Information Flows and Stock Market Volatility

Chew Lian Chua and Sarantis Tsiaplias*

Abstract

This study examines how news is distributed across stocks. A model is developed that categorizes a stock's latent news into normal and non-normal news, and allows both types of news to be filtered through to other stocks. This is achieved by formulating a model that jointly incorporates a multivariate lognormal-Poisson jump process (for non-normal news) and a multivariate GARCH process (for normal news), in addition to a news (or shock) transmission mechanism that allows the shocks from both processes to impact inter-temporally on all stocks in the system. The relationship between news and the expected volatility surface is explored and a unique news impact surface is derived that depends on time, news magnitude and news type. We find that the effect of non-normal news on volatility expectations typically builds up before dissipating, with the news transmission mechanism effectively crowding-out normal news and crowding-in non-normal news. Moreover, in contrast to the standard approach for measuring leverage effects using asymmetric GARCH models, we find that leverage effects stem predominantly from non-normal news. Finally, we find that the capacity to identify positively or negatively correlated stock returns is ambiguous in the short term, and depends heavily on the behaviour of the non-normal news component.

Keywords: news, stock market, volatility, jumps.

JEL codes: G10, C58, C32.

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1002/ias.2649

Chua is from the School of Economics, University of Nottingham, Ningbo, China. Tsiaplias is from the Melbourne Institute of Applied Economic and Social Research, University of Melbourne, Melbourne, Australia. Corresponding author: Sarantis Tsiaplias, Melbourne Institute of Applied Economic and Social Research, University of Melbourne, Melbourne, Australia, 3010. E-mail: stsiaplias@unimelb.edu.au. We thank three anonymous referees and the editor Eric Ghysels for their many helpful comments and suggestions.

1. Introduction

Asset returns are primarily driven by the inter-relationship between the arrival of new information and the underlying mechanism that processes this information into the asset price (Glosten and Milgrom, 1985; Andersen, 1996). This paper seeks to model this mechanism in order to better understand the news transmission process, whether the transmission process for non-normal or unusual news differs from that of normal or ‘typical’ news, and how the normal and non-normal news of other stocks impact on a given stock’s expected volatility. This is important in terms of better characterizing and measuring financial volatility, and in improving our understanding of financial market behaviour.

The link between volatility and information (or news) flows is well established, with theory indicating that the volatility of an asset’s return is related to the flow of information about the asset (Ross, 1989). In other words, an asset’s volatility can be used to glean information about the impact of news on that asset. In this respect, a number of papers consider the impact of information flows for financial assets and, in some cases, how an asset’s own news can be decomposed into different news types (Campbell and Hentschel, 1992; Maheu and McCurdy, 2004; Chen and Ghysels, 2011; Patton and Verardo, 2012; Giglio and Shue, 2014).

Further research is, however, required on the important issue of information cross-overs and the inter-relationship between the information flows of sets of assets. In practice, the normal and non-normal news in univariate models is likely to reflect a combination of public and private news about the particular asset and other assets (Engle, Hansen, and Lunde, 2012; Patton and Verardo, 2012); yet, it is typically not possible to disentangle this complex web of information in order to examine how the (normal and non-normal) information flows of other assets influence the volatility of a particular asset, and to learn about the impact of an asset’s own normal news relative to the impact of non-normal news from other assets.

We explore these issues by formulating and estimating a generalized model of news transmission and by deriving an inter-temporal news impact surface that is able to express an asset’s expected volatility in terms of its own normal and non-normal news and the (normal and non-normal) news of other assets. The model jointly decomposes shocks into those that can be explained by a multivariate lognormal-Poisson jump process and a multivariate GARCH (mGARCH) process, and allows both shocks to be transmitted across the asset space by reference to a news transmission factor. We follow Maheu and McCurdy’s (2004) terminology in interpreting the shocks in the jump process as reflecting non-normal news and the shocks in the GARCH process as reflecting normal news.¹ Accordingly, the news transmission mechanism filters both types of news and

¹News can be about ‘real’ events or nothing more than speculation (or noise) (see, further, DeLong

disseminates them across the asset space such that the volatility of an asset is potentially influenced by its own news and the normal and non-normal news of other assets.

The paper explores three particular features of asset returns, being: spillovers and interdependencies between asset returns, volatility clustering and leverage effects (Bollerslev, 1986; Engle and Kroner, 1995; Karolyi, 1995; Engle, 2002; Chen and Ghysels, 2011; Kuhnén, 2015). In particular, we generalize the notion of a spillover by allowing both the normal and non-normal news processes to have their own covariance structures. A news transmission factor then acts as a conduit for filtering and disseminating normal and non-normal news throughout the asset space. The model allows for two forms of multivariate clustering (being clustering associated with both normal and non-normal information flows), and complements the research on asymmetric effects by incorporating good versus bad response asymmetries into the non-normal news process. This is achieved by conditioning the jump intensities on the sign of the non-normal news, thereby allowing each asset's jump intensity to respond differently to negative non-normal news. The resulting model is unique in allowing for two potential sources of leverage effects that can stem from normal or non-normal news.

The paper also derives analytical news impact surfaces that are novel in explicitly accounting for news heterogeneity in order to better understand the impact of information flows on asset volatilities. The surfaces incorporate new features that are important in determining the contemporaneous and inter-temporal information cross-overs between assets and understanding how information flows generate the observed heteroscedasticity and leptokurtosis of asset returns. In this respect, the expected volatility of an asset's return potentially responds to its own normal and non-normal news, and to the news of other assets. Expected volatility is also influenced by whether the normal or non-normal news is 'good' or 'bad' and this produces interesting results, particularly when 'bad' news from a set of assets results in lower volatility expectations for a second set of assets that can therefore be treated as substitutes. To the best of our knowledge, this type of surface is explored for the first time in this paper.

We apply our model to two US stock market datasets. The first represents stock returns for four major technology stocks listed on the NASDAQ exchange (Apple, Google, Intel and Microsoft) and is used to identify the impact of information flows for a set of fundamentally inter-related, well-known stocks issued by businesses that either compete with each other on numerous fronts (eg. search engine platforms, operating systems, tablets and smart phones, and software development) or are technologically inter-dependent (particularly in terms of the inter-dependence between software and hardware for both the PC and tablet/smart-phone markets). We also look at the impact of news across different segments of the universe of investable US stocks by considering the information flows

et al., 1990). Its characterisation as normal or non-normal is objective in that it is based on the impact of the news rather than its factual or speculative basis.

underpinning movement in the S&P 500, S&P (MidCap) 400 and the S&P (SmallCap) 600 indexes which account for about 90 per cent of US stock market capitalization.

The paper provides a number of significant new findings regarding the relationship between news transmission and asset prices. We find that, in contrast to the standard approach adopted for measuring leverage effects, such effects predominantly stem from non-normal news and are driven by expectations regarding the probability of a jump. In this respect, we observe that if the expected probability of non-normal news is close to zero, the level of asymmetry in the response of expected volatility to (good or bad) normal news declines substantially.

The empirical findings to date have also given rise to the stylized fact that large shocks have a less persistent impact on expected volatility than normal shocks. Schwert (1990), Engle and Mustafa (1992) and Maheu and McCurdy (2004) all find, perhaps counter-intuitively, that the persistence effect on expected volatility from large shocks is smaller than that of normal shocks. Although the theoretical justification for this result is debatable, this empirical property is widely held to be correct. We provide evidence that the opposite, in fact, holds, with the effect of normal news on expected volatility dissipating at a faster pace than the effect of non-normal news. In particular, the effect of non-normal news on volatility expectations often builds up before dissipating. We show that this finding is intrinsically linked to information cross-overs, with the news transmission mechanism playing an important role in effectively crowding-out normal news and crowding-in non-normal news into the expected volatility equation.

Moreover, we find that the expected covariance surface between stock returns is relatively flat (viz. the correlation between any pair of assets is relatively stable if we remove non-normal news), but can change sharply with non-normal news. It is shown that the notion of positively or negatively correlated stock returns is ambiguous in the short term, with the sign of the expected correlation depending on the magnitude of non-normal news. Importantly, the sign of the expected correlation between stocks does not appear to vary with changes in normal news. However, the expected correlation for two stocks that are normally positively correlated can switch from being positive to negative in the presence of non-normal news. This has significant ramifications for determining whether stocks are substitutes or complements.

Our results indicate that accounting for information spillovers is fundamental in determining the expected volatility of an asset, and that the impact of information cross-overs on expected volatility differs significantly across assets in a manner that is not well explained by asset pricing theory. The findings therefore relate directly to the measurement and management of financial risk but also have implications for hedging and portfolio management purposes, particularly since they indicate that there is no need to hedge against the normal news of other stocks. Instead, the benefits of hedging are maximized by explicitly identifying stocks whose expected volatility declines in response

to non-normal news associated with the stock for which insurance is being sought.

Finally, we corroborate the empirical performance of the model using an out-of-sample forecasting exercise that compares the model's volatility forecasts with those from a set of alternative multivariate and univariate GARCH volatility models. We also limit the out-of-sample volatility forecasting exercise to periods with high volatility and to periods following a large negative return. In nearly all cases, the news transmission model outperforms the alternative models.

This paper is organized as follows. Section 2 presents a generalized model of news spillovers that can account for the transmission of normal and non-normal news in a vector of asset returns. Section 2 also describes the news impact surfaces associated with the news transmission model which are used to better understand the cross-sectional and inter-temporal transmission of normal and non-normal news. Section 3 describes the model estimation procedure, while Section 4 describes the technology stocks and S&P index data used in this study. The results are presented in Section 5, with Section 6 evaluating the out-of-sample volatility forecasting performance of the model. Section 7 concludes the study.

2. A generalized model of news spillovers

Consider the following general process for asset returns $y_t \in \mathbb{R}^N$

$$y_t = BE_{n_{t-1}|I_{t-1}}[f_{t-1}] + u_t + e_t \quad (1)$$

where u_t and e_t are innovations subject to $E(u_t|I_{t-1}) = E(e_t|I_{t-1}) = 0$ and $u_t \perp e_t$. I_{t-1} is a σ -field representing information up to time $t-1$ and contains $\{y_1, y_2, \dots, y_{t-1}\}$. B is an $(N \times M)$ matrix of coefficients.

The $M < N$ -dimensional vector f_{t-1} acts as an unobserved news transmission mechanism. This mechanism depends on the unobserved asset-specific jump dynamics represented by the N -dimensional vector n_t . The notation $E_{n_{t-1}|I_{t-1}}[f_{t-1}]$ denotes the dependence of f_{t-1} on the conditional jump dynamics $n_{t-1}|I_{t-1}$. At each time period agents form a set of conditional probabilities $n_t|I_t$ (these are denoted as $p(n_t|I_t)$ which is described in Section 3). These probabilities are used as weights for $f_{t|n_t}$ (being the value of the news transmission factor given n_t , see Eq. (12)) thereby yielding $E_{n_t|I_t}[f_t]$.

The use of this conditional expectation is important for a number of reasons. First, due to the presence of jump dynamics, f_t exhibits a path dependency issue whereby derivation of the news transmission process at time t would otherwise require knowledge of the entire history of n_t up to time t . The reason for this is that we allow the news transmission factor to depend on its past values. If $\{n_1, n_2, \dots, n_t\}$ were observed then the derivation of f_t would be straightforward. Since the jump dynamics are unobserved, however, we observe

a range of values of f_t conditional on the possible values of $\{n_1, n_2, \dots, n_t\}$. This renders the likelihood function intractable for any reasonable size t . The use of $E_{n_{t-1}|I_{t-1}}[f_{t-1}]$ follows Kim (1994) and avoids the intractability stemming from the adoption of f_{t-1} rather than its conditional expectation.

The vector u_t represents unusual or non-normal news, whereas e_t represents normal news. In line with Maheu and McCurdy (2004), normal news is assumed to produce smooth changes in the conditional variance, whereas non-normal news is associated with a jump in returns. Our approach is also consistent with Andersen (1996) who posits the existence of two or more types of information arrival processes for stock prices and volumes. We deliberately postpone any discussion on the distribution of u_t and e_t .

The transmission mechanism is capable of transmitting normal and non-normal news (e_t and u_t respectively) across assets i, j irrespective of the conditional covariance between the innovations to the i th and j th assets, being $cov(u_{it}, u_{jt}|I_{t-1}) + cov(e_{it}, e_{jt}|I_{t-1})$. Accordingly, news spillover between assets i and j can occur in three ways.

First, a non-zero covariance between e_{it} and e_{jt} whereby normal news for asset i contemporaneously affects asset j . Similarly, a non-zero covariance between u_{it} and u_{jt} implies that non-normal news for asset i contemporaneously affects asset j . Both these forms of spillover are direct in the sense that they require a non-zero covariance between the innovations to assets i and j .

The third spillover form accommodates the indirect spillover of news between assets i and j via the news transmission mechanism. This form of transmission is not contingent on any covariance between e_{it} and e_{jt} or u_{it} and u_{jt} and occurs intertemporally. Moreover, it is consistent with the notion of cross-predictability, whereby the current return of a stock predicts future returns of related stocks (Schlusche and Scherbina, 2015). In this respect, if the normal and non-normal shocks associated with asset j at time $t - 1$ (being $e_{j,t-1}$ and $u_{j,t-1}$) are embedded into the k th transmission mechanism f_{kt} , the shocks will be transmitted to any asset with a non-zero exposure to f_{kt} . Consequently, if $B_{ik} \neq 0$, the time $t - 1$ shocks $e_{j,t-1}$ and $u_{j,t-1}$ can be transmitted to the i th asset at time t irrespective of any contemporaneous covariance between the innovations. The mechanisms that facilitate these spillover forms are described in further detail below.

2.1 Accounting for non-normal news

We assume the following structure for non-normal news. In period t , the i th asset has n_{it} jumps where

$$n_{it} \sim \text{Poisson}(\theta_{it}). \quad (2)$$

A log-normal distribution is adopted for the jump intensities such that

$$\theta_t \sim LN(\lambda_t, \Sigma)$$

where $\theta_t = \begin{bmatrix} \theta_{1t} & \theta_{2t} & \cdots & \theta_{Nt} \end{bmatrix}'$, $\lambda_t = \begin{bmatrix} \lambda_{1t} & \lambda_{2t} & \cdots & \lambda_{Nt} \end{bmatrix}'$ and Σ is a positive-definite variance-covariance matrix.

The non-normal news process therefore accommodates co-jumps which are known to be important for the pricing of risk (Bandi and Reno, 2016; Li, Todorov and Tauchen, 2017). Furthermore, the co-jumps may be positively or negatively correlated. We do not impose the restriction $\lambda_{it} = \lambda_i$ thereby avoiding the assumption that the expected number of jumps for asset i is necessarily time-invariant (see, further, Christoffersen, Jacobs and Ornathanalai, 2012). Given the log-normal distribution of θ_{it} , we are also able to estimate the model without imposing any positivity restrictions on λ_{it} .

We consider two types of jump sizes in the model. In the first case, we assume that jump sizes are given by the vector N -vector μ_x and are therefore constant. In the second case, we allow normal news to influence the jump size; particularly in terms of negative normal news potentially resulting in a larger jump size in the following period. In this case, the jump size at time t (conditional on I_{t-1}) is defined as

$$x_t = \mu_x + Ge_{t-1}^* \quad (3)$$

where $x_t = \begin{bmatrix} x_{1,t} & x_{2,t} & \cdots & x_{N,t} \end{bmatrix}'$, μ_x is the mean jump size, G is a diagonal matrix, and $e_{t-1}^* = E_{n_{t-1}|I_{t-1}}[e_{t-1}]$ is the conditional expectation of normal news at time $t-1$ after accounting for the jump dynamics $n_{t-1}|I_{t-1}$. This expectation is derived in the same manner as for $E_{n_{t-1}|I_{t-1}}[f_{t-1}]$ and is described further in Section 3.

The non-normal news component u_t is defined as the conditionally zero-mean process

$$u_t = \tilde{x}_t - E(\tilde{x}_t|I_{t-1}) = (n_t - \lambda_t^*) \odot x_t \quad (4)$$

where $\tilde{x}_t = n_t \odot x_t$ and $\lambda_t^* = \begin{bmatrix} \lambda_{1t}^* & \lambda_{2t}^* & \cdots & \lambda_{Nt}^* \end{bmatrix}'$.

The following autoregressive conditional parameterisation is adopted for λ_{it}

$$\lambda_{it} = \omega_i + \psi_i \lambda_{i,t-1} + \gamma_{i1} \xi_{i,t-1} + \gamma_{i2} \xi_{i,t-1}^- \quad (5)$$

$$\xi_{i,t} = E(n_{it}|I_t) - \lambda_{it}^* \quad (6)$$

$$\lambda_{it}^* = \exp(\lambda_{it} + \Sigma_{ii}/2) \quad (7)$$

where $\xi_{i,t-1}^-$ is equal to $\xi_{i,t-1}$ if $\xi_{i,t-1} < 0$ and zero otherwise, and Σ_{ii} is the i, i th element of Σ .

The parameter λ_{it}^* is the conditional (on I_{t-1}) mean of n_{it} such that the conditional expectation of ξ_{it} is zero (see, further, Chan and Maheu, 2002).² The introduction of $\xi_{i,t-1}^-$ into the conditional expectation of the log jump intensity allows us to estimate

²Since λ_{it}^* is necessarily positive, we do not need to impose any positivity restriction on λ_{it} or on the jump intensity parameters.

the asymmetric impact of non-normal news on volatility. As such, it provides a natural counterpart to the asymmetric component often incorporated in GARCH models and allows us to identify asymmetric volatility effects induced by both normal and non-normal news.

The conditional variance of n_t is also related to λ_t^* and is denoted by Σ_t^λ with individual elements

$$\Sigma_{ii,t}^\lambda = \lambda_{it}^{*2} \times (\exp(\Sigma_{ii}) - 1) + \lambda_{it}^*, \quad i = 1, \dots, N \quad (8)$$

$$\Sigma_{ij,t}^\lambda = \lambda_{it}^* \lambda_{jt}^* \times (\exp(\Sigma_{ij}) - 1), \quad i \neq j. \quad (9)$$

2.2 Accounting for normal news

We assume that normal news e_t is conditionally multivariate normal, $e_t|I_{t-1} \sim N(0_N, H_t)$, where 0_N represents an N -dimensional vector of zeros. Since almost all positive definite VEC-based volatility models can be represented using the BEKK representation (Engle and Kroner, 1995), we adopt a BEKK model for H_t . Given the preponderance of evidence in support of a leverage effect, we also augment the BEKK form to account for leverage effects stemming from normal news such that

$$H_t = A_1 + A_2 e_{t-1}^* e_{t-1}' A_2' + A_3 H_{t-1} A_3' + A_4 e_{t-1}^{*-} e_{t-1}' A_4' \quad (10)$$

where $e_{t-1}^* = E_{n_{t-1}|I_{t-1}}[e_{t-1}]$ and A_1, A_2, A_3 and A_4 are $N \times N$ matrices. A_1 is estimated as $A_1^* A_1'^*$ (with A_1^* being a lower diagonal matrix of rank N) to ensure it is always positive definite.³ e_{t-1}^{*-} is a N -vector with i th element equal to e_{t-1}^* if $e_{t-1}^* < 0$ and zero otherwise.

Since the number of parameters in the BEKK representation increases sharply with N , we follow Engle and Kroner (1995) for estimation purposes and adopt the diagonal-BEKK. Accordingly, A_2, A_3 and A_4 are restricted to their diagonal forms.

Pursuant to (10), the impact of e_t^* on the conditional variance of the normal news component depends on A_2 and A_4 , with the latter term capturing potential asymmetry associated with the response to ‘bad’ or negative normal news. This normal news asymmetry complements the asymmetry allowed in the non-normal news component. The model therefore provides two avenues for leverage effects, which may stem from normal news or non-normal news (or both). This issue is explored further in Section 5D.

2.3 The news transmission component

The transmission process f_t incorporates the news observed at time t and transmits it across the asset space. The resulting cross-flow of information allows the normal or non-normal news of a given asset to influence the expected volatility of any of the assets in

³For simplicity, we assume a first order BEKK form.

the asset space. Investors can, for example, use news for a given firm to inter-temporally revise their expectations about the profitability of other firms (Patton and Verardo, 2012). In particular, the transmission process identifies and accepts the asset specific non-normal and normal news at time t , and then filters the news using Π_A and Π_N respectively before transmitting it to the remaining assets

$$f_t = \mu + \Phi E_{n_{t-1}|I_{t-1}}[f_{t-1}] + \Pi_A u_t + \Pi_N e_t \quad (11)$$

where μ is a vector of length M , and Π_A, Π_N are $M \times N$ matrices.

We also allow for the transmission of news expectations observed at time $t - 1$ and these are filtered using Φ . As such, just as expected returns depend on $E_{n_{t-1}|I_{t-1}}[f_{t-1}]$, the evolution of the news transmission process also depends on $E_{n_{t-1}|I_{t-1}}[f_{t-1}]$. Note that replacing f_{t-1} with $E_{n_{t-1}|I_{t-1}}[f_{t-1}]$ avoids having to observe the entire history of the jump variable (from $t = 0$ up to time $t - 1$) in order to determine the value of the news transmission factor at time t .

To better understand the transmission component we can expand f_t as follows

$$f_t = \mu + \Phi E_{n_{t-1}|I_{t-1}}[f_{t-1}] + \Pi_A ((n_t - \lambda_t^*) \odot x_t) + \Pi_N e_t. \quad (12)$$

The expansion renders it clear that the transmission component embeds information about historical expectations and current news, including the vector of non-normal news intensity shocks $(n_t - \lambda_t^*)$, the size of non-normal news x_t , and normal news e_t . This information is filtered using the Π_A and Π_N matrices thereby allowing us to establish hypothesis tests based on $\Pi_A = 0$ or $\Pi_N = 0$ in order to examine whether (and what) information is being filtered out by the news transmission process.

2.4 The model's news spillover dynamics and their comparison to other models

The conditional variance of y_{t+k} given information up to time t , $V(y_{t+k}|I_t)$, may be determined recursively as

$$V(y_{t+k}|I_t) = B P_{t+k-1} B' + H_{t+k} + \Lambda_{t+k} \quad (13)$$

$$P_{t+k} = \Phi P_{t+k-1} \Phi' + \Pi_N H_{t+k} \Pi_N' + \Pi_A \Lambda_{t+k} \Pi_A' \quad (14)$$

$$\Lambda_{ij,t+k} = (\mu_{i,x} \mu_{j,x} + G_i H_{t+k-1} G_j') \Sigma_{ij,t+k}^\lambda \quad (15)$$

where G_i is the i th row of G , and $\Lambda_{ij,t+k}, \Sigma_{ij,t+k}^\lambda$ are the (i,j) th elements of Λ_{t+k} and Σ_{t+k}^λ respectively. H_{t+k} and Λ_{t+k} are the conditional expectations of the covariances for normal news e_t and non-normal news u_t respectively.

In the case of $k = 1$, $P_t = 0$ since the common transmission f_t and the jump size x_t are conditionally observed. Consequently, the covariance matrix for y_{t+1} is determined as the sum of the covariances associated with normal and non-normal news, $V(y_{t+1}|I_t) = H_{t+1} + \Lambda_{t+1}$.

We can expand the recursive process (13) to obtain the general result

$$V(y_{t+k}|I_t) = \underbrace{H_{t+k}}_{\text{normal}} + \underbrace{\Lambda_{t+k}}_{\text{non-normal}} + \underbrace{BS_{t+k}^N B'}_{\text{normal spillover}} + \underbrace{BS_{t+k}^A B'}_{\text{non-normal spillover}} \quad (16)$$

where $S_{t+k}^N = \sum_{l=1}^{k-1} \Phi^{l-1} \Pi_N H_{t+k-l} \Pi_N' \Phi^{l-1'}$ and $S_{t+k}^A = \sum_{l=1}^{k-1} \Phi^{l-1} \Pi_A \Lambda_{t+k-l} \Pi_A' \Phi^{l-1'}$.

The model therefore generates news dynamics for expected volatility that are readily interpretable in terms of four distinct components: normal volatility H_{t+k} , non-normal volatility Λ_{t+k} , the intertemporal transmission of normal news across assets S_{t+k}^N , and the intertemporal transmission of non-normal news S_{t+k}^A . In this respect, we note that the identification of Φ , Π_A and Π_N requires that $\{S_{t+k}^A, S_{t+k}^N\} \neq 0$ for at least one $k > 0$.

An interesting property of the model's expected volatility is that it nests a range of popular models that describe the interaction between the news arrival process and the volatility of a set of assets. In the absence of any indirect volatility spillover or jump effects the conditional variance collapses to the basic mGARCH model $V(y_{t+k}|I_t) = H_{t+k}$ such that volatility is characterized solely by contemporaneous normal news. If, however, there are jump effects but no indirect spillovers, the additional presence of non-normal news results in $V(y_{t+k}|I_t) = H_{t+k} + \Lambda_{t+k}$ such that the conditional volatility of y_{t+k} is determined by contemporaneous normal and non-normal news.

The presence of information spillovers produces additional dynamics for the expected variance-covariance of y_{t+k} . In the case where $\Pi_A = S_{t+k}^A = 0$, $V(y_{t+k}|I_t) = H_{t+k} + BS_{t+k}^N B'$ such that the news transmission process is limited to normal news. Nevertheless, the news transmission process is significantly more flexible than the ordinary mGARCH case ($V(y_{t+k}|I_t) = H_{t+k}$). In this respect, note that notwithstanding the adoption of a diagonal BEKK form for H_{t+k} , the presence of S_{t+k}^N implies that normal news for asset i can affect the conditional variance of asset j (instead of only the conditional covariance between asset i and j). This holds even where $H_{ij,t}$ is zero. In such a scenario, normal news regarding asset i at time t can affect asset j at time $t+k$ via asset l ($l \neq i, j$). Clearly, asset dynamics of this kind require three or more assets.

Finally, the result $S_{t+k}^N \neq 0, S_{t+k}^A \neq 0$ implies the full volatility structure (16) and accommodates a range of characteristics that can be present in the volatility spillover process. Volatility can be transmitted between all assets irrespective of the contemporaneous correlation observed between the assets' normal or non-normal news; this is facilitated by the news transmission mechanism which can transmit normal news (if $\Pi_N \neq 0$),

non-normal news (if $\Pi_A \neq 0$) or both.⁴ Accordingly, the expected volatility of $y_{i,t+k}$ is potentially a function of the normal and non-normal news of *any* of the N assets.

2.5 Deriving a heterogeneous news impact surface

The estimation of normal and non-normal news components allows us to generalise the news impact curve, which typically defines the relationship between news in the current period and expected volatility in the next period (Engle and Ng, 1993). In this respect, we analytically derive the impact of contemporaneous normal and non-normal shocks on expected volatility, thereby producing what we term a ‘generalised’ news impact surface.

The news impact surface for the i th asset is denoted Δ_i and represents the step-ahead impact of uncertainty regarding ordinary news e and the unexpected number of jumps ξ on the expected variance.

$$\Delta_i = Q_{ii} + C_{1,i} \times (\mu_{ix} + G_{(i,i)}e_{i,0})^2 \times \exp(\gamma_{i1}\xi_{i,0} + \gamma_{i2}\xi_{i,0}^-), \quad i = 1, \dots, N \quad (17)$$

$$Q_{ij} = C_{0,ij} \times (\mu_{ix} + G_{(i,i)}e_{i,0}) (\mu_{jx} + G_{(j,j)}e_{j,0}) \times (C_{1,i}C_{1,j}\tilde{\xi}_{ij}) + C_{2,ij} + \tilde{H}_{ij} \quad (18)$$

$$\tilde{\xi}_{ij} = \exp(\gamma_{i1}\xi_{i,0} + \gamma_{j1}\xi_{j,0} + \gamma_{i2}\xi_{i,0}^- + \gamma_{j2}\xi_{j,0}^-) \quad (19)$$

$$\tilde{H}_{ij} = A_{(i,i),2}A_{(j,j),2}e_{i,0}e_{j,0} + A_{(i,i),4}A_{(j,j),4}(e_{i,0}^-e_{j,0}^-) \quad (20)$$

where $C_{0,ij} = \exp(\Sigma_{ij}) - 1$, $C_{1,i} = \exp(\omega_i + \psi_i\bar{\lambda}_i + \Sigma_{ii}/2)$, $C_{1,j} = \exp(\omega_j + \psi_j\bar{\lambda}_j + \Sigma_{jj}/2)$ and $C_{2,ij} = A_{(i,j),1} + A_{(i,i),3}A_{(j,j),3}\bar{H}_{ij}$.

The terms $\bar{\lambda}_i$, $\bar{\lambda}_j$ and $\bar{H}_{ij}$ represent pivot points for the time-varying processes λ_{it} , λ_{jt} and $H_{ij,t}$ (being the (i, j) th element of H_t). In our setting, we adopt the unconditional expectation of the relevant time-varying parameters as the pivot points. The shocks $e_{i,0}$, $\xi_{i,0}$ are the i th elements of the N -dimensional shocks e_0 , ξ_0 (respectively), while $G_{(i,i)}$, $A_{(i,i),1}$ and $A_{(i,i),3}$ are the i th diagonal elements of G , A_1 and A_3 respectively. Note that in the case $i = j$, the proportionalities $C_{0,ii}$, $C_{1,i}$ and $C_{2,ii}$ are non-negative such that V_i is non-negative for any $e_{i,0}$, $\xi_{i,0}$. The *covariance* surface for assets i, j is given by Q_{ij} and represents the impact of the shocks $e_{i,0}$, $\xi_{i,0}$ on the step-ahead expected covariance. Supplementary information for deriving the news impact surface is provided in the online Appendix A.

The dynamics of the model change markedly when we move from period 1 to period 2. Although Δ_i , Q_{ij} are only influenced by the period 0 shocks to assets i and j , the entire set of shocks is relevant in period 2. It is therefore instructive to consider the news impact surface in period 2 stemming from a period 0 shock. We denote the period 2 impact as Δ_i^+ , Q_{ij}^+ . The sequence $\{\Delta_i, \Delta_i^+\}$ therefore represents the period 1 and 2 impact of the

⁴The Φ matrix determines the extent to which news disseminated by the news transmission mechanism impacts on expected volatility. In particular, the impact of news on the news transmission component of expected volatility ($BS_{t+k}^N B' + BS_{t+k}^A B'$) decays at the rate $\Phi^{2(l-1)}$ where l represents the number of periods prior to $t+k$.

period 0 arrival of N dimensions of normal news (e_0) and N dimensions of non-normal news (ξ_0) on the expected variance of asset i . Similarly, $\{Q_{ij}, Q_{ij}^+\}$ represents the period 1 and 2 impact of e_0, ξ_0 on the expected covariance between assets i and j .

$$\Delta_i^+ = Q_{ii}^+ + C_{1,i}^+ \times \left(\mu_{ix}^2 + G_{(i,i)}^2 \left[C_{2,ii} + A_{(i,i),2}^2 e_{i,0}^2 + A_{(i,i),4}^2 (e_{i,0}^-)^2 \right] \right) \quad (21)$$

$$\times \exp(\psi_i [\gamma_{i1} \xi_{i,0} + \gamma_{i2} \xi_{i,0}^-]) \quad (22)$$

$$Q_{ij}^+ = C_{0,ij} \times \mu_{ix} \mu_{jx} \times \left(C_{1,i}^+ C_{1,j}^+ \tilde{\xi}_{ij}^+ \right) + C_{2,ij}^+ + A_{(i,i),3} A_{(j,j),3} \times \tilde{H}_{ij} + B_i S^N B_j' + B_i S^A B_j' \quad (23)$$

$$S^N = \Pi_N (A_1 + A_2 e_0 e_0' A_2' + A_3 \bar{H} A_3' + A_4 e_0^- e_0'^- A_4') \Pi_N' \quad (24)$$

$$S^A = \Pi_A ((\mu_x + G e_0) (\mu_x + G e_0)' \odot \Sigma^\lambda) \Pi_A' \quad (25)$$

$$\tilde{\xi}_{ij}^+ = \exp(\psi_i [\gamma_{i1} \xi_{i,0} + \gamma_{i2} \xi_{i,0}^-] + \psi_j [\gamma_{j1} \xi_{j,0} + \gamma_{j2} \xi_{j,0}^-]) \quad (26)$$

where $C_{1,i}^+ = \exp(\omega_i + \psi_i \lambda_{i0} + \psi_i^2 \bar{\lambda}_i + \Sigma_{ii}/2)$, $C_{1,j}^+ = \exp(\omega_j + \psi_j \lambda_{j0} + \psi_j^2 \bar{\lambda}_j + \Sigma_{jj}/2)$, and $C_{2,ij}^+ = A_{(i,j),1} + A_{(i,i),3} A_{(j,j),3} \times A_{(i,j),1} + (A_{(i,i),3} A_{(j,j),3})^2 \bar{H}_{i,j}$.

The N -dimensional matrix Σ^λ used to derive S^A is the contemporaneous covariance of the jumps and its individual elements are given by

$$\Sigma_{ii}^\lambda = C_{0,ii} \times \exp(\lambda_i^* + \Sigma_{ii}/2)^2 + \exp(\lambda_i^* + \Sigma_{ii}/2), \quad i = 1, \dots, N \quad (27)$$

$$\Sigma_{ij}^\lambda = C_{0,ij} \times \exp(\lambda_i^* + \Sigma_{ii}/2) \times \exp(\lambda_j^* + \Sigma_{jj}/2), \quad i \neq j \quad (28)$$

where $\lambda_i^* = \lambda_{i0} + \psi_i \bar{\lambda}_i + \gamma_{i1} \xi_{i,0} + \gamma_{i2} \xi_{i,0}^-$.

The change in dynamics from period 1 to period 2 following a shock highlights the inter-temporal propagation of shocks implied by the model (readers interested in the steady state news impact surface Δ_i^* are referred to the online Appendix A). In period 0, consider that asset i is subject to either normal news $e_{i,0}$ or news leading to an unexpected jump $\xi_{i,0}$. In the initial period 1, this impacts on both the volatility of asset i and on its covariance with the other assets. In period 2, however, the news propagates across all assets such that the news pertaining to asset i impacts not only on its covariance with other assets but directly on the variance of the other assets as well. This can be caused by, for example, asset managers changing their portfolio mix following the arrival of news e_0, ξ_0 .

The nature of the propagation that takes place, however, can depend on whether the news is normal or non-normal. As such, non-normal news pertaining to asset i can have a different impact on the expected variance of other assets than normal news. This unique distinction ends up being critical in our empirical assessment of the news transmission that takes place between assets and on its ramifications for both expected volatility and the expected covariance between assets.

3. Estimation

The model requires estimation of $\Omega = \{\mu, B, \Phi, \Pi_A, \Pi_N, G, A_1, A_2, A_3, A_4, \omega, \psi, \gamma, \Sigma\}$. This implies the estimation of $N(N+1)/2 + 3N$ parameters for the diagonal-BEKK representation of the mGARCH, $N(N+1)/2 + 4N$ parameters for the lognormal-Poisson representation of non-normal news, $2N$ parameters for the size of the non-normal news, $M + M^2 + 2MN$ parameters for the news transmission process and $NM - M^2$ parameters for the loadings on the news factor.

3.1 Likelihood function

The model's conditional likelihood function is

$$L(y|\Omega, n_1, \dots, n_T) = \prod_{t=1}^T L(y_t|\Omega, n_t, \lambda_t, E_{n_{t-1}|I_{t-1}}[f_{t-1}], x_t) \quad (29)$$

$$= \prod_{t=1}^T f_N(y_t - BE_{n_{t-1}|I_{t-1}}[f_{t-1}] - (n_t - \lambda_t^*) \odot x_t | 0_N, H_t) \quad (30)$$

$$= \prod_{t=1}^T f_N(e_{t|n_t} | 0_N, H_t) \quad (31)$$

where $y = (y_1, y_2, \dots, y_T)'$, f_N is the multivariate normal density and $e_{t|n_t}$ denotes the deviation of y_t from its conditional expectation $E(y_t|\Omega, n_t, \lambda_t, E_{n_{t-1}|I_{t-1}}[f_{t-1}], x_t)$.

Similar to Hamilton (1989; 1990), the likelihood function is complicated by the presence of an unobserved state (in our case, e_t is conditional on the unobserved n_t rather than an unobserved Markovian regime s_t). To estimate the model, we therefore require the likelihood after integrating out the jump component n_t

$$L(y_t|\Omega, I_{t-1}) = \sum_{n_{1t}=0}^{\infty} \dots \sum_{n_{Nt}=0}^{\infty} p(n_t|\lambda_t, \Sigma) L(y_t|\Omega, n_t, \lambda_t, E_{n_{t-1}|I_{t-1}}[f_{t-1}], x_t) \quad (32)$$

with the full likelihood being $L(y|\Omega) = \prod_{t=1}^T L(y_t|\Omega, I_{t-1})$. The conditional probability of n_t is $p(n_t|\lambda_t, \Sigma)$ and involves integrating over the N -dimensional positive subspace. The estimation of this conditional probability is discussed in Section 3.2 below which considers the scalability of the model.

Given $L(y_t|\Omega, I_{t-1})$, we are able to update the conditional probability $p(n_t|\lambda_t, \Sigma)$ with information up to time t thereby yielding the filtered probability of n_t

$$p(n_t|I_t) = \frac{p(n_t|\lambda_t, \Sigma) L(y_t|\Omega, n_t, \lambda_t, E_{n_{t-1}|I_{t-1}}[f_{t-1}], x_t)}{L(y_t|\Omega, I_{t-1})}. \quad (33)$$

This probability is used to obtain $\xi_{i,t} = E(n_{i,t}|I_t) - \lambda_{i,t}^*$ thereby providing the next pe-

riod's λ_{t+1} (based on Eq. (5)). It is also used to obtain $f_{t|n_t}$ which is necessary for the evaluation of the likelihood function. In this respect, conditional on $\{\Omega, n_t, \lambda_t, E_{n_{t-1}|I_{t-1}}[f_{t-1}], x_t\}$, the news transmission factor is given by

$$f_{t|n_t} = \Phi E_{n_{t-1}|I_{t-1}}[f_{t-1}] + \Pi_A ((n_t - \lambda_t^*) \odot x_t) + \Pi_N e_{t|n_t}. \quad (34)$$

To facilitate the iteration we require the conditional expectations $E_{n_t|I_t}[f_t]$ and $E_{n_t|I_t}[e_t]$. The former is used as an input into the conditional expectation of y_{t+1} and in next period's new transmission factor $f_{t+1|n_{t+1}}$. The latter expectation $E_{n_t|I_t}[e_t]$ serves as an input into the next period jump size x_{t+1} and normal news volatility H_{t+1} . We estimate these values in a similar manner to Eq. (32) which is used to obtain $L(y_t|\Omega, I_{t-1})$. This involves weighting $f_{t|n_t}$ and $e_{t|n_t}$ using $p(n_t|I_t)$ as follows

$$E_{n_t|I_t}[f_t] = \sum_{n_{1t}=0}^{\infty} \dots \sum_{n_{Nt}=0}^{\infty} f_{t|n_t} p(n_t|I_t) \quad (35)$$

$$E_{n_t|I_t}[e_t] = \sum_{n_{1t}=0}^{\infty} \dots \sum_{n_{Nt}=0}^{\infty} e_{t|n_t} p(n_t|I_t). \quad (36)$$

In practice, the summation in Eqs. (35)-(36) (and in Eq. (32)) is undertaken up to point c such that $p(c|\lambda_t, \Sigma) \rightarrow 0$. For the daily datasets used in this paper, the number of jumps is rarely greater than 2 such that setting c beyond 2 or 3 extends the estimation time with little difference in the parameter estimates.

The individual steps involved in forming the likelihood function $L(y|\Omega)$ are detailed in the online Appendix B. Maximum likelihood estimates of the parameters Ω are obtained by maximizing $L(y|\Omega)$ with respect to the parameter set.⁵ To implement the relative pricing of assets we place M unity restrictions on B , in addition to the restriction that all elements of B are non-negative. This implies that, if μ is positive, the expected return will also converge to a positive value. Finally, stationarity restrictions are placed on λ_t and H_t (see Engle and Kroner (1995) for further details regarding the properties of H_t).

3.2 Scalability

Given that mGARCH models are typically parameter intensive, it is useful to compare the parameters required to estimate the proposed model relative to those of competing models. We focus on the full BEKK mGARCH model which requires the estimation of $(N(N+1))/2 + 2N^2$ parameters but also allows for rich volatility dynamics.⁶ In terms of scalability, when $N > 8$, the full BEKK model requires the estimation of a greater

⁵For further information regarding the maximum likelihood estimation of models in the single source of error family see Ord, Koehler and Snyder (1997).

⁶To maintain consistency with our specification for normal news volatility in Section 2.2, we specify the number of parameters for a first-order BEKK model.

number of parameters than the model proposed here (assuming $M = 1$). In particular, with $N = 10$, the full BEKK model requires the estimation of an additional 24 parameters, rising to an additional 267 parameters for $N = 15$. The additional number of parameters required increases further if the BEKK model is augmented to account for asymmetry in the volatility response.

In terms of estimating the model using larger values of N (for example, $N = 20$), the primary issue for the model we propose is the dimensionality of the jump component n_t . The reason for this is that model estimation requires summation over the conditional probabilities associated with the possible combinations of the N -dimensional vector n_t . This becomes prohibitive from a time perspective such that, for larger values of N , it is sensible to reduce the dimensionality of n_t . Consider, for example, a model with $N = 20$ assets across five asset classes or sectors with a jump process specified for each of the five asset classes or sectors rather than for each individual asset. An effective way to proceed is to specify such common jumps yet continue to allow individual jump sizes for each asset. In this situation, a rich dynamic process allowing for normal and non-normal news can be estimated for relatively large N , with both normal and non-normal news continuing to be transmitted across individual assets.

A related difficulty is that obtaining the conditional probability of n_t (denoted $p(n_t|\lambda_t, \Sigma)$ in Eq. (32)) involves integrating over the N -dimensional positive subspace

$$p(n_t|\lambda_t, \Sigma) = \int_{\mathbb{R}_+^N} \prod_{i=1}^N p_0(n_{it}|\theta_{it}) f_{LN}(\theta_t|\lambda_t, \Sigma) d\theta_t \quad (37)$$

where $\theta_t = (\theta_{1t}, \theta_{2t}, \dots, \theta_{Nt})'$ is an N -dimensional vector, $p_0(n_{it}|\theta_{it})$ is the Poisson density evaluated at n_{it} conditional on the time-varying intensity θ_{it} , and $f_{LN}(\theta_t|\lambda_t, \Sigma)$ is the N -variate log-normal density with location and scale parameters λ_t, Σ respectively.

For $N \leq 5$, we find that the grid-based numerical estimation method proposed in Aitchison and Ho (1989) is fast and accurate. For larger values of N , however, we propose a simulation-based approach to estimating $p(n_t|\lambda_t, \Sigma)$. The proposed approach yields estimates that are effectively identical to those based on the grid-method but with a substantial improvement in computation time. In particular, we propose drawing K^* values of θ from the log-normal density $f_{LN}(\theta_t|\lambda_t, \Sigma)$ and then evaluating the Poisson density $\prod_{i=1}^N p_0(n_{it}|\theta_{it})$ for each draw of θ .⁷ The evaluation of the latter is trivial as $\prod_{i=1}^N p_0(n_{it}|\theta_{it})$ is simply the product of univariate Poisson densities. The resulting estimate of $p(n_t|\lambda_t, \Sigma)$ is given by

$$\frac{1}{K^*} \sum_{k=1}^{K^*} \prod_{i=1}^N p_0(n_{it}|\theta_i^{(k)}) \quad (38)$$

⁷We observe little change in the estimated conditional probabilities when setting K^* greater than 1000.

where $\theta^{(k)} = (\theta_1^{(k)}, \theta_2^{(k)}, \dots, \theta_N^{(k)})'$ is the k th draw from the density $f_{LN}(\theta_t | \lambda_t, \Sigma)$.

4. Data

We estimate the model parameters using daily returns data for four major technology stocks: Google (GOOGL), Apple (AAPL), Intel (INTC) and Microsoft (MSFT). The four stocks are chosen as they potentially exhibit technological inter-dependencies (for example, partial reliance on Intel processors or the Microsoft Windows OS) and offer a competing range of products (such as tablets and smart-phones). The stocks therefore provide an interesting basis for examining the impact of information cross-overs on volatility.

The returns are obtained from Yahoo Finance and are computed as 100 times the first difference of log closing prices over the period 24 March 2005 to 11 December 2013 for a total of $T = 2,196$ observations for each stock.⁸ Descriptive statistics for the returns data are presented in the online Appendix C and indicate that Apple and Google have higher daily returns and volatility than Intel or Microsoft. In turn, the stocks differ in terms of their skewness and kurtosis, with Google and Microsoft exhibiting greater skewness and kurtosis than Apple or Intel.

The model is also estimated on the S&P 400, S&P 500 and S&P 600 index returns data. The data are obtained from the Zacks Equity Prices backtesting database and span the period September 10, 2007 to August 28, 2013 for a total of $T = 1500$ observations per index. In line with the summary statistics, the indexes exhibit substantially greater correlation than the technology stocks but smaller levels of volatility and kurtosis.

5. Results

A. Estimates for Google, Apple, Intel and Microsoft

We estimate several variants of the model on both the technology and S&P datasets based on the adoption of a single news transmission factor ($M = 1$), and find that restricting λ_0 in (5) and G in (3) to zero has little difference on the findings for the technology stocks.⁹ In contrast, the restrictions have implications for the S&P index data. On this basis, we present results for the more parsimonious restricted model ($\lambda_0 = 0$ and $G = 0$) for the technology stocks and for the unrestricted model on the index returns. The parameter estimates for the technology stocks are provided in the online Appendix D. To

⁸We estimate the model using adjusted and unadjusted closing prices but find little difference in the parameter estimates. To ensure comparability with previous studies (eg. Maheu and McCurdy, 2004), however, the forecasting exercise is based on closing values that are adjusted for dividends and splits.

⁹To identify the news transmission factor, we set the first element of B to unity.

preserve space, particularly since the results for the index returns are consistent those of the technology stocks, readers interested in the results and associated discussion for the index stocks are referred to the online Appendix E. We note that diagnostic tests on the standardized residuals, $e_{i,t}/H_{(i,i),t}^{1/2}$ and $u_{i,t}/(x_{i,t}^2 \Sigma_{ii,t}^\lambda)^{1/2}$, indicate that there remains no statistically significant autocorrelation in either the level or squared standardized residuals.¹⁰

Prior to discussing the model's key results we examine briefly the time-varying jump intensities $\lambda_{i,t}$ which (together with Σ) determine the volatility of the non-normal news component, and the diagonal elements of H_t which determine the volatility of each stock's normal news component. The jump shocks $\xi_{i,t}$ are also examined as they are critical to determining time-variation in the jump intensities (see, further, Patton and Sheppard, 2015). The figures for these variables are presented in the online Appendix F and show that the jump intensities vary substantially over time for Google and Intel, indicating that the volatility of non-normal news stemming from these stocks has historically exhibited sharp changes (Figure F1 in the online Appendix F). In contrast, Apple and Microsoft exhibit relatively smaller time variation in their jump intensities indicating that, at least historically, non-normal news for both of these stocks has been less volatile than that of Google or Intel.

It is clear from Figure F2 that it is the presence of relatively large jump shocks $\xi_{i,t}$ that is responsible for the sharp changes in the jump intensities of Google, Intel and Microsoft. The shocks are typically positive indicating that when a jump shock occurs it is usually the result of a greater number of jumps than expected. Since the expected jump *size* for Google is positive ($\mu_{x,1} = 4.7072$), the positive shocks in Figure F2 indicate that non-normal news from Google has tended to be 'good news'. Conversely, the expected jump size for Intel and Microsoft are similar but negative ($\mu_{x,3} = -4.8561$ and $\mu_{x,4} = -5.5497$), indicating that jump shocks for these two stocks tend to be the result of 'bad news'. Apple, on the other hand, has a small (albeit statistical significant) negative expected jump size ($\mu_{x,2} = -0.32035$) indicating that non-normal news from Apple has been neither particularly good or bad (at least, relative to Google, Intel or Microsoft), although it is slightly skewed to the negative.

Figure F3 shows the volatility of each stock's normal news component. The volatility of normal news rises during the GFC, indicating general investor uncertainty about the post-GFC 'status-quo'. The volatility is greatest for Apple indicating significant uncertainty about the post-GFC Apple stock price, but endures longest for Intel (with normal

¹⁰The Engle-LM test of additional ARCH effects in the standardised normal residuals returned p-values that were not distinguishable from those for an *iid* random draw (being 0.7236, 0.5438, 7687, 0.8414 for Google, Apple, Intel and Microsoft respectively when applying the test at 5 lags). The relevant p-values for lags smaller or greater than 5 (up to 10) were similarly insignificant, as were the p-values based on the Ljung-Box Q statistic on the squared standardised residuals (being 0.7604, 0.5669, 0.7798, 0.8421).

volatility not stabilizing until mid-2009).

B. What news is being transmitted?

Table 1 presents likelihood ratio tests of the parameters relevant to determining whether $cov(u_{it}, u_{jt}) = 0$ and $cov(e_{it}, e_{jt}) = 0$, and of the parameters in the transmission component. Although the focus is on the latter parameters, it is first noted that likelihood ratio tests of the null hypotheses of a diagonal Σ and A_1 indicate rejection of the null at the .05 level. Rejection of the diagonality of Σ implies a rejection of $cov(u_{it}, u_{jt}) = 0$ and indicates that, to some extent, the transmission of non-normal news between the technology stocks is contemporaneous. Similarly, the rejection of a diagonal A_1 indicates a rejection of $cov(e_{it}, e_{jt}) = 0$ (irrespective of the parameter values for A_2, A_3 or A_4) with an analogous interpretation to that for non-normal news. It is noted that although the rejection of $cov(e_{it}, e_{jt}) = 0$ is well known, the additional rejection of $cov(u_{it}, u_{jt}) = 0$ indicates that the contemporaneous transmission of non-normal news is distinct to that of normal news.

Table 1: Likelihood ratio tests of the transmission component

	LR statistic	df	95% Crit.
$\mu = 0$	56.268	1	3.8415
$\Phi = 0$	82.858	1	3.8415
$\Pi_A = 0$	90.509	4	9.4877
$\Pi_N = 0$	49.18	4	9.4877
$\Pi_A = \Pi_N$	94.555	4	9.4877
$B = 0$	66.625	3	7.8147
$B = 1$	93.467	3	7.8147
$\Sigma_{ij} = 0, i \neq j$	15.812	6	12.592
$A_{1,ij} = 0, i \neq j$	514.72	6	12.592

B.1. The significance of the news transmission process

To determine the significance of the news transmission process f_t we examine the transmission of non-normal news by reference to Π_A and the transmission of normal news based on Π_N . We also examine the persistence of the transmission component Φ which provides information on the rate of news decay.

At the outset, the rejection of the hypothesis $B = 0$ in Table 1 indicates that news pertaining to any of the technology stocks is *potentially* embedded in the expected volatility function of each of the remaining stocks.¹¹ To identify the news being transmitted we

¹¹The test $B = 0$ (or $B = 1$) pertains to the elements of B other the first element which is restricted to unity for identification purposes.

require the additional testing of the hypotheses $\Pi_A = 0$, $\Pi_N = 0$. Both hypotheses are clearly rejected, implying the transmission of both normal and non-normal news across the asset space. Moreover, the rejection of $\Pi_A = \Pi_N$ indicates the presence of asymmetries in the news transmission process, with the transmission of non-normal news differing significantly from that of normal news. This asymmetry may result in the dampening of non-normal news (if Π_A is smaller than Π_N) or the amplification of non-normal news, with research pointing heavily to the former outcome (whereby the impact of non-normal news on expected volatility is found to decay at a faster rate than that of normal news) (Schwert, 1990; Engle and Mustafa, 1992; Maheu and McCurdy, 2004).

We find, however, that Π_A is substantially larger than Π_N thereby highlighting the relative importance of non-normal news to the transmission process. In particular, univariate tests of each of the parameters in Π_A and Π_N indicate that greater weight is attached to the non-normal news of Google, Apple and Microsoft, with the weight being attached to the non-normal news from Intel not being significantly different to the weight attached to Intel's normal news. As such, we cannot observe $\Pi_{N,i} > \Pi_{A,i}$ for any i .

The results provide strong evidence that the transmission of normal news is dampened at a substantially heavier rate than that of non-normal news. To better understand why we observe converse results to previous studies it is necessary to consider the implications of setting either $\Pi_A = \Pi_N = 0$ or $\Pi_A = \Pi_N$. These restrictions generate dynamics that are similar to those in previous studies. Under the former restriction ($\Pi_A = \Pi_N = 0$), expected volatility is determined as the sum of the expected volatility stemming from u_t and e_t . Given the latter restriction ($\Pi_A = \Pi_N$), both normal and non-normal news are transmitted in the same manner. In both cases, the dynamics underlying expected volatility are determined solely by the mGARCH process H_t and the jump intensity λ_t . Since the persistence associated with H_t is typically greater than that for λ_t , previous studies generally conclude that the impact of non-normal news on expected volatility decays at a faster rate than the impact of normal news. In this study, however, we show that both $\Pi_A = \Pi_N = 0$ and $\Pi_A = \Pi_N$ are clearly rejected, with Π_A being significantly greater than Π_N , such that the aforementioned outcome no longer holds.

B.2. Estimating the conditional variance using the full BEKK mGARCH model

In the absence of a non-normal component, it is in principle clear that even a parameter rich model such as the full BEKK will be unable to capture the heterogeneous normal and non-normal dynamics observed in the technology dataset. In practice, however, it is unclear how the full BEKK (or a comparable model) will approximate the volatility dynamics associated with heterogeneous news. To examine this, we simulate 100 datasets of length $T = 2,196$ from a DGP process based on the model proposed in Section 2 using the estimated parameters for Google, Apple, Intel and Microsoft. We then obtain parameter estimates for the datasets using both the full BEKK model and the model

proposed in this paper. To ensure, however, that the results are not peculiar to the full BEKK, we also estimate the model on the asymmetric full BEKK and the (symmetric and asymmetric) VEC-based mGARCH models with consistently similar results.¹²

The results show unambiguously that the full BEKK model produces parameter estimates that amplify the one-step ahead conditional volatilities above their actual values. This holds for all four stocks. In particular, the autoregressive coefficients in the full BEKK model produce a substantially larger impact of normal news on expected volatility than that implied by the underlying DGP. This essentially yields a model that attaches a greater weight to normal news than that warranted by the underlying DGP.

On average, the full BEKK produces estimates of conditional variances that are approximately 27 per cent greater than the conditional variance of the normal news component in the model.¹³ This holds irrespective of whether the conditional variance of the normal news component is constructed using parameter estimates or by using the actual parameters in the DGP. Moreover, there are numerous periods where the surge in expected volatility predicted by the full BEKK is two or three times larger than the conditional volatility of the normal news component. As noted above, similar results hold for the asymmetric full BEKK and the VEC-based mGARCH models.

The problem of inflated conditional volatilities continues to exist even if we compare the step-ahead conditional volatilities of the full BEKK with the total volatility implied by the DGP process (viz. the sum of normal and non-normal volatilities). In such a scenario, the full BEKK continues to produce conditional volatility estimates that are - on average - 17 per cent greater than the actual one-step ahead total volatilities. The forecasting implications of these results for the full BEKK model are explored further in Section 6.

C. Understanding news spillovers using news impact surfaces

The statistical significance of the news transmission process provides some insight into the complexity of the news transmission process. To better appreciate this complexity, however, we use the news impact surfaces derived in Section 2.5 to examine the relationship between news and volatility expectations. In this respect, the news impact surfaces allow us to examine the impact of normal and non-normal news on expected volatility taking into account the entire set of model parameters.

To examine the relationship between news and volatility expectations, we first construct the surfaces contingent on an environment where there is no non-normal news. This provides us with a basis for examining how important the normal news of asset j is

¹²The parameter estimates for the alternative models were obtained using Kevin Sheppard's MFE toolbox.

¹³The average ratios of the conditional volatility estimates stemming from the full BEKK relative to the estimates implied by the underlying DGP are: 1.34 for Google; 1.46 for Apple; 1.15 for Intel; and 1.13 for Microsoft.

in determining the expected variance of asset i . We then introduce non-normal news into the financial environment in order to determine how expected volatility changes in the joint presence of normal and non-normal news. Since, in our news model, the expected volatility of an asset is a function of both its own news and news flows from other assets, we are able to examine an asset's volatility response to its own news and the news of other assets. We focus on the $k = 1$ and $k = 2$ day ahead impact of news since these two periods are sufficient in order to reflect differences in the news transmission process that are a function of k .

The results indicate that information spillovers are fundamental in determining the expected volatility of an asset. Although the impact of information cross-overs on expected volatility differs across assets, there are several clear commonalities in our findings that shed new light on the relationship between expected volatility and news. The commonalities relate to the relative importance of normal news and to the persistence of non-normal news. They also relate to the capacity to interpret asset price correlations in the context of normal and non-normal information flows between assets. Each of these is discussed, in turn, below.

C.1. How important is normal news?

In the absence of any non-normal news, the expected volatility of a stock's return is dominated by its own normal news. The impact of normal news stemming from other stocks is relatively minor, even when we assume that the normal shock is large.¹⁴ Effectively, the normal news of other assets is largely filtered out in determining the expected volatility of a given asset.

The impact surface changes dramatically with the introduction of non-normal news, with the expected volatility of each of the technology stocks being sensitive to the non-normal news of at least one of the other technology stocks. Given the presence of normal news for asset i and non-normal news for asset j , we observe two types of surfaces. The first type is where the expected volatility of asset i is significantly influenced by its own normal news and the non-normal news of another asset *if* the non-normal news is of a sufficiently large magnitude. Consider, for example, Figure 1 which indicates that Google's expected volatility is clearly influenced by large non-normal shocks from Intel and Microsoft. Depending on the magnitude of the non-normal news, news from Intel may increase or decrease Google's expected volatility. Intel can therefore act as *both* a substitute or a complement in terms of minimizing the volatility of an investor holding Google stock.

The second type of response is a collapse in the asset's sensitivity to its own normal

¹⁴Figure G1 in the online Appendix G shows the impact of normal news on the expected volatility of Google and Microsoft, and it is clear that the normal news of other stocks has a negligible impact on the *expected* volatility of either asset. This result also holds for Apple and Intel.

news. In such a scenario, the expected volatility of asset i is largely determined by the non-normal news of asset j . In Figure 1, for example, we observe that Intel’s expected volatility is fairly flat but increases sharply in response to non-normal news from Apple or Google. The contrasting response of Google’s expected volatility to non-normal Intel news (with Intel’s expected volatility to non-normal Google news) highlights the lack of symmetry in news responses. The identity of the news source and recipient appear to be critical in determining the volatility response to news.

Figure 1: Expected variance surface conditional on normal (or ‘typical’) and non-normal (or ‘atypical’) news. The expected variance is at $t = 2$ following the initial ($t = 0$) observation of normal and non-normal news (denoted as e_0 and ξ_0 respectively in Section 2.5).

C.2. The impact of non-normal news on the persistence of volatility

The empirical findings to date indicate that large shocks tend to have a less persistent impact on expected volatility than normal shocks (eg. Engle and Mustafa, 1992; Maheu and McCurdy, 2004). According to this stylized fact, the persistence effect on expected volatility from normal news is greater than that of non-normal news (or, in other words, that non-normal news has a significant immediate impact that appears to be quickly discounted in terms of the stock’s expected volatility). Indeed, if we restrict our attention to the persistence of the volatility associated with normal and non-normal news (which

are captured by equations (10) and (5) respectively) our findings effectively corroborate this result.

We propose, however, to augment the likelihood ratio tests in Section B.1 by using the news impact surface to examine the persistence of non-normal news. We note that in Section B.1 it is shown that the propagation or transmission mechanism is statistically significant, and that the propagation of non-normal news differs significantly from its normal counterpart (with $\Pi_A \geq \Pi_N$ where Π_A propagates non-normal news and Π_N propagates normal news).

The news impact surface, which accounts for the individual persistence of each type of news and the inter-temporal propagation of the different news across the asset space, is used to visualize the impact of the propagation mechanism and the asymmetry $\Pi_A \geq \Pi_N$ on expected volatility. The propagation mechanism is, in fact, key to measuring the persistence of news and its omission results in a measurement of persistence that fails to account for an important characteristic of the returns data.

To measure the relative persistence of non-normal news, we consider the news impact surface for asset i in the absence of any news from other assets. This provides a depiction of the relative importance of normal news on expected volatility whilst avoiding the need to consider the way in which news from other stocks obfuscates the impact surface. The surfaces observed at successive time periods examine the impact on expected volatility of an initial (time 0) set of non-normal and normal news.

The news impact surfaces show that the 1-period ahead expected impact of normal news induces a curvature in the expected volatility surface that expands with the magnitude of the normal news (see, further, Figure G2 in the online Appendix G). On the other hand, the 1-period ahead expected impact of non-normal news tends to be asymmetric and induces a fairly monotonic increase in expected volatility if the observed number of non-normal shocks exceeds the expected number (such that $\xi_0 > 0$).

In the following period, the curvature induced by the period 0 normal news declines sharply for all four stocks and the surface is almost fully characterized by steepness stemming from the period 0 non-normal news. Importantly, the maximal 2-step ahead volatility *exceeds* its 1-period ahead counterpart for all four stocks (in some cases by a substantial margin). The maximal volatility begins to decline at 3 or 4 steps ahead and asymptotes exponentially to the limiting surface described in the online Appendix A.

It is clear that the news transmission mechanism effectively *filters out* normal news and *filters in* non-normal news into the expected volatility equation. Through this filtering process, the impact of non-normal news expands whilst the impact of normal news declines even in the presence of highly persistent GARCH volatilities that characterize the volatility of normal news.

C.3. How does news affect asset return correlations?

The expected covariance surfaces depend heavily on the extent to which the period 0 news is normal or non-normal. When news is limited to the latter, all the covariance pairs are defined by hyperbolic paraboloid surfaces that are similar to a saddle. This is depicted in the online Appendix G (Figure G3). Saddle points are clearly observed that depict the ‘news-free’ covariance between the stock pairs; beginning at the saddle point, the expected covariance increases in the presence of normal (period 0) news from any two stocks that have the same sign (ie. two stocks that both have ‘good’ normal news or ‘bad’ normal news) and declines when the news is of opposing sign. This result suggests that the information content of no-news described in Giglio and Shue (2014) can be understood as the saddle point that the covariance function ‘converges’ to in the absence of any news. Put another way, the absence of news can be considered ‘news’ in itself in the sense that it establishes convergence towards the news-free saddle point.

It is interesting that the Microsoft-Google and Microsoft-Apple covariance surfaces contain negative values (implying that normal news of opposing sign can shift the sign of the expected covariance between the asset pairs) whereas the Microsoft-Intel surface is always positive. This may reflect the long-standing collaboration between Microsoft and Intel in the development of multiple generations of the Windows OS that have been optimized for use on Intel processors (in contrast, the Microsoft-Google and Microsoft-Apple pairings involve competing organizations).

The covariance surfaces become significantly more heterogeneous in the presence of non-normal news, and a common shape is no longer apparent (Figure 2). Most importantly, this can result in the expected covariance changing sign and becoming negative. As such, the expected covariance function’s dependence on non-normal news implies that even the sign of the correlation between stock returns is ambiguous in the short term. In particular, the expected correlation for two stocks that are normally positively correlated can switch from being positive to negative in the presence of significant non-normal news (even if the contemporaneous covariance in the jump components for the two stocks, Σ_{ij} , is positive).

In contrast, normal news alone generally does not change the sign of the expected covariance function for any of the asset pairings. These two results are particularly important for hedging and portfolio management purposes, indicating that there is effectively no need to hedge against the normal news of other stocks with the portfolio variance minimization problem being largely a function of identifying stocks with inversely related non-normal news (in effect, identifying sets of stocks whose covariance surface has large negative subspaces).

Figure 2: Expected covariance surface conditional on normal (or ‘typical’) and non-normal (or ‘atypical’) news. The expected covariance is presented at $t = 1$ and $t = 2$ and is conditional on the set of $t = 0$ normal and non-normal news for the depicted pair of assets. Note that although the figures are pairwise the depicted covariance also depends on the other assets in the model.

D. Leverage effects and asymmetries

The asymmetric response to negative normal news is widely held to describe a leverage effect associated with the devaluation of the equity component of a firm’s capital. The leverage effects are extended here with the incorporation of sign-dependent asymmetric effects in the non-normal news component thereby allowing for two possible leverage effects: an effect associated with devaluation attributed to normal news and a second effect associated with devaluations caused by non-normal news.

The leverage effects are obtained by allowing for asymmetric responses to news about the number of jumps ξ_t and to normal news e_t .¹⁵ The latter asymmetry follows a standard form (eg. Glosten, Jagannathan and Runkle, 1993), while the former asymmetry is based on the premise that the expected probability of n jumps in the next period is influenced by whether investors over- or under- estimate the number of jumps in the current period. Both asymmetries are relevant in determining the stock’s expected volatility; normal news e_t impacts directly but asymmetrically on the volatility of the normal news component

¹⁵The number of jumps can be defined as the standardised non-normal news component $\xi_{i,t} = u_{i,t}/x_{i,t}$, where $u_{i,t}$ is the overall non-normal or atypical news and $x_{i,t}$ is the conditional expectation of the jump size.

via H_t (Eq. (10)) whereas ξ_t impacts non-linearly on the volatility of the non-normal component through $\Sigma_{ii,t}^\lambda$ (Eq. (8)).

Likelihood ratio tests are used to evaluate whether the two types of asymmetries are statistically significant (formally, to test the null hypotheses $A_4 = 0$, $\gamma_2 = 0$, for the normal and non-normal news components respectively). As expected, the results for normal news are consistent with a long line of research indicating the presence of asymmetric effects in the GARCH equation (Table 2). However, we find that there are also significant asymmetric effects associated with non-normal news such that the expected volatility response to bad news (irrespective of its magnitude) differs depending on whether the devaluation stems from normal or non-normal news.

Table 2a,b: (a) Likelihood ratio tests of the asymmetric components (γ_2, A_4), and (b) the impact of asymmetries on non-normal and normal news .

(a)				(b)	Google	Apple	Intel	Microsoft
	LR statistic	df	95% Crit.	$\xi_t \geq 0$	1.157	1.737	1.213	0.143
$\gamma_2 = 0$	133.74	4	9.4877	$\xi_t < 0$	5.961	2.036	3.182	5.005
$A_4 = 0$	313.45	4	9.4877	$e_t \geq 0$	0.088	0.075	0.027	0.056
				$e_t < 0$	0.088	0.159	0.043	0.059

Note: The individual elements in $(\xi_t \geq 0, \xi_t < 0)$ and $(e_t \geq 0, e_t < 0)$ represent the marginal impact of asymmetries on the non-normal news intensity (5) and on the GARCH equation (10) respectively.

The unambiguous rejection of the null hypotheses suggests that there are two distinct sources of leverage effects. Closer inspection of the impact of the two types of news, however, renders it clear that leverage effects predominantly stem from non-normal news and appear to be driven by expectations regarding the probability of a jump. Table 2(b) shows the marginal impact of ξ_t and e_t on the step-ahead values of the jump intensity λ_t and the GARCH volatility H_t . Negative normal news has a relatively small incremental impact on the volatility of the normal news H_t . In contrast, there is a large differential impact when $\xi_t < 0$ relative to $\xi_t \geq 0$ resulting in an expected news intensity λ_{it} that depends heavily on the sign of ξ_t for all four stocks.

The findings indicate that expectations regarding non-normal news (in particular, the extent to which investors have over- or under- estimated jump probabilities) are key to the measurement of the leverage effect, with standard asymmetric volatility models failing to reflect the additional volatility implications stemming from non-normal news asymmetries.

6. Volatility forecasting performance

We compare the out of sample volatility forecasts from the news transmission model with those from a set of competing models for the 516 days from the end of our sample period to the end of December 2015. The competing models are the full BEKK, diagonal BEKK, scalar BEKK, the GARCH(1,1)-DCC, the GJR-GARCH-t model and a univariate GARCH-jump model. The first three models belong to the BEKK family of multivariate volatility models, whereas the GARCH(1,1)-DCC is a popular multivariate volatility model. The GJR-GARCH-t (Glosten, Jagannathan and Runkle, 1993) is a widely used univariate volatility model designed to account for asymmetric (or leverage) effects (we note that the other competing models are also estimated subject to the potential presence of leverage effects). The univariate GARCH-jump model is based on Chan and Maheu (2002) and incorporates both GARCH and jump processes.¹⁶

In order to assess model performance under potential parameter uncertainty, we do not re-estimate the model parameters during the forecasting period. This yields information about model performance under potentially different conditions to those used to estimate the model parameters. We follow Maheu and McCurdy (2004) in comparing the forecasting accuracy of the period-ahead volatility forecast made in day $t - 1$ with the actual log-range for day t .¹⁷ To examine model performance, we consider three forecasting exercises. The first covers the entire forecasting period, the second covers periods of high volatility and the third covers the 10-day period following large negative returns. High volatility is defined by reference to periods where the stock's range is greater than its mean plus one standard deviation. Large negative returns are when $r_t < -2.0$. Table 3 presents the MAE and RMSE values for forecasts produced by the competing models, in addition to the p-values associated with the application of the Model Confidence Set test of equal predictive ability proposed by Hansen, Lunde and Nason (2011).

The results provide strong evidence in favour of the news transmission model, which is in the 90 per cent model confidence set in almost all cases and is clearly the dominant model across the four stocks. In addition to being the preferred model in all three forecasting exercises, the forecasts from the news transmission model also produce smaller RMSE and MSE values than the competing models.

We note, in particular, the news transmission model's significant forecast improvement on the diagonal BEKK model. Since the diagonal BEKK forms the basis of the normal news innovations in the proposed model, any improvement of the model on the diagonal BEKK indicates the importance of accounting for the transmission of non-normal news.

¹⁶Chan and Maheu (2002) find that the out-of-sample performance of the model with constant jump intensity is similar to that of the model with time-varying jump intensities. We therefore estimate the simpler model based on a constant jump intensity.

¹⁷The log-range (being the log of the ratio of the intraday high to the intraday low prices) is scaled so that its sample mean accords with the sample standard deviation of the relevant return.

Table 3. Forecasting evaluation of the news transmission model

	MAE				RMSE				pMCS						
	Google	Apple	Intel	Microsoft	All	Google	Apple	Intel	Microsoft	All	Google	Apple	Intel	Microsoft	All
All data															
News model	0.892	0.734	0.819	0.963	0.853	1.290	1.2530	1.3320	1.4168	1.324	1.000*	1.000*	1.000*	0.000	1.000*
Full-BEKK	1.007	0.793	0.863	0.942	0.902	1.434	1.3999	1.3889	1.3969	1.405	0.000	0.012	0.000	0.000	0.000
Diag-BEKK	1.034	0.799	0.847	0.920	0.900	1.459	1.3916	1.3723	1.3772	1.400	0.000	0.001	0.137*	0.000	0.000
Scalar-BEKK	0.923	0.785	0.869	0.910	0.872	1.343	1.3928	1.3969	1.3676	1.375	0.057	0.012	0.000	0.879*	0.032
DCC	0.931	0.766	0.872	0.926	0.874	1.347	1.3276	1.3680	1.3616	1.351	0.048	0.027	0.223*	1.000*	0.052
GJR-GARCH-t	0.968	0.772	0.896	0.945	0.895	1.374	1.3367	1.3885	1.3679	1.368	0.000	0.012	0.024	0.331*	0.002
GARCH-Jump	0.974	0.802	0.895	0.972	0.911	1.372	1.377	1.407	1.415	1.393	0.003	0.010	0.000	0.000	0.000
Obs	516	516	516	516											
High volatility periods															
News model	2.773	2.845	3.132	3.147	2.972	2.988	3.681	3.720	3.373	3.408	1.000*	1.000*	1.000*	0.000	1.000*
Full-BEKK	3.111	3.394	3.321	3.162	3.228	3.299	4.202	3.891	3.363	3.644	0.000	0.001	0.000	0.000	0.000
Diag-BEKK	3.142	3.389	3.287	3.131	3.219	3.325	4.148	3.863	3.336	3.626	0.000	0.001	0.000	0.000	0.000
Scalar-BEKK	2.957	3.368	3.341	3.113	3.167	3.157	4.196	3.910	3.321	3.599	0.005	0.014	0.000	0.261*	0.000
DCC	2.948	3.148	3.134	3.052	3.052	3.147	3.939	3.750	3.253	3.481	0.005	0.014	0.689*	0.470*	0.055
GJR-GARCH-t	2.946	3.135	3.127	3.041	3.050	3.147	3.957	3.755	3.245	3.484	0.002	0.014	0.689*	1.000*	0.055
GARCH-Jump	2.858	3.241	3.277	3.087	3.116	3.091	4.041	3.850	3.334	3.579	0.075	0.001	0.000	0.261*	0.000
Obs	64	44	46	63											
10-day periods following large negative returns (<-2.0)															
News model	1.251	0.994	1.209	1.173	1.150	1.732	1.927	2.134	1.783	1.884	1.000*	1.000*	1.000*	0.173*	1.000*
Full-BEKK	1.512	1.193	1.380	1.317	1.351	1.987	2.242	2.293	1.881	2.106	0.000	0.046	0.000	0.000	0.000
Diag-BEKK	1.542	1.193	1.353	1.294	1.350	2.007	2.199	2.273	1.868	2.092	0.000	0.045	0.000	0.000	0.000
Scalar-BEKK	1.391	1.176	1.394	1.280	1.305	1.879	2.240	2.308	1.859	2.074	0.000	0.046	0.000	0.081	0.000
DCC	1.335	1.053	1.292	1.114	1.197	1.819	2.058	2.144	1.693	1.934	0.001	0.046	0.923*	0.421*	0.001
GJR-GARCH-t	1.360	1.064	1.309	1.111	1.210	1.835	2.069	2.146	1.676	1.939	0.000	0.046	0.923*	1.000*	0.001
GARCH-Jump	1.376	1.075	1.310	1.135	1.224	1.841	2.085	2.135	1.740	1.939	0.000	0.046	0.000	0.421*	0.000
Obs	370	365	237	257											

Note: The full BEKK, diagonal BEKK, scalar BEKK, and DCC models are estimated with leverage effects. The conditional variances for the DCC model are based on (and therefore equivalent to) the GARCH(1,1) model with leverage effects. The GARCH-Jump is based on Chan and Maheu (2002) with constant jump intensity. pMCS is the probability associated with the model confidence set on the squared forecast errors. The block lengths used in the bootstrapping procedure (based on 10,000 resamples) for the pMCS test are determined using the maximum number of significant parameters obtained when fitting an AR process to the squared errors (Hansen, Lunde and Nason, 2011). The superscript * denotes that the relevant model is in the 90 per cent model confidence set. The column denoted 'All' is based on the vector of model-specific errors for all four stocks (eg. the vector of model-specific forecast errors with length 4×516 when evaluating the MAE, RMSE and pMCS using all forecast observations). High volatility periods are defined as those where the stock's range is greater than a standard deviation above its mean.

In this respect, a key issue with the diagonal BEKK model (or, indeed, the full or scalar BEKK model) is that it restricts all news to have the same impact. Consequently, it is not possible to distinguish between the impact of normal and non-normal news on the conditional variance of an asset return. The significant forecast improvement yielded by the news transmission model indicates that the latter restriction can have substantial ramifications for forecasting accuracy. Overall, the results provide clear support for our general finding that both news type and the news transmission process are important for determining an asset's volatility.

7. Conclusion

We examine the impact of information cross-flows on the behavior of asset prices, including how information from a given stock affects the volatility of other stocks. In particular, we develop a model that categorizes an asset's news into normal and non-normal news, and allows both types of news to influence the expected volatility of other assets. Each asset therefore potentially produces news that can be normal or non-normal, with the two types of news having their own unique distribution. This news is then filtered through to other assets such that the asset's expected volatility is a function of its own news and the news of other assets.

The model is used to consider how the transmission of normal news differs from that of non-normal news, to examine the relative importance of an asset's own news relative to the news of other assets, and to establish the type and magnitude of news from other assets that is sufficient to 'dominate' an asset's own news. To address these questions, we derive a news impact surface that characterizes an asset's expected variance (and its covariance with other assets) in terms of the asset's own (normal and non-normal) news and the news of other assets. The surface provides insights into the dynamics underlying the impact of news on volatility expectations, and allows us to reach a number of important conclusions regarding the inter-relationship between news and expectations.

We find that the expected volatility surface changes dramatically in the presence of non-normal news such that a stock in a given portfolio can *reduce or increase* the portfolio's volatility even if the stock's own volatility rises following the news.

In exploring the impact of news type on the expected volatility surface, we observe clearly different asymmetries associated with normal and non-normal news. These asymmetries allow us to identify two distinct sources of leverage effects: a leverage effect for normal news and a (more important) leverage effect for non-normal news. Accordingly, estimates of the leverage effect based on GARCH or stochastic volatility models will typically underestimate the extent of the effect in stock returns data.

We also show that news transmission tends to crowd-in non-normal news and crowd-out normal news, such that the effect of non-normal news on volatility expectations often

builds up before dissipating. This produces a shift in the persistence of the impact of non-normal and normal news that negates the stylized fact that larger shocks have a less persistent impact on expected volatility than normal shocks.

Finally, we provide evidence that non-normal news is primarily responsible for the negative subspace in the covariance surface of the stocks that we consider. For these types of stocks, non-normal news can impact significantly on their expected correlation such that the ability to determine even the sign of their correlation is ambiguous in the short term.

References

- [1] Aitchison J. and Ho, C.H., 1989, The multivariate Poisson-Log normal distribution, *Biometrika* 76(4), 643-653.
- [2] Andersen, T.G., 1996, Return volatility and trading volume, *Journal of Finance* 51, 169-204.
- [3] Bae, K-H., Ozogus, A., Tan, H., and Wirjanto, T.S., 2012, Do foreigners facilitate information transmission in emerging markets?, *Journal of Financial Economics* 105(1), 209-227.
- [4] Bandi, F.M., and Reno, R., 2016, Price and volatility co-jumps, *Journal of Financial Economics* 119(1), 107-146..
- [5] Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31(3), 307-327.
- [6] Campbell, J.Y. and Hentschel, L. 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281-318.
- [7] Chan, W.H. and Maheu, J.M., 2002, Conditional jump dynamics in stock market returns, *Journal of Business and Economic Statistics* 20, 377-389.
- [8] Chen, X., and Ghysels, E., 2011, News—good or bad—and its impact on volatility predictions over multiple horizons, *Review of Financial Studies* 24, 46-81.
- [9] Christoffersen, P., Jacobs, K., and Ornathanalai, C., 2012, Dynamic jump intensities and risk premiums: Evidence from S&P500 returns and options, *Journal of Financial Economics* 106(3), 447-472.
- [10] DeLong, B.J., Shleifer, A., Summers, L. and Waldmann, R.J., 1990, Noise trader risk in financial markets, *Journal of Political Economy* 98 (4), 703-738.
- [11] Engle, R.F., 2002, Dynamic conditional correlation—a simple class of multivariate GARCH models, *Journal of Business and Economic Statistics* 20, 339-350.
- [12] Engle, R.F., and Kroner, K.F., 1995, Multivariate simultaneous generalized ARCH, *Econometric Theory* 11, 122-150.

- [13] Engle, R.F., and Mustafa, C., 1992, Implied ARCH models from options prices, *Journal of Econometrics* 52, 289 – 311.
- [14] Engle, R.F., Hansen, M.K. and Lunde, A. 2012, And now, the rest of the news: Volatility and firm specific news arrival, CREATES Research Paper 2012-56.
- [15] Engle, R.F., and Ng, V.K. 1993, Measuring and testing the impact of news on volatility, *Journal of Finance* 48(5), 1749-1778.
- [16] Giglio, S., and Shue, K., 2014, No news Is news: Do markets underreact to nothing?, *Review of Financial Studies* 27(12), 3389-3440.
- [17] Glosten, L.R., Jagannathan, R., and Runkle, D.E., 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779-1801.
- [18] Glosten, L.R., and Milgrom, P.R., 1985, Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71-100.
- [19] Hansen, P.R., Lunde, A., and Nason, J.M., 2011, The model confidence set, *Econometrica* 79(2), 453–497.
- [20] Hamilton, J.D., 1989, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57, 357-384.
- [21] Hamilton, J.D., 1990, Analysis of time series subject to changes in regime, *Journal of Econometrics* 45, 39-70.
- [22] Karolyi, G.A., 1995, A multivariate GARCH model of international transmissions of stocks returns and volatility: The case of United States and Canada, *Journal of Business and Economic Statistics* 13, 11-25.
- [23] Kim, C-J., 1994, Dynamic linear models with Markov-switching, *Journal of Econometrics* 60(1-2), 1-22.
- [24] Kuhnen, C.M., 2015, Asymmetric learning from financial information, *Journal of Finance* 70(5), 2029-2062.
- [25] Li, J., Todorov, V., and Tauchen, G. 2017, Jump regressions, *Econometrica* 85(1), 173-195.
- [26] Maheu, J.M., and McCurdy, T.H., 2004, News arrival, jump dynamics, and volatility components for individual stock returns, *Journal of Finance* 59(2), 755-793.
- [27] Ord, J.K., Koehler, A.B, and Snyder, R.D., 1997, Estimation and prediction for a class of dynamic nonlinear statistical models, *Journal of the American Statistical Association* 92, 1621-1629.

- [28] Patton, A. and Verardo, V., 2012, Does beta move with news? Firm specific information flows and learning about profitability, *Review of Financial Studies* 25; 2789-2839.
- [29] Patton, A. and Sheppard, K., 2015, Good volatility, bad volatility: Signed jumps and the persistence of volatility, *Review of Economics and Statistics* 97(3), 683-697.
- [30] Ross, S.A., 1989, Information and volatility: The no-arbitrage martingale approach to timing and resolution irrelevancy, *Journal of Finance* 44, 1-17.
- [31] Scherbina, A. and Schlusche, B., 2015, Cross-firm information flows and the predictability of stock returns, mimeo.
- [32] Schwert, G.W., 1990, Stock market volatility and the crash of '87, *Review of Financial Studies* 3, 77-102.

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Tsiaplias, S;Chua, CL

Title:

Information flows and stock market volatility

Date:

2019-02

Citation:

Tsiaplias, S. & Chua, C. L. (2019). Information flows and stock market volatility. Journal of Applied Econometrics, 34 (1), pp.129-148. <https://doi.org/10.1002/jae.2649>.

Persistent Link:

<http://hdl.handle.net/11343/284059>