What Types of Events Provide the Strongest Evidence that the Stock Market is Affected by Company Specific News?

Calum Robertson
School of Software Engineering and Data Communications
Faculty of Information Technology
Queensland University of Technology
Level 7, S Block, 2 George Street
Brisbane, Queensland Australia 4000
cs.robertson@qut.edu.au

Shlomo Geva
School of Software Engineering and Data Communications
Faculty of Information Technology
Queensland University of Technology
Level 7, S Block, 2 George Street
Brisbane, Queensland Australia 4000
s.geva@qut.edu.au

Rodney Wolff
School of Economics and Finance
Faculty of Business
Queensland University of Technology
Level 8, Z Block, 2 George Street
Brisbane, Queensland Australia 4000
r.wolff@qut.edu.au

Abstract

The efficient market hypothesis states that an efficient market immediately incorporates all available information into the price of the traded entity. It is well established that the stock market is not an efficient market as it consists of numerous traders with differing strategies and interpretations of information. However there is substantial evidence to suggest that the stock market does incorporate new information into prices. Unfortunately little research has focussed on the high frequency effect of real time news, across a broad base of assets. This paper investigates how the US, UK, and Australian markets incorporate all real time news, not just Press Announcements, Annual Reports, etc. We find that there is strong evidence to suggest that the markets do incorporate news quickly.

Keywords. Stock Market, News, Return, Volatility, Market Reaction.

1. Introduction

A plethora of research is available which shows that the occurrence of news does effect the market, with the majority of research focusing on macroeconomic news, which provides an indication of the state of the economy (Almeida, Goodhart and Payne 1998, Bomfim 2003, Brannas and De Gooijer 2004, Ederington and Lee 1993, 1995, 2001, Ewing 2002, Graham, Nikkinen and Sahlstrom 2003, Han and Ozocak 2002, Hess 2004, Kim 1998, 2003, Kim, McKenzie and Faff 2004, Nikkinen and Sahlstrom 2004a, 2004b, Nofsinger and Prucyk 2003, Simpson and Ramchander 2004, Sun and Sutcliffe 2003, Tse 1999). However macroeconomic news is relatively infrequent compared to asset specific information, e.g. Simpson and Ramchander (2004) state that the United States of America releases 23 macroeconomic reports regularly, usually monthly, whilst Fung, Yu and Wai (2003) found an average over 373 news articles per asset per month. Not only is asset specific information more frequent but it has been shown to have a noticeable effect on the given asset (Chan 2003, Donders and Vorst 1996, Dungey, Fry and Martin 2004, Fung, Yu and Wai 2003, Goodhart 1989, Goodhart and Figliuoli 1992, Goodhart, Hall, Henry and Pesaran 1993, Hong, Lim and Stein 2000, Melvin and Yin 2000, Michaely and Womack 1999, Mitchell and Mulherin 1994, Mittermayer 2004, Roll 1984, Womack 1996, Wuthrich, Permunetilleke, Leung, Cho, Zhang and Lam 1998).

Previous research has shown that the market reacts quickly to macroeconomic news (Ederington and Lee 1993, 1995, 2001, Han and Ozocak 2002, Nofsinger and Prucyk 2003). However as macroeconomic news is scheduled (i.e. the market is aware exactly when the news is released) the market anticipates the content of news and can react quickly based on whether the actual news matches analyst forecasts. It is not known how rapidly the stock market reacts to non-macroeconomic news and therefore it would be interesting to determine if the market responds in a timely manner, if at all, to non-macroeconomic news.

Most research to date which investigates the intraday effect (reaction of the market on the day which the news was released) of news has focussed on the Foreign Exchange markets (Almeida, Goodhart and Payne 1998, Bollerslev and Domowitz 1993, Ederington and Lee 1993, 1995, 2001, Goodhart 1989, Goodhart and Figliuoli 1991, 1992, Goodhart, Hall, Henry and Pesaran 1993, Han and Ozocak 2002, Melvin and Yin 2000, Peiers 1997), or Futures markets (Hess 2004, Tse 1999), whilst little has focussed on the Stock market (Mittermayer 2004, Nofsinger and Prucyk 2003). Nofsinger and Prucyk (2003) investigated the intraday effect of macroeconomic news on the S&P 100 Index Option and found that bad news with high information surprise is responsible for most abnormal volume associated with macroeconomic news. Mittermayer (2004) investigated the effect of Press Announcements on stocks on the New York Stock Exchange and NASDAQ and found evidence to suggest that the market does react to the news, and furthermore the content of the news which triggered the reaction. However Press Announcements are relatively rare compared to other types of news available from real time news providers, so it bears further investigation.

The purpose of this paper is to identify events which occur with a high correlation to the occurrence of real time news, such that they can be used to identify "interesting" news articles. Furthermore this paper aims to investigate the percentage of news articles which the market appears to find significant, as the sheer volume of news available would prohibit an individual from reading all available news.

2. Data

All data for this research was obtained using the Bloomberg Professional® service. The dataset consists of stocks which were in the S&P 100, FTSE 100, and ASX 100 indices as at the 1st of July 2005 and continued to trade through to the 1st of September 2006, which is a total of 286 stocks. For each stock the Trading Data, and News were collected for the period beginning 1st of May 2005 through to and including the 31st of August 2006.

The set defined in Eq. (1) consists of each distinct minute where trading occurred for the stock (s), within all minutes for the period of data collection (T_A). For each minute the average price, the volume, and the number of ticks (number of trades) for trades during that minute are also stored. However we are only interested in the business time scale (minutes which occurred during business hours for the market on which the stock trades). Furthermore we want a heterogeneous time series (i.e. an entry for every business trading minute for the stock, regardless of whether any trading occurred). Therefore we produce the date, price, volume, and tick time series for all minutes in the business time scale (T_B) with the definitions in Eqs. (2)-(5). We define the price at time t as the price of the last actual trade for the stock prior to or at the given time. We set the volume and ticks equal to 0 if there wasn't a trade at time t. Note that if the stock was suspended from trading for a whole day then the day is excluded from T_B .

$$I_{(s)} = \{I_1, I_2, \ldots, I_m\} \mid I_{(s,z)} = \left(d_{(s,z)}, p_{(s,z)}, v_{(s,z)}, k_{(s,z)}\right) \wedge z \in \mathrm{T}_A \tag{1}$$

$$D_{(s)} = \{D_1, D_2, \ldots, D_n\} \mid D_{(s,t)} > D_{(s,t-1)} \wedge D_{(s,t)} \in \mathrm{T}_B \wedge \mathrm{T}_B \subseteq \mathrm{T}_A \tag{2}$$

$$P_{(s)} = \{P_1, P_2, \ldots, P_n\} \mid P_{(s,t)} = \left(p_{(s,z)} \mid z = \max\left(z \mid d_{(s,z)} \le D_{(s,t)}\right)\right) \tag{3}$$

$$V_{(s)} = \{V_1, V_2, \ldots, V_n\} \mid V_{(s,t)} = \left(\exists! d_{(s,z)} = D_{(s,t)} \ ? \ v_{(s,z)} \ : \ 0\right) \tag{4}$$

$$T_{(s)} = \{T_1, T_2, \ldots, T_n\} \mid T_{(s,t)} = \left(\exists! d_{(s,z)} = D_{(s,t)} \ ? \ k_{(s,z)} \ : \ 0\right) \tag{5}$$

The news search facility within the Bloomberg Professional® service was used to download all relevant articles for each stock within the dataset. These articles include Press Announcements, Annual Reports, Analyst Recommendations and general news which Bloomberg has sourced from over 200 different news providers. The set defined in Eq. (6) consists of each distinct news article for the stock and contains the time and content of the article. However we are only interested in the business time scale and are only concerned whether news occurred at the given time. Therefore we produce the news time series defined in Eq. (7) such that each business trading minute for the stock contains the count of the articles which occurred during it. If an article occurs after hours then it is stored in the first trading minute of the next trading day, as defined in Eq. (7).

$$A_{(s)} = \{A_1, A_2, \ldots, A_p\} \mid A_{(o)} = \left(d_o, c_o\right) \wedge s \in \mathrm{T}_A \tag{6}$$

$$N_{(s)} = \{N_1, N_2, \ldots, N_q\} \mid N_{(s,t)} = \prod \left\{\forall A_{(s)} \mid D_{(t-1)} < d_o \le D_{(t)}\right\} \tag{7}$$

3. Methodology

The return time series for a stock (s), defined in Eq. (8), is formed by taking the difference in the log prices from the trading data defined in Eq. (3) over the period Δt. The return time series identifies periods of high return which may indicate that the market is reacting to news.

The change in volume time series for a stock (s), defined in Eq. (9), is formed by taking the difference in the log of the average volume, over n minutes, between the time t and t-Δt. The volume defined in Eq. (4) is averaged over n minutes to limit the effect of trading minutes where no trade occurred. The conditional log of the average volume is used such that in the case where no trade occurs during the given n minutes, the function still produces an answer. The change in volume time series detects periods where there is a sudden increase in the volume of the stock traded, which might suggest that the market is reacting to news.

$$R_{(s,\Delta t)} = \{R_1, \ldots, R_m\} \mid R_{(s,t,n,\Delta t)} = \log\left(P_{(s,t)}\right) - \log\left(P_{(s,t-\Delta t)}\right) \tag{8}$$

$$CV_{(s,n,\Delta t)} = \{CV_1, \ldots, CV_m\} \mid CV_{(s,t,n,\Delta t)} = L\left(V_{(s,t,n)}\right) - L\left(V_{(s,t-\Delta t,n)}\right) \\ \wedge V_{(s,t,n)} = \frac{1}{n}\sum_{i=0}^{n-1} V_{(s,t-i)} \wedge L(x) = \left((x>0) \ ? \ \log(x) \ : \ \log(0.5)\right) \tag{9}$$

$$CT_{(s,n,\Delta t)} = \{CT_1, \ldots, CT_m\} \mid CT_{(s,t,n,\Delta t)} = L\left(T_{(s,t,n)}\right) - L\left(T_{(s,t-\Delta t,n)}\right) \\ \wedge T_{(s,t,n)} = \frac{1}{n}\sum_{i=0}^{n-1} T_{(s,t-i)} \wedge L(x) = \left((x>0) \ ? \ \log(x) \ : \ \log(0.5)\right) \tag{10}$$

$$\upsilon_{(s,n,\Delta t)} = \{\upsilon_1, \ldots, \upsilon_m\} \mid \upsilon_{(s,t,n,\Delta t)} = \sqrt{\frac{y}{\Delta t}} \times \left[\frac{1}{n}\sum_{i=0}^{n-1}\left(R_{(s,t-i,n,\Delta t)} - M_{(s,t,n,\Delta t)}\right)^p\right]^{\frac{1}{p}} \\ \wedge M_{(s,t,n,\Delta t)} = \frac{1}{n}\sum_{i=0}^{n-1} R_{(s,t-i,\Delta t)} \wedge y = 250 \times \begin{pmatrix} 390 \mid s \in US \\ \vee \ 510 \mid s \in UK \\ \vee \ 360 \mid s \in AU \end{pmatrix} \tag{11}$$

The change in ticks time series for a stock (s), defined in Eq. (10), is formed by taking the difference in the log of the average ticks, over n minutes, between the time t and t-Δt. The ticks defined in Eq. (5) are averaged, and the conditional log is used for the same reasons as for the volume. The change in ticks time series pinpoints periods where there is a sudden increase in the number of trades, which implies that the market is reacting to news.

The volatility time series for a stock (s), defined in Eq. (11), calculates the annualised volatility of the stock. When p is set to 2 it is simply the annualised variance of the return of the stock. The y value is calculated by multiplying the number of trading days per year (generally set to 250), by the number of trading minutes for the day (390 minutes for the US, 510 for the UK, and 360 for Australia). The volatility is annualised such that the results of each country can be directly compared, as suggested by Dacorogna, Gencay, Muller, Olsen and Pictet (2001). The volatility time series discovers points where the stock price changes rapidly, which could insinuate that the market is reacting to news.

The stocks are grouped together as per Eq. (12) so that we can examine the effect of news on an individual country. We divide the trading day into equally sized time windows ΔT, as defined in Eq. (13), in order to examine the intraday effect of the news. Note that the first period is ignored because we don't want the after hours news and market behaviour to skew results. The first period is the larger of ΔT and Δt.

We define a generalised time series F, where F is one of the return, volume, tick, and volatility time series. We use the generalised time series to define the Event Point Process (EPP) in Eq. (14). In this point process a point value of 1 indicates that the generalised time series for the given stock exceeded the specified threshold (x), which we will refer to as an event. It should be noted that the return, volume and tick time series are log values about 0, so a threshold of 10% means that the value should be ≥ to log(11/10) or ≤ to log(10/11). The volatility time series is always positive when p is even so the ≤ condition is ignored.

$$S = \{S_1, S_2, \ldots, S_h\} \mid h \geq 1 \tag{12}$$

$$\begin{aligned} &W = \{W_1, W_2, \ldots, W_g\} \mid \\ &W_{(w,\Delta t,\Delta \mathrm{T})} = \{\forall \mathrm{T}_B \mid t_0 \leq time(T_{(B,\ i)}) < t_0 + \Delta \mathrm{T} \ \wedge \ t_0 = t_f + (w-1) \times \Delta \mathrm{T}\} \\ &\wedge t_f = \min(time(T_B)) + \max(\Delta \mathrm{T}, \Delta t) \end{aligned} \tag{13}$$

$$\begin{aligned} &E_{(w,\Delta \mathrm{T},S,n,\Delta t,x)} = \{E_1, \ldots, E_m\} \mid E_{(w,\Delta \mathrm{T},s,t,n,\Delta t,x)} = \left(F_{(s,t,n,\Delta t)} \begin{pmatrix} \geq x \\ \vee \leq -x \end{pmatrix} \right) ?1:0 \\ &\wedge s \in S \ \wedge d_{(s,t)} \in W_{(w,\Delta t,\Delta \mathrm{T})} \end{aligned} \tag{14}$$

$$\begin{aligned} &EN_{(w,\Delta \mathrm{T},S,n,\Delta t,x,\Delta\tau)} = \{EN_1, \ldots, EN_m\} \mid EN_{(w,\Delta \mathrm{T},s,t,n,\Delta t,x,\Delta\tau)} = \\ &\left((E_{(w,\Delta \mathrm{T},s,t,n,\Delta t,x)} = 1) \wedge \left(\prod_{i=1}^{\Delta\tau} A_{(s,\ t-i)} > 0 \right) \right) \ ?\ 1 \ :\ 0 \end{aligned} \tag{15}$$

$$\begin{aligned} &E\overline{N}_{(w,\Delta \mathrm{T},S,n,\Delta t,x,\Delta\tau)} = \{E\overline{N}_1, \ldots, E\overline{N}_m\} \mid E\overline{N}_{(w,\Delta \mathrm{T},s,t,n,\Delta t,x,\Delta\tau)} = \\ &\left((E_{(w,\Delta \mathrm{T},s,t,n,\Delta t,x)} = 1) \wedge \left(\prod_{i=1}^{\Delta\tau} A_{(s,\ t-i)} = 0 \right) \right) \ ?\ 1 \ :\ 0 \end{aligned} \tag{16}$$

The Event given News Point Process (ENPP) is defined in Eq. (15), where a point value of 1 symbolises that an event occurred at time t for the stock s in the EPP and at least one company specific news article arrived between t-Δτ and t-1. The Event Without news Point Process (EWPP) defined in Eq. (16) has a point value of 1 when an event occurred at time t for the stock s in the EPP and no company specific news arrived between t-Δτ and t-1.

The Ratio of Events Related to News to Events (RERNE) defined in Eq. (17) indicates the percentage of events which are preceded by news. A high RERNE value suggests that most events for the given parameters are preceded by news, which would imply that news is responsible for these events. A low RERNE value can denote that the market takes longer than the specified Δτ time to react to news, or that the events are caused by other factors, or that the events themselves are merely noise.

The Benchmark defined in Eq. (18) provides a measure of the likelihood of news arriving within the specified Δτ time. This is achieved by calculating a return of 0% in Eq. (14), which produces a point process where every point has a value of 1, and therefore the point process in Eq. (15) simply indicates the points when news occurs within the specified Δτ time.

This Benchmark is then used to calculate the Likelihood that Events are Related to News (LERN) in Eq. (19). A high LERN value implies that it is more likely for news to occur prior to an event that it is normally. A LERN value equal to 100% indicates that it is just as likely for news to occur before an event as it is to occur at any other time. A low LERN value signifies that it is less likely for news to occur prior to an event than normal, which would imply that news isn't responsible for the event.

$$RERNE_{(w,\Delta \mathrm{T},S,n,\Delta t,x,\Delta\tau)} = \frac{\sum EN_{(w,\Delta \mathrm{T},S,n,\Delta t,x,\Delta\tau)}}{\sum E_{(w,\Delta \mathrm{T},S,n,\Delta t,x)}} \tag{17}$$

$$B_{(w,\Delta \mathrm{T},S,n,\Delta t,\Delta\tau)} = \frac{\Sigma}{\Pi} EN_{(w,\Delta \mathrm{T},S,n,\Delta t,0,\Delta\tau)} \tag{18}$$

$$LERN_{(w,\Delta \mathrm{T},S,n,\Delta t,x,\Delta\tau)} = \frac{RERNE_{(w,\Delta \mathrm{T},S,n,\Delta t,x,\Delta\tau)}}{B_{(w,\Delta \mathrm{T},S,n,\Delta\tau)}} \tag{19}$$

The Event T-Test (ETT) defined in Eq. (20) performs a Student t-Test on the period distribution of the chance of an event occurring with news versus the chance of an event occurring without news during each period. The purpose of the ETT is to test the null hypothesis that the occurrence of events is not influenced by the occurrence of news.

The News T-Test (NTT) defined in Eq. (21) performs a Student t-Test on the period distribution of the chance of news occurring prior to an event versus the chance of news occurring during each period. The intent of the NTT is to test the null hypothesis that the occurrence of news before events is the same as the occurrence of news normally.

$$\begin{aligned} &ETT_{(W,\Delta T,S,n,\Delta t,x,\Delta\tau)} = tTest\left(\left\{ \forall \frac{\Sigma}{\Pi} EN_{(w,\Delta T,S,n,\Delta t,x,\Delta\tau)} \mid w \in W \right\}, \right. \\ &\left. \left\{ \forall \frac{\Sigma}{\Pi} E\overline{N}_{(w,\Delta \mathrm{T},S,n,\Delta t,x,\Delta\tau)} \mid w \in W \right\} \right) \end{aligned} \tag{20}$$

$$\begin{aligned} &NTT_{(W,\Delta T,S,n,\Delta t,x,\Delta\tau)} = tTest\left(\left\{ \forall \frac{\Sigma}{\Pi} EN_{(w,\Delta T,S,n,\Delta t,x,\Delta\tau)} \mid w \in W \right\}, \right. \\ &\left. \{\forall B_{(w,\Delta \mathrm{T},S,n,\Delta t,\Delta\tau)} \mid w \in W\} \right) \end{aligned} \tag{21}$$

4. Results

Table 1 shows some of the characteristics of the dataset. The number of trading minutes during standard business days, week, month, and year can be used by the reader to appreciate how frequently events occur in the Table 2.

The Average Minutes without a Trade gives an indication of how many minutes within normal trading hours have no trades for each country. Clearly trading on the US market is more frequent than on the others. Bearing in mind that the Australian market is far smaller than the other two, it shouldn't be a surprise that there is less activity than on the others.

The News Articles in Dataset gives the reader an idea of the frequency of news in the different markets.

Variable	US	UK	AU
Trading Minutes Per Business Day	390	510	360
Trading Minutes Per Typical Business Week	1,950	2,550	1,800
Trading Minutes Per Typical Business Month	8,125	10,625	7,500
Trading Minutes Per Typical Business Year	97,500	127,500	90,000
Average Minutes without a Trade (%)	2.36%	37.65%	50.68%
News Articles in Dataset	293,416	136,627	130,988
Average After Hours News Articles (%)	57.76%	43.22%	76.61%

Table 1. This table shows some characteristics of the dataset for each country to help the reader appreciate later results.

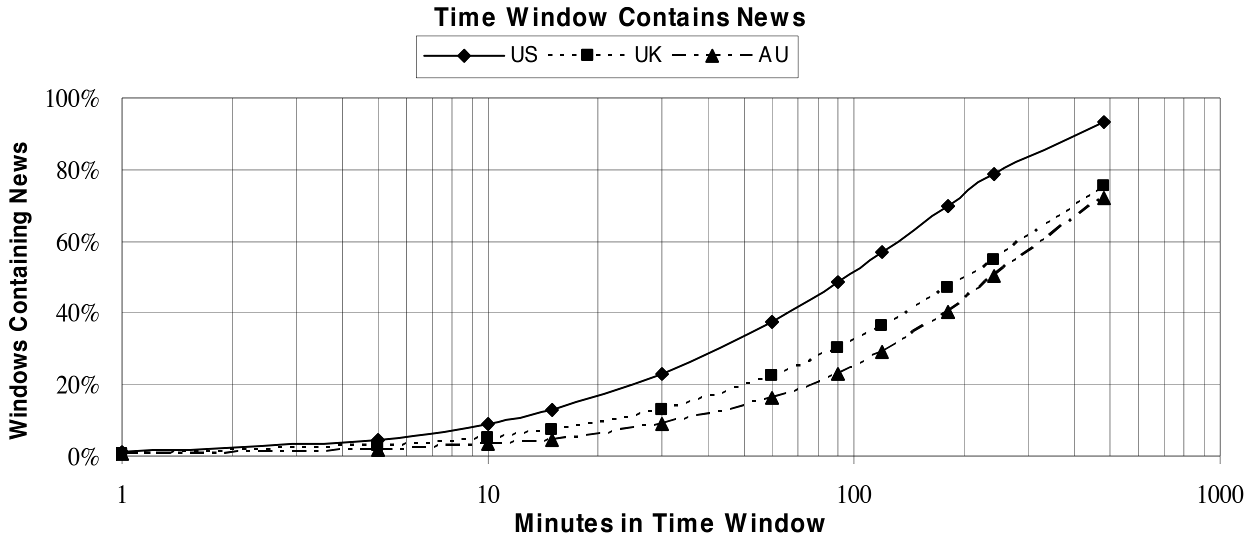

Fig. 1. As the size of the time window is increased the percentage of windows which contain news also does. Clearly there is more news distributed in the US market than the others.

Obviously the US receives far more information than the other markets, which means it should be easier to identify events linked to news articles within that market.

The Average After Hours News articles give the percentage of the articles which occur outside business hours in the different markets. The UK market has the lowest ratio but this could be a factor of longer trading hours than the other markets. The Australian market has the least business hours, and the highest after hours news ratio, though it is a far smaller market than the other two, so it shouldn't be a surprise that less information is made available during the trading day.

The choice of a time window $\Delta\tau$ in which news can be found is significant. If the percentage of minutes which contain a news article within the time window is too high it is difficult to establish whether the news was responsible for the event, or whether it was a coincidence. The Benchmark results using Eq. (18), shown in Fig. 1, identify the chance of finding news within the given time window. The figure shows that a reasonable amount of minutes are preceded by news within 60 minutes ($\Delta\tau$=60), without every minute being preceded by news.

It is logical to assume that if news causes the event then the reaction should be within the same time window as that for the news, and therefore the values $\mathbf{n=\Delta t=\Delta\tau=60}$ are used for all tests in this paper. Furthermore we use **p=2** for the volatility tests, which means that we are calculating the annualised variance. Finally we set $\mathbf{\Delta T=30}$ in order to analyse the intraday effect of news, which means that we ignore the first 60 minutes of the trading day. Further tests using different variations may yield more intriguing results but the aim of this paper is to identify the types of events which appear to be linked to news, rather than finding the ideal parameters.

The results in Table 2 show the average number of minutes between events during the first period (60 minutes), and the rest of the day. Return and volatility events are rarer during the rest of the day than during the first period, with return events a lot rarer. However volume and tick events tend to be rarer in the first period than during the rest of the day in the UK and Australia. This is probably due to the number of minutes without trades in these countries, which limits the average volume and number of ticks. Alternatively it could indicate that these markets exhibit a fairly steady rate of trade in the opening hour of business.

The results of the RERNE and LERN tests are shown in Table 3 where bolded LERN values highlight results where the value exceeds 100%. As the return threshold is increased to 5% there is an increase in both RERNE and LERN values for all countries. Only the US fails to show a further increase by the 10% threshold. Whilst the 10% return results have higher values for the UK and Australian markets than the 5% threshold it should be

Event Type	Threshold	First Period			Rest of Day		
		US	UK	AU	US	UK	AU
Return	0.1%	1.14	1.16	1.14	1.36	1.41	1.40
	0.2%	1.32	1.35	1.25	1.88	2.02	1.81
	0.5%	2.16	2.30	1.85	5.20	6.02	4.90
	1.0%	5.27	5.93	3.63	23.34	30.96	20.47
	2.0%	23.20	28.64	13.18	187.38	296.73	176.53
	5.0%	172.61	452.16	190.78	2,986.46	6,669.99	4,177.65
	10.0%	644.54	3,102.39	1,674.15	30,892.04	57,065.47	37,419.38
Volume	10%	1.1	1.1	1.1	1.2	1.1	1.1
	20%	1.2	1.2	1.2	1.4	1.2	1.2
	50%	1.7	1.7	1.7	2.6	1.6	1.4
	100%	3.6	2.7	2.7	7.0	2.5	2.0
	200%	13.9	5.3	6.0	35.7	4.8	3.7
	500%	80.8	17.4	24.0	269.9	15.0	10.9
	1000%	142.8	44.9	74.7	561.4	35.7	27.5
Ticks	10%	1.2	1.2	1.1	1.3	1.2	1.1
	20%	1.6	1.4	1.3	1.8	1.4	1.3
	50%	3.6	2.4	2.0	5.2	2.5	1.9
	100%	15.2	4.8	4.8	30.8	5.8	4.1
	200%	63.0	13.8	22.1	191.4	21.0	15.4
	500%	147.8	97.6	280.5	433.9	144.0	117.7
	1000%	167.5	533.9	1,417.3	513.2	572.1	360.7
Volatility	1%	1.0	1.0	1.0	1.0	1.0	1.0
	2%	1.0	1.0	1.0	1.0	1.0	1.0
	5%	1.1	1.0	1.0	1.3	1.2	1.2
	10%	1.8	1.4	1.4	3.0	2.5	2.4
	20%	5.8	4.2	3.7	12.9	11.1	9.9
	50%	52.4	53.4	40.5	149.9	191.4	148.8
	100%	264.7	501.3	401.6	936.3	1,832.4	1,516.9

Table 2. This table shows the average number of minutes between events of each type during the first 60 minute period, and the rest of the day for the given thresholds.

noted that these are over 8.5 times rarer than 5% return events when excluding the first 60 minutes, and therefore the 5% return values are the most interesting. These results signify a strong correlation between the advent of news and subsequent return events.

There appears to be a steady decrease in correlation between news and events as the threshold for volume events is increased. Only the 10-100% threshold range for LERN tests for the US show that there is an increased likelihood of news prior to a volume event. None of these has a value sufficiently high to suggest that there might be some correlation between the arrival of news followed by a volume event. Therefore there appears to be little correlation between news and volume events.

The UK market reveals that there is a stable decline in correlation between news and events as the threshold for tick events is increased. The US market shows a slight rise to the 100% threshold and then a stable decline afterwards. The 10%-200% threshold range for LERN tests for the US, and the 10%-20% threshold range, and 500%-1000% threshold range for LERN tests for the Australian markets imply that there is some correlation between news and tick events. However only the 500%-1000% threshold tests for the Australian market have values high enough to denote that could be a link between news and tick events. Therefore, whilst tick events appear to be a better indicator than volume events they don't appear to be too reliable.

All bar the 1-10% threshold range for LERN results for the Australian market and the 2%-5% threshold range for the UK suggest that volatility events are linked to news. There is a steady increase in correlation for all countries as the threshold is increased to 100%. Therefore it appears that there is a strong correlation between the arrival of news and later volatility events, with the 50% threshold providing the strongest evidence.

The results in Table 4 show the p-values for the ETT and NTT tests, where ΔT is set to 30 minutes, and therefore

Event Type	Threshold	RERNE			LERN		
		US	UK	AU	US	UK	AU
Return	0.1%	29.19%	17.52%	7.01%	**101.06%**	**100.25%**	**101.84%**
	0.2%	29.44%	17.80%	6.84%	**101.93%**	**101.87%**	99.35%
	0.5%	30.64%	18.83%	7.49%	**106.08%**	**107.77%**	**108.74%**
	1.0%	34.42%	22.12%	9.80%	**119.15%**	**126.57%**	**142.24%**
	2.0%	45.80%	30.90%	19.37%	**158.54%**	**176.86%**	**281.25%**
	5.0%	66.35%	55.29%	67.82%	**229.66%**	**316.41%**	**984.76%**
	10.0%	60.64%	79.76%	73.44%	**209.91%**	**456.48%**	**1066.41%**
Volume	10%	29.04%	17.37%	6.86%	**100.52%**	99.41%	99.64%
	20%	29.15%	17.26%	6.83%	**100.92%**	98.79%	99.16%
	50%	29.35%	16.88%	6.74%	**101.61%**	96.60%	97.93%
	100%	29.03%	16.23%	6.50%	**100.50%**	92.91%	94.45%
	200%	27.85%	15.08%	5.93%	96.39%	86.32%	86.06%
	500%	25.76%	13.64%	5.04%	89.16%	78.07%	73.18%
	1000%	22.63%	13.83%	4.83%	78.33%	79.16%	70.19%
Ticks	10%	29.19%	17.30%	6.91%	**101.04%**	98.99%	**100.36%**
	20%	29.51%	17.07%	6.94%	**102.15%**	97.68%	**100.81%**
	50%	30.51%	16.28%	6.84%	**105.62%**	93.15%	99.35%
	100%	33.03%	14.78%	6.29%	**114.33%**	84.56%	91.34%
	200%	32.75%	12.05%	5.70%	**113.35%**	68.94%	82.73%
	500%	21.22%	10.82%	8.24%	73.47%	61.93%	**119.61%**
	1000%	18.65%	11.80%	13.58%	64.55%	67.55%	**197.24%**
Volatility	1%	28.90%	17.47%	6.88%	**100.05%**	**100.00%**	99.89%
	2%	28.96%	17.40%	6.84%	**100.24%**	99.59%	99.37%
	5%	29.48%	17.39%	6.51%	**102.06%**	99.54%	94.53%
	10%	31.01%	17.81%	6.87%	**107.35%**	**101.93%**	99.79%
	20%	36.36%	21.03%	10.39%	**125.87%**	**120.35%**	**150.87%**
	50%	57.31%	37.62%	26.55%	**198.37%**	**215.32%**	**385.51%**
	100%	71.13%	61.24%	47.25%	**246.23%**	**350.47%**	**686.17%**

Table 3. This table shows the results of the RERNE, and LERN tests for each type for the given thresholds. The LERN results in bold have values over 100% which indicate that news is more likely prior to the given event type and threshold, than it is normally.

the first 60 minutes of the day are ignored. The tests in which we have at least 95% confidence to reject the null hypothesis are in bold. The 5% threshold range for every country and the 2% and 10% thresholds for the UK and the Australian markets are the only ETT results where we can reject the null hypothesis that return events are not influenced by news.

The 2%-5% threshold range for every country, and the 1%-2% threshold range and the 10% threshold for the UK and Australian markets are the only NTT results where we can reject the null hypothesis that the arrival of news before return events is the same as the arrival of news normally. Combining the two indicates that the 5% threshold for every country and the 2% and 10% thresholds for the UK and Australian markets have a strong correlation between news and return events. This concurs with the results of the LERN and RERNE tests, which imply that return events are linked to the arrival of news.

No ETT test results can be used to reject the null hypothesis that volume events are not linked to the arrival of news. Furthermore only the 500-1000% threshold range for the Australian market have NTT results which can reject the null hypothesis that the occurrence of news prior to a volume event is the same as news normally. Therefore these results appear to confirm the RERNE, and LERN tests that volume events do not imply that the market has reacted to news.

Only the 1000% threshold for the Australian market provides ETT test results which can reject the null hypothesis that tick events are not related to the advent of news. Furthermore only the 200% thresholds for the US, and 1000% threshold for the Australian market for the NTT test provide enough evidence to reject the null hypothesis that tick events do not imply that the market has reacted to news. This suggests that there is weak evidence that tick events are related to the arrival of news. However low RERNE and LERN values reveal

Event Type	Threshold	ETT			NTT		
		US	UK	AU	US	UK	AU
Return	0.1%	87.29%	85.15%	88.81%	97.22%	94.47%	94.46%
	0.2%	93.88%	62.19%	66.29%	97.31%	71.68%	71.26%
	0.5%	86.30%	42.14%	86.73%	80.35%	14.90%	62.05%
	1.0%	55.00%	16.31%	38.16%	8.07%	**0.01%**	**0.00%**
	2.0%	7.34%	**2.02%**	**4.22%**	**0.00%**	**0.00%**	**0.01%**
	5.0%	**2.15%**	**0.07%**	**0.18%**	**0.00%**	**0.00%**	**0.00%**
	10.0%	25.29%	**0.17%**	**1.17%**	5.09%	**0.03%**	**0.38%**
Volume	10%	95.32%	32.23%	88.29%	99.24%	93.65%	98.02%
	20%	98.17%	32.21%	83.90%	99.47%	87.16%	94.69%
	50%	89.48%	27.67%	84.73%	91.03%	65.75%	86.28%
	100%	76.17%	24.78%	73.12%	71.78%	39.02%	55.93%
	200%	58.37%	21.54%	51.10%	50.65%	14.75%	8.25%
	500%	75.37%	20.21%	25.12%	50.69%	6.16%	**0.23%**
	1000%	77.80%	20.64%	23.18%	86.15%	10.79%	**0.28%**
Ticks	10%	61.02%	49.98%	86.53%	92.69%	87.54%	96.66%
	20%	60.42%	47.79%	84.08%	85.10%	72.26%	92.54%
	50%	38.56%	47.83%	94.86%	55.47%	34.38%	92.37%
	100%	14.21%	47.48%	57.02%	5.52%	12.18%	39.67%
	200%	40.53%	45.23%	42.76%	**1.84%**	11.68%	20.74%
	500%	58.37%	50.59%	39.09%	8.91%	97.50%	5.57%
	1000%	49.22%	57.04%	**2.54%**	34.37%	92.87%	**0.84%**
Volatility	1%	36.25%	65.22%	7.26%	99.56%	99.96%	98.78%
	2%	47.01%	**0.12%**	11.14%	98.55%	95.00%	92.48%
	5%	93.38%	99.25%	18.48%	95.70%	98.90%	38.24%
	10%	99.90%	69.93%	71.52%	99.18%	39.44%	32.14%
	20%	63.53%	34.59%	52.36%	**2.05%**	**0.02%**	**0.00%**
	50%	10.56%	10.01%	9.26%	**0.00%**	**0.00%**	**0.00%**
	100%	13.92%	**4.74%**	**4.57%**	**0.01%**	**0.00%**	**0.00%**

Table 4. This table shows the p-values of the ETT, and NTT tests for each type for the given thresholds. The results in bold indicate that there is 95% confidence that the null hypothesis can be rejected for the given test, country, and parameters.

Event Type	Threshold	US	UK	AU
Return	2%	3.58%	3.94%	6.68%
	5%	0.37%	0.35%	1.29%
	10%	0.03%	0.06%	0.24%
Volatility	100%	0.62%	0.50%	1.07%

Table 5. This table shows the percentage of articles which occur outside the first period of the day which correlate to the given event type and threshold.

that, whilst tick events are more reliable than volume events, tick events aren't strongly correlated to news.

The 2% and 100% thresholds for the UK, and the 100% threshold for the Australian market for the ETT test provide evidence to reject the null hypothesis that volatility events are not affected by news. The 20-100% threshold range for all countries for the NTT test can be used to reject the null hypothesis that the advent of news prior to a volatility event are the same as the advent of news normally. It may bear further investigation into the values n, Δt, and $\Delta\tau$ for the US, as the RERNE, LERN, and NTT tests all strongly imply that news is correlated to volatility events. It could be that a shorter period is required to obtain better ETT results. Apart from the US ETT test, it appears that there is a strong correlation between the arrival of news and subsequent volatility events. This is strongly supported by the evidence of the RERNE and LERN tests, and therefore we conclude that volatility events are linked to news.

Finally Table 5 shows that only a fraction of articles which occur during the trading day, excluding the first 60 minutes, correlate with the given return and volatility events.

5. Conclusions

We have found strong evidence to suggest that the stock market does react to real time news. Return and volatility appear to give the most compelling evidence. However only the 5% threshold for all countries, and the 2% and 10% thresholds for the UK and Australian markets have return events which are supported by all of the RERNE, LERN, ETT, and NTT tests. This implies that further research is required to determine if volatility events occur differently when the market reaction period is changed. However there appears to be some weak evidence that news affects volume and tick events. Furthermore the most significant tests, shown in Table 5 imply that only a fraction of news is responsible for the most significant market reactions. This reveals that the market interprets news differently, and considers some news more significant than others.

Further research into when events and news occur is necessary to establish if the market behaves in a uniform manner. Furthermore the content of news which the market reacts to should be investigated, as Fung, Yu and Wai (2003) and Mittermayer (2004) have studied, in order to highlight potentially significant news articles for investors.

6. References

Almeida, A., Goodhart, C. A. E. and Payne, R. (1998): The Effects of Macroeconomic News on High Frequency Exchange Rate Behavior. *Journal of Financial & Quantitative Analysis,* **33**(3):383-408.

Bollerslev, T. and Domowitz, I. (1993): Trading Patterns and Prices in the Interbank Foreign Exchange Market. *Journal of Finance,* **48**(4):1421-1443.

Bomfim, A. N. (2003): Pre-announcement Effects, News Effects, and Volatility: Monetary Policy and the Stock Market. *Journal of Banking and Finance,* **27**(1):133-151.

Brannas, K. and De Gooijer, J. G. (2004): Asymmetries in conditional mean and variance: modelling stock returns by asMA-asQGARCH. *Journal of Forecasting,* **23**(3):155-171.

Chan, W. S. (2003): Stock Price Reaction to News and No-news: Drift and Reversal After Headlines. *Journal of Financial Economics,* **70**:223-260.

Dacorogna, M. M., Gencay, R., Muller, U., Olsen, R. B. and Pictet, O. V. (2001) *An Introduction to High-Frequency Finance,* Academic Press, London.

Donders, M. W. M. and Vorst, T. C. F. (1996): The Impact of Firm Specific News on Implied Volatilities. *Journal of Banking and Finance,* **20**(9):1447-1461.

Dungey, M., Fry, R. and Martin, V. L. (2004): Currency Market Contagion in the Asia-Pacific Region. *Australian Economic Papers,* **43**(4):379-95.

Ederington, L. H. and Lee, J. H. (1993): How markets process information: News releases and volatility. *Journal of Finance,* **48**(4):1161-1191.

Ederington, L. H. and Lee, J. H. (1995): The short-run dynamics of the price adjustment to new information. *Journal of Financial & Quantitative Analysis,* **30**(1):117-134.

Ederington, L. H. and Lee, J. H. (2001): Intraday Volatility in Interest-Rate and Foreign-Exchange Markets: ARCH, Announcement, and Seasonality Effects. *Journal of Futures Markets,* **21**(6):517-552.

Ewing, B. T. (2002): Macroeconomic news and the returns of financial companies. *Managerial and Decision Economics,* **23**(8):439-446.

Fung, G. P. C., Yu, J. X. and Wai, L. (2003): Stock Prediction: Integrating Text Mining Approach using Real-Time News. *Proc. IEEE International Conference on Computational Intelligence for Financial Engineering*, Hong Kong, 395-402.

Goodhart, C. A. E. (1989): News and the foreign exchange market. *Proc. Manchester Statistical Society*, 1-79.

Goodhart, C. A. E. and Figliuoli, L. (1991): Every Minute Counts in Financial Markets. *Journal of International Money and Finance,* **10**(1):23-52.

Goodhart, C. A. E. and Figliuoli, L. (1992): The Geographical Location of the Foreign Exchange Market: A Test of an 'Island' Hypothesis. *Journal of International and Comparative Economics,* **1**:13-27.

Goodhart, C. A. E., Hall, S. G., Henry, S. G. B. and Pesaran, B. (1993): News Effects in a High-Frequency Model of the Sterling-Dollar Exchange Rate. *Journal of Applied Econometrics,* **8**:1-13.

Graham, M., Nikkinen, J. and Sahlstrom, P. (2003): Relative Importance of Scheduled Macroeconomic News for Stock Market Investors. *Journal of Economics and Finance,* **27**(2):153-165.

Han, L.-M. and Ozocak, O. (2002): Risk-return relationships in foreign-currency futures following macroeconomic announcements. *Journal of Futures Markets,* **22**(8):729-764.

Hess, D. (2004): Determinants of the Relative Price Impact of Unanticipated Information in U.S. Macroeconomic Releases. *Journal of Futures Markets,* **24**(7):609-629.

Hong, H., Lim, T. and Stein, J. C. (2000): Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies. *Journal of Finance,* **55**(1):265-95.

Kim, S.-J. (1998): Do Australian and the US Macroeconomic News Announcements Affect the USD/AUD Exchange Rate? Some Evidence from E-GARCH Estimations. *Journal of Multinational Financial Management,* **8**(2-3):233-248.

Kim, S.-J. (2003): The Spillover Effects of US and Japanese Public Information News in Advanced Asia-Pacific Stock Markets. *Pacific-Basin Finance Journal,* **11**(5):611-630.

Kim, S.-J., McKenzie, M. D. and Faff, R. W. (2004): Macroeconomic News Announcements and the Role of Expectations: Evidence for US Bond, Stock and Foreign Exchange Markets. *Journal of Multinational Financial Management,* **14**(3):217-232.

Melvin, M. and Yin, X. (2000): Public Information Arrival, Exchange Rate Volatility, and Quote Frequency. *Economic Journal,* **110**(465):644-661.

Michaely, R. and Womack, K. L. (1999): Conflict of Interest and the Credibility of Underwriter Analyst Recommendations. *Review of Financial Studies,* **12**(4):653-86.

Mitchell, M. L. and Mulherin, J. H. (1994): The Impact of Public Information on the Stock Market. *Journal of Finance,* **49**(3):923-50.

Mittermayer, M.-A. (2004): Forecasting Intraday Stock Price Trends with Text Mining Techniques. *Proc. 37th Annual Hawaii International Conference on System Sciences (HICSS'04)*, Big Island, Hawaii, 30064b.

Nikkinen, J. and Sahlstrom, P. (2004a): Impact of the Federal Open Market Committee's Meetings and Scheduled Macroeconomic News on Stock Market Uncertainty. *International Review of Financial Analysis,* **13**(1):1-12.

Nikkinen, J. and Sahlstrom, P. (2004b): Scheduled Domestic and US Macroeconomic News and Stock Valuation in Europe. *Journal of Multinational Financial Management,* **14**(3):201-215.

Nofsinger, J. R. and Prucyk, B. (2003): Option volume and volatility response to scheduled economic news releases. *Journal of Futures Markets,* **23**(4):315-345.

Peiers, B. (1997): Informed Traders, Intervention, and Price Leadership: A Deeper View of the Microstructure of the Foreign Exchange Market. *Journal of Finance,* **52**(4):1589-1614.

Roll, R. (1984): Orange Juice and Weather. *American Economic Review,* **74**(5):861-80.

Simpson, M. W. and Ramchander, S. (2004): An examination of the impact of macroeconomic news on the spot and futures treasuries markets. *Journal of Futures Markets,* **24**(5):453-478.

Sun, P. and Sutcliffe, C. (2003): Scheduled announcements and volatility patterns: The effects of monetary policy committee announcements on LIBOR and short sterling futures and options. *Journal of Futures Markets,* **23**(8):773-797.

Tse, Y. (1999): Market microstructure of FT-SE 100 index futures: An intraday empirical analysis. *Journal of Futures Markets,* **19**(1):31-58.

Womack, K. L. (1996): Do Brokerage Analysts' Recommendations Have Investment Value? *Journal of Finance,* **51**(3):137-67.

Wuthrich, B., Permunetilleke, D., Leung, S., Cho, V., Zhang, J. and Lam, W. (1998): Daily Stock Market Forecast from Textual Web Data. *Proc. IEEE International Conference on Systems, Man, and Cybernetics*, 2720-2725.