A Comparison of Spillover Effects before, during and after the 2008 Financial Crisis

Alethea Rea[1,2], William Rea[1,3], Marco Reale[1], Carl Scarrott[1],

1. Department of Mathematics and Statistics, University of Canterbury, New Zealand

2. Data Analysis Australia, Perth

3. Department of Economics and Finance, University of Canterbury, New Zealand

November 20, 2013

Abstract

This paper applies graphical modelling to the S&P 500, Nikkei 225 and FTSE 100 stock market indices to trace the spillover of returns and volatility between these three major world stock market indices before, during and after the 2008 financial crisis. We find that the depth of market integration changed significantly between the pre-crisis period and the crisis and post-crisis period. Graphical models of both return and volatility spillovers are presented for each period. We conclude that graphical models are a useful tool in the analysis of multivariate time series where tracing the flow of causality is important.

Keywords: Volatility spillover, graphical modelling, financial crisis, causality

1 Introduction

The world's financial markets are becoming increasingly integrated through the use of high speed telecommunications and computer networks both for the dissemination of financial information about assets traded and for trading in these

markets. Thus traders must be aware not only of direct influences on their domestic markets but events in foreign markets which may be transmitted to their domestic markets, the so-called contagion or spillover effects.

A mechanism for spillover effects proposed by [1] is that often the underlying information which drives prices may not be immediately available to a trader but pricing information itself can be obtained in near real time. Thus in the absence of such information, the prices which other traders are willing to pay for an asset may be used as a proxy for the missing information. For example, a trader engaged in buying or selling on a London exchange may believe prices of similar assets traded in New York or Tokyo are a proxy for relevant information which is not directly available to him. Such an explanation is consistent with models of rational expectations equilibrium in which market prices reveal all relevant information, see [2]. But such an explanation also encompasses the case where mistakes or idiosyncratic changes in one market are transmitted to other markets thus increasing volatility. Because volatility is a key element in pricing derivatives such as options, understanding the influence of volatility in foreign markets on a trader's domestic market is important for the implementation of trading strategies, independent of whether the volatility is driven by new, but currently unavailable, information or for other reasons.

In order to understand return and volatility transmission between assets traded in financial markets a multivariate model is essential for multiple markets. There is the AutoRegressive Conditional Heteroskedastic (ARCH) family of models and a graphical modelling approach. We build upon the work of [3] and others, however, often these works did not address the issue that not all coefficients in a vector model are statistically significant, thus a model so identified may well be over-specified. Our contribution by using graphical modelling is to highlight a framework within which the statistically significant variables and lags may be identified.

Generalised ARCH (GARCH) models assume no shift in volatility occurs in the sample period, as noted in [4]. This leads to models which overestimate the persistence of the volatility - the so-called long memory effect (see [5] for a discussion in a univariate context) - and reduced effectiveness as forecasting tools. As a consequence, a two-step modelling process is required. First, any structural breaks in the data must be identified. Second, return and volatility transmission is modelled within the identified regimes. In this paper we advocate checking for structural breaks in the data in both returns and volatility and then modelling the regimes identified.

As indicated above, underlying the modelling of return and volatility transmission is the assumption that the indices represent a summary statistic of all currently available price sensitive information. This assumption allows the modelling of

return and volatility transmission by only examining market returns and volatilities without the need to have access to these information flows or to quantify their effects.

This paper studies return and volatility transmission taking into account structural breaks and using graphical modelling to analyse each identified regime. Graphical modelling is a multivariate technique which is widely applied in other branches of statistics where identifying the structure of the relationships and the flow of causality between variables is important. A graphical model of stock market returns or volatilities obtained from their indices objectively tests the potential influences on an index from its own past and other indices, including contemporaneous relationships.

In other work authors who used standard vector autoregressions to model volatility spillover often did not address the issue that not all coefficients in a vector model are statistically significant, thus a model so identified may well be over-specified. Graphical modelling provides a framework within which the significant variables and lags may be identified.

The remainder of the paper is structured as follows. Section (2.1) reviews literature on volatility spillover while Section (2.2) briefly reviews graphical modeling. Section (3) outlines the use of graphical modeling in the context of financial time series analysis. Section (4) presents an application to both return and volatility spillover among the Standard and Poor's Composite 500, FTSE 100 and the Nikkei 225 stock market indices. Section (5) contains the discussion and Section (6) the conclusions.

2 Literature Review

2.1 Volatility Spillover

In order to understand volatility transmission between assets traded in financial markets a multivariate model is essential. Previous investigations into spillover effects commonly used models from the Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) [6] family. For example, [7] used a GARCH(1,1) model while [8] used an M-GARCH model. These models offer insights into volatility within and between markets.

[4] noted that the GARCH models assume that no shift in volatility occurs in the sample period. When studying volatility transmission the most common of these are the multivariate AutoRegressive Conditional Heteroskedastic (ARCH) family

3

of models. The ARCH model was originally proposed as a univariate method by [9] and extended by [6] to the Generalized ARCH (GARCH) model. Subsequent developments lead to a range of extensions to include, among others, exponential GARCH [10], fractionally integrated GARCH [11] and a number of multivariate models, see [12] for a discussion of a number of these.

As noted in the introduction, assuming no shift in volatility leads to models which overestimate the persistence of the volatility, the so-called long memory effect [5], and consequently reduced effectiveness as forecasting tools.

Studies in the area of volatility spillover have reported evidence of (sometimes bidirectional) return and volatility spillover from major to minor markets and between major markets. In all cases the authors were dealing with multivariate data and either explicitly state or implicitly assume that the spillover from one market to another was causal. The dual consideration of multivariate data and the direction of causation makes the use of graphical modelling an ideal tool.

These previous studies include [1] who investigated volatility transmission between the Dow Jones (US), Nikkei (Japan) and FT30 (London) indices in the period surrounding the stock market crash in October 1987 to study the validity of the so-called "contagion" model of price movements. They concluded that the influence of volatility in one stock market's price, such as occurred on the New York exchange in the October 1987 crash, increases as the level of integration between markets increases and hence results in greater volatility transmission between markets. Furthermore they included [13] who reported short run interdependence of the S&P 500, Nikkei, and FTSE indices. They reported spillover effects from the S&P 500 and FTSE to the Nikkei in the study period but not from the Nikkei to the S&P 500 and FTSE. They also reported the strength of the spillover effects varied with time. A suite of other papers ([3], [14], [8], [15], [7] and [16]) provided further insights into volatility spillover between at least one of these three major markets and other financial centers. In a study of the stock markets of Bahrain, Kuwait and Saudi Arabia, [17] claim to make the first application of structural time-series modeling to volatility spillover, an approach related to that which we take here.

As indicated above, underlying the modelling of return and volatility transmission is the assumption that the indices represent a summary statistic of all currently available price sensitive information. This assumption allows the modelling of return and volatility transmission by only examining market returns and volatilities without the need to have access to these information flows or to quantify their effects.

4

2.2   Graphical Models

Graphical models are an important tool for analyzing multivariate data. Statisticians often ignore issues of causality preferring instead to leave such matters to subject specialists. However, any model constructed for the purpose of prediction or forecasting (as many time series models are) implicitly assumes that either the variables used for prediction or forecasting directly measure the causal mechanism(s) or that they are sufficiently good proxies that they can be used for prediction without undue caution. Graphical models provide an excellent framework for dealing with issues of causal relationships. The roots of such graphs can be traced as least as far back as [18]. Much of the large body of research literature has been summarized in the recent monographs and texts of [19], [20], and [21]. In these works the basic notation is developed and an overview of the different methods are presented. These works do not include time series data. Two approaches to using graphical modelling with time series data have been presented; a frequency domain approach by [22], and a time domain approach by [23], and [24, 25].

Here we briefly outline the important concepts of a conditional independence graph (CIG), a directed acyclic graph (DAG) and the process of *moralization*.

In graph theory terminology a graph is a pair $G = (V, E)$ where the elements of V are called vertices (or nodes) and the elements of E are called edges (or lines). In graphical modeling the vertices represent variables and the edges represent relationships between the variables. The graph $G = (V, E)$ in Figure 1 has a set of three vertices, $V = \{A, B, C\}$ and a set of two directed edges, $E = \{AC, BC\}$. Vertices A and B are called parents of C while C is called the child of A and B. Figure 1 is a directed acyclic graph because all edges are directed but there is no cycle.

It is often the case with highly correlated sets of variables, that some variables do not make a significant contribution to prediction in the presence of other predictors, although they are correlated with the predicted variable. Because of this we now introduce conditional independence. Statistically, if A, B and C are random variables and A and B are conditionally independent given C, which we write as $A \perp\!\!\!\perp B|C$, then the probability can be factorized

$$f_{a,b|c}(a, b|c) = f_{a|c}(a|c) f_{b|c}(b|c) \tag{1}$$

Conditional independence between A and B given C is seen in a graph when A and B are connected by (directed or undirected) edges to C but not to each other as in Figure 1.

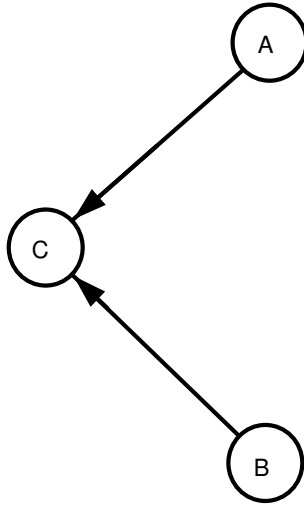Graphical modelling creates a conditional independence graph (CIG). A CIG is

Figure 1: A simple graph with variables A, B and C and two directed edges.

a graph with only the edges which represent the significant partial correlations. The zero partial correlations indicate that the the two variables are independent given all of the other variables.

A simple example will be used to explain this. If we have two series and the order of the vector autoregressive model (explained in Section 3) is one then we allow the graph to have edges between $t - 1$ and t. To extend the model to include contemporaneous relationships a link between Series 1 and Series 2 at time t is allowed.
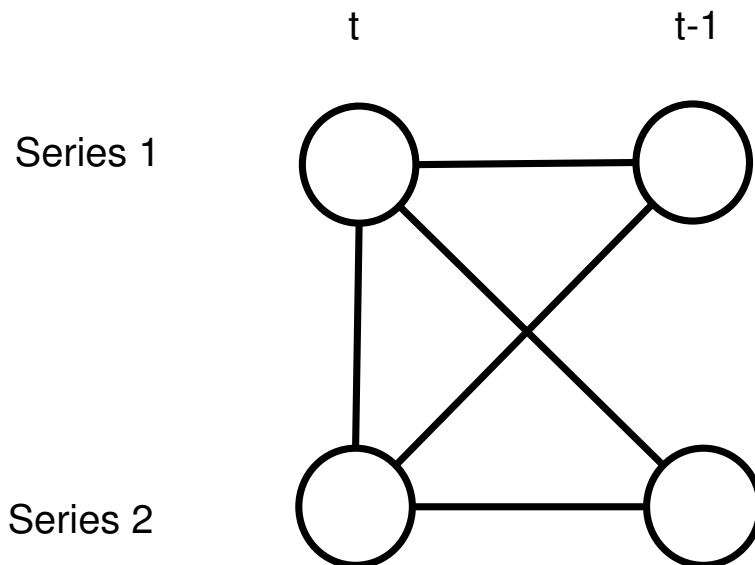

Figure 2: A simple graph with two series and a single time lag.

The model in Figure 2 is saturated. This means the model has all the admissible edges. Admissible edges are those which connect vertices representing the past with today and all contemporaneous edges.

The *moral graph* associated with the directed graph $G^\prec = (V, E^\prec)$ is the undirected graph $G^m = (V, E^m)$ on the same vertex set but with an edge set which includes all the edges in $E^\prec$ plus all necessary edges required to eliminate so-called forbidden Wermuth configurations from $G^\prec$. In essence it forbids sub-graphs of the type in Figure 1. These are called a moral graphs because they "marry" parent nodes, a term due to [26].

Further details on CIGs and DAGs can be found in [23], [20], and [21].

3 Graphical modelling for financial time series

Graphical modelling in a time series context seeks to find causal links between past and present observations. In addition, it also allows the study of causality among contemporaneous variables. Graphical modelling applies to vector autoregressive moving average (VARMA) models of the form

$$x_t = c + \Phi_1 x_{t-1} + \ldots + \Phi_p x_{t-p} + \Psi_1 \epsilon_{t-1} + \ldots + \Psi_q \epsilon_{t-q} + \epsilon_t \tag{2}$$

where x_t is a $m \times 1$ vector of variables measured at time t. Also, ϵ_t is assumed to be normally and identically distributed with a mean of zero and general covariance matrix, Ω.

To allow for contemporaneous relationships both sides of Equation (2) must be multiplied by Φ_0 as follows

$$\Phi_0 x_t = d + \Phi_1^* x_{t-1} + \ldots + \Phi_p^* x_{t-p} + \Psi_1^* \epsilon_{t-1} + \ldots + \Psi_q^* \epsilon_{t-q} + a_t. \tag{3}$$

Considering only the autoregressive components Equation (3) reduces to

$$\Phi_0 x_t = d + \Phi_1^* x_{t-1} + \ldots + \Phi_p^* x_{t-p} + a_t. \tag{4}$$

Two restrictions apply to Equation (4). The first is that the variance matrix of $a_t = \Phi_0 \epsilon_t$ is diagonal and the second is that Φ_0 is upper triangular with a unit diagonal. Φ_0 represents the causal dependence of each variable on its contemporaneous counterparts.

Graphical modelling involves firstly finding the conditional independence graph (CIG) and secondly finding the directed acyclic graph (DAG). Determining the

CIG involves three steps: (1) calculating the pairwise correlations or preliminary Φ and Ψ coefficients, (2) determining the statistically significant coefficients (3) and defining the graph.

The CIG is determined by nodes, indexed by the series and the lag, and edges representing statistically significant relationships. When applying graphical modelling to time series the first step is to determine the order, or number of lags, in the model. This defines the nodes of the graph. The set of admissible edges contains only those edges from the lagged nodes to the present nodes and all possible contemporaneous relationships. A preliminary set of significant edges is given by the non-zero partial correlations. A partial correlation between two variables is equivalent to their correlation with the linear dependence of both of them and the remaining variables subtracted. The set of statistically significant edges are those whose partial correlation differs significantly from zero.

A CIG is a statement about a single joint distribution. A CIG does not allow one to make statements about causality, that is, one cannot make statements about which events have directly influenced others. DAG's do allow such statements to be made. The edges of a DAG contain arrows from the cause nodes to the effect nodes. A DAG represents marginal conditional relationships and results from inferring causality. Creating a set of marginal conditional relationships from a joint distribution is not unique meaning that a single CIG can give rise to many DAGs.

Assuming for the moment that we know the DAG we can determine the CIG required to represent the relationships. In each scenario involving two 'cause' or parent nodes and a single 'effect' or child node the DAG has two edges with the arrows pointing towards the child node. To represent this without directed edges the parental nodes must be connected. This process is called moralisation.

Given the CIG created by the procedure outlined above the final step in graphical modelling is to convert the CIG to a DAG; a process called demoralisation. For time series applications causality is a direct consequence of time because the past influences the present and not the other way around; this determines the direction of the arrows. With contemporaneous relationships the direction of the arrows is determined by an information criterion such as the AIC [27]. The resulting DAG may have some moral links remaining.

In a financial context the VAR model is very similar to the unconstrained multivariate autoregressive conditional heteroskedasticity (ARCH). The form of the unconditional ARCH(m) model is

$$
\begin{aligned}
E(Y_t|Y_{t-1}) &= 0 \\
Var(Y_t|Y_{t-1}) &= H_t
\end{aligned}
$$

where

$$H_t = c + A_1 Y_{t-1} Y'_{t-1} + \ldots + A_p Y_{t-p} Y'_{t-p}$$

$$= c + \sum_{i=1}^{p} A_i Y_{t-i} Y'_{t-i}, \tag{5}$$

while the form of the VAR model is by

$$x_t = c + \Phi_1 x_{t-1} + \ldots + \Phi_p x_{t-p} + \epsilon_t. \tag{6}$$

The VAR(p) model and the ARCH(m) have the same form as can be seen by setting $\Phi_i = A_i$ and $x_{t-1} = Y_{t-i} Y'_{t-i}$ in Equation (3). While ARCH(m) is parameter rich, graphical modelling chooses the best model according to an information criterion. The VAR(p) as fitted by graphical modelling is usually relatively sparse.

In a financial context the VARMA model of Equation (3) does not correspond directly to a commonly used ARCH or generalised ARCH (GARCH) type model. The model deals only with variances and their associated estimation error from the previous time lag and as such is not a GARCH model because it does not model covolatilities. The model is however more than a collection of ARCH models because the observed variances are modelled based on the observed variances of all the series under consideration.

4 Financial Integration Example

In this section we use graphical modelling as a tool for studying spillover effects. The analysis has three phases;

1. A visual inspection of the three time series, Section (4.2).

2. The determination the structural breaks and consequently the regimes, Section (4.3).

3. An analysis of each regime, Section (4.4).

We begin with a description of the data set.

4.1   The Data

To investigate spillover effects three stock market indices were used, namely: Standard & Poor's Composite 500 for the USA, FTSE 100 for the UK and the Nikkei 225 for Japan. The data were downloaded from Datastream for the period 1 January 2001 to 22 August 2011. These three stock indices are ideal as they are widely followed and over a 24 hour period there is little overlap in their trading hours. The London stock exchange opens at 4am Eastern Standard Time (EST) and closes at 12noon EST. The New York stock exchange opens at 9:30am EST and closes at 4pm. The Tokyo stock exchange opens at 7pm and closes at 1am EST but here we must note that Japan is on the opposite side of the date line from New York and London. Therefore there is an overlap of two and a half hours between the London and New York exchanges. By calendar day the Japanese market is the first to open.

4.2   Visual Inspection

A plot of the values of the three indices, from 1 January 2001 until 22 August 2011, is presented in Figure 3. This graph has the breakpoints which bound our three study periods marked and labelled. See our breakpoint analysis in Section 4.3 below. A visual inspection suggests they are not three independent time series. The FTSE 100 and S&P 500 indices in particular are remarkably similar in appearance. Visually the plots show four distinct periods. In the initial period the three indices were all in decline. This was followed by period of increase for the three indices from 2003 until 2007, followed by a period of steady decline. The final period was one of relative stability for the Japanese index, while for the UK and USA there was time of increase until they were close to their early 2007 levels.

4.3   Structural Break Analysis

We used atheoretical regression trees (ART) [28] to investigate evidence for any structural breaks in the mean of the absolute values of the returns. Regression trees are widely used in many branches of statistical analysis as a non-parametric regression method, see [29] for a detailed description. A regression tree will model the relationship between the response variable and the covariates, which in time series analysis is the single variable time, by fitting piece-wise constant functions to the data. In univariate time series analysis the points at which these piece-wise constant functions change are interpreted as candidate breakpoints.
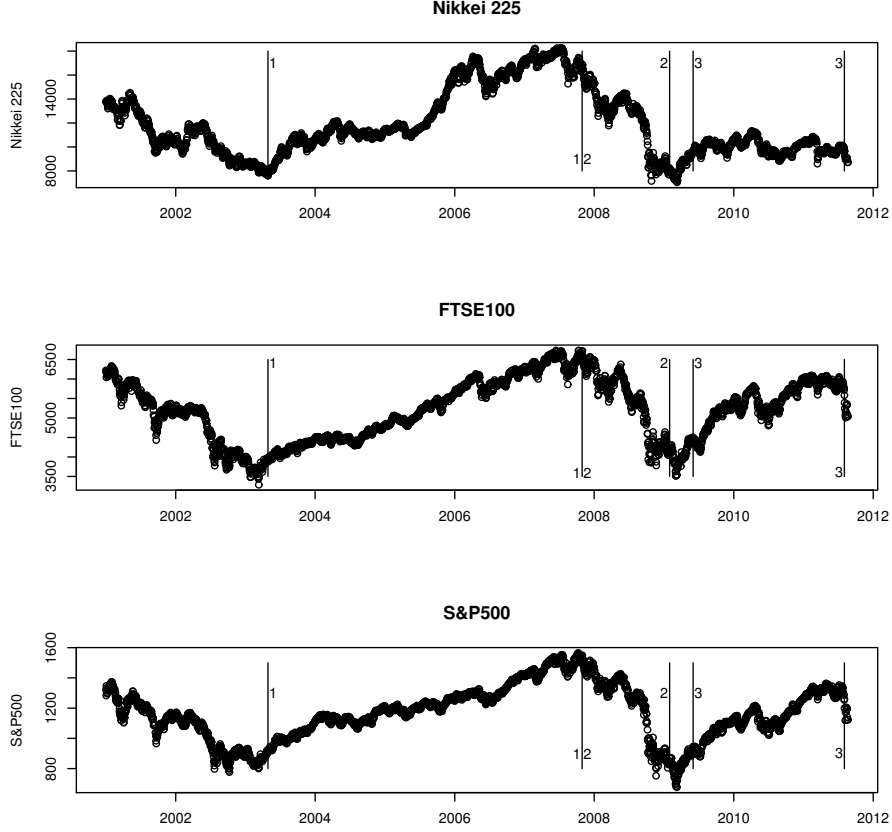
Figure 3: A plot of the values of the S&P500, FTSE100, and Nikkei 225 indices between 1 January 2001 and 22 August 2011.

We decided to use the S&P 500 series as the master series because numerous authorities consider the American markets to be the source of volatility which then spills over into other markets. The structural breaks reported by ART for the S&P 500 were used to identify the study periods.

We used the `tree` [30] package in R [31] to implement ART. A plot of the regression tree is presented in Figure (4). ART reported several breaks in the absolute values of the log return series for the S&P 500 index. Examining both the index series (Figure 3) and the absolute value of the log return series (Figure 5) together with the regression tree (Figure 4) we chose study period one to be 28 April 2003 to 29 October 2007, as this was before a noticeable rise in volatility leading into the financial crisis of 2008. ART reported two structural breaks, yielding three regimes, during the market decline and initial recovery in the period 12 September 2008 to 31 May 2009. We chose study period two to be 30 October 2007 to 11 September 2008 because this was the longest of the three regimes within this

11

period. The other two regimes contained too few data points to yield a useful graphical model of the spillover effects. We chose study period three to be 1 June 2009 to 2 August 2011, which was after the markets had experienced a significant decline and before the period of volatility associated with the credit downgrade of US Government debt began in August 2011. Within these three selected periods there were no reported structural breaks.
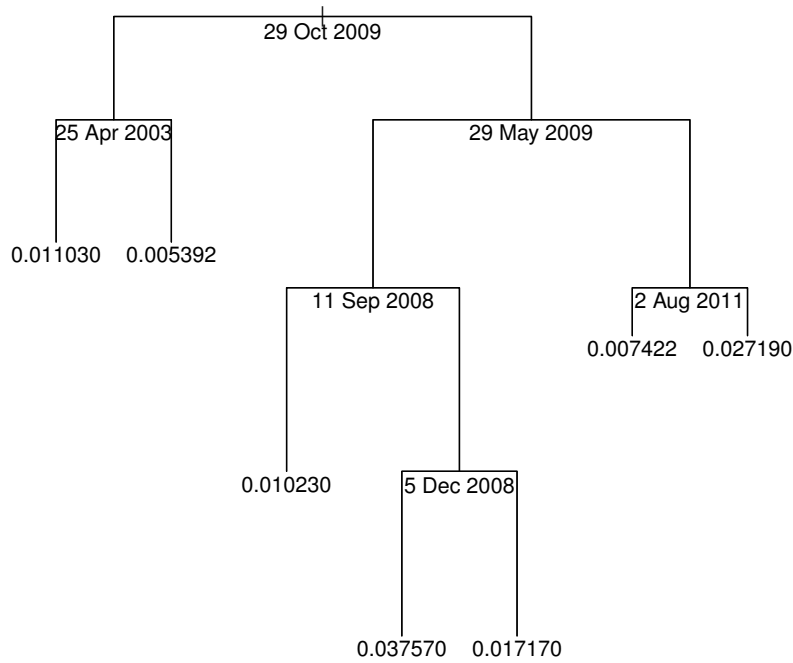

Figure 4: The regression tree for the absolute values of the returns of the S&P500 stock index.

4.4 VAR models

In this section we describe how to fit an vector autoregressive model of order p (VAR(p)) using the three stock indices as our example.

The first step is to find the order of the VAR(p) model, that is estimate the number of lags p. For each value of p considered, a VAR model of order p is fitted
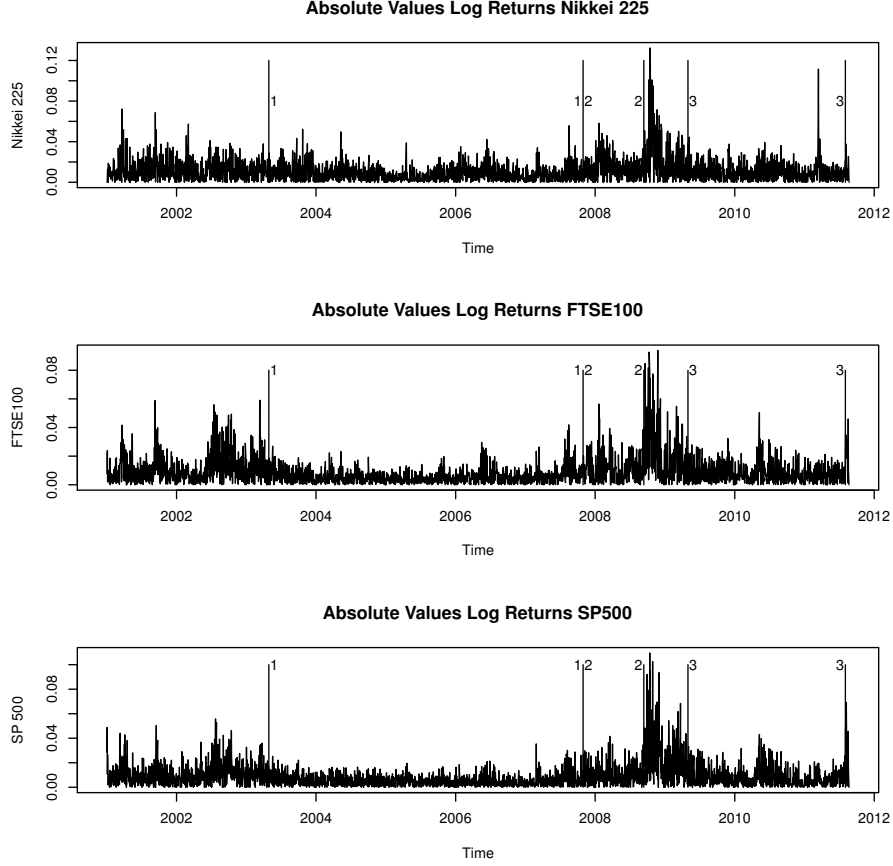
Figure 5: Plots of the absolute values of the log returns of the Nikkei 225, FTSE100, and S&P500 stock indices together with the locations of the identified structural breaks.

and its parsimony vs explanatory power trade-off was evaluated using several information criterion; Akaike's information criterion (AIC) [27] , the corrected Akaike information criterion (AICC) [32], Hannan information criterion (HIC) [33] and the Schwarz information criterion (SIC) [34]. The most reasonable p value was chosen.

We fitted graphical models to the VAR(p) models selected in the previous step for the log returns and the squared log returns lags. The squared log returns provide a measure of stock market volatility, hence are used to provide insight into volatility spillover effects.

The reported partial correlations correspond to the conditional independence graph (CIG). CIG's are converted to directed acyclic graphs (DAG's) by determining the causal relationships between the variables. In this case the causal

13

relationship is determined because of the innate temporal ordering because only the past can influence the present. In this case even the contemporaneous variables have causal relationships based on the closing times varying though the day.

For the log returns, study periods one, two and three had the same graphical model structure. This common DAG structure is presented in Figure (6). Thus the spillover of log returns did not change in structure between any of the three periods studied.

When considering the squared log return series as a proxy for volatility, for period one, the various information criteria reported optimal orders ranging from three to five lags. An order three model was selected for parsimony reasons as it should reveal the most important spillover effects. With three indices and three lags there are 30 edges in the saturated model (that is the model with all possible edges). The number of statistically significant partial correlation coefficients was 19. The DAG of this model is presented in Figure (7).

For both study periods two and three, each of the information criteria reported an optimal model order of one lag. With three indices and just one lag there are 12 edges in the saturated model. The model for study period two reported just six statistically significant partial correlation coefficients, The DAG for these six significant partial correlations is presented in Figure (8). The model for study period three reported seven statistically significant partial correlations. The DAG for these seven significant partial correlations is presented in Figure (9).

The VAR models were fitted using MATLAB. The first program requires the data from three time series as an input and returns the order of the model selected by the range of information criteria. The second program fits the CIG. It requires the same data, the model order selected in the previous step and a user-chosen t-value corresponding to the desired alpha level. We choose alpha to be 0.05 and the corresponding t-value is 1.96. Both programs were written by one of the authors and are available on request.

5 Discussion

Spillover effects are widely regarded in the literature as the result of market integration. The phenomenon refers to the general tendency for a market to move in the same direction as other markets. This ought to be particularly important for these three large markets, if the theory is correct.

A return or a volatility channel is where we have evidence of spillover effects. In

a graphical modelling context, the channel A-B is considered active if there is at least one parent-child link between the lagged nodes of market A and the present node of market B. For these three series the lag zero nodes are ordered by time.

The order of the graphical model defines the number of statistically significant lags for which there is market integration, the statistically significant partial correlations are used to define which channels are active. Partial correlations are used to define the CIG which is then converted to a DAG, as described above. A channel is considered active if one or more direct links are present from one stock index to another.

5.1 Return Spillovers

In the graphical models of return spillover all study periods exhibited the same structure, see Figure (6). It is well-known that returns are less predictable than volatilities. In fact, often univariate stock market return data are indistinguishable from noise. Our results showed that returns at lag one from the FTSE 100 and S&P 500 directly influenced, or spilled over, to the returns in all three markets studied whereas the Nikkei 225 at lag one only influenced itself and did not spillover into either of the other two major markets in any of the study periods. When considering same day returns the Nikkei 225 returns spilled over to the FTSE 100 which in turn spilled over to the S&P 500. However, the Nikkei 225 to S&P 500 channel was not open, thus the same day returns on the Nikkei 225 and the S&P 500 were independent given the returns on the FTSE 100. This differs from one of the volatility spillover models (see Figure 8 and discussion below) in which the Nikkei 225 and the S&P 500 were not independent given the FTSE 100.

So although stock market returns may appear to be unpredictable when considered in a univariate context, when considered in a multivariate context graphical modelling has identified active spillover channels indicating that returns have a degree of predictability when returns on other major markets are known. It is perhaps somewhat surprising that the structure of the graphical model did not change between study periods.

In the graphical model for the returns only the FTSE 100 and S&P 500 channels were saturated. Only two of the possible five edges originating from the Nikkei 225 nodes had statistically significant partial correlations; the contemporaneous edge from the Nikkei 225 to the FTSE 100 and the Nikkei 225 at lag one to itself. The fact that one of the two statistically significant edges is from the Nikkei 225 at lag one to itself indicates that returns on the Nikkei 225 have limited relevance to the other two major markets during the study period.

15

5.2 Volatility Spillover

For the graphical models of volatility spillover, study period one had three statistically significant lags, and consequently the largest number of lags for which there is evidence of market integration (see Figure 7). In study periods two and three the number of lags for which there is evidence of market integration was much shorter at just a single day (see Figures 8 and 9, respectively). This implies that before the 2008 financial crisis the spillover effects were longer lived and the indices responded to recent events over a greater time period than in the crisis and post-crisis periods.

With three indices and three lags a saturated model would have thirty edges among the twelve nodes, but in Figure (7) we have 19 edges representing the statistically significant partial correlations. Because of their ordering in time the Nikkei 225 could have 11 directed edges pointing away from its four nodes (three edges from each of the lagged nodes and two from the contemporaneous node), the FTSE100 could have 10 edges and the S&P 500 nine. Of these, six of the 11 possible directed edges from the Nikkei 225 are statistically significant, six of the 10 possible directed edges from the FTSE100 are statistically significant while seven of the nine possible directed edges from the S&P 500 are statistically significant. This structure confirms conventional wisdom that the US market is the most influential of the three during the first study period.

The graphical model reported for period two (see Figure 8) was of order one and would have 12 edges in the saturated model (three from each of the lagged nodes, two from the contemporaneous Nikkei 225 and one from the contemporaneous FTSE 100). Because of the ordering in time the Nikkei 225 could have five directed edges, the FTSE 100 four directed edges and the S&P 500 three. Of these, two were statistically significant for the Nikkei 225 (both contemporaneous channels), three for the FTSE 100 and one for the S&P 500.

The second study period was a period of market turmoil and decline and the graphical model indicates that only the most immediate market information from each market was relevant with one exception. That is the directed edge from the FTSE 100 at lag one to the S&P 500. (Note that the directed edge from lag-1 S&P500 to the Nikkei is the most recent data from the S&P 500 available when the Tokyo Market opens.) The directed edge from the S&P 500 at lag one to the Nikkei 225 represents the most recent information available from the S&P 500. This model confirms street wisdom that traders in periods of market turmoil to have concentrate only on the most recent events. Fast breaking news is assimilated quickly.

The second study period has the fewest statistically significant partial correla-

tions. With six statistically significant partial correlations, there is no directed edge in the model from the S&P 500 to the FTSE 100. This may be regarded surprising as the US market is generally considered to be one of the most influential markets (see Figure 8). However, one should not interpret this graph to mean that the volatility S&P 500 at lag one did not spill over to either the FTSE 100 or the S&P 500 the following day but rather that the spillover was mediated via the Nikkei 225. Thus during the second study period the volatility of the current trading day of the FTSE 100 (and the S&P 500) was independent of the volatility of the S&P 500 at lag one given the contemporaneous volatility of the Nikkei 225. If the information from the Nikkei 225 were unavailable, for example the market was closed for a holiday, then a volatility spillover model different from the one presented here would be required.

The graphical model reported for period three is also of order one. The saturated model would again have 12 directed edges. Of these, seven statistically significant partial correlations were present. Three of the possible five directed edges from the Nikkei 225, one of the possible four directed edges from the FTSE 100, while all three of the possible directed edges from the S&P 500 were present in study period three. This again confirms conventional wisdom that the American markets are the most influential. There is no spillover from the FTSE 100 to the Nikkei 225 (see Figure 9).

The time period over which market integration occurs is very different. As discussed above, before the financial crisis the market responded to movements from up to three trading days prior to the current day. As the market was declining during turmoil of the 2008 financial crisis the period of market integration or spillover effects were confined to at most one previous trading day. Three of the six edges are contemporaneous, that is within a trading day. Therefore, during this period of decline half the spillover effects occurred within 24 hours or less. After the market stabilised in 2009 the effects mostly took a calendar day which is longer than within the period of decline but much shorter than before the financial crisis.

6 Conclusions

Graphical modelling is a quick and efficient way to study financial integration. It investigates the casual structure both between multiple time series and within each individual time series.

We investigated the stock market integration between the US, UK and Japan using the indices S&P 500, FTSE 100 and Nikkei 225 respectively. Using structural

break analysis three periods of interest were defined, before the financial crisis of 2008, during the crisis and after the worst of the crisis.

The novelty of our approach is that we first used ART, a structural breakpoint detection method, to determine suitable regimes and then used graphical modelling to analyse the causal spillover effects within each period.

Our key findings were that the period of market integration was much longer before the crisis of 2008 and most of the time, most of the spillover effects channels are active. None of the models presented here were saturated.

The study has some limitations. Firstly, we have not studied the relative levels of activity of the channels and secondly very small periods can not be analysed (consequently we did not analyse every regime only the three largest). Periods smaller than 300 observations are too small too analyse.

This approach has the potential to be more widely applicable to market integration analysis and to multivariate time series analysis where tracing the flow of causality among the variables is important.

7   Acknowledgements

References

[1] M. A. King and S. Wadhwani, "Transmission of Volatility between Stock Markets," *The Review of Financial Studies*, vol. 3, no. 1, pp. 5–33, 1990.

[2] M. Bray, "Rational Expectations, Information and Asset Markets: An Introduction," *Oxford Economic Papers*, vol. 37, no. 2, pp. 161–195, 1985.

[3] J. Park and A. M. Fatemi, "The Linkages between the Equity Markets of Pacific-Basic Countries and thos of the U.S., U.K., and Japan: A Vector Autoregression Analysis," *Global Finance Journal*, vol. 4, no. 1, pp. 49–64, 1993.

[4] S. H. Kang, C. Cheong, and S.-M. Yoon, "Structural Changes and Volatility Transmission in Crude Oil Markets," *Physica A*, 2011.

[5] C. G. Lamoureux and W. D. Lastrapes, "Persistance in Variance, Structural Change, and the GARCH Model," *Journal of Business and Economic Statistics*, vol. 8, no. 2, pp. 225–234, 1990.

[6] T. Bollerslev, "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, vol. 31, pp. 307–327, 1986.

[7] M. S. Mulyadi, "Volatility Spillover in Indonesia, USA, and Japan Capital Markets," MPRA Paper 16914, University Library of Munich, Germany, July 2009.

[8] M. Martins and S.-H. Poon, "Returns Synchronization and Daily Correlation Dynamics Between International Stock Markets," *Journal of Banking and Finance*, vol. 25, pp. 1805–1827, 2001.

[9] R. F. Engle, "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation," *Econometrica*, vol. 50, pp. 987–1008, 1982.

[10] D. B. Nelson, "Stationarity and Persistence in the GARCH(1,1) Model," *Econometrica*, vol. 59, pp. 318–334, 1991.

[11] R. T. Baillie, T. Bollerslev, and H. O. Mikkelson, "Fractionally Integrated Generalised Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, vol. 74, pp. 3–30, 1996.

[12] R. S. Tsay, *Analysis of Financial Time Series*. John Wiley & Sons, Inc., 2002.

[13] Y. Hamao, R. W. Masulis, and V. Ng, "Correlations in Prices Changes and Volatility across International Stock Markets," *The Review of Financial Studies*, vol. 3, no. 2, pp. 281–307, 1990.

[14] A. Ng, "Volatility Spillover Effects from Japan and the US to the Pacific-Basin," *Journal of International Money and Finance*, vol. 19, pp. 207–233, 2000.

[15] C. Savva, D. R. Osborn, and L. Gill, "Volatility, spillover Effects and Correlations in US and Major European Markets," Money Macro and Finance (MMF) Research Group Conference 2005 23, Money Macro and Finance Research Group, Sept. 2005.

[16] P. Singh, B. Kumar, and A. Pandey, "Price and Volatility Spillovers across North American, European and Asian Stock Markets," *International Review of Financial Analysis*, vol. 19, pp. 55–64, 2010.

[17] T. Al-Deehani and I. A. Moosa, "Volatility Spillover in Regional Emerging Stock Markets: A Structural Time-Series Approach," *Emerging Markets Finance & Trade*, vol. 42, no. 4, pp. 78–89, 2006.

[18] S. Wright, "Correlation and Causation," *Journal of Agricultural Research*, vol. 20, pp. 557–585, 1921.

[19] S. L. Lauritzen, *Graphical Models*. Oxford University Press, 1996.

[20] D. Edwards, *Introduction to Graphical Modelling*. Springer, second ed., 2000.

[21] J. Whittaker, *Graphical Models in Applied Multivariate Statistics*. Wiley, 2009.

[22] R. Dahlhaus, "Graphical Interaction Models for Multivariate Time Series," *Metrika*, vol. 51, pp. 157–172, 2000.

[23] M. Reale, *A Graphical Modelling Approach to Time Series*. PhD thesis, Lancaster University, 1998.

[24] M. Reale and G. Tunnicliffe-Wilson, "Identification of Vector AR Models with Recursive Structural Errors using Conditional Independence Graphs," *Statistical Methods and Applications*, vol. 10, pp. 49–65, 2001.

[25] M. Reale and G. Tunnicliffe-Wilson, "The Sampling Properties of Conditional Independence Graphs for Structural Vector Autoregressions," *Biometrika*, vol. 89, no. 2, pp. 457–461, 2002.

[26] S. L. Lauritzen and D. J. Spiegelhalter, "Local Computations with Probabilities on Graphical Structures and Their Application to Expert Systems," *Journal of the Royal Statistical Society. Series B*, vol. 50, no. 2, pp. 157–224, 1988.

[27] H. Akaike, "Information theory and an extension of the maximum likelihood principle," in *2nd International Symposium on Information Theory* (B. Petrov and F. Csaki, eds.), pp. 267–281, Akademia Kadio, Budapest, 1973.

[28] C. Cappelli, R. N. Penny, W. S. Rea, and M. Reale, "Detecting Multiple Mean Breaks at Unknown Points with Atheoretical Regression Trees," *Mathematics and Computers in Simulation*, vol. 78, no. 2-3, pp. 351–356, 2008.

[29] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*. Chapman & Hall/CRC, 1993.

[30] B. Ripley, *tree: Classification and Regression Trees*, 2011. R package version 1.0-28.

[31] R Development Core Team, *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009. ISBN 3-900051-07-0.

[32] C. M. Hurvich, J. S. Simonoff, and C. Tsai, "Smoothing parameter selection in non-parametric regression using an improved Akaike information criterion," *Journal of the Royal Statistical Society Series B*, vol. 60, pp. 271–293, 1998.

[33] E. J. Hannan and B. G. Quinn, "The Determination of the Order of an AutoRegression," *Journal of the Royal Statistical Society Series B*, vol. 41, pp. 190–195, 1979.

[34] G. Schwarz, "Estimating the Dimension of a Model," *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.
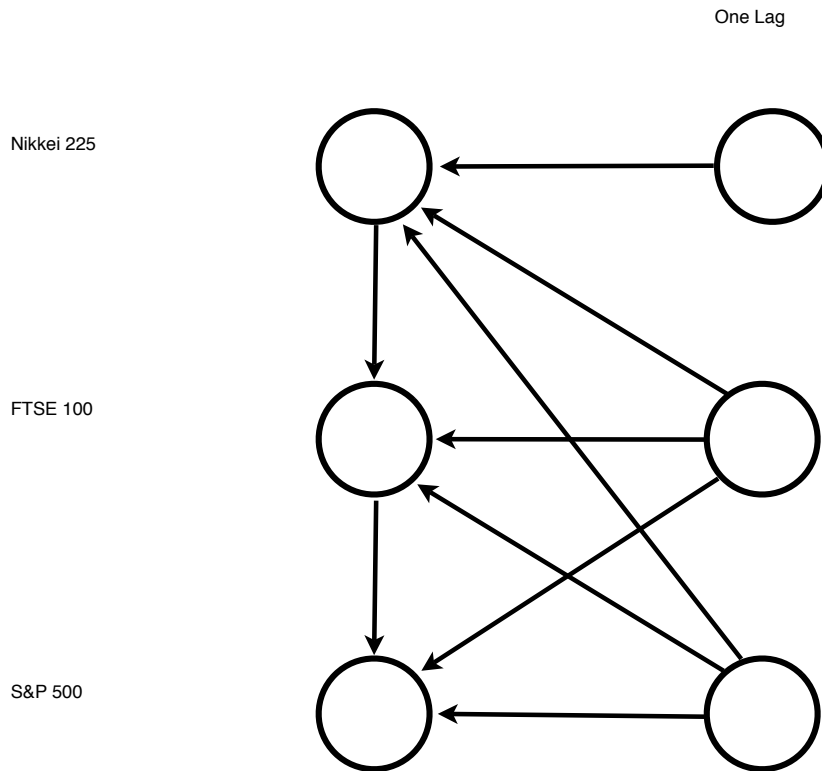
Figure 6: A graphical model of the log returns fitted to the lags. Period one (28 April 2003 to 29 October 2007), two (30 October 2007 to 11 September 2008) and three (1 June 2009 to 2 August 2011) log returns had identical models.
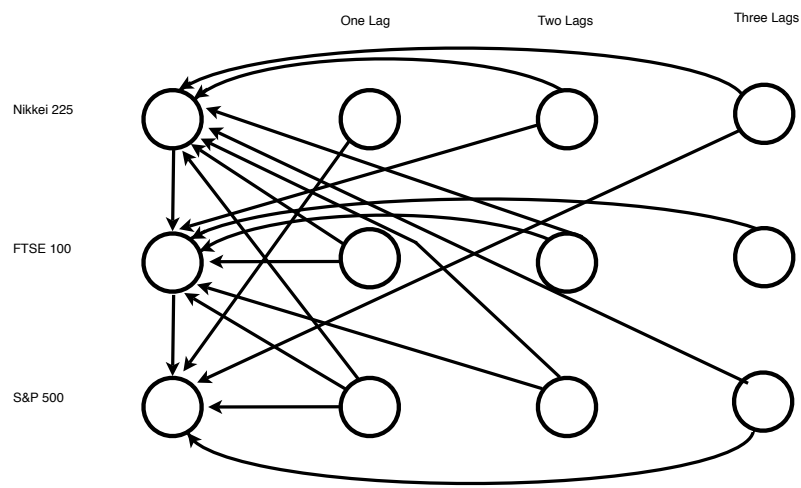
Figure 7: A graphical model fitted to the squared log returns for study period one – 28 April 2003 to 29 October 2007.
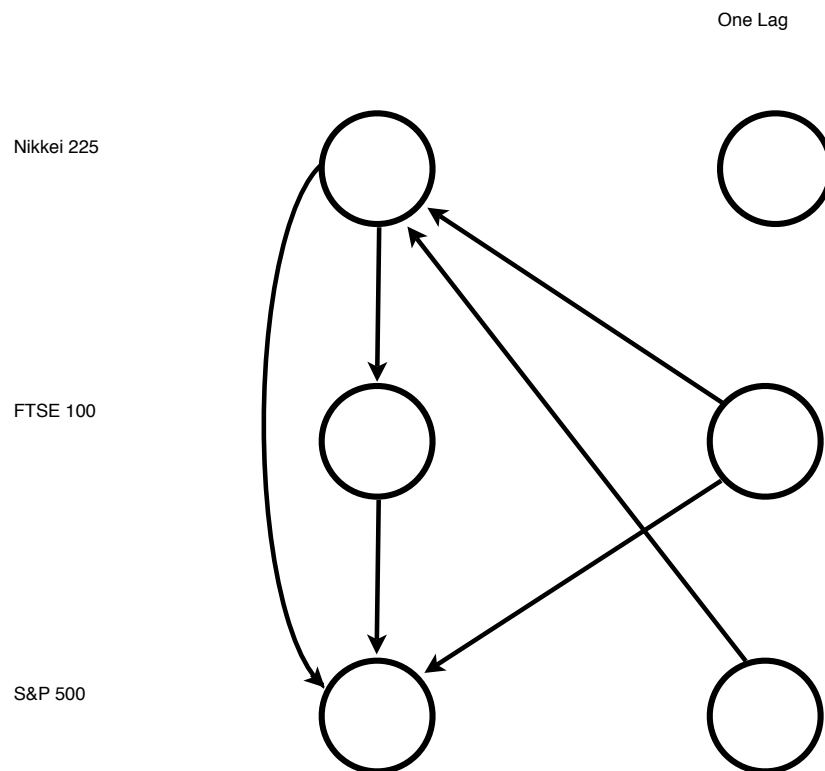
Figure 8: A graphical model fitted to the squared log returns for study period two – 30 October 2007 to 11 September 2008.
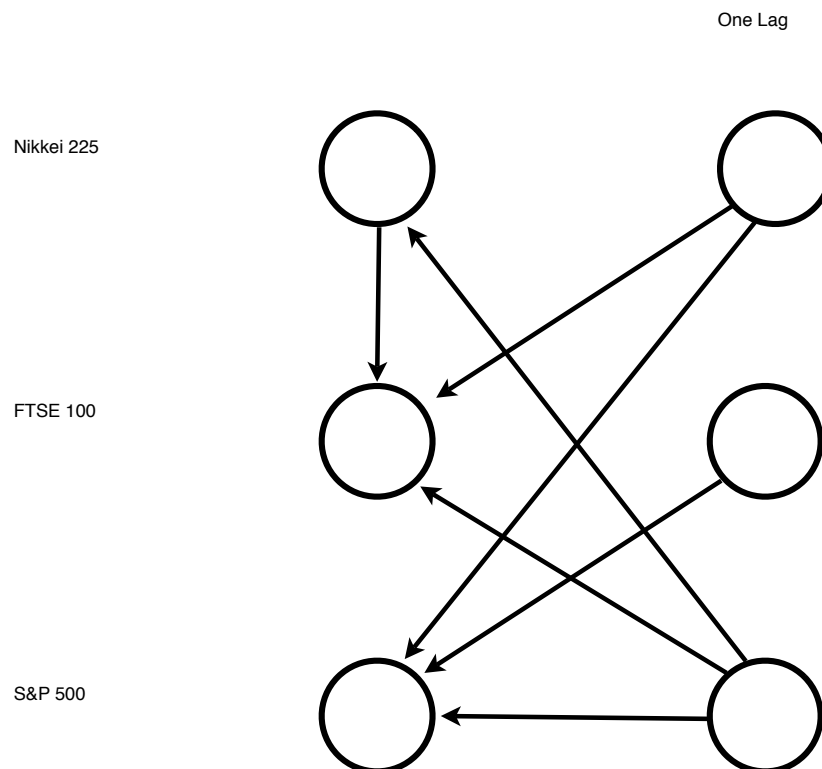
Figure 9: A graphical model fitted to the squared log returns for study period three – 1 June 2009 to 2 August 2011.