates. If the Hessian is positive semi-definite, where all eigenvalues are non negative, the function is said to be convex. Slaters condition states that the feasible region of a convex optimisation problem must have an interior point [82]. This requires all non linear functions to be convex where nonlinear constraints are able to be satisfied with strict inequalities, and affine/linear equalities are simply satisfied. A MINLP requires these conditions hold for all feasible combinations of integer values [80].

To demonstrate convexity in the negative objective function, $-\pi$ first note the sum of convex functions is itself convex. This can help simplify the objective function into suitable smaller functions, where convexity is more easily assessed.

$$-\pi = f^* (\mathrm{S}_i) + g\left(\mathrm{S}_i, I_i^{(To)}, \mathrm{store}_i^{(To)}, D_i\right) + h(W_i), \tag{4.20}$$

$$f^* = \frac{-1}{\mathrm{PET}^{(To)}} \left(\mathrm{PET}_1 f(S_1) + \cdots + \mathrm{PET}_T f(S_T)\right), \tag{4.21}$$

$$g = I_1^{(To)} + \cdots + I_T^{(To)} + \mathrm{store}_1^{(To)} + \cdots + \mathrm{store}_T^{(To)} + D_1 + \cdots + D_T, \tag{4.22}$$

$$h = W_1 + \cdots + W_T. \tag{4.23}$$

The g component of the objective function is a linear combination of decision variables and so has non negative second order derivates. Considering the first function f^*, it is possible to further break this down to the components of $\mathrm{PET}_k f(\mathrm{S}_k)$, where $f(\mathrm{S}_k)$ represents the reduction function equation 3.8. The non zero entry in the Hessian matrix will be given by the second derivative of the reduction function with respect to the S_k variable. This can be evaluated by setting the soil parameter to $r = 10$, as is used in the framework, and field capacity of 100 mm. This is given as,

$$\begin{aligned}
-\frac{\partial^2 f(\mathrm{S}, 10)}{\partial \mathrm{S}^2} &= \frac{1}{50}\left(\exp\left(-\frac{x}{10}\right) + \exp\left(-\frac{x}{10} - 10\right)\right) A^{-2} \\
&\quad - \frac{1}{5} A \exp\left(-\frac{x}{10}\right)\left(\frac{1}{5}\exp\left(-\frac{x}{10}\right) - \frac{1}{5}\exp\left(\frac{x}{10} - 10\right)\right) A^{-4}, \\
A &= 1 - 2\exp(-10) + \exp\left(-\frac{x}{10}\right).
\end{aligned}$$

Looking at where the graph of this function is positive will give an associated convex region. This can be seen as the positive real line, and so for all values of soil moisture with a soil parameter $r = 10$, the reduction function is convex. Therefore f^* is a convex function.

This leaves the sum of water loss functions to deal with, h. On any day the water loss function depends on both the soil moisture and the irrigation amount. This requires a Hessian matrix with respect to both S and $I^{(To)}$. Here only the total irrigation, $I^{(To)}$ is considered as this is a linear sum of irrigation variables and thus convexity with respect to the total irrigation will imply convexity with respect to each individual irrigation variable, $I,\ I^{(st)},\ I^{(b)}$. Looking at the function,

$$
\begin{aligned}
W &= (1-\text{eff})I^{(To)} \\
&= \left(1 - a + b\frac{I^{(To)}}{F_c - \mathrm{S}}\right) I^{(To)} \\
&= I^{(To)}(1-a) + b\frac{\left(I^{(To)}\right)^2}{F_c - \mathrm{S}},
\end{aligned} \tag{4.24}
$$

S	current soil moisture level [mm],
F_c	Field capacity [mm],
$I^{(To)}$	total irrigation applied [mm],
eff	efficiency of irrigation event,
a, b	positive fitted coefficients,

the appropriate Hessian of this function can then be evaluated as,

$$
\mathbf{H}_W = \frac{2b}{F_c - \mathrm{S}} \begin{bmatrix} 1 & \frac{I^{(To)}}{F_c - \mathrm{S}} \\ \frac{I^{(To)}}{F_c - \mathrm{S}} & \left(\frac{I^{(To)}}{F_c - \mathrm{S}}\right)^2 \end{bmatrix}. \tag{4.25}
$$

The eigenvalues of this matrix are

$$
\lambda_1 = 0, \quad \lambda_2 = 2b\frac{\left(I^{(To)}\right)^2 + (F_c - \mathrm{S})}{(F_c - \mathrm{S})^3}.
$$

Since $b > 0$ and $0 \leq \mathrm{S} \leq F_c$ at all times, the Hessian is positive semi-definite, and so the water loss function is convex. This implies the negative of the objective function is convex under minimisation, and therefore the original objective function is concave under maximisation as required. The concavity of constraints now needs to be confirmed. Since the ET reduction function and water loss function are part of the objective function, these have already been shown to meet the convexity requirement. Since these are the only nonlinear constraints it can be concluded that the original problem is convex under minimisation and therefore meets the concavity requirement of the decomposition method.

4.3.3 Mathematical Constraints

The constraints in the optimisation problem can now be formalised in the appropriate form to meet the assumptions outlined above. Considering the full problem with the inclusion of the bulk and on-farm storages, the constraints on some day t which relate the availability of water and irrigation events are:

$$\text{store}_t \leq \text{pump capacity} \times \text{availability of water source}_t, \tag{4.26}$$
$$I_t \leq \text{irrigation system limit} \times \text{availability of water source}_t, \tag{4.27}$$
$$I_t + \text{store}_t \leq \text{pump capacity}, \tag{4.28}$$
$$I_t^{(\text{of})} \leq V_t, \tag{4.29}$$
$$I_t + I_t^{(\text{b})} + I_t^{(\text{of})} \leq \text{irrigation decision}_t \times \text{irrigation system limit}, \tag{4.30}$$
$$I_t + I_t^{(\text{b})} + I_t^{(\text{of})} \leq 2\,(F_c - \text{S}_t). \tag{4.31}$$
$$\tag{4.32}$$

Constraints for the efficiency, drainage and water loss, and the ET can be defined as,

$$\text{eff}_t \leq a - b\frac{I_t + I_t^{(\text{b})} + I_t^{(\text{of})}}{F_c - \text{S}_t}, \tag{4.33}$$
$$\text{eff}_t \leq 1.0, \tag{4.34}$$
$$W_t \leq -(1 - \text{eff}_t)\left(I_t + I_t^{(\text{b})} + I_t^{(\text{of})}\right), \tag{4.35}$$
$$D_t \leq 0.0, \tag{4.36}$$
$$\text{AET}_t \leq \text{PET}_t \times f(\text{S}_t). \tag{4.37}$$

The definition of these constraints allows the nonlinear constraints (4.35, 4.37) to be satisfied with strict inequalities, satisfying the final assumption associated with Slaters condition. By restricting the efficiency, the linear approximation to the application dynamics of the efficiency curve can be used. This is without needing to also manage the 100% efficient region. Finally the soil moisture can be constrained to not exceed the field capacity:

$$\text{S}_t \leq F_c. \tag{4.38}$$

This final constraint helps implicitly define the drainage amount. If the field capacity is exceeded, the drainage will be set to the excess in order to meet the soil moisture constraint. All together this allows a definition of the daily update constraints which are equality constraints for the soil moisture and storage level. These are given as,

$$V_{t+1} = V_t + \text{store}_t + \text{store}_t^{(\text{b})} - I_t^{(\text{of})}, \tag{4.39}$$
$$\text{S}_{t+1} = \text{S}_t + R_t + \left(I_t + I_t^{(\text{b})} + I_t^{(\text{of})}\right) + D_t + W_t - \text{AET}_t. \tag{4.40}$$

The objective function is also modified to maintain convexity with respect to the drainage and water loss. This is given as,

$$\begin{aligned}\pi = c_1 Y_0 &\left(1 - K_s\left(1 - \frac{\text{AET}}{\text{PET}}\right)\right) \\ &+ c_2\,(D + W) - c_3\,(I + \text{Store}) \\ &- c4\left(I^{(\text{b})} + \text{Store}^{(\text{b})}\right) \;-\; c_5 I^{(\text{of})}.\end{aligned} \tag{4.41}$$

A *python* code listing of the complete complied optimisation program can be found in Appendix A.

4.3.4 Sub Problem Solvers

MindtPy allows the choice of solvers for each of the resulting sub problems. For the MILP problem a Branch and Bound enumeration method is chosen. This is implemented in a solver called Coin-or Branch and Cut (Cbc) [83]. While the choice for the NLP is a interior point method encoded in the Interior Point Optimisation solver (IPOPT) [84]. Cbc was selected due to fast convergence in initial testing, while IPOPT was chosen for its extensive configurability, implementation of adaptive barrier strategies [85], and ability to use second order corrections to reduce infeasibility in solutions.

Branch and Bound methods systematically enumerate through valid combinations of integers. At each iteration the integer values are fixed and the resulting continuous sub problem solved [86]. Cbc is a branch and cut and bound implementation, that makes use of cutting planes to tighten the relaxations on the original problem. The branch and bound method works by initially relaxing the integer variable restrictions. The Linear Program is then solved to get a lower bound on the original MILP objective function. If the lower bound found by the relaxation equals the upper bound found by a feasible solution of the MILP, the solution is optimal. This is due to a feasible solution of the MILP providing a valid upper bound to the objective function. If these bounds do not meet, there exists a non-integer value in place of an integer variable. The approach then branches this variable and the nodes are added to a search tree of integers. This search tree works by choosing a particular node and creating a new relaxed LP and solving it. This can then generate additional cutting planes which serve to tighten the relaxations on the relaxed node LP. This solution is then investigated for infeasibility or if it violates the current upper bound. Both options result in the node being pruned from the tree. If the solution is feasible and improves the upper bound, it replaces the best known MILP solution. If the node is unable to be pruned or the best known MILP solution updated, it will be branched again and the process repeated [81].

IPOPT makes use of a interior point line search filter algorithm. The interior point method is related to the simple primal-dual iteration. To solve the NLP, IPOPT incorporates a filter line-search method to construct a primal-dual interior-point algorithm. A brief overview of the method is presented below. For complete details of this algorithm, the original paper on the implementation of this method *On the implementation of an interior-point filter line-search* by A Wächter and L Biegler [84] is recommended.

Defining a simple NLP of the form:

$$\begin{aligned} \min_x \quad & f(x), \\ \text{s.t.} \quad & c(x) = 0, \\ & x \geq 0, \end{aligned} \tag{4.42}$$

where $x \in \mathbb{R}^n,\ f : \mathbb{R}^n \to \mathbb{R},\ c : \mathbb{R}^n \to \mathbb{R}^m,$ s.t. $m \leq n$. This can then be converted to an appropriate barrier problem,

$$\begin{aligned} \min_x \ \varphi_\mu(x) = f(x) - \mu \sum_{i=1}^{n} \ln x^{(i)}, \\ \text{s.t. } c(x) = 0, \end{aligned} \tag{4.43}$$

where $\mu > 0$ is the barrier parameter. With the use of the Karush-Kuhn-Tucker (KKT) conditions it is possible to rewrite this with corresponding Lagrangian multipliers, λ, z on the constraints $c(x) = 0$, $x \geq 0$. This then becomes:

$$\begin{aligned} \nabla f(x) + \nabla c(x)\lambda - z = 0, \\ c(x) = 0, \\ \mathrm{X}z - \mu e = 0, \end{aligned} \tag{4.44}$$

where e is a vector of ones to ensure dimensionality remains correct, and X represents the diagonal matrix of the corresponding x vector. For the case where $\mu = 0$ the requirement that $x \geq 0$ and $z \geq 0$ defines the KKT conditions for the original problem, equation 4.42. Solutions are then able to be found through applications of Newtons method.

Chapter 5

Verification of Model Design

With the optimisation framework developed in Chapter 4 a baseline was established to compare the optimisation approach with a current rule based model. This was simulated by the simple no storage model, where the only water supply was the available source water. This was to verify the optimisation model against irricalc's rule based modelling approach. It also allowed insight into how different parameters and data inputs affect the basic soil moisture dynamic. From these simulations the storages options were added to the modelling. The initial condition on the simulation was 100 mm of plant available water that was set to be full at the start of the season. The start and end of the simulation was the typical irrigation season as recommended by Aqualinc: the first of September until the end of March. This date range is to ensure a drought is not split over multiple irrigation seasons.

Simulation results were able able match similar dynamics as the rule based modelling. The optimisation was also able to improve on the amount of drainage that occurred. Soil moisture levels were kept optimal, and the ET demand was met as well as the supply allowed. Irrigation decisions through the peak demand of the season remained under an increased cost of irrigation.

5.1 Initial Simulations

The initial model output from optimisation is compared with the rule based irrigation modelling. The rule system operates by soil moisture levels, where a drop below a particular trigger point will cause an irrigation event. If the return period of the irrigation allows an irrigation to be scheduled, water will be applied at a fixed amount of 15 mm. The rule based system can still only apply water when the supply is available; rainfall can help increase the soil moisture level.

The 1997/98 irrigation season was simulated first. The water supply for this period is maintained at 100% reliability, allowing irrigation as and when needed. This provides information on how the optimal approach compares to irrigating on a trigger, when there is supply to irrigate as required. The rainfall for this season is constant, with a large rainfall event towards the end of the season. A large rainfall event has the ability to supplement the soil moisture towards the end of the season. These larger rainfall events can bring unwanted drainage however, and an optimal situation should use this rainfall to an advantage. A simulation on this season was computed, where the cost to irrigate is set to \$1/mm/Ha. The cost of drainage is set similarly to \$1/mm/Ha. These costings put emphases on maintaining soil moisture levels with low repercussions in excessive drainage. Practically this represents an approach to maintain outright productivity of a crop, at the potential expense of excess water and drainage. Profit is \$0.4/Kg/Ha with a maximum potential yield of 16000 kg/Ha. Irrigation is fixed at a return period of three days and a maximum application amount of 15 mm.

The trajectory of the soil moisture level, as well as the difference in potential ET and actual ET are shown for this season at Figure 5.1. The soil moisture for the optimal based output is similar to the higher 80% trigger point rule for irricalc. Towards the later part of the season this dips to similar levels of the lower 60% trigger. These trigger points represent a percent of the field capacity at which the soil moisture will trigger an irrigation event. This allows the soil moisture to make full use of the larger rainfall event previously mentioned. The degree to which the evapotranspiration load is met is sufficient, as is expected with 100% supply availability. If this demand is not able to be met there is a decrease in crop production, leading to a decreased final yield. Despite the decline in soil moisture towards the end of the season, the demand is still able to be met, largely in part due to the reduced ET demand at the tail end of the season. Looking specifically at the irrigation events, Figure 5.2, with a high supply reliability, irrigation events are regular but not at the minimum return period restriction. This is in contrast to the higher 80% trigger where the high demand section of the season has very frequent irrigation events at three day intervals. Furthermore with the ET taper at the tail of the season, a large amount of water is not used. This shows it is possible to maintain a very similar soil moisture level to a higher irrigation trigger level with an average soil moisture level of 84 mm. This was achieved with less total irrigation water. The high trigger level used 540 mm of irrigation, compared to an optimal 440 mm. Considering the drainage loss in Figure 5.3, less drainage occurs with the optimal approach. Less drainage is also observed with the lower trigger point, due to the ability to take up more rainfall. The start of the season sees similar water loss across all approaches, but the lower soil moisture level is able to make use of more rainfall events, later in the season. This helps to reduce drainage and requires less irrigation. Across all three methods the objective function for the optimisation and equivalent function for the rule based methods produced a similar scaled value. When normalised by the maximum possible profit, the optimisation approach produced a 0.92 value. The two rule based methods produced 0.91 and 0.90 for the lower and high trigger respectively.

Simulating the 2000/2001 season there is a significant restriction to the supply for just over a month, starting around the middle of February. The same parameters as the 1997/98 season are again used. The trajectory of the soil moisture for each scenario is shown at Figure 5.4. Through peak summer months the soil moisture tracks similarly to the higher trigger level. Once the restriction begins, all moisture levels drop, and the demand is not able to be met. Such a drop in satisfying the demand will have adverse effects on crop growth, if not resulting in the death of the crop. Across all three approaches the deficit loss

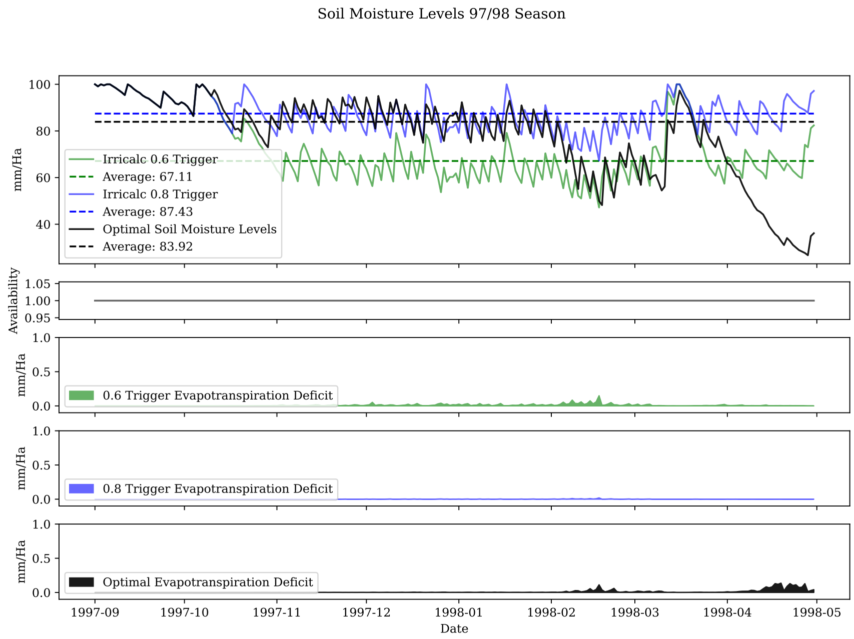

Figure 5.1: Soil moisture trajectories for the 1997/98 season using 60% and 80% soil moisture triggers, alongside the optimal based soil moisture output. Also shown is the consistent supply for the season. The bottom three plots show the deficit in the AET compared to the PET.

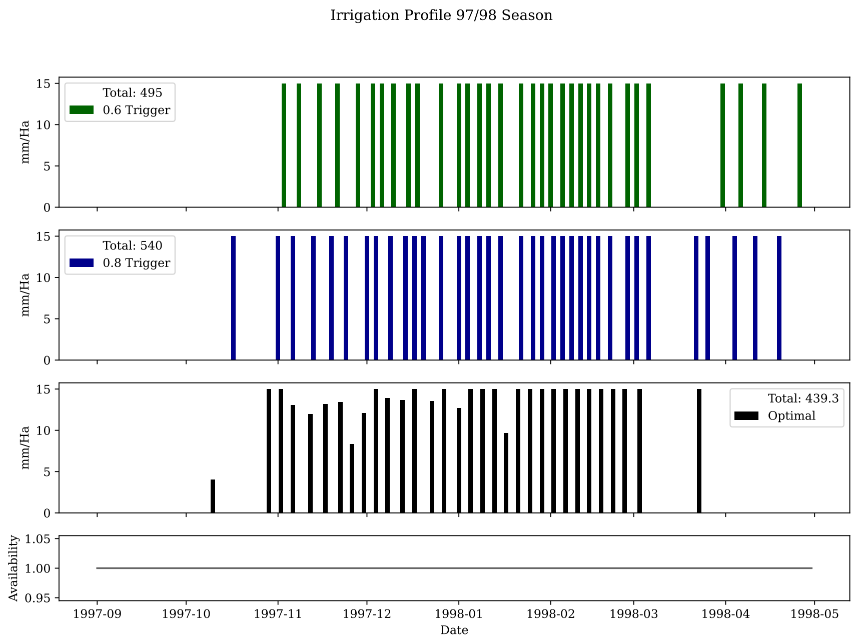

Figure 5.2: Irrigation events for the 1997/98 season, note the decline in events towards the end of the season in the optimal approach.

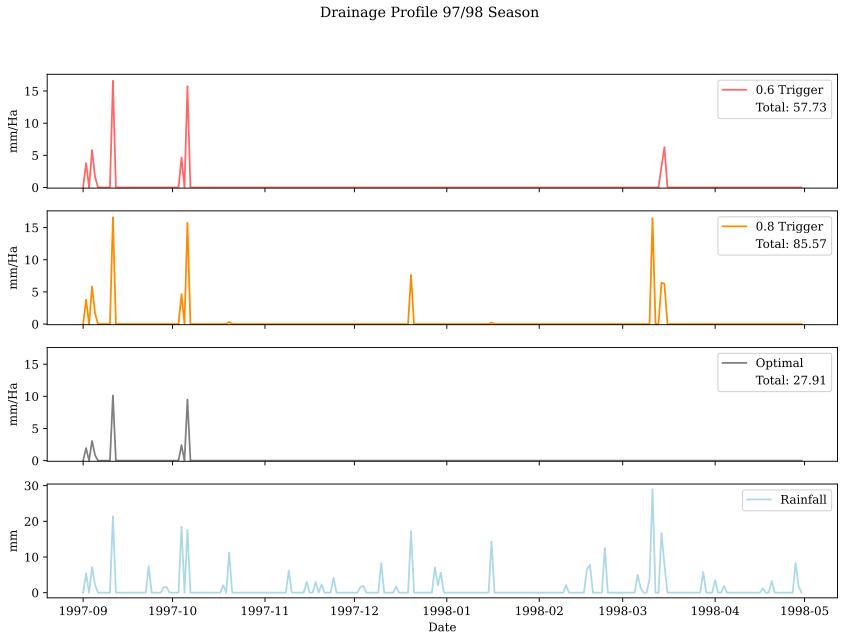

Figure 5.3: Drainage amounts for the 1997/98 season. The higher trigger of 80% experiences more drainage due to the higher overall soil moisture level. The optimal approach still experiences drainage where the soil moisture level is too high to take full advantage of the rainfall; this can be expected at the start of the season where spring rainfall combined with a high soil moisture level lead to drainage.

is very similar. The lower trigger point spends more time at compromising moisture levels, while the other two experience a similar daily maximum in inability to meet the ET load. The total amount of optimal irrigation falls between the two trigger levels, and mainly differs through the month leading up to the restriction. The irrigation is shown at Figure 5.5. Similar results are observed with the drainage (Figure 5.6), where the optimal approach experiences less drainage about the middle of the season. The high trigger rules maintain a moisture level that is too high to take in all the rainfall, and the 60% trigger also has a irrigation coinciding with a large rainfall event. The amount of rainfall tapers off towards the end of the season, and this, coupled with the heavy restriction, results in no drainage at the end of the season. This season fits the expected behaviour where a heavy restriction in supply greatly reduce the ability to meet the ET demand. Varying levels of irrigation and different starting soil moisture levels before the restriction, have little effect on the demand not being met. When starting the restriction with less moisture, the trajectory spends longer in a deficit. In both cases it is expected that the crop would suffer heavily, as the soil moisture tends very close to the permanent wilting point with zero plant available water. The scaled objective functions here showed a similar result again. The optimal approach improves on both the trigger rules with a value of 0.86. This large reduction from the 1997/98 season is due to unmet demand. The two rules achieved 0.83 for the low trigger and 0.85 for the higher trigger.

The last season to be simulated was 2005/2006. This season is of interest as it has a

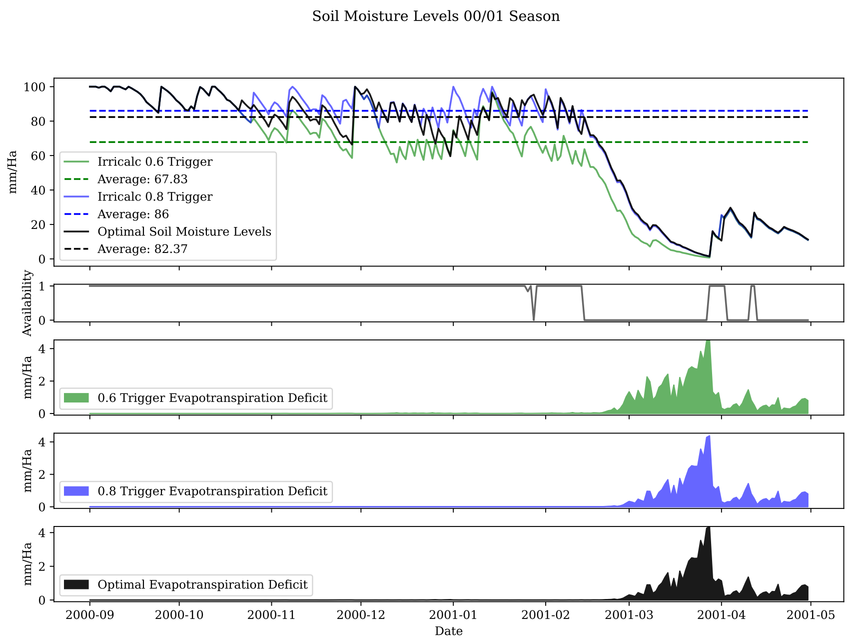

Figure 5.4: Soil moisture trajectories for the 2000/01 season. A higher initial soil moisture before restriction, allows slightly more of the ET load to be met. Ultimately the final result of an extremely low moisture level is unavoidable. The ET load that is unmet over the restriction period is very close to the assumed daily demand of five mm per day. This will have a severe impact on the crop.

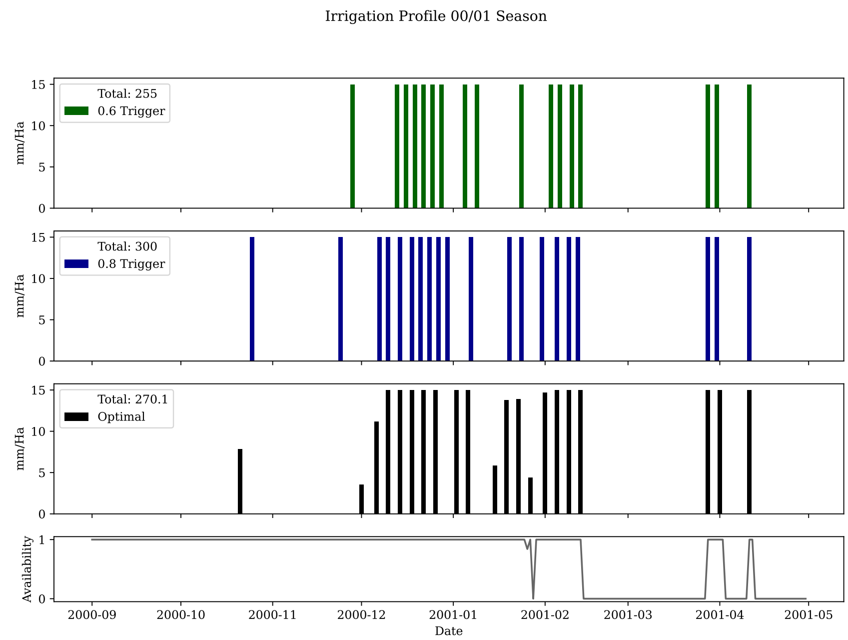

Figure 5.5: Irrigation events for the 2000/01 season. Very similar across the three methods when supply is available, the optimal approach varies the amount of water per irrigation event, but overall maintains a similar moisture level as the higher trigger method, as seen in Figure 5.4.

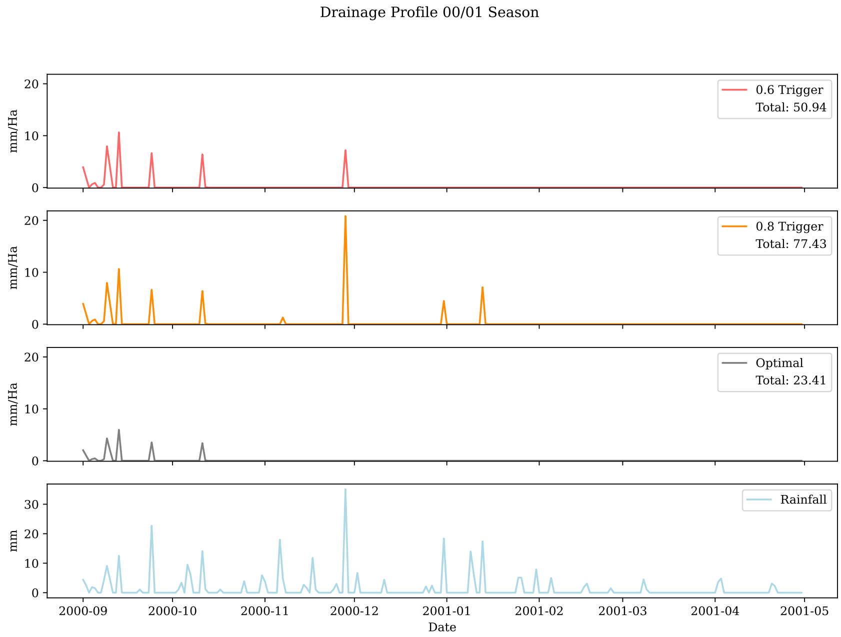

Figure 5.6: Drainage amounts for the three methods applied to the 2000/01 season. The optimal approach is again able to maintain less drainage through the main portion of the season. Drainage at the beginning is very similar across methods.

high level of undulating restriction throughout. These restrictions vary from a few days to almost a month. The rainfall for this season is quite varied with dry periods replenished by large rain events. From the soil moisture plot at Figure 5.7 the soil moisture level is observed to be undulating in response to the varied source availability. The average soil moisture level sits very close to that of the 80% trigger at 77 mm. The effects of the varying supply are more pronounced towards the end of the season, where the soil moisture is unable to recover as well after each successive restriction. This causes the soil moisture to fall lower and lower, and the ability to meet the ET load decrease. The optimal approach satisfied the demand similarly to the higher trigger, which can be accredited to the soil moistures following very similar paths. The lower trigger experiences more deficit to the demand earlier on in the season. For this particular scenario maintaining a higher soil moisture level is advantageous, as the rainfall events tend to signify a return to the supply availability. There is then ample time for the soil moisture to decrease and be able to take up the extra water during the restricted period. The irrigation events for this season (Figure 5.8) show the higher trigger level and the optimal approach begin irrigating earlier in the season. This allows the soil moisture to start at a higher point before each restriction. Through the center of the season irrigation events are very similar between the three methods. At the tail end of the season there is again a taper in the optimal irrigation amount. This works well with the reduced ET demand and there is only see a small decrease in the met demand. The drainage of each method shows similar results as previously (Figure 5.9). There is drainage across all methods at the start of the season, followed by no losses for the optimal approach. There is little drainage for the low trigger and slightly more for the higher trigger. This season is

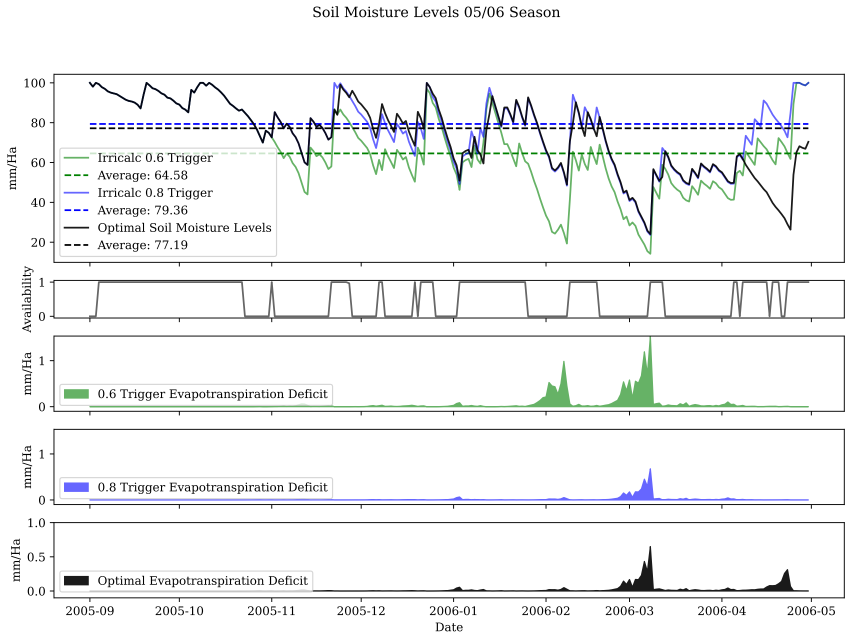

Figure 5.7: Soil moisture levels for the 20005/06 season. This season experiences a number of restrictions and the soil moisture levels reflect that. The lower trigger level sees greater overall dips in soil moisture due to the poorer starting position when entering the restricted period. A high soil moisture level appears advantageous here as it prevents a compounding of water deficits. This is further seen in the ability to meet the ET load. The optimal approach and higher trigger level have similar deficits in meeting the demand, while the lower trigger experiences a greater inability to meet the demand.

a good insight into variable supply reliability. The peak demand of the season has similar irrigation amounts but correspondingly different soil moisture levels. The trajectory of the soil moisture is dependent on the position the moisture levels are initially at before the restriction. This can be further seen by the optimal approach using less water than both trigger levels, yet managing the restrictions with a better overall outcome. Attained objective function values are very similar to that of the 1997/98. The optimal approach is able to achieve a value of 0.95, while the two rule methods were approximately 0.93.

5.2 Comparison with Irricalc Modelling

The three simulated seasons represent different supply availability profiles. These different supply profiles cover a wide range of situations, and in general it can be expected that a water supply is represented by one of these examples, either with less serve restrictions or less restrictions over all. If the results above are assumed to represent most real world supply situations, the optimal approach and the rule based method can be holistically compared.

The optimal approach falls between the two trigger levels for the average soil moisture, yet for both the 100% reliable season (1997/98) and the undulating season (2005/06), the total amount of irrigation is around 40 mm less than the lower trigger level. This can be

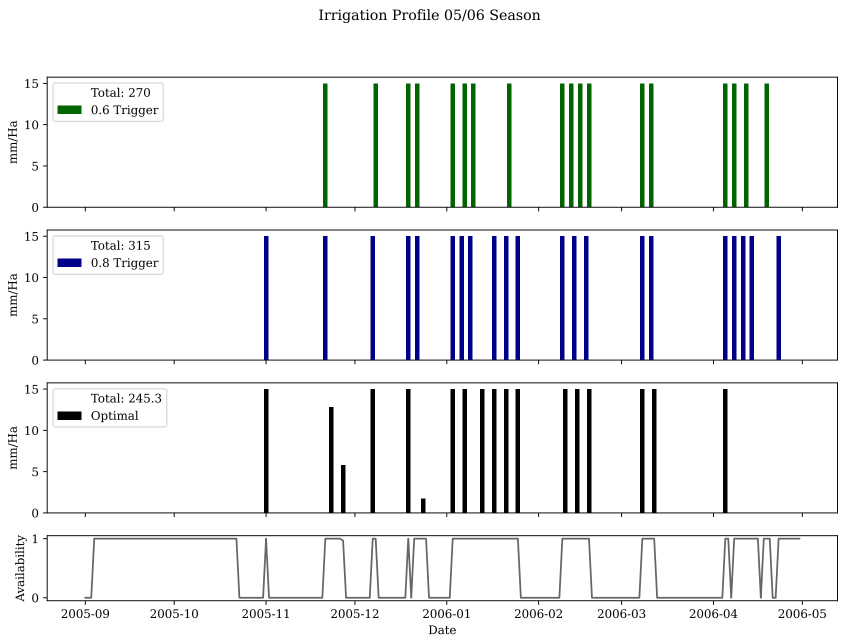

Figure 5.8: Irrigation events for the three approaches of the 2005/06 season. The irrigation amounts are very similar between methods throughout the middle of the season. This was seen to translate to different soil moisture levels. A lower trigger level meant the soil moisture may have gone into a restriction with an already lower soil moisture level. This manifested as the low trigger having a much lower soil moisture average, despite having similar water use.

Figure 5.9: Drainage amounts for the 2005/06 season. The optimal approach has less drainage, while the higher trigger level again experiences the greatest amount of water loss.

observed in part due to the decrease in irrigation events towards the end of the season. As the demand for transpiration decreases there is no longer a need to maintain a higher soil moisture level. For the season with the large restriction (2000/01), all methods apply similar amounts of irrigation at the end of the season, due to the extreme drop in soil moisture level. The optimal approach is able to improve the scheduling of irrigation to maximise the net benefits across the three scenarios.

The drainage and evapotranspiration demand show a similar result across the simulated seasons. The optimal approach matches the ET demand the same as the higher trigger, as would be expected with a higher average soil moisture level. Drainage is also able to be minimised by keeping a fluctuating moisture level to account for rainfall events. This perhaps highlights a single trigger level across the entire season is not the best approach to manage soil moisture levels and ET demand. Instead the output from the optimisation suggests a lower trigger point at the start and tail end of a season (shoulders) would provide better irrigation management. This would also cater to managing drainage levels by allowing more free soil moisture to take up rainfall during wetter parts of the season. This is a change that would be able to be implemented in the rule based approach by a time series of preset trigger levels.

5.3 Higher Cost Emphasis

The above simulations are a more productivity focused parameter set up. The low costs encourage irrigation when needed and do not heavily penalise water loss and drainage. To consider a more cost efficient approach the price of water was increased to \$3/mm/Ha. By keeping the drainage at the same cost as before, emphasis is shifted to be more conservative with water use. Starting with the heavy restriction of the 2000/01 season there is little change in the ability to meet the ET demand, and only a few mm difference in irrigation amounts (Figures 5.10, 5.11). This suggests that difference in water costings are over shadowed by the severity of the restriction, and the soil moisture levels before the restriction must remain similar. This allows recovery of some ET demand towards the end of the restriction.

The 2005/06 provides an interesting insight with an increased water price. There is a slight decrease in the total irrigation used (Figure 5.12), but the average soil moisture is increased compared to the higher trigger point (Figure 5.13). This is due to an increased soil moisture level before restrictions begin. Through the middle of the season the irrigation remains the same. This again shows having a higher soil moisture before periods of restriction is beneficial for maintaining overall ET demand. This particular example also shows that it is possible to achieve high soil moisture levels with better timing of irrigation events, even when restrictions are frequent.

Changing the cost of drainage was also simulated. This placed emphasis on minimising the total drainage. The cost of drainage was increased to \$8/mm/Ha and irrigation was returned to the cheaper cost of \$1/mm/Ha. This did not invoke any different irrigation decisions compared to the above cases, and the overall drainage amounts remain unchanged.

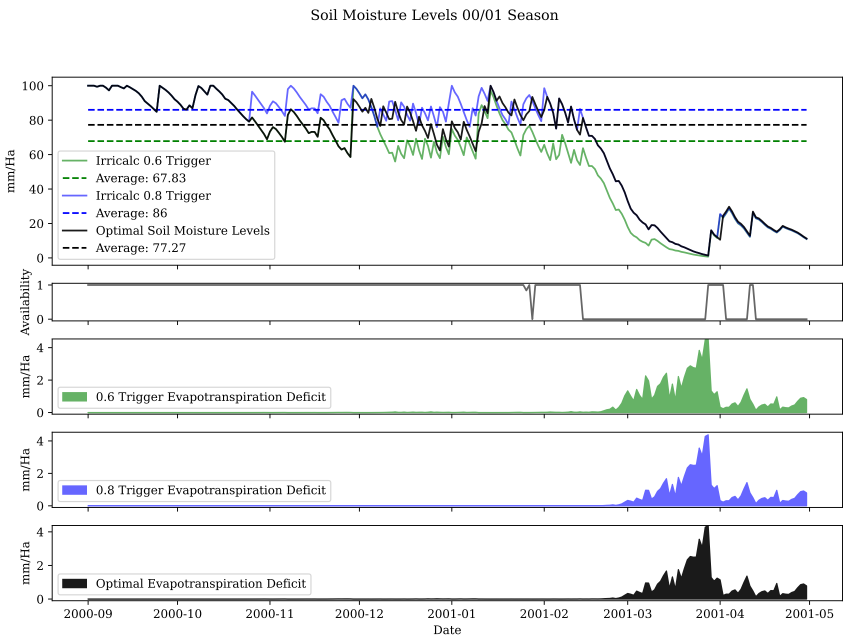

Figure 5.10: Soil moisture dynamics with increased irrigation cost for the 2000/01 season. There is a more pronounced reduction of irrigation at the start of the season, before reaching a higher soil moisture level as the restriction begins. The ability to meet the ET demand remains poor as the restriction takes hold.

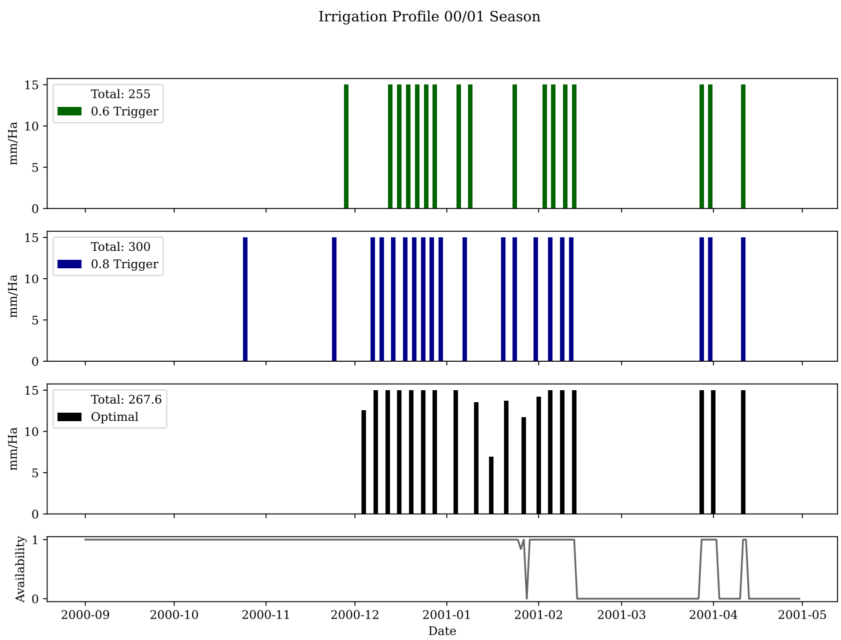

Figure 5.11: Irrigation events with the increased irrigation cost for the 2000/01 season. There is less initial irrigation, before ramping up towards the heavy restriction. Overall there is only a slight decrease in the total irrigation used, compared to the cheaper costing.

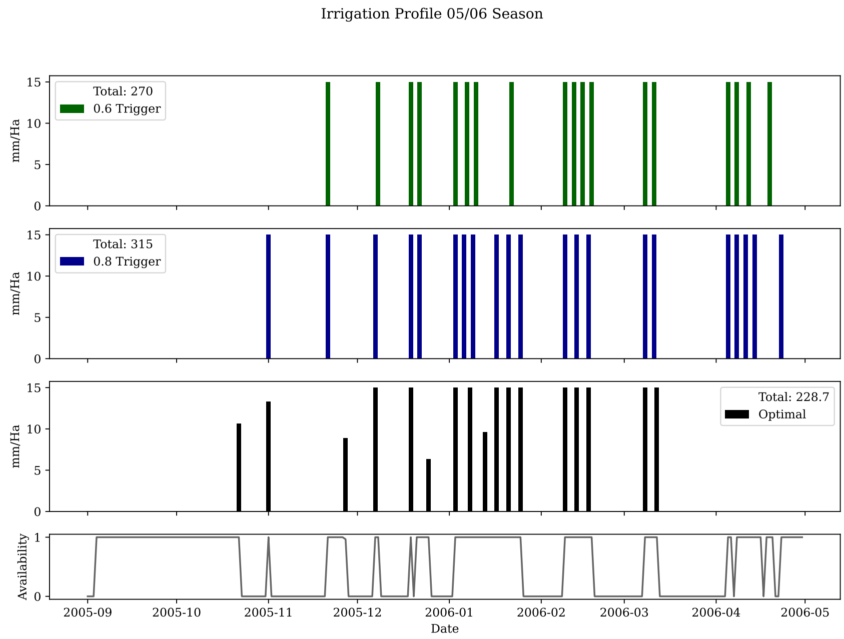

Figure 5.12: Irrigation events for the 2005/06 season at an increased irrigation cost of three dollars per mm per hectare. There is a slight increase in the irrigation earlier in the season.

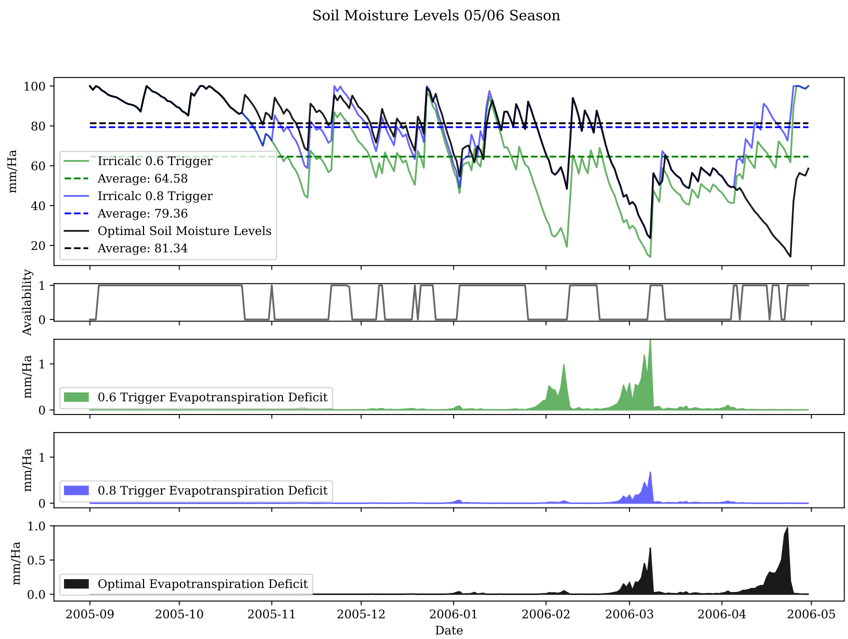

Figure 5.13: 2005/06 season soil moisture with increased irrigation cost. A higher average soil moisture is achieved through more irrigation events earlier in the season. This further shows through better timing of irrigation events, the effects of restrictions can be managed.

The initial drainage spike at the start of the season remained due to the assumption that the soil moisture starts the season at field capacity. Assuming the starting soil moisture is at field capacity makes sense in principal, as there should be ample moisture in the soil profile after the winter months of rain. However this also assumes that a farmer will begin utilising this soil moisture level at the beginning of the season, and not when they believe they need to start irrigating. A more realistic starting soil moisture could be 80% of field capacity. Reducing the initial moisture level would allow the model to better work through irrigation decisions, without being impeded by the unavoidable drainage. This should allow more of a focus on optimising the distribution of water between storages.

5.4 Rainfall Events and Modelled Drainage

With different supply profiles and parameter combinations simulated, some conclusions can be drawn about how the optimisation method is working and where dependencies lie. When comparing the rule based trigger to the optimisation, it is important to note that the optimisation method has the entire season's data available at once. This is contrary to the rule method which looks day by day. Relating back to the real world, the best case for a farmer is a long range forecast to base irrigation decisions from. Therefore this optimisation approach makes sense not as a forecasting tool, but as a tool that can inform on overarching goals or investments. It does however show that it is possible to achieve a more optimal irrigation strategy than a single rule based trigger. As previously mentioned, a varying trigger level seems most appropriate to minimise drainage through the wetter shoulders of the season, while maintaining a high soil moisture buffer through potential restrictions.

The rainfall observed in a season can play an important role in resetting the soil moisture conditions. Larger rainfall events tend to allow supply availability to return after restrictions. This appears especially true during a season with a large number of restrictions, as opposed to a longer single restriction. This can be investigated by running the above simulations with the initial focus on maintaining crop production, but with a quarter of the rainfall amount. The 2000/01 (Figure 5.14) and 2005/06 (Figure 5.15) seasons were considered. The multiple restrictions of the 2005/06 season with the lower rainfall did not allow the soil moisture to recover as well before the next restrictive period. In contrast, the single longer restriction of the 2000/01 season behaves quite similarly to the original simulation, but with more irrigation to compensate. This highlights the role rainfall plays in the soil moisture dynamics. With less rainfall it becomes harder to manage the soil moisture levels through restricted periods.

5.5 Summary

The modelling framework was shown to follow the expected dynamics of the state of the art rule based model. This gives confidence that the simulation results are able to represent physical situations. Through the initial production focused simulations the optimal

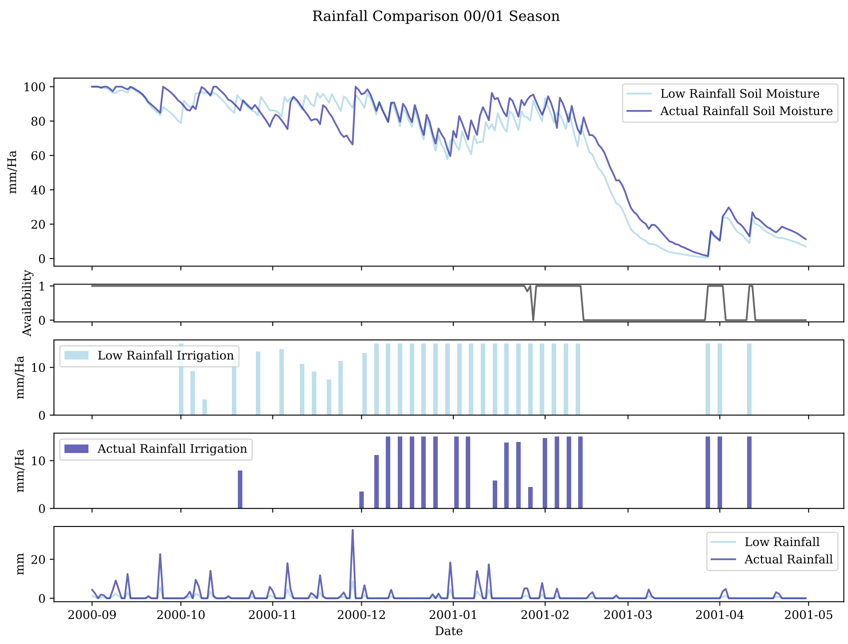

Figure 5.14: Comparing soil moisture levels and irrigation amounts for the 2000/01 season when different rainfall amounts are used as input series. The low rainfall is a quarter of the regular amount. There is little difference observed in the soil moisture levels between the two simulations, but the lower rainfall requires more irrigation to account for less moisture.

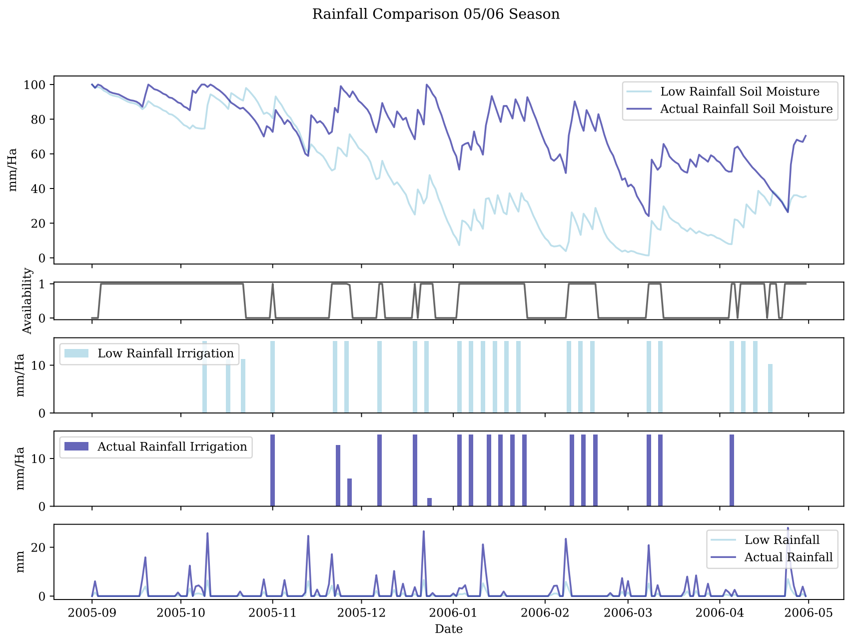

Figure 5.15: Soil moisture and irrigation for different rainfall profiles on the 2005/06 season. Much lower soil moisture recovery is observed with less rainfall when supply is available. This is despite similar amount of irrigation. This shows there is some dependency on additional rainfall in managing the soil moisture levels.

approach was able to schedule irrigation events that maximised the ability to meet the evapotranspiration demand. This was reflected in both the average soil moistures, and the attained objective functions, improving on the equivalent rule based function. It was shown that a time varying trigger point could offer improved drainage and irrigation amounts in rule based modelling. Increasing costs did not greatly impact the dynamics. This increase was instead reflected in the attained objective function values. This was seen to decrease by 0.08 in the 2000/01 and 2005/06 season, and by 0.14 in the 1997/98 season. With a decrease in rainfall, the optimisation was not able to manage the frequent restrictions of 2005/06, and more irrigation was required in both low rainfall simulations to compensate. This again shows the dependency on rainfall to assist irrigation in restoring the soil moisture levels.

Chapter 6

Extending the Framework for Water Storage

With the initial results presented above in Chapter 5, water storages were implemented into the simulations. This gives the ability to utilise the reliable bulk storage at the increased price per mm of water, as well as access to an on-farm storage. The goal here was to gain an understanding into what role these external stores of water played for maintaining crop production, while maximising the possible profit. The inclusion of these storages is in an effort to understand the potential benefits these can bring to a farm and how they could be managed in the scheme of the framework. This in contrast to finding an optimal irrigation schedule as has already been developed in previous research [67, 8, 9, 10] For these simulations the full model and objective function, detailed at the end of Chapter 4, was used.

6.1 Water Storage Rules

From the results of the initial testing, the initial soil moisture level was set to 80%. It is possible during extreme drought the bulk storage may become restricted, but to better analyse the distribution of water between these stores it was assumed that the bulk storage is 100% reliable. The on-farm storage pond was assumed to be able to hold enough water for two irrigation events, 30 mm of water per hectare. It was initially assumed this storage was at capacity before the start of the irrigation season. The cost to irrigate from the on-farm storage is less than the cost to directly irrigate. This essentially give a store of cheap water to access when full. If the storage is refilled, it costs the same as directly irrigating. This gives a relation between the cost parameters for the three possible irrigation paths as

$$\text{source irrigation cost} < \text{storage cost} + \text{storage irrigation cost} < \text{bulk irrigation cost} \tag{6.1}$$

It was also assumed that there is sufficient capacity in the pump system to be able to meet the demand of irrigating and storing water. This decision was based on the assumption that a farmer will want to maximise the potential availability for water, and so would be have the capacity to irrigate and store water at the same time. This becomes especially important when using water that becomes briefly available during fluctuating restrictions. The cost of drainage was also increased over the initial simulations. This increased to $2/mm/Ha. With little difference between the cheap and expensive drainage, this cost is set to be between the two costs of irrigation that will be used, one and three dollars. The goal here is to maintain a regulation on the amount of drainage. Since there is an increase in the opportunity to irrigate, the cost of drainage should increase as well.

6.2 Simulation Results

The 1997/98 season is not discussed with the inclusion of different water storages. With an already reliable source of water, the behaviour of an on-farm storage depends on the initial volume stored. If it is empty, there is no need for any water to be stored as the source is cheaper and reliable. If there is the assumption that the on-farm storage is full, the irrigation from the storage is spread out across the season. This is used to reduce the cost of irrigation from the source due to the cheaper cost to irrigate from the on-farm storage. The 2000/01 season was simulated first and an output of this season is displayed at Figure 6.1

[1].

For this simulation there was a focus on scheduling that will maintain production, with cheap pricing for all forms of irrigation. In contrast to the initial simulations, a high soil moisture level is able to be maintained for roughly half of the large restriction. To achieve this irrigation, is drawn from the on-farm storage pond as well as the bulk storage. Because of the sunk cost to have the on-farm storage filled, the optimal strategy is to use the stored water to supplement the bulk storage irrigation. This allows an additional five irrigation events, adding an extra 75mm of water to the soil. The season is finished by additional irrigation from the reinstated source, to reduce the falloff of ET demand. The ability to meet the ET demand through the restriction is greatly improved. There is a decline towards the end of the restriction, but at worst, this is less than 0.5 mm per day of unmet demand. Comparing the objective function value found with the initial no storage simulation, there is an improvement despite the increased cost of irrigation. It is

[1]The plots have been developed to show all relevant details of the modelling output. The top plot represents the total available water level at any time. This is ordered by cost to access the water, where soil moisture is immediately accessible, followed by the water source/supply. Then there is access to the on-farm storage ponds and finally the reliable bulk storage has 15 mm available at all times. This represents the potential water that is available on the supply side. It does not necessarily represent the water that is used in irrigation. Below this the availability of the source and on-farm storage are shown. The on-farm storage has a capacity of 30 mm per hectare. This is an availability percentage of the possible take. The actual total water take is found by scaling the availability by the pump capacities. The middle plots shows the irrigation events and colour represents which supply is drawn from. Source represents the main water supply and the appropriate costings in $/mm/Ha for each irrigation is shown in the legend. Direct from the source is lightest in colour, followed by the on-farm in a darker shade and finally the bulk storage irrigation in the darkest shade. Below this is the difference between the actual ET and the maximum potential ET. Finally the rainfall across the season is shown. This allows reference to irrigation decisions where large rainfall events occur. Also shown in the plot title is the normalised objective function value.

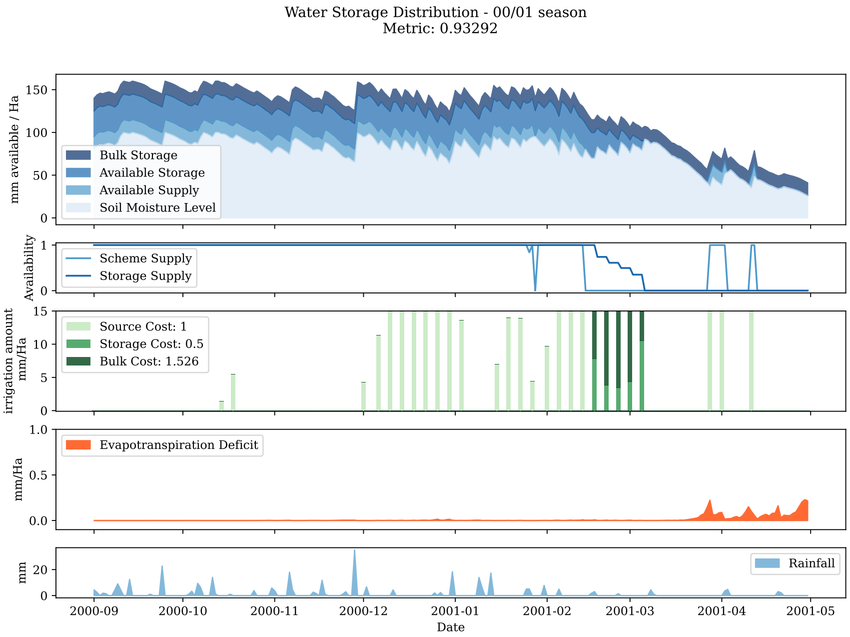

Figure 6.1: Simulation results for the 2000/01 season with access to on-farm and bulk storages. Towards the end of the season the actual ET falls away from the maximum potential, as expected with less water availability. Finally the rainfall over the season is shown in the last plot.

possible to normalise the objective value by the maximum potential profit. With only the soil moisture available, the normalised objective function is 0.86 for the initial simulation of the 2000/01 season. In comparison this is improved to 0.93 with the inclusion of additional storages. This increase is due to the increase in soil moisture though the restriction, keeping the AET high.

Increasing the cost of irrigation to $3/mm/Ha, a similar pattern in the soil moisture levels, irrigation, and meeting the ET demand (Figure 6.2) was observed. The biggest change is the lower normalised objective function value of 0.84. This is attributed to the increase in irrigation cost. The final irrigation event of the season also no longer occurs, and there is an increase in the ET deficit as a result. This shows once again that it is perhaps beneficial to reduce the amount of irrigation towards the end of the season to conserve costs. Furthermore, the high demand at the middle of the season unavoidably requires constant irrigation.

The 2005/06 season has frequent restrictions to manage (Figure 6.3). With the same initial conditions as the first simulation above, a different response to the restrictions was observed. The soil moisture level was more undulating in response to the varying availability of the source. Furthermore, with the restrictive periods being far shorter in length, there was no need to utilise the bulk storage. On-farm storage is used to bridge one of the longer restrictions at the tail end of the season to ensure the soil moisture level does not fall far from field capacity. Unlike the large restriction of the 2000/01 season there was not as much fall off in the ability to meet the ET demand. This can be attributed to the ability to keep the soil moisture at a higher level throughout the shorter restrictions. With

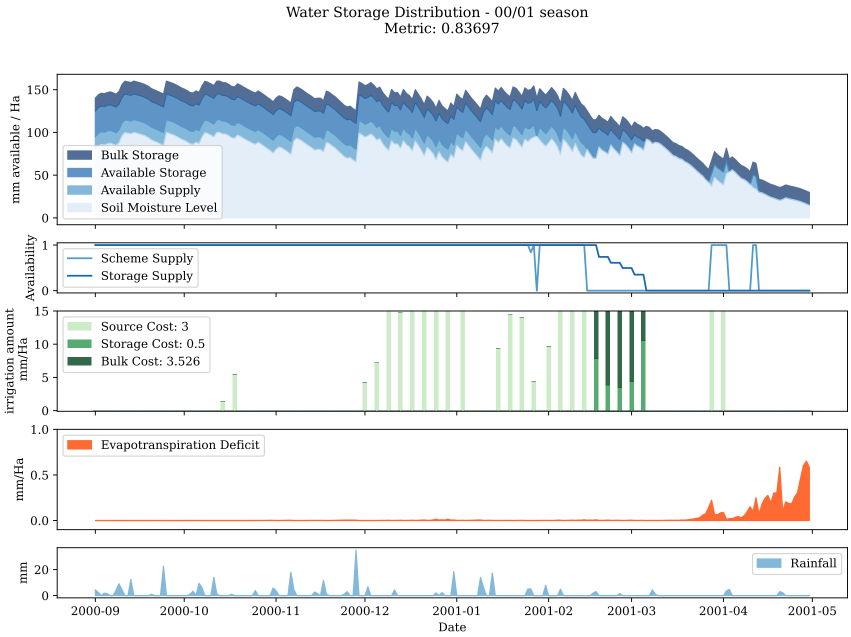

Figure 6.2: Output from the 2000/01 season with an increased cost to irrigate from the supply and bulk storage. An almost almost identical results as the cheaper cost simulation is observed. Irrigation stops earlier and there is more unmet ET demand at the very tail of the season. The normalised objective function is also reduced as a result, down to 0.84 from 0.93.

the addition of access to on-farm storages the objective function is maintained at 0.95, despite the slight cost increase with this access. Perhaps this more notably shows that with optimal irrigation management, short restrictive periods can be managed. This allows the ET demand to be met without the need for on-farm storage. Looking more closely at the amount of rainfall there is a large number of events applying around 20 mm of water each. It is possible this irrigation pattern would not be achievable with less rainfall, and more overall irrigation would be required with the use of on-farm and bulk storages.

6.2.1 Initial on-farm Storage Levels

In the previous simulations it was assumed the cost of filling the on-farm storage up had occurred outside the immediate irrigation season. If instead the storage was required to be filled during the season, the initial storage level can be set to zero. This was then simulated with he initial soil moisture at 80%. Initially the same cheaper cost scenario was used. Running the simulation for the 2005/06 season only enough for 5 mm of irrigation is put into the on-farm storage (Figure 6.4). This shows given the current restrictions and rainfall associated with this season, it is possible to maintain a high level of crop production without additional access to water storage. This is further reflected in the scaled objective function value, which remains very similar to the previously discussed 0.95. The ET is satisfied to a similar amount as the initial results with no storage available. The additional 5 mm that is applied through storage irrigation is done so at the first significant restriction period during the middle of the season. The timing of this will be to maintain a high soil

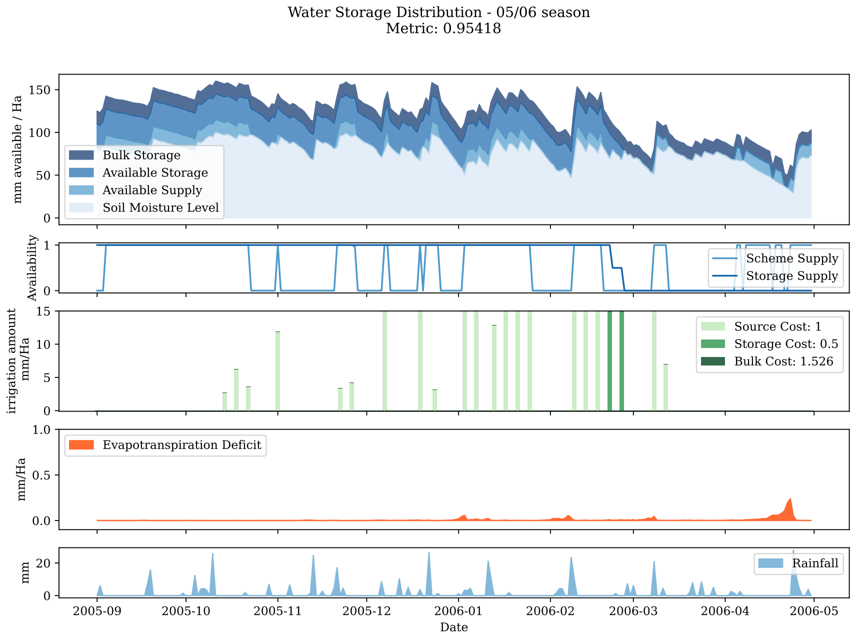

Figure 6.3: Simulation results for the 2005/06 season. Overall the available water is kept high, allowing the ET demand to be met. The on-farm storage is used over two irrigation events during a longer restriction towards the tail of the season. For this season there was a number of high rainfall events, it is possible that without these events the ability to meet the ET demand would require more storage input.

moisture before additional longer restrictions begin.

6.2.2 Rainfall Dependency

With no initially stored water in the on-farm storage, the amount of rainfall can be reduced to considered how this influences the decision to irrigate. The effect rainfall has on the total amount of irrigation needed in a season is shown in Figure 6.5. This was achieved by running the simulation over 30 years of data from 1980 to 2010. The total optimal irrigation applied in a season was then plotted against the total rainfall. Clearly there exists a relationship between the two variables, where more rainfall implies less total irrigation needed. This makes sense physically, as increased rainfall refills the soil moisture more often.

The 2000/01 season was then simulated with the rainfall reduced to a quarter of the original amount. The output of this simulation is shown at Figure 6.6 and it is immediately noticeable that additional irrigation is needed to maintain the ET. Irrigation for the season starts earlier and an attempt to stretch the on farm storage is made. There are now seven irrigation events during the restriction, and all of these draw from both the on-farm storage and the bulk storage. The increased cost reduces the normalised objective function to 0.90, but the ability to meet the ET demand is successfully maintained.

Reducing the rainfall over the 2005/06 season produced a more interesting insight into

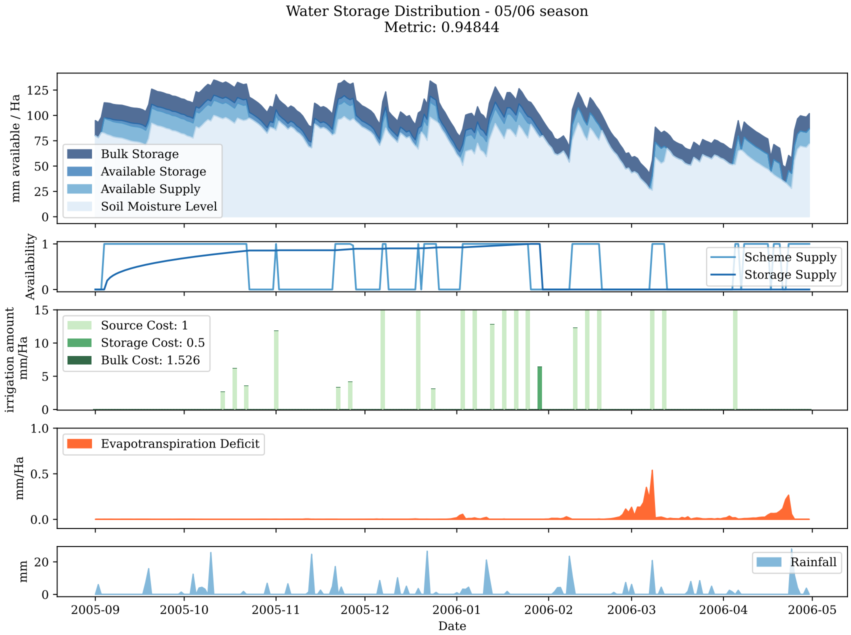

Figure 6.4: Simulation of the 2005/06 season with no initially stored water in the on-farm storage. The results show a similar pattern to the previous 2005/06 season simulations in terms of soil moisture and irrigation events. The storage pond is only filled up to 5 mm to allow for a single irrigation event. This irrigation bridges the first main restriction period during the middle of the season. The ET demand is satisfied similarly to the no storage case, and the attained normalised objective function value is very similar.

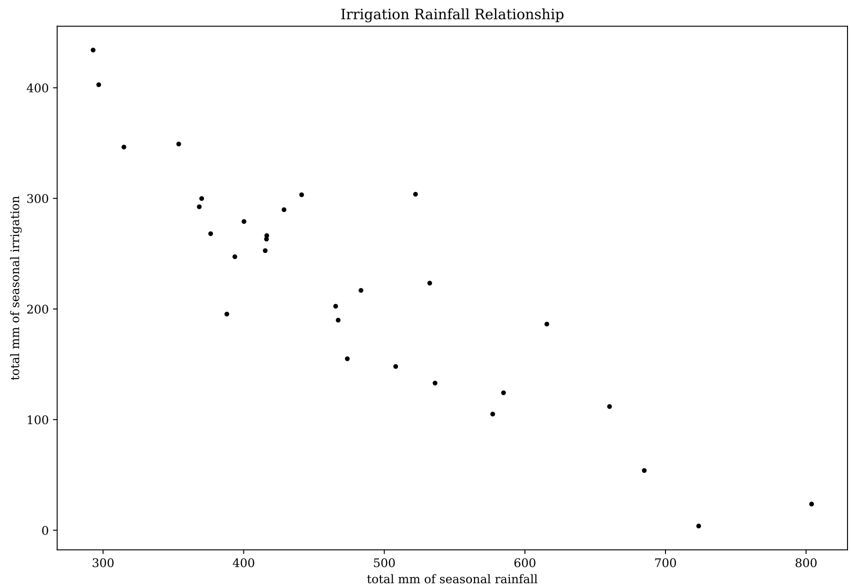

Figure 6.5: Dependency between total rainfall and optimal irrigation over 30 years. More rainfall refilling the soil moisture requires less irrigation.

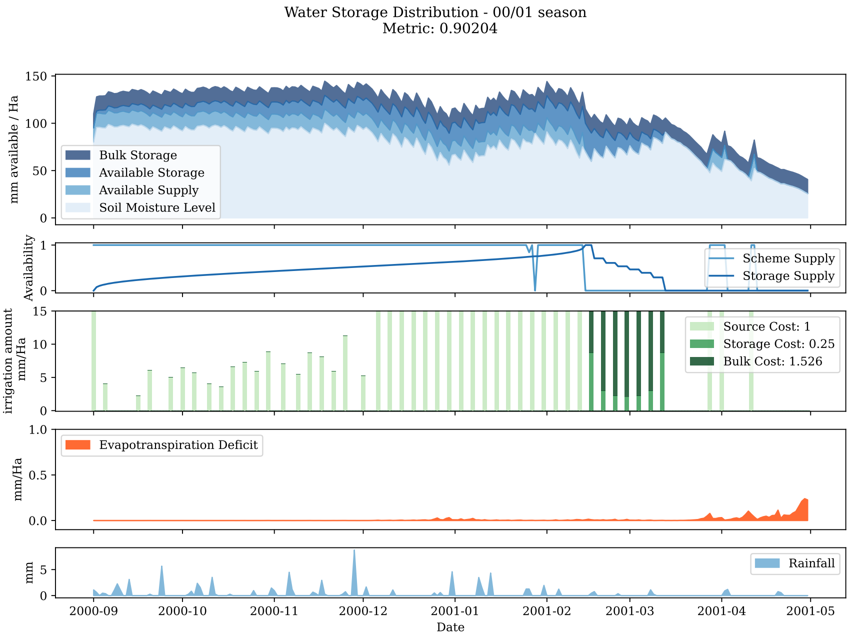

Figure 6.6: Simulation of the 2000/01 season with reduced rainfall. The rainfall was scaled to a quarter of the original amount. A greater number of irrigation events compensate the reduced rainfall, and there are more irrigation events during the restriction. These irrigations are a combination of the on-farm storage and the bulk storage. This is in an effort to maintain ET in a cost effective way. The reduced rainfall peaks at roughly 8 mm.

how the on-farm storage is refilled and used for irrigation. With the same parameters again, and the rainfall a quarter of the original amount, the simulation produced the output at Figure 6.7. In this scenario the on-farm storage is regularly refilled and used to bridge the restrictions. The bulk storage is used to stretch the on-farm storage to three irrigation events during the first larger restriction in the middle of the season. Overall the total irrigation is increased, while the ET demand is met the same as with the original rainfall. Through the middle of the season the irrigation remains as consistent as possible. At the beginning and tail end the irrigation is again diminished as the ET demand is not as significant. The price of irrigation was then increased to $3/mm/Ha and a similar pattern observed. Figure 6.8 shows the results of this parameter combination. The increase in price removes the use of the bulk storage to stretch the on-farm storage. There are now only two irrigation events during the large restriction. As with previous cost increases, the irrigation drops off towards the tail of the season and the actual ET falls away from the maximum potential.

6.3 Sensitivity to Parameters

The effects of the different costings for storages were simulated. Bulk storage irrigation and the bundled cost to irrigate from the on-farm storage were simulated to see the effects on different components of the modelling dynamic. Again only the 2000/01 and 2005/06 seasons were considered as the optimal irrigation for 1997/98 remains unchanged

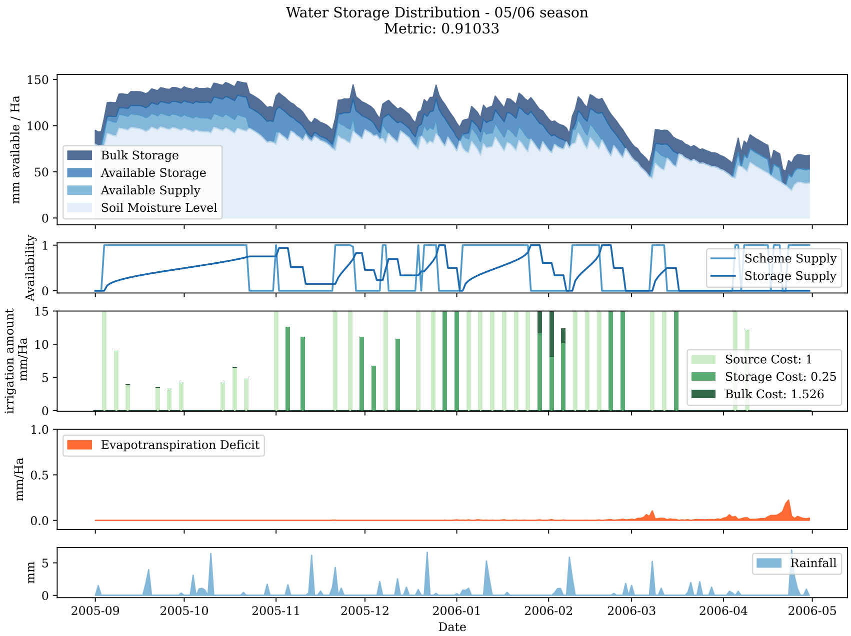

Figure 6.7: Low rainfall for the 2005/06 season. The on-farm storage is empty at the start of the season and this is refilled and then irrigated from during most restrictions. The ET demand is able to be met as well as with the original rainfall. Once again the bulk storage is used to stretch the on-farm stored water during irrigation. This increases the number of irrigation events during one of the longer restricted periods.

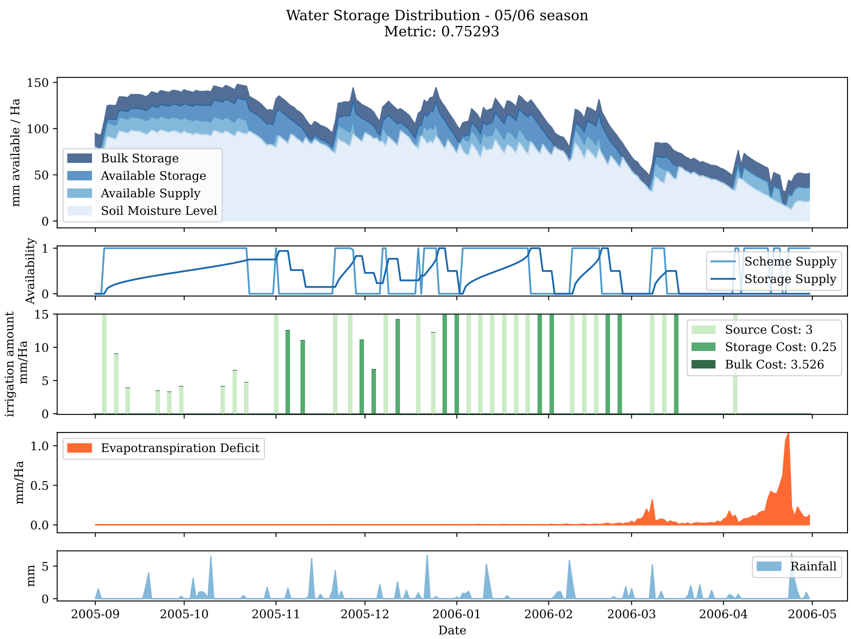

Figure 6.8: 2005/06 season simulated with low rainfall, an empty on-farm storage, and higher irrigation costs. The changes are a lack of bulk storage being used, and an earlier decrease in irrigation at the tail of the season. The decrease in soil moisture towards the end of the season also reduces the ability to meet all of the ET demand, and there is an increase in the deficit between the maximum potential and AET.

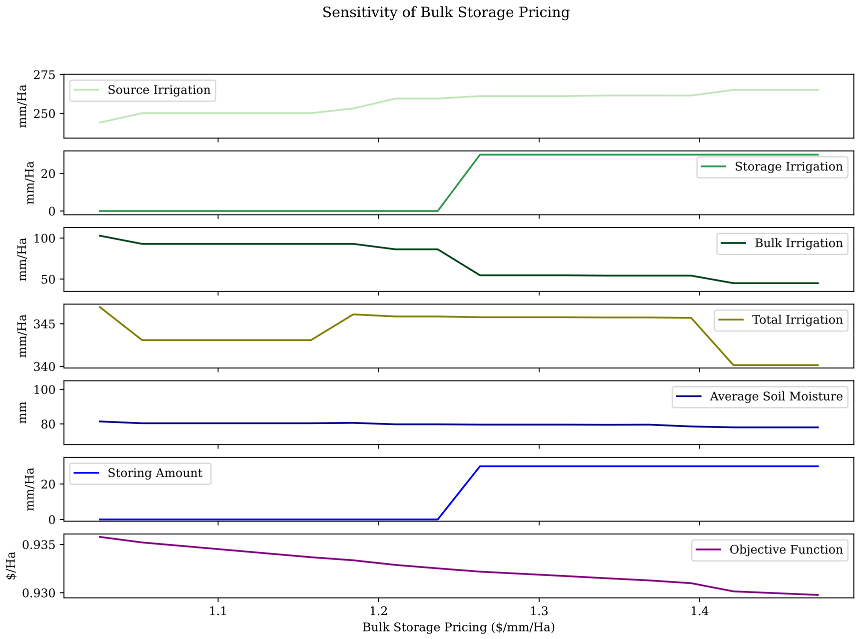

Figure 6.9: Increasing cost of bulk storage irrigation for the 2000/01 season. Irrigation is plotted as total millimeters applied during the season against the respective cost of bulk storage. Average soil moisture is computed as the seasonal average of daily soil moisture levels. The objective function is normalised with respect to the maximum possible profit. There is initially no water in the on-farm storage. When the cost of the bulk storage exceeds the cost to fill and irrigate from the on-farm store, the irrigation split between the three supplies reshuffles to accommodate the lowest cost combination.

with modification to the storage costs due to the reliability. The cost of storage irrigation is fixed at $0.25/mm/Ha. The cost to irrigate and store from the bulk storage is increased from the cost of the source up to $0.5 more expensive. Starting with the cheaper irrigation cost of $1/mm/Ha chosen metrics are plotted against the changing cost of the bulk supply, Figure 6.9. As the cost to use the bulk storage increases, there is a decline in the amount of bulk storage irrigation. At $1.25 the cost of the bulk supply is greater than the source plus the cost to irrigate from the on-farm storage. There is a further decline in the amount of bulk supply used from here. Since the 2000/01 season has the large restriction, there is still a need to use the bulk supply, despite the increased cost. It is also observed that source irrigation increases as the bulk supply irrigation decreases, keeping the total irrigation relativity constant. This is in an effort to maintain more optimal soil moisture levels at other times during the season. This keeps the profit high enough to compensate for the increasing cost. Over the price increase the objective function dips by 0.005. This indicates the ET demand is met equally well across the price increase.

Similarly, the cost to irrigate from the on-farm storage was changed. For this the bulk supply was fixed at a cost of $1.5/mm/Ha, Figure 6.10. There is a similar total irrigation throughout the cost increase. Irrigation from each supply increases or decreases to meet this total irrigation at higher costs. At its cheapest, the storage irrigation is used and refilled with a total irrigation of 45 mm/Ha. As the cost increases, this levels off to irrigation

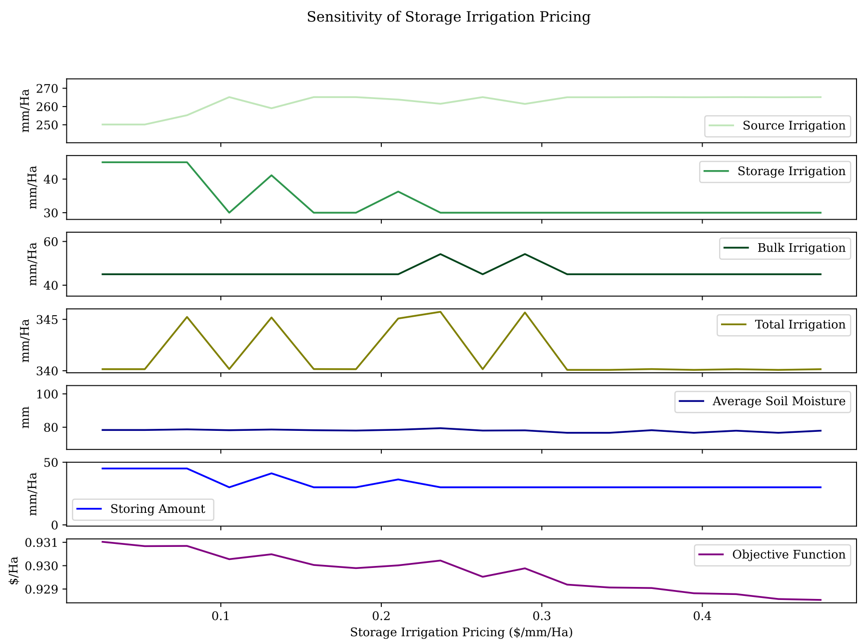

Figure 6.10: Increasing irrigation from on-farm storage cost for 2000/01 season. Total irrigation is again very similar with the distribution between the three supplies managing the increasing cost. Objective function decreases as a response to increasing price.

amounts that have been seen in previous simulations. The average soil moisture for both variations of cost increase remains steady around 80 mm.

The 2005/06 season provides a more interesting case when increasing the cost of the bulk storage, Figure 6.11. As previously, the total irrigation remains reasonably constant throughout the cost increase. As the cost of bulk storage increases to the cost of on-farm storage, the amount of irrigation between the two storages exchanges. Furthermore, the amount of bulk storage irrigation decreases slightly with the increasing cost, and this is accounted for by the source irrigation. The objective function decreases very slightly, again implying the ET demand is able to be met.

Finally the price of all irrigation options was increased to observe the total effects on the dynamics. The cost to irrigate from the source was increased and the price differentials between the source, bulk and on-farm storages maintained. The price of source irrigation was increased from $1/mm/Ha to $12/mm/Ha, where bulk storage cost an extra $0.5/mm/Ha over the source. The on-farm storage cost an extra $0.25/mm/Ha above the price of the source cost. This maintained the relationship outlined in equation 6.1. Figure 6.12 shows the simulation of the 2000/01 season. This shows a cut off point in irrigation around $9/mm/Ha. At costs higher than this cut off point the cost to irrigate outweighs the net gain in irrigation. This leaves the soil moisture to be determined by the rainfall that occurs. Although the average soil moisture dips just below 50 mm, the soil moisture will almost certainly have dropped to zero. Again there is a reasonably constant level of total irrigation, with a slight downward trend until the cut off. The objective function steadily tracks down with the increasing costs, but levels off at the maximum no irrigation profit. The

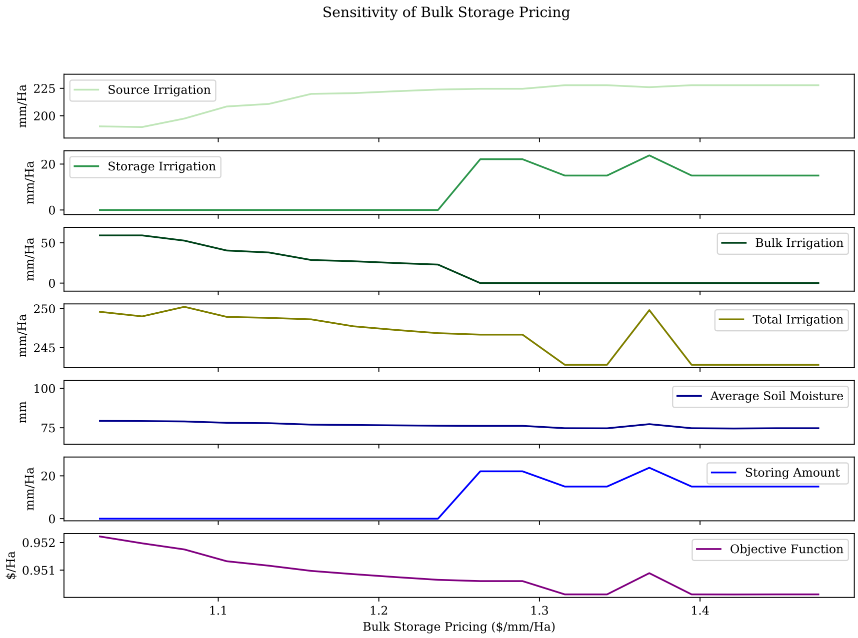

Figure 6.11: Increasing cost of Bulk storage access for the 2005/06 season. As with the 2000/01 season the total irrigation remains reasonably constant throughout, as does the soil moisture. This perhaps implies the total irrigation for a season is optimal and the distribution between irrigation is one that minimises cost.

2005/06 season produces a similar result. Here the cut off point is slightly higher around $10/mm/Ha. Interestingly there is a spike in the storage irrigation as the costs increase. This is mostly likely due to a net benefit in increasing soil moisture during restrictions, that can overcome the increased cost of source irrigation. This spike drops away again, and there is a gradual decline in the objective function to match. There is no bulk storage irrigation used to improve the later longer restrictions in this season. Over both the 2000/01 and 2005/06 season, the average soil moisture remains constant, until the cut off points.

6.4 Benefit of Storage Access

To test the benefit of additional storage access, the initial modelling detailed in Chapter 5 was compared with the simulation results presented in this Chapter. Analysing the 2000/01 season under regular rainfall, Figure 6.14 shows the benefit storage access brings. With no storage access the ability to meet the ET demand diminishes as the supply remains restricted. Towards the end of the restriction, there is around 5 mm of unmet ET demand. Contrast with access to storage, where the demand is able to be met throughout the restriction. Comparing scaled objective function values, without storage access the maximum profit is 0.85. When storage is available, this is higher at 0.92. Such a large restriction benefits from the additional storage. In contrast, the 2005/06 season (figure 6.15) does not see a noticeable decline in the objective function without storage access. Here the objective function attains 0.95 with storage, and 0.94 without.

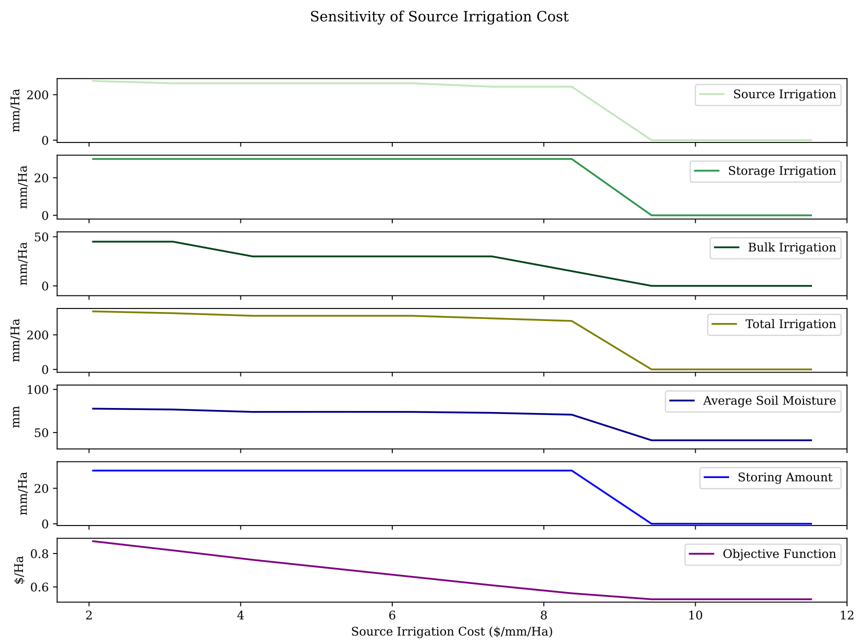

Figure 6.12: Increasing cost of source irrigation and storage for the 2000/01 season. The bulk storage is priced as 50 cents higher than the source, and on-farm storage irrigation is fixed at $0.25/mm/Ha, with an increasing cost to store. Irrigation amounts hold constant until around nine dollars, at which the cost to irrigate to sufficient soil moisture levels out weigh the profit generated.

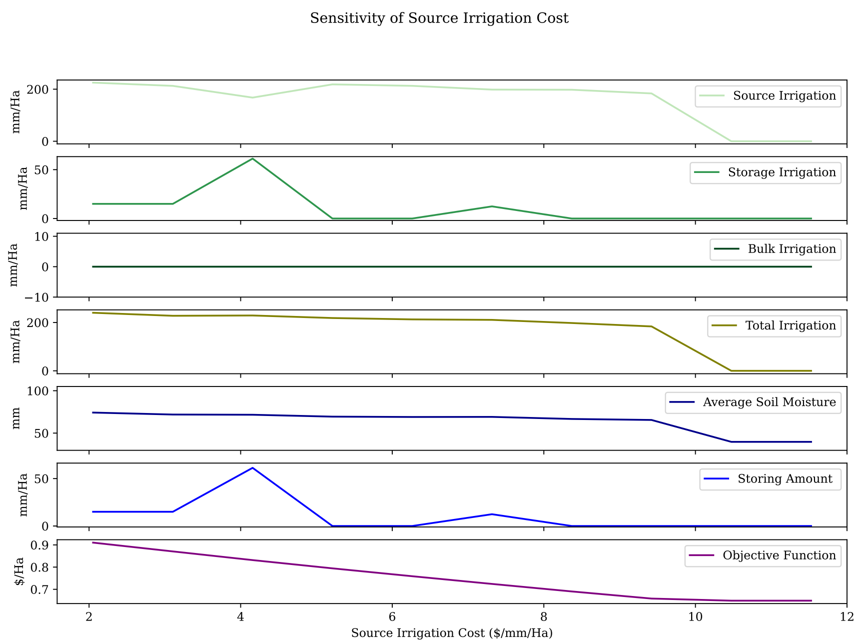

Figure 6.13: Source irrigation cost increase for the 2005/06 season. As with the 2000/01 season, the irrigation levels hold constant until a cut off around $10/mm/Ha. Around four dollars there is an increase in the amount of on-farm storage irrigation. No bulk storage irrigation is used with the shorter restriction periods found in the 2005/06 season.

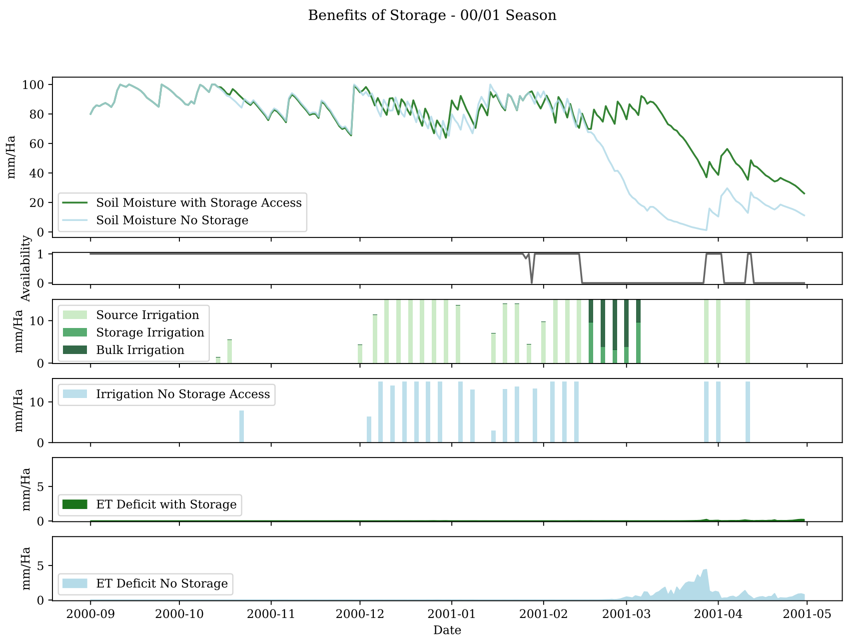

Figure 6.14: Comparison of the 2000/01 season with and without access to additional storage. Top plot represents the soil moisture trajectories of the two cases. The middle plots show the irrigation, and in the additional storage case, which supply the irrigation is coming from. Finally at the bottom is the unmet ET demand. With the additional irrigation made possible with the storage access, the ET demand is met as required.

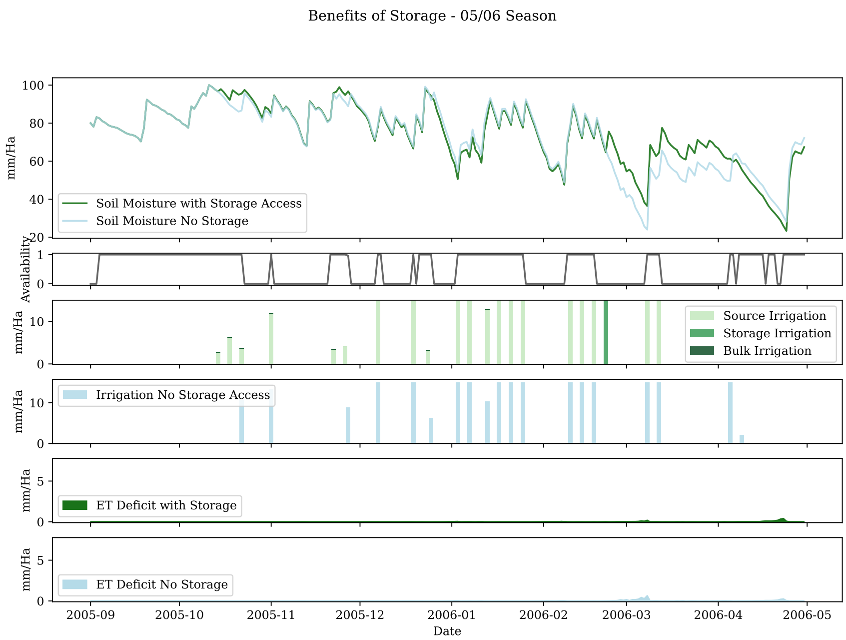

Figure 6.15: Comparison for the 2005/06 season. The additional supply with storage does not present a significant improvement. Both cases are able to meet the ET demand well.

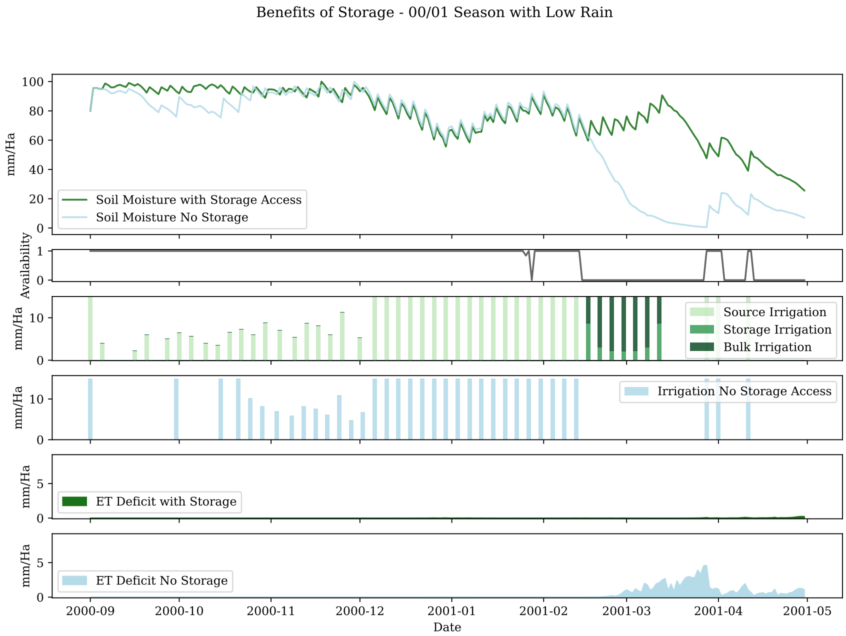

Figure 6.16: Contrast between storage access for the 2000/01 season with greatly reduced rainfall. Similar dynamics are observed, but with a greater amount of irrigation across both cases.

The results between the two different supply availabilities in each season show how the need for storage access can depend on the type of restriction. If there is extended restrictive periods in the supply, storage access adds a clear advantage. With multiple shorter restrictions this advantage is not present, and the soil moisture is able to cope without relying on additional supply. Altering the rainfall received over the season changes how the dynamics respond. This reduced rainfall is assumed to maintain the same supply availability profile. When reducing the rainfall the benefits of the storage access remain similar for the 2000/01 season, as shown in Figure 6.16. The 2005/06 sees a large change in the dynamics, Figure 6.17. With reduced rainfall, the soil moisture for 2005/06 is unable to both sustain, and recover to, sufficient levels needed to meet the ET demand. With access to the additional storage, the demand is able to be sufficiently met. This is further shown in the objective function values, where without storage the attained value is only 0.72 for 2005/06. With storage access the objective value increases to 0.91.

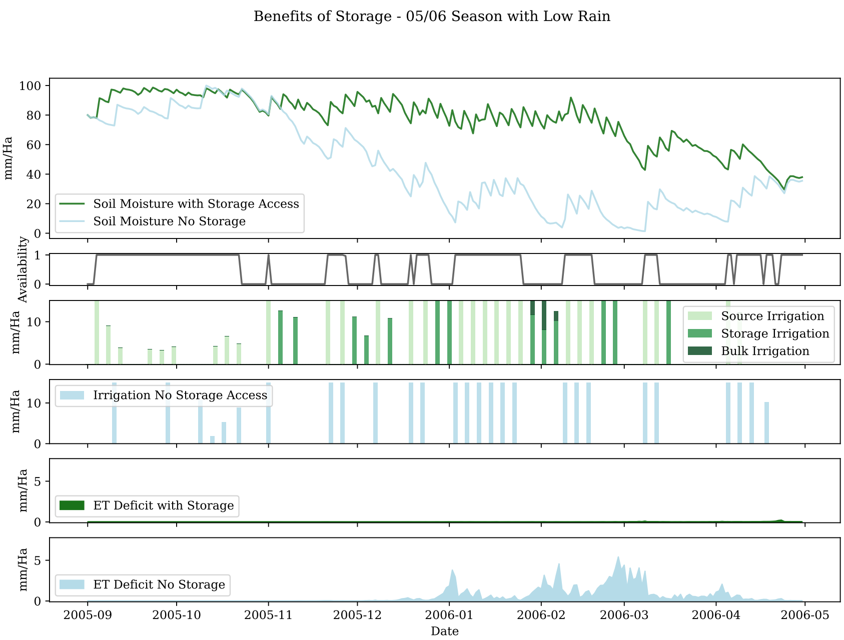

Figure 6.17: Contrast between 2005/06 season with and without storage as well as reduced rainfall. Here the observed dynamics greatly differ. Without access to additional storage, the soil moisture is unable to maintain or recover as well as before. This leads to severe drop off in the ability to meet the demand. By contrast, with additional storage the demand is able to be met as well as with the original amount of rainfall.

Chapter 7

Discussion

Analysing the results of Chapters 5 and 6 the developed framework and subsequent simulation results can be related back to the real world operation of a farm. The initial results compared the rule based irrigation scheduling with the developed MINLP model. These initial simulations showed good agreement with the rule based approach in overall soil moisture dynamics. This gave confidence that the developed optimisation model was behaving as expected, and would allow a physical interpretation of the outputs. The initial results showed that it was possible to improve on the rule based approach, through more optimal timing of irrigation events and utilisation of rainfall. This also reduced the amount of drainage, while meeting the evapotranspiration demand to an equal or improved amount. The improvement of soil moisture dynamics with the developed optimisation framework was inline optimisation results in the literature [71, 67, 8, 72, 9].

With the initial model verified against the rule based approach, storage access was able to be implemented. With access to water storage included in the model, the ability to meet the evapotranspiration demand increased significantly, as Section 6.4 showed. The exact use of the storage was dependent on the cost structure surrounding the relevant dynamics. If the on-farm storage was required to be refilled, the degree to which it was utilised depended on the supply availability and the cost. The large restriction in 2000/01 required the full use of the available storage, while the frequent restrictions of 2005/06 used very little. The question surrounding the use of these storages now becomes one of practical terms. An assumption that the on-farm storage capacity held enough for two irrigation events was made. This was based on the practical assumption that for an on-farm storage to be viable, it should hold enough water to be able to reduce the impact of a restriction. This would be in an effort to minimise the risk to crop production. Two irrigation events, with a return period of three days and an application of 15 mm, are able to meet the demand of 5 mm of evapotranspiration per day, over six days. This is based on an assumed demand of 5 mm per day that must be met. During the peak of the season, the maximum ET demand was seen to reach 7 mm per day. With no water input this could cause the soil moisture to drop to compromising levels over the same six day period. This storage size was useful for the extended restriction for the 2000/01 season. Here it was used with supplemental water from the bulk storage to provide an extra 75 mm of water application. Two irrigation events appears to be a manageable size of on-farm storage, but this is equivalent to having a storage pond of 3000 m^3/Ha. This storage amount can be contrast with a real world farm that has an available on-farm storage volume of 407000 m^3. The total amount of farming land the model on-farm storage could supply, is about 135 Ha. This particular real world farm is actually 1500 Ha and the 407000 m^3 of storage would be able to serve 3mm/Ha of storage. These two storage amounts are very different. However, it is possible a farm will rarely irrigate all hectares at once, and so it is possible the instantaneous storage volume available is higher. This becomes an important factor to consider, as the proportionality of mm of storage per hectare does not necessarily directly scale to a real world situation. There also exists a problem with on-farm storage where potentially profitable land must be given up to construct the storage. The feasibility of a storage size ultimately comes down to the whether this land is available and can be converted. While the inclusion of storages improve the soil moisture dynamics, there are still additional considerations on a farm scale that must be accounted for.

7.1 Sensitivity to Pricing

The sensitivity of the model to cost increases showed the optimal solution relied on a particular amount of irrigation and soil moisture levels. An increasing cost served to reduce water, only where it was strictly possible, with minimal effect on the trajectory of the soil moisture. This was almost always towards the end of the season. The total amount of irrigation remained steady throughout all cost increases as shown in Section 6.3. When referencing the irrigation amounts to the soil moisture level, there exists an optimal balance between the evapotranspiration load and the soil moisture that maximises the net profit of the objective function. This makes sense when considering how the evapotranspiration demand relates directly to the production and profit of the crop. The respective irrigation amounts between the three different supplies was seen to reshuffle to represent the lowest combined cost, while maintaining required irrigation amounts during restrictions. It would seem that provided the costings are reasonably similar, there is always a net benefit to maintaining irrigation and hence maintaining the soil moisture. In both seasons investigated this held true until the cost of all irrigation methods was increased to \$9 or \$10 per hectare. This cost was too high to sustain profit, and the model instead relied only on rainfall. Furthermore, this showed that the optimal solution is relativity robust to pricing increases. The biggest observed difference was the cost effective way to utilise the different supply options.

7.2 Strategies

The results of the simulations allow conclusions to be drawn on optimal strategies to manage the water under restricted supplies. When the supply is 100% available, an optimal irrigation strategy appears to have a varying trigger point across the irrigation season. A lower trigger level on the shoulders of the season means less water is used, as well as less drainage occurring from a lower soil moisture level. The evapotranspiration demand is also able to be met as there is less on the shoulders. During the middle of the season across all simulations, a higher soil moisture level is more beneficial. The higher evapotranspiration demand requires a higher soil moisture supply to maintain optimal crop growth. This varying trigger level also matches experimental design by Aqualinc, where modifications to the irrigation trigger point across a season was found to reduce drainage while maintaining crop production.

The use on-farm and bulk storage is done in one of two ways. If the restriction is long, the best strategy appears to be a combination of both the on-farm and bulk storages. This allows the cheaper on-farm storage to offset the higher cost of the bulk storage, while still maximising the amount of water applied. This strategy appears whenever the bulk storage is used for irrigation. If the cost of the bulk storage is too high, only the on-farm storage is used during irrigation events that occur under a restricted source. There is a general pattern regarding the irrigation leading up to a restriction. During this period the soil moisture is held as high as possible. A high soil moisture level allows the evapotranspiration demand to be met for a greater duration of time without irrigation. Across the restrictions in the simulated seasons, the level of soil moisture before a restriction did not reach field capacity. Despite sub maximal soil moisture levels, the ability to meet evapotranspiration demand was not significantly impacted. This showed there is little benefit to having the soil moisture at field capacity, prior to a restriction. This also decreases the chance of excess drainage. Without access to storage, the ability to meet the evapotranspiration demand is heavily impaired across the season. This was even more apparent with the lower levels of rainfall.

When rainfall is greatly reduced, the system requires an increase in the total amount of irrigation to compensate. Under this circumstance, the middle of the irrigation season requires a constant application of the maximum possible irrigation rate. It is under these conditions that the use of storage becomes critical, otherwise it is not possible to maintain a minimum soil moisture level, and the evapotranspiration demand is not met. During a restriction period the storage is used as frequently as possible until depleted. Maintaining a full on-farm storage throughout the peak of the irrigation season appears to be paramount, especially under limited rainfall.

7.3 Viability of Soil Moisture as a Storage

The viability of using soil moisture more actively as a store of water is a key idea that needs addressing in the soil-crop-atmosphere system. This involves relying on the soil moisture to provide sufficient water to meet evapotranspiration demand. For the soil to be a viable store of moisture, without additional storage access, there should be low risk of frequent, or extended restrictions. This is seen in the 1997/98 season, where the demand is able to be sufficiently met only relying on source irrigation. During the restriction periods analysed here, soil moisture alone is not sufficient to the demand during the extended restriction during the 2000/01 season. The frequent restrictions of 2005/06 allow for more recovery of the soil moisture after each restriction. This enabled the evapotranspiration demand to be met without access to storage and the soil moisture on its own was a sufficient store of water.

As Bras and Cordova note [71], the extent to which rainfall plays a part in the dynamics must be understood. When reducing the amount of rainfall to a quarter of the original, the above dynamics are changed, most notably in the case of frequent restrictions. The large restriction of 2000/01 remained relativity unchanged, with just an increase in the amount of irrigation to maintain the same soil moisture levels. The results from the 2005/06 season show a drastic change where the model is unable to maintain the evapotranspiration demand with the lower rainfall amount. For this season, the soil moisture no longer provides a viable supply of water to the crop. Under situations where there are back to back restrictions, the role of on-farm and bulk storage access becomes important with lower rainfall. These restrictions result in substantial decreases in soil moisture, and with irrigation only possible when supply returns, there is a high dependency on the rainfall to supplement the soil moisture.

Even with modifications to the data presented in one season, the ability of the soil moisture to hold enough water to supply a crop during a restriction varies. This is further complicated by differences in crop sensitivity to water stress. The limited simulations investigated here show that there are benefits from on-farm storage supplementing the water supply, but this is not a universal conclusion. Crops that are more resilient to dry conditions, may be able to utilise the soil moisture as the only supply of water. In contrast, highly sensitive crops may require the constant use of on-farm and bulk storages, to maintain an optimal soil moisture level and meet the evapotranspiration demand.

7.4 Limitations and Future Work

The framework developed here has been constructed on existing hydrological and agricultural models. These have been incorporated in such a way that allows for additional extension. This allows the framework to simulate a variety of situations through the modification of input parameters and data. One current limitation is the way the model makes use of a per hectare basis. The assumption of a per hectare basis does not necessarily allow a real representation of a farming system. Farms will often have multiple crops and more than one irrigator. For simplicity the framework has been used here with a single crop and irrigator. This treats a farm as a single paddock. Furthermore, with larger farms there is a high chance the soil 'r' parameter will vary. This gives the model limited scope in investigating dynamics across multiple crop types and irrigator set ups in the current configuration.

Another limitation comes in the form of being able to make general claims about strategies and soil moisture dynamics. As has been previously noted, the large number of input variables and high variance across different seasons make it difficult to determine concise conclusions about the dynamics for a general farm. This is further complicated by a high dependency on rainfall. Instead this framework is better suited to case by case investigation of optimal strategies or storage options, given a fixed set of parameters. This is what has been simulated through the results section in this thesis. Based on these results, it was possible to conclude particular trends and optimal improvements for the set of input data and parameters. This shows a framework of this type is best suited as a tool to be utilised with specific data, and strategies can be inferred from there. This could be extended in the future by running a set of fixed parameters over multiple seasons. This would allow greater insight into decisions surrounding storage options and management, as these farming structures would are in use for multiple seasons.

With more time, the framework presented here could be expanded with greater detail. The per hectare basis could be removed, and replaced with a more complete representation of the farm wanting to be simulated. Under the `pyomo` optimisation module this would be possible through increasing the dimensionality of the vectors associated with the decision variables. Such a change would allow a per crop or per irrigator optimisation, where global constraints covering each crop type or irrigator, would represent the overall available storage or supply access across the farm. It is not known exactly how the computation time would react to such changes, but it is expected that this would increase in some capacity. This would be due to the increased number of integer combinations that would need to be searched though. This extension would allow more accurate investigations of how to optimally utilise storages, as well as planning for on-farm storage sizes. It would also allow cost pricing structures to be included in the simulation for different storage options. This would also remove some assumptions regarding the initial volume and price of stored water, contributing to a more accurate farm model. Furthermore, these changes would allow an assessment of pricing structures to bulk storages, and under what costs they show a net benefit in paying to access. As well as the bulk storage, it would be possible to establish an optimal on-farm storage size that balances land required with the water benefits.

Chapter 8

Conclusion

This thesis developed a framework to be able to investigate the supply and distribution of water between multiple storages, and how that water could be utilised to maximise a profit function. One of the objectives with this framework was to develop a tool that allowed for investigation of particular soil-crop-atmosphere dynamics. This meant historic input data was utilised to assess an irrigation season as a whole. This is in contrast to a forecasting model where time series inputs for the current season, would be used to determine the optimal strategy forward in time. The model developed here allows for a comparison between existing strategies, and ones that are more optimal with respect to desired traits, such as reducing drainage. The framework presented combined various soil hydrology models, with models developed for relating a crop's yield to its evapotranspiration demand. A shift in the supply-demand relationship moved the soil moisture from the demand to the supply side. This allowed a focus on crop evapotranspiration, as opposed to maximising soil moisture levels. With a more active management possible on the soil moisture, a reduction in the total amount of drainage was observed. The potential to use less irrigation, while maintaining crop production was also shown.

The framework was built around the use of a soil moisture mass balance model. This required a representation of how water entered, and was lost from, the soil profile. To represent irrigation, supply availability, return period constraints, and a variable application amount were utilised. Water was lost from the soil through evapotranspiration, and from previous research [37], a relationship between the maximum potential ET and the actual ET was incorporated. This relationship was linked to the current soil moisture level. A drainage model for water lost during irrigation events was also introduced. This was reconstructed from work by Bright [67], and was presented in detail to confirm the statistical methods used. These models acted as regulators in the optimisation framework, and tied the dynamics together over the season.

A Mixed Integer Non-Linear Program optimisation methods was used to represent the models in an optimisation framework. This allowed a binary decision variable to enforce the return period of the irrigator. This was critical to getting an accurate representation of the irrigation scheduling. Due to the complications solving a MINLP, a solver based off Outer Approximation was used. This algorithm found solutions by alternating through

MILP and NLP sub problems, each with their own respective solver. This required a reformulation of some of the models used to ensure the assumptions of the solver were satisfied. In particular, the drainage model was approximated with a linear function, and other models were reworked to ensure concavity under maximisation.

The optimisation was initially compared to a rule based scheduling model, irricalc. This had no additional access to storage, and irrigation was limited by supply availability. Optimisation results showed similar dynamics to the rule based model, giving confidence that the results were inline with current modelling methods. With optimisation, drainage was able to be significantly reduced. In a period leading up to a restriction, the soil moisture did not attain field capacity, but was generally able to meet the evapotranspiration demand of the crop. The exception to this was the extended restriction of 2000/01. This season saw the soil moisture drop to almost zero, and would almost certainly have resulted in a severe reduction of production, or crop loss. Through the use of the optimisation model, it was also observed that with better timing of irrigation events, crop demand could be met with less irrigation used. This was attributed to using the evapotranspiration as the demand, instead of the soil moisture as a proxy. Such a shift allowed lower soil moisture levels on the shoulders of the season, while still maintaining maximum evapotranspiration.

Storages were then added into the simulations. This showed an increase in the ability to meet the evapotranspiration demand, under extended restrictions. In particular, during the 2000/01 season soil moisture levels were kept around the optimum and the attained objective function value was greatly improved. Strategies for using the storages involved maintaining a full on-farm storage throughout the first half and middle of the season. The on-farm storage was also used to supplement the bulk storage irrigation. This allowed the full irrigation amount to be applied, at a slightly lower cost than the bulk storage on its own. The ability of soil moisture to hold enough water to meet the demands of a crop was found to depend on the timing, duration, and frequency of restrictions, as well as the amount of rainfall. Low rainfall amounts, coupled with frequent restrictions, performed the worst with no storage access. By contrast, frequent restrictions with moderate amounts of rainfall were still able to have the demand managed.

The simulations showed the benefits in optimal irrigation scheduling, the use of storages, and also strategies associated with maintaining optimal soil moisture levels. Bulk storage was used to stretch the volume of on-farm storage across multiple irrigation events, totalling more then the storage volume of 30 mm. Such a strategy was seen to add an extra 75 mm of water onto the soil during the extended restriction of 2000/01. This strategy vastly improved both the limited capacity of the on-farm storage and the attained objective function during this season. Throughout the simulations, the total amount of irrigation for a particular season was seen to remain relativity constant, implying an optimal total irrigation amount. The proportion of the total irrigation that came from each supply changed as modifications to the pricing structure were made. This was to represent the lowest cost combination, while maintaining the same total irrigation amount.

If time allowed modifications could be made to the framework that incorporated different crops, irrigators, and soil types. Such a modification would be possible with the `pyomo` and MindtPy modelling components, but it is not known how adversely this would affect the computation time, due to additional integer combinations being added to the MILP search tree. This would allow specific modelling of a farm to investigate possible storage options. Additionally, this could inform in better detail how improvements could be made to fit a more optimal scheduling of irrigation.

Bibliography

[1] D. SCHOLEFIELD, K. C. TYSON, E. A. GARWOOD, A. C. ARMSTRONG, J. HAWKINS, and A. C. STONE. Nitrate leaching from grazed grassland lysimeters: effects of fertilizer input, field drainage, age of sward and patterns of weather. *Journal of Soil Science*, 44:601–613, 12 1993.

[2] N. Beaudoin, J.K. Saad, C. Van Laethem, J.M. Machet, J. Maucorps, and B. Mary. Nitrate leaching in intensive agriculture in northern france: Effect of farming practices, soils and crop rotations. *Agriculture, Ecosystems & Environment*, 111:292–310, 12 2005.

[3] N. Vázquez, A. Pardo, M.L. Suso, and M. Quemada. Drainage and nitrate leaching under processing tomato growth with drip irrigation and plastic mulching. *Agriculture, Ecosystems & Environment*, 112:313–323, 3 2006.

[4] Lincoln Zotarelli, Johannes M. Scholberg, Michael D. Dukes, and Rafael Muñoz Carpena. Monitoring of nitrate leaching in sandy soils. *Journal of Environmental Quality*, 36:953–962, 7 2007.

[5] M. Quemada, M. Baranski, M.N.J. Nobel-de Lange, A. Vallejo, and J.M. Cooper. Meta-analysis of strategies to control nitrate leaching in irrigated agricultural systems and their effects on crop yield. *Agriculture, Ecosystems & Environment*, 174:1–10, 7 2013.

[6] Olga M. Grant, Michael J. Davies, Helen Longbottom, and Christopher J. Atkinson. Irrigation scheduling and irrigation systems: optimising irrigation efficiency for container ornamental shrubs. *Irrigation Science*, 27:139–153, 1 2009.

[7] J. G. Arnold and C. O. Stockle. Simulation of supplemental irrigation from on-farm ponds. *Journal of Irrigation and Drainage Engineering*, 117:408–424, 5 1991.

[8] N.K. Garg and Sushmita M. Dadhich. Integrated non-linear model for optimal cropping pattern and irrigation scheduling under deficit irrigation. *Agricultural Water Management*, 140:1–13, 7 2014.

[9] S. Anvari, S.J. Mousavi, and S. Morid. Stochastic dynamic programming-based approach for optimal irrigation scheduling under restricted water availability conditions. *Irrigation and Drainage*, 66:492–500, 10 2017.

[10] Francesco Galioto and Adriano Battilani. Agro-economic simulation for day by day irrigation scheduling optimisation. *Agricultural Water Management*, 248:106761, 4 2021.

[11] Getachew Belaineh, Richard C. Peralta, and Trevor C. Hughes. Simulation/optimization modeling for water resources management. *Journal of Water Resources Planning and Management*, 125:154–161, 5 1999.

[12] N. Jovanovic, L.S. Pereira, P. Paredes, I. Pôças, V. Cantore, and M. Todorovic. A review of strategies, methods and technologies to reduce non-beneficial consumptive water use on farms considering the fao56 methods. *Agricultural Water Management*, 239:106267, 9 2020.

[13] Björn Guse, Matthias Pfannerstill, and Nicola Fohrer. Dynamic modelling of land use change impacts on nitrate loads in rivers. *Environmental Processes*, 2:575–592, 12 2015.

[14] Zhanpeng Pan, Yu Tong, Jingming Hou, Jian Zheng, Yongde Kang, Yan Wang, and Changchong Cao. Hole irrigation process simulation using a soil water dynamical model with parameter inversion method. *Agricultural Water Management*, 245:106542, 2 2021.

[15] John G Ndiritu. Maximising water supply system yield subject to multiple reliability constraints via simulation-optimisation. *Water SA*, 31, 9 2006.

[16] Enrique Playán and Luciano Mateos. Modernization and optimization of irrigation systems to increase water productivity. *Agricultural Water Management*, 80:100–116, 2 2006.

[17] David Molden, Theib Oweis, Pasquale Steduto, Prem Bindraban, Munir A. Hanjra, and Jacob Kijne. Improving agricultural water productivity: Between optimism and caution. *Agricultural Water Management*, 97:528–535, 4 2010.

[18] Zhiyu Wang, Guangcheng Shao, Jia Lu, Kun Zhang, Yang Gao, and Jihui Ding. Effects of controlled drainage on crop yield, drainage water quantity and quality: A meta-analysis. *Agricultural Water Management*, 239:106253, 9 2020.

[19] Yongchul Shin, Binayak P. Mohanty, and Amor V. M. Ines. Soil hydraulic properties in one-dimensional layered soil profile using layer-specific soil moisture assimilation scheme. *Water Resources Research*, 48, 6 2012.

[20] Andrew J. Guswa, M. A. Celia, and I. Rodriguez-Iturbe. Models of soil moisture dynamics in ecohydrology: A comparative study. *Water Resources Research*, 38:5–1–5–15, 9 2002.

[21] Susha Lekshmi S.U., D.N. Singh, and Maryam Shojaei Baghini. A critical review of soil moisture measurement. *Measurement*, 54:92–105, 8 2014.

[22] B. S. Minhas, K. S. Parikh, and T. N. Srinivasan. Toward the structure of a production function for wheat yields with dated inputs of irrigation water. *Water Resources Research*, 10:383–393, 6 1974.

[23] Neil C. Turner and John E. Begg. *Plant-water relations and adaptation to stress*, pages 97–131. Springer Netherlands, 1981.

[24] Elfatih A. B. Eltahir. A soil moisture-rainfall feedback mechanism: 1. theory and observations. *Water Resources Research*, 34:765–776, 4 1998.

[25] H. Vereecken, J. A. Huisman, H. Bogena, J. Vanderborght, J. A. Vrugt, and J. W. Hopmans. On the value of soil moisture measurements in vadose zone hydrology: A review. *Water Resources Research*, 44, 4 2008.

[26] Gayathri K. Devia, B.P. Ganasri, and G.S. Dwarakish. A review on hydrological models. *Aquatic Procedia*, 4:1001–1007, 2015.

[27] Julian Koch, Thomas Cornelissen, Zhufeng Fang, Heye Bogena, Bernd Diekkrüger, Stefan Kollet, and Simon Stisen. Inter-comparison of three distributed hydrological models with respect to seasonal variability of soil moisture patterns at a small forested catchment. *Journal of Hydrology*, 533:234–249, 2 2016.

[28] C. Albergel, C. Rüdiger, T. Pellarin, J.-C. Calvet, N. Fritz, F. Froissard, D. Suquia, A. Petitpa, B. Piguet, and E. Martin. From near-surface to root-zone soil moisture using an exponential filter: an assessment of the method based on in-situ observations and model simulations. *Hydrology and Earth System Sciences*, 12:1323–1337, 12 2008.

[29] S. J. Schymanski, M. Sivapalan, M. L. Roderick, J. Beringer, and L. B. Hutley. An optimality-based model of the coupled soil moisture and root dynamics. *Hydrology and Earth System Sciences*, 12:913–932, 6 2008.

[30] Ángel González-Zamora, Nilda Sánchez, José Martínez-Fernández, and Wolfgang Wagner. Root-zone plant available water estimation using the smos-derived soil water index. *Advances in Water Resources*, 96:339–353, 10 2016.

[31] Michael Blaschek, Pierre Roudier, Matteo Poggio, and Carolyn B. Hedley. Prediction of soil available water-holding capacity from visible near-infrared reflectance spectra. *Scientific Reports*, 9, 12 2019.

[32] Irrigation New Zealand. Soil texture and water. https://www.irrigationnz.co.nz/news-resources/irrigation-resources/soil-texture-water/Attachment?Action=Download&Attachment_id=103.

[33] Willem Verstraeten, Frank Veroustraete, and Jan Feyen. Assessment of evapotranspiration and soil moisture content across different scales of observation. *Sensors*, 8:70–117, 1 2008.

[34] Lingling Zhao, Jun Xia, Chong-yu Xu, Zhonggen Wang, Leszek Sobkowiak, and Cangrui Long. Evapotranspiration estimation methods in hydrological models. *Journal of Geographical Sciences*, 23:359–369, 4 2013.

[35] I.W. Eames, N.J. Marr, and H. Sabir. The evaporation coefficient of water: a review. *International Journal of Heat and Mass Transfer*, 40:2963–2973, 8 1997.

[36] Rohitashw Kumar, Vijay Shankar, and Mahesh Kumar Jat. Evaluation of root water uptake models – a review. *ISH Journal of Hydraulic Engineering*, 21:115–124, 5 2015.

[37] Richard Allen, L. Pereira, D. Raes, and M. Smith. Fao irrigation and drainage paper no. 56. *Rome: Food and Agriculture Organization of the United Nations*, 56:26–40, 01 1998.

[38] Shohreh Didari and Seyed Hamid Ahmadi. Calibration and evaluation of the fao56-penman-monteith, fao24-radiation, and priestly-taylor reference evapotranspiration models using the spatially measured solar radiation across a large arid and semi-arid area in southern iran. *Theoretical and Applied Climatology*, 136:441–455, 4 2019.

[39] Mario Córdova, Galo Carrillo-Rojas, Patricio Crespo, Bradford Wilcox, and Rolando Célleri. Evaluation of the penman-monteith (fao 56 pm) method for calculating reference evapotranspiration using limited data. *Mountain Research and Development*, 35:230, 8 2015.

[40] Andrew dark. Nation irrigated land spatial dataset 2020 update. techreport, Aqualinc, 09 2020.

[41] E. Fereres and M. A. Soriano. Deficit irrigation for reducing agricultural water use. *Journal of Experimental Botany*, 58:147–159, 11 2006.

[42] Henry E. Igbadun, Baanda A. Salim, Andrew K. P. R. Tarimo, and Henry F. Mahoo. Effects of deficit irrigation scheduling on yields and soil water balance of irrigated maize. *Irrigation Science*, 27:11–23, 11 2008.

[43] J. T. Waddell, S. C. Gupta, J. F. Moncrief, C. J. Rosen, and D. D. Steele. Irrigation- and nitrogen-management impacts on nitrate leaching under potato. *Journal of Environmental Quality*, 29:251–261, 1 2000.

[44] A.I. Gärdenäs, J.W. Hopmans, B.R. Hanson, and J. Šimůnek. Two-dimensional modeling of nitrate leaching for various fertigation scenarios under micro-irrigation. *Agricultural Water Management*, 74:219–242, 6 2005.

[45] T. P. Gaines and S. T. Gaines. Soil texture effect on nitrate leaching in soil percolates. *Communications in Soil Science and Plant Analysis*, 25:2561–2570, 8 1994.

[46] S. E. Simmelsgaard. The effect of crop, n-level, soil type and drainage on nitrate leaching from danish soil. *Soil Use and Management*, 14:30–36, 3 1998.

[47] T. W. Andraski, L. G. Bundy, and K. R. Brye. Crop management and corn nitrogen rate effects on nitrate leaching. *Journal of Environmental Quality*, 29:1095–1103, 7 2000.

[48] Roy F. Spalding, Darrell G. Watts, James S. Schepers, Mark E. Burbach, Mary E. Exner, Robert J. Poreda, and Glen E. Martin. Controlling nitrate leaching in irrigated agriculture. *Journal of Environmental Quality*, 30:1184–1194, 7 2001.

[49] R.B. Thompson, C. Martínez-Gaitan, M. Gallardo, C. Giménez, and M.D. Fernández. Identification of irrigation and n management practices that contribute to nitrate leaching loss from an intensive vegetable production system by use of a comprehensive survey. *Agricultural Water Management*, 89:261–274, 5 2007.

[50] Jordi Marsal, Mercè Mata, Amadeu Arbonés, Josep Rufat, and Joan Girona. Regulated deficit irrigation and rectification of irrigation scheduling in young pear trees: an evaluation based on vegetative and productive response. *European Journal of Agronomy*, 17:111–122, 9 2002.

[51] Stephen N. Ngigi, Hubert H.G. Savenije, Josephine N. Thome, Johan Rockström, and F.W.T. Penning de Vries. Agro-hydrological evaluation of on-farm rainwater storage systems for supplemental irrigation in laikipia district, kenya. *Agricultural Water Management*, 73:21–41, 4 2005.

[52] I.R. Calder, R.J. Harding, and P.T.W. Rosier. An objective assessment of soil-moisture deficit models. *Journal of Hydrology*, 60:329–355, 1 1983.

[53] Manijeh Mahmoudzadeh Varzi. Crop water production functions—a review of available mathematical method. *Journal of Agricultural Science*, 8:76, 3 2016.

[54] Sander J. Zwart and Wim G.M. Bastiaanssen. Review of measured crop water productivity values for irrigated wheat, rice, cotton and maize. *Agricultural Water Management*, 69:115–133, 9 2004.

[55] G.P. Tsakiris. A method for applying crop sensitivity factors in irrigation scheduling. *Agricultural Water Management*, 5:335–343, 12 1982.

[56] Henry E. Igbadun, Andrew K.P.R. Tarimo, Baanda A. Salim, and Henry F. Mahoo. Evaluation of selected crop water production functions for an irrigated maize crop. *Agricultural Water Management*, 94:1–10, 12 2007.

[57] Hayrettin Kuşçu. Mathematical modelling of crop water productivity for processing tomato. *Tarım Bilimleri Dergisi*, 25:137–146, 6 2019.

[58] Marvin E Jensen. Water consumption by agricultural plants (chapter1). 1968.

[59] J. Doorenbos and Amir Kassam. Fao irrigation and drainage paper 33. *Yield Response to Water*, 1979.

[60] Pasquale Steduto, Theodore C. Hsiao, Ellias Fereres, and Dirk Raes. Crop yield response to water. *Food and Agriculture Organization of the United Nations*, 66, 2012.

[61] Thomas R. Sinclair. "natural evaporation from open water, bare soil and grass" by harold l. penman, proceedings of the royal society of london (1948) a193:120-146. *Crop Science*, 59:2297–2299, 11 2019.

[62] John Lennox Monteith. Evaporation and environment. *Symposia of the Society for Experimental Biolsogy*, pages 205–234, 1965.

[63] Luis S. Pereira, Richard G. Allen, Martin Smith, and Dirk Raes. Crop evapotranspiration estimation with fao56: Past and future. *Agricultural Water Management*, 147:4–20, 1 2015.

[64] Gideon Peri, Donald I. Norum, and William E. Hart. Optimal irrigation depths — a method of analysis. *Journal of the Irrigation and Drainage Division*, 105:341–355, 12 1979.

[65] None Chairatana Seniwongse, None l-pai Wu, , and None W. N. Reynolds. Skewness and kurtosis influence on uniformity coefficient, and application to sprinkler irrigation design. *Transactions of the ASAE*, 15:0266–0271, 1972.

[66] E. Maroufpoor, A. Faryabi, H. Ghamarnia, and G.Y. Moshrefi. Evaluation of uniformity coefficients for sprinkler irrigation systems under different field conditions in kurdistan province (northwest of iran). *Soil and Water Research*, 5:139–145, 12 2010.

[67] John Charles Bright. Optimal control of irrigation systems; an analysis of water allocation rules. 1986.

[68] Mohammed Mainuddin, Mac Kirby, and M. Ejaz Qureshi. Integrated hydrologic–economic modelling for analyzing water acquisition strategies in the murray river basin. *Agricultural Water Management*, 93:123–135, 11 2007.

[69] Mahdi Moradi-Jalal, Omid Bozorg Haddad, Bryan W. Karney, and Miguel A. Mariño. Reservoir operation in assigning optimal multi-crop irrigation areas. *Agricultural Water Management*, 90:149–159, 5 2007.

[70] Biju George, Hector Malano, Brian Davidson, Petra Hellegers, Luna Bharati, and Sylvain Massuel. An integrated hydro-economic modelling framework to evaluate water allocation strategies i: Model development. *Agricultural Water Management*, 98:733–746, 3 2011.

[71] Rafael L. Bras and Jose R. Cordova. Intraseasonal water allocation in deficit irrigation. *Water Resources Research*, 17:866–874, 8 1981.

[72] Raphael Linker, Ilya Ioslovich, Georgios Sylaios, Finn Plauborg, and Adriano Battilani. Optimal model-based deficit irrigation scheduling using aquacrop: A simulation study with cotton, potato and tomato. *Agricultural Water Management*, 163:236–243, 1 2016.

[73] William E. Hart, Jean-Paul Watson, and David L. Woodruff. Pyomo: modeling and solving mathematical programs in python. *Mathematical Programming Computation*, 3:219–260, 9 2011.

[74] David E. Bernal, Qi Chen, Felicity Gong, and Ignacio E. Grossmann. *Mixed-Integer Nonlinear Decomposition Toolbox for Pyomo (MindtPy)*, pages 895–900. Elsevier, 2018.

[75] Marco A. Duran and Ignacio E. Grossmann. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Mathematical Programming*, 39:337–337, 10 1987.

[76] Tapio Westerlund and Frank Pettersson. An extended cutting plane method for solving convex minlp problems. *Computers & Chemical Engineering*, 19:131–136, 6 1995.

[77] Filippo Pecci, Edo Abraham, and Ivan Stoianov. Model reduction and outer approximation for optimizing the placement of control valves in complex water networks. *Journal of Water Resources Planning and Management*, 145:04019014, 5 2019.

[78] Maria P. Papadopoulou, George F. Pinder, and G. P. Karatzas. Enhancement of the outer approximation method for the solution of concentration-constrained optimal-design groundwater-remediation problems. *Water Resources Research*, 39, 7 2003.

[79] Alexander A. Spiliotopoulos, George P. Karatzas, and George F. Pinder. A multiperiod approach to the solution of groundwater management problems using an outer approximation method. *European Journal of Operational Research*, 157:514–525, 9 2004.

[80] Jan Kronqvist, David E. Bernal, and Ignacio E. Grossmann. Using regularization and second order information in outer approximation for convex minlp. *Mathematical Programming*, 180:285–310, 3 2020.

[81] Laurence A. Wolsey. Mixed integer programming. *Wiley Encyclopedia of Computer Science and Engineering*, 9 2008.

[82] Morton Slater. *Lagrange Multipliers Revisited*, pages 293–306. Springer Basel, 7 2014.

[83] John Forrest, Ted Ralphs, Stefan Vigerske, LouHafer, Bjarni Kristjansson, jpfasano, EdwinStraver, Miles Lubin, Haroldo Gambini Santos, rlougee, and Matthew Saltzman. coin-or/cbc: Version 2.9.9, jul 2018.

[84] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 3 2006.

[85] Jorge Nocedal, Andreas Wächter, and Richard A. Waltz. Adaptive barrier update strategies for nonlinear interior methods. *SIAM Journal on Optimization*, 19:1674–1693, 1 2009.

[86] E. L. Lawler and D. E. Wood. Branch-and-bound methods: A survey. *Operations Research*, 14:699–719, 8 1966.

Appendix A

Constraint Formalism

The *Python* optimisation program is presented below. Listing A.1 shows the initialisation of the variables used. The set up of the constraint structure is then shown at listing A.2. These were incorporated in a larger `class` object that handled data processing and plotting.

Listing A.1: Initialisation of Variables and Model Structure

```python
import pyomo.environ as env

# initialize the model
m = env.ConcreteModel()

m.t = env.Set(initialize= ic.seasonal_data.index)
t_i = m.t[1]
t_f = m.t[-1]

# model parameters; data passed into the model to initialize the variables
rain_data = dict((time_stamp, ic.seasonal_data['Rain'][time_stamp]) for time_stamp in m.t)
m.rainfall = env.Param(m.t, initialize=rain_data)
avaliable_1 = dict((time_stamp,(ic.seasonal_data['Flow_1'][time_stamp]/
                                max(ic.seasonal_data['Flow_1']))) for time_stamp in m.t)
m.supply_1 = env.Param(m.t,initialize= avaliable_1)
avaliable_2 = dict((time_stamp,(ic.seasonal_data['Flow_2'][time_stamp]/
                            max(ic.seasonal_data['Flow_2']))) for time_stamp in m.t)
m.supply_2 = env.Param(m.t,initialize= avaliable_2)
crop_coefficient = dict((time_stamp, ic.seasonal_data['Crop_Factor'][time_stamp])
                        for time_stamp in m.t)
m.crop_factor = env.Param(m.t, initialize=crop_coefficient)
# crop factor is built into the potential ET
et_data = dict((time_stamp, ic.seasonal_data['Crop_Factor'][time_stamp]*
                ic.seasonal_data['ref_ET'][time_stamp]) for time_stamp in m.t)
m.et_o = env.Param(m.t, initialize=et_data)

# ------ variables -----
m.soil_moisture = env.Var(m.t, bounds=(0, ic.field_capacity))
m.drainage = env.Var(m.t, bounds=(-2*ic.field_capacity, 0))
m.water_loss = env.Var(m.t, bounds=(-ic.irrigation_system_limit, 0))
m.efficiency = env.Var(m.t, bounds=(0.0,1.0))
m.evapotranspiration = env.Var(m.t, bounds=(0, ic.field_capacity))
m.storage_size = env.Var(m.t, bounds=(0, 10*ic.storage_size))
m.storage_volume = env.Var(m.t, bounds=(0, ic.storage_size))
m.store_1 = env.Var(m.t, bounds=(0, ic.storage_size))
m.store_2 = env.Var(m.t, bounds=(0, ic.storage_size))
m.irrigation_1 = env.Var(m.t, bounds=(0,2*ic.irrigation_system_limit))
m.irrigation_2 = env.Var(m.t, bounds=(0,2*ic.irrigation_system_limit))
m.storage_irrigation = env.Var(m.t, bounds=(0, 2*ic.irrigation_system_limit))
```

Listing A.2: Constraint Structure

```python
# -------  Binary decision to irigate or not -----
m.irrigation_decision = env.Var(m.t, domain=env.Binary)
# --- retern period constraint ---
m.cooldown = env.Constraint(m.t,
   rule=lambda m, k: sum(m.irrigation_decision[k + datetime.timedelta(days=s)] \
      for s in range(ic.return_period+1) if k + datetime.timedelta(days=s) <= t_f) <= 1)
# ----------- on-farm storage ------------
# m.store_1_zero = env.Constraint(m.t, rule=lambda m, k: m.store_1[k] >= 0.0)
# m.store_2_zero = env.Constraint(m.t, rule=lambda m, k: m.store_2[k] >= 0.0)
m.storage_volume_limit = env.Constraint(m.t,
    rule=lambda m, k: m.storage_volume[k] <= ic.storage_size)
m.daily_store_limit_1 = env.Constraint(m.t,
    rule=lambda m, k: m.store_1[k] <= ic.storage_system_limit_1 * m.supply_1[k])
m.daily_store_limit_2 = env.Constraint(m.t,
    rule=lambda m, k: m.store_2[k] <= ic.storage_system_limit_2 * m.supply_2[k])

# ----------- system limits ---------------
# ----------- irrigation and restrictive constraints --------------
m.daily_irrigation_upper_1 = env.Constraint(m.t,
    rule=lambda m, k: m.irrigation_1[k] <= ic.irrigation_system_limit*m.supply_1[k])

m.daily_irrigation_upper_2 = env.Constraint(m.t,
    rule=lambda m, k: m.irrigation_2[k] <= ic.irrigation_system_limit*m.supply_2[k])

# m.daily_irrigation_2 = env.Constraint(m.t,
    # rule=lambda m, k: m.irrigation_2[k] <= ic.irrigation_system_limit*m.bulk_buy_in_decision[k])

m.irrigation_system = env.Constraint(m.t,
    rule=lambda m, k: (m.irrigation_1[k] + m.irrigation_2[k] +
                      m.storage_irrigation[k]) <= m.irrigation_decision[k]* ic.irrigation_system_limit)

# ----------- storage irrigation --------------
m.daily_storage_irrigation_upper = env.Constraint(m.t,
    rule=lambda m, k: m.storage_irrigation[k] <= m.storage_volume[k])

m.ratio_cap = env.Constraint(m.t,
    rule=lambda m, k: (m.irrigation_1[k] + m.irrigation_2[k] +
                      m.storage_irrigation[k]) <= 2*(ic.field_capacity - m.soil_moisture[k]+0.01))

m.pump_1_limits = env.Constraint(m.t,
    rule=lambda m, k: m.store_1[k] + m.irrigation_1[k] <= ic.storage_system_limit_1*m.supply_1[k])
m.pump_2_limits = env.Constraint(m.t,
    rule=lambda m, k: m.store_2[k] + m.irrigation_2[k] <= ic.storage_system_limit_2* m.supply_2[k])

# ----------- update storage volume ------------
m.storage_volume[t_i].fix(ic.initially_stored)
def storage_update(m, k):
    if k < t_f:
        return m.storage_volume[k+datetime.timedelta(days=1)] == m.storage_volume[k] +\
                                m.store_1[k] + m.store_2[k] - m.storage_irrigation[k]
    else:
        return env.Constraint.Skip
m.daily_pond = env.Constraint(m.t,
    rule=storage_update)

# ---------- water loss ---------
def efficiency_calculation(m, k):
    deficit = ((m.irrigation_1[k] + m.irrigation_2[k] + m.storage_irrigation[k])\
              /(ic.field_capacity - m.soil_moisture[k]+0.01))
    return m.efficiency[k]  <= (dist.poly(deficit, *ic.popt))
m.efficiency_constraint = env.Constraint(m.t,
    rule=efficiency_calculation)
m.water_loss_calculation = env.Constraint(m.t,
    rule=lambda m, k: m.water_loss[k] <= -(1-m.efficiency[k])*
        (m.irrigation_1[k] + m.irrigation_2[k] + m.storage_irrigation[k]))

# ----------- drainage -----------
m.drainage_calculation = env.Constraint(m.t,
    rule=lambda m, k: m.drainage[k] <= 0.0)

# ----------- evapotranspiration ----------
def evapotranspiration(m, k):
    reduction = (1.0 - env.exp(-ic.soil_r * ((m.soil_moisture[k]+0.1) / ic.field_capacity))) / \
        (1.0 - 2.0 * env.exp(-ic.soil_r) + env.exp(-ic.soil_r * ((m.soil_moisture[k]+0.1) /\
        ic.field_capacity)))
    return m.evapotranspiration[k] <=  (m.et_0[k]) * (reduction)
m.daily_et = env.Constraint(m.t,
    rule=evapotranspiration)

# ----------- soil moisture ----------------
m.soil_moisture[t_i].fix(ic.initial_soil_moisture)
def irricalc(m, k):
    if k < t_f:
        return m.soil_moisture[k + datetime.timedelta(days=1)] == m.soil_moisture[k] +\
        ((m.storage_irrigation[k] + m.irrigation_1[k] + m.irrigation_2[k])) +\
                    m.rainfall[k] + m.drainage[k] - m.evapotranspiration[k] + m.water_loss[k]
    else:  # termination condition
        return env.Constraint.Skip
m.sm_balance = env.Constraint(m.t,
    rule=irricalc)
```

```
# ----------- model objective ------------
m.pi = env.Objective(
    rule=lambda m:
    ic.yield_profit*ic.potential_yield*
    (1-ic.crop_sensitivity*(1-(sum(m.evapotranspiration[k] for k in m.t)/\
        (sum(m.et_0[k] for k in m.t)+0.0001))))+\
            ic.drainage_cost*sum(m.drainage[k] for k in m.t) +\
            ic.drainage_cost*sum(m.water_loss[k] for k in m.t) -\
                ic.irrigation_cost_1*sum(m.irrigation_1[k] for k in m.t) -\
                ic.irrigation_cost_2*sum(m.irrigation_2[k] for k in m.t) -\
                ic.stored_irrigation_cost*sum(m.storage_irrigation[k] for k in m.t) -\
                    ic.store_cost_1*sum(m.store_1[k] for k in m.t) -\
                    m.storage_size[t_i] -\
                    ic.store_cost_2*sum(m.store_2[k] for k in m.t),
        sense=env.maximize)

# ----------- extra constraints -------
m.field_capacity_limit = env.Constraint(m.t,
    rule=lambda m, k: m.soil_moisture[k] <= ic.field_capacity)
```