Real-time Diffusion of Information on Twitter and the Financial Markets

Financial firms and academic researchers have recently begun to study the predictive value of information gathered from social media (e.g. Bollen et al. 2011, Hristidis et al. 2012). The study and tracking of chatter on large-scale social networks such as Twitter has spawned new industries that aim to harness valuable information about popular and consumer sentiment (Chozick and Perlroth 2013). Social networks also provide channels for widespread chatter and speculation about the financial viability and success of firms on the stock market.

Although researchers have begun to study the relationship between patterns observed on Twitter and stock market prices at daily levels of aggregation, we do not yet have much understanding of intra-day responses of the stock market in relation to the spread of news on Twitter. Such effects may be transitory and dissipate within hours or minutes, and it is precisely these effects that many algorithmic traders would like to exploit in real time. Moreover, interesting details about the rates at which signals propagate and dissipate on the Twitter social network and in the financial markets can be observed at time-resolutions of minutes; these may be missed entirely when the data is aggregated at the daily level.

We study the real-time relationship between chatter on Twitter and the trading volume and stock price of Nasdaq 100 firms during 165 days of trading in the period from May 21, 2012 to June 3, 2013. We collected more than nine thousand observations for 96 firms (we excluded the four most common house-hold names), resulting in a total of 923,925 observations comprising a panel time-series dataset. We adopted a quasi-experimental design, in which we identified observations featuring firm-specific spikes in Twitter activity, and randomly assigned each observation to a ten-minute increment matching on the firm symbol and a number of repeating time indicators. We examine the extent that unusual levels of chatter on Twitter about a firm portend an oncoming surge of trading of its stock within the hour, over and above what would normally be expected for the stock in a given day of week or time of day.

Our results suggest that, through monitoring of chatter on Twitter about firms listed on the Nasdaq 100, observing spikes of chatter affords a reliable and non-trivial amount of foresight into oncoming surges in trading volume. Foresight into changes in stock price appears to be much more limited, but our preliminary results indicate that further investigation may be worthwhile. We do not posit a causal relationship between chatter on Twitter and movements in financial markets; although recently a hoax that spread on Twitter, claiming that President Obama was injured by an explosion at the White House, reportedly caused a temporary drop of 150 points in the Dow Jones industrial average (Chozick and Perlroth 2013). However, we believe that monitoring of chatter on Twitter can be potentially useful for modest improvements in real-time predictions of oncoming surges in trading activity. Anomalies of chatter on Twitter can also reveal certain competitive dynamics within industries; for example, when product announcements of one firm impact their suppliers or rivals, as we discuss below with specific examples related to Garmin and Akamai in June 2012.

This study has two primary objectives as a research contribution. First, the study represents a microscope upon the diffusion of information in social networks that becomes observable at a resolution of minutes. Second, the approach allows a better understanding of how social networks and financial markets simultaneously respond in real-time to external events, drawing contrasts in the speed in which information is propagated in both types of spaces. Unlike traditional time-series approaches that consider spikes as anomalies in the data that need to be removed, we treat spikes as central to the analysis because they represent real reactions to news and have tangible impacts that should not be ignored. We ask the following research questions: 1) To what extent is there a predictive relationship between the spread of new information on Twitter about a firm, and the reaction in the financial markets? 2) Is there a measurable and predictable difference in the time it takes for new information to spread in Twitter and the time it takes for that information to be absorbed in the financial markets?

2. Data

We study the real-time relationship between chatter on Twitter and the trading volume and stock price of Nasdaq 100 firms during 165 days of trading in the period from May 21, 2012 to June 3, 2013.[1] Since our emphasis is on investor rather than consumer sentiment, we excluded the most common household names listed in the Nasdaq 100 from this study: Facebook, Microsoft, Intel, Google and Apple. [2] Among Nasdaq 100 firms, these particular firms have a dominant presence on Twitter in terms of their frequency of being mentioned. The comments we observed on Twitter mentioning these household names are seldom directly related to their financial performance and more often represent consumer sentiment.

During the same 165 trading days, we collected a stream of continuous Twitter feeds of messages that mention the common names of firms in the Nasdaq 100 index.[3] We obtain counts of the number of Twitter messages in which each firm is mentioned in each ten-minute period. Using an automated screen scraper, we also gathered Yahoo! Finance data at the beginning and end of each ten-minute period, in particular for the price and trading volume of each stock.

Overall, we collected more than nine thousand observations for each firm, with a total of 923,925 observations comprising a panel time-series dataset for 96 firms. Twenty of the Nasdaq 100 firms were rarely mentioned on Twitter by their common names during the study period,

---

[1] Data was not collected during weekends and holidays, when financial markets were closed. One author operated the data-collection program on a daily basis, but on some days had to forgo data collection during extended travel or due to disruptions in internet connectivity.

[2] We used the Nasdaq 100 list as of May 1, 2012. Facebook had not yet been listed.

[3] For example, "Expedia" for "Expedia, Inc.", "Lam Research" for "Lam Research Corporation"

and we excluded those firms from the final sample. We also restricted the sample to the middle hours of trading from 10 am to 2 pm eastern standard time (EST), reducing the number of observations to 390,248.

Adopting a quasi-experimental design, we identified observations featuring firm-specific spikes in Twitter activity, and assigned each observation to a randomly selected ten-minute increment matching on the firm and a number of repeating time indicators. This resulted in a final data sample of 58,509 observations in 27,409 treatment-control group pairings for 76 firms. Table 1 shows an example of a treatment-control group pairing taken directly from our final data sample. On July 13, 2012, our algorithm determines the 80th and 90th percentile levels for the firm Adobe to be 0.14 and 0.17 Twitter mentions per second (TPS), respectively. Between 2:45 PM and 2:55 PM of that day, we observe 1000 Twitter mentions corresponding to 1.67 TPS, a level that exceeds the current 90th percentile threshold for this firm. Thus, this period represents a Twitter spike event that we define as an instance of the treatment. As July 13, 2012 fell on a Friday afternoon, our algorithm selects randomly among all other observations for Adobe that occur on some other Friday afternoon between 2:30 and 3:00 pm. In this case, the selected control instance was Friday May 17, 2013 beginning at 2:56 pm; where the observed Twitter mentions per second (TPS) is below Adobe's 80th percentile TPS threshold for the training-window period that ends prior to that week.

Table 1: Example of a treatment-control group in our data sample

Symbol	Treatment-group ID	Treatment	Start	End	Twitter Mentions (twitter)	twitter/second (TPS)	TPS 80th pctile	TPS 90th pctile
ADBE	4	0	5/17/13 2:56 PM	5/17/13 3:06 PM	114	0.19	0.21	0.29
ADBE	4	1	7/13/12 2:45 PM	7/13/12 2:55 PM	1000	1.67	0.14	0.17

3. Case Studies in the Data

To provide insight into the measurable effects contained in the data, we consider some case studies that we extracted from the data early in the study period in June 2012. These examples, among many others, provide some guidance in operationalizing constructs related to Twitter activity and financial markets reactions at the appropriate levels of granularity in time.

3a. Akamai: June 5, 2012

On June 4, 2012, NetFlix announced that it would develop its own content delivery network (CDN) that it called “Open Connect” (openconnect.netflix.com). The news had a sudden impact upon the price and trading volume of stock for Akamai Technologies (AKAM). Akamai is a leading provider of content delivery network services to NetFlix (Larmer 2012). Although it will take several years before Open Connect is fully operational to handle all of NetFlix data, the prospect that NetFlix will eventually become less reliant on Akamai’s services triggered a burst of trading activity, momentarily bringing down the price of the stock by more than 10%. Within hours, the stock price recovered entirely and the volume of trading resumed to its normal daily cyclical pattern. Several features of the associated Twitter and trading spikes are noteworthy. The spikes of Twitter activity (in blue) are relatively sharp in that they occur over a compressed time frame, as seen in Figure 1. The first spike in Twitter activity precedes the change in stock price or trading volume by approximately ten minutes. A larger spike appears within the hour, after which another flurry of trading activity takes place that restores Akamai to about its original trading price prior to the NetFlix announcement. In comparison to Twitter spikes, the reaction of the financial markets is more gradual and appears to have required a larger amount of time to process and react to new information. Since the reaction of the stock market is transitory, it can be missed entirely at the aggregation level of daily financial returns.

Figure 1: Twitter mentions, stock price, and trading volume for ***Akamai***, June 3-6, 2012.

Fig. 1A Stock price vs. Twitter mentions at ten-minute intervals

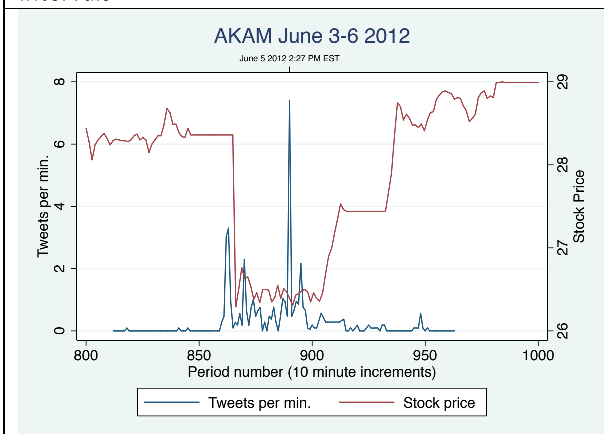

Fig. 1B Trading Volume vs. Twitter mentions at ten-minute intervals

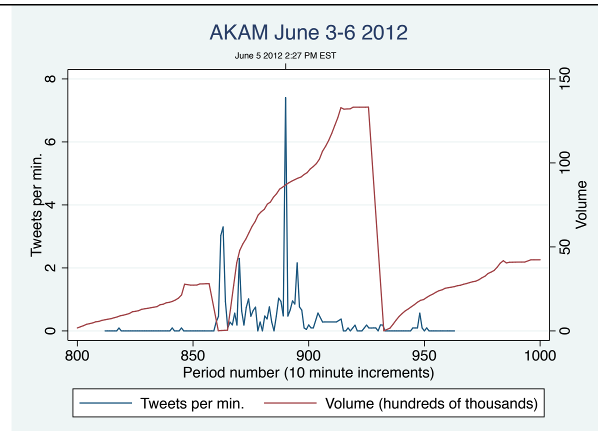

4b. Garmin: June 11, 2012

On June 11, 2012, at approximately 2:40 pm Eastern time, Apple announced the launch of new mapping software for its iO6 devices, leading to speculation that its next versions of the iPhone would be equipped with its own voice-enabled GPS service with turn-by-turn navigation (Zack's Equity Research 2012). Within minutes, this had a direct impact on Garmin's stock price and trading volume. Just as in the Akamai case, it is worthwhile to note the difference in the duration and timing between the spike in Twitter mentions and the spikes representing stock price and trading volume. The spike in Twitter activity mentioning Garmin occurs shortly after the announcement and is compressed in a relatively small duration of time. The frequency of Twitter mentions for this stock quickly returns to normal levels. The reaction in the financial markets takes approximately ten minutes, and it takes at least an hour for the reaction to this news to fully register in both trading volume and price of that stock.

Figure 2: Twitter mentions, stock price, and trading volume for ***Garmin***, June 10-12, 2012.

Fig. 2A Stock price vs. Twitter mentions at ten-minute intervals	**Fig. 2B Trading Volume** vs. Twitter mentions at ten-minute intervals
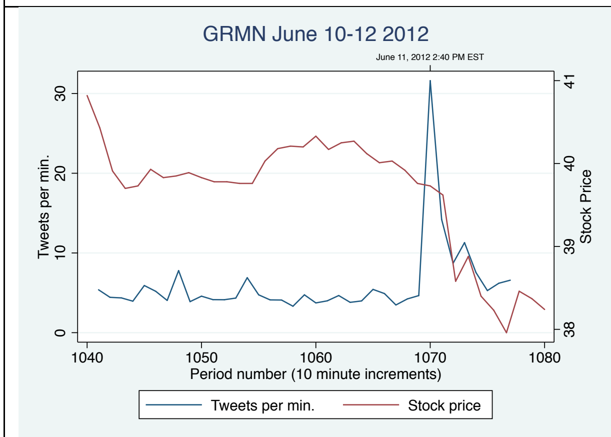	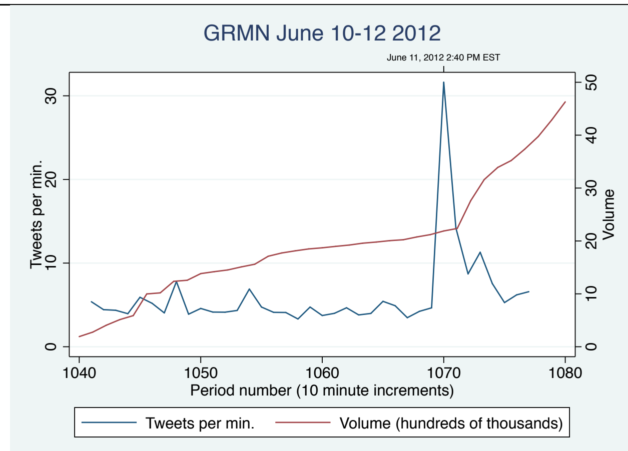

4. Empirical Model

The two case examples highlighted above are among many in the dataset that we studied to gain a better understanding of the effects that are present in the data. Exploration of the data suggests that spikes of Twitter activity are much more pronounced over smaller durations of time, whereas the financial market reactions take longer to register the effects and result in more gradual slopes.

Guided by visual exploration of the data, we define a spike in Twitter activity for a firm as the 90th percentile in the mentions of the common name of the firm per minute, which is calculated with reference to an expanding training-period window that ends prior to the week of measurement. We consider stock market reactions in the forty minutes immediately after a spike in Twitter activity. To capture trading volume spikes, we identify instances in which trading volume reaches an 80th percentile level, a threshold defined using past trading data. To capture stock price spikes, we identify the 80th percentile in the absolute value of the slope of the change in the price of the stock, based on past stock price data. We illustrate these definitions in Figure 3.

Figure 3: Measurement of spikes in Twitter activity, stock price and trading volume for a firm. Each period represents a 10-minute increment of time. The straight dashed-line represents the change in trading volume (which we denote as $\Delta TradingVolume_{40min}$) in the 40-minute period subsequent to the period in which the Twitter spike event is observed.

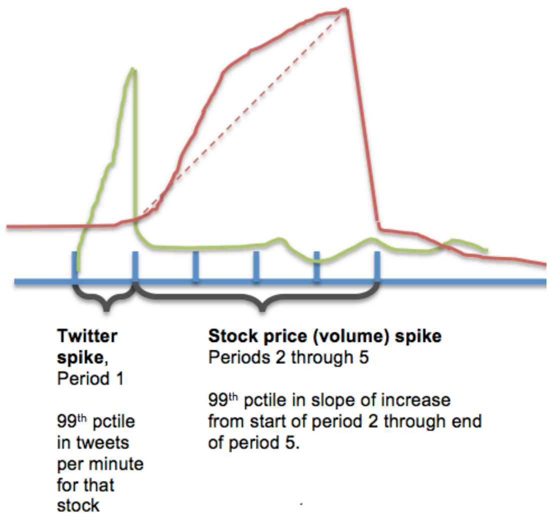

Figure 4: Training and determination periods for 99 percentile spike events

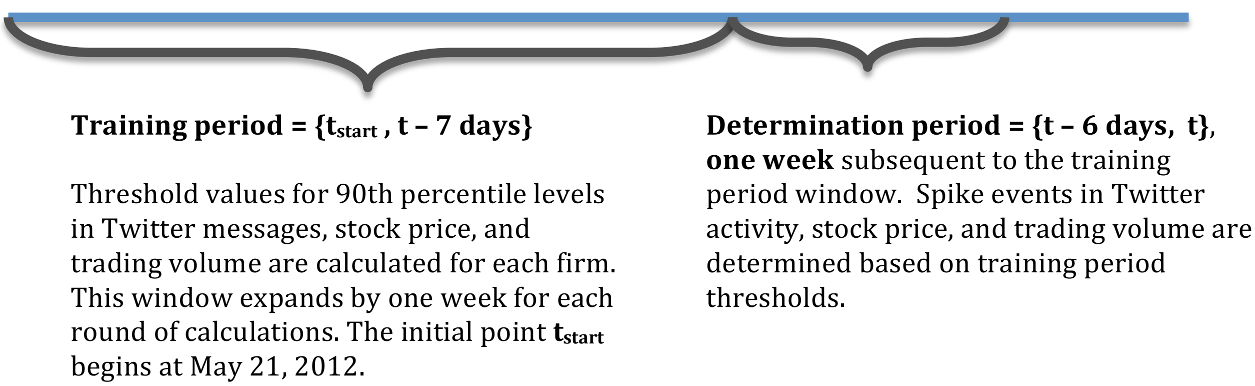

To identify spikes in Twitter activity or changes in trading volume or stock price, we used a procedure to determine the percentile thresholds beginning with an initial training period comprised of the first month of data. These thresholds are firm-specific; that is, our algorithm determines a unique threshold level for each firm that changes each week of the study period based on the history of data for that firm ending in the prior week. The training window then expands one week at a time, to determine the percentile threshold levels for each subsequent week. This method ensures that we use only past data to determine the threshold values for spikes in Twitter as well as future financial trading activity (see Figure 4).

We define changes in trading volume over the subsequent forty-minutes, which we define in four distinct units of ten-minute increments as in Fig. 1. Thus, the subsequent change in trading volume is represented by the following formula:

$$\Delta TradingVolume_{40min} = TradingVolume_{t+4} - TradingVolume_{t} \quad (1)$$

Similarly, we defined the quantities $\Delta TradingVolume_{30min}$, $\Delta TradingVolume_{20min}$, and $\Delta TradingVolume_{10min}$ as changes in trading volume over the subsequent three, two, and single 10-minute increments respectively. Likewise, we defined quantities representing the absolute value of stock price changes over the subsequent forty, thirty, twenty, and ten minutes, respectively, as $Abs(\Delta StockPrice)_{40min}$, $Abs(\Delta StockPrice)_{30min}$, $Abs(\Delta StockPrice)_{20min}$, and $Abs(\Delta StockPrice)_{10min.}$ We use the absolute value of stock price changes over the stated time periods in order to capture changes in magnitude, rather than direction, of stock price. Thus, our approach is agnostic to positive or negative sentiment, as we are interested in how the quantity of chatter pertaining to any firm will portend the volume of trading or magnitude of changes in stock price in the immediate future.

Our general estimation model is a within-estimator (i.e. fixed-effects) to implement a differences-in-differences analysis, in which groups are defined through random treatment-control assignment. To create treatment-control groups, we identified Twitter spike events as treatments and randomly assigned each such observation to a control observation in which Twitter activity was observed to be under the 80th percentile for the firm. Control observations were drawn from past data, and the method for determining the 80th percentile limits for eligible controls is based on the expanding training-period window that ends in the week prior to each observation. Random control selections were done conditional on matching of the firm symbol, as well as multiple repeating time indicators such as day of week, hour of day, and half hour. Our four-step procedure for identification of treatment-control groups is outlined in Table 2. This procedure resulted in identification of 31,126 Twitter spike events randomly assigned (with replacement) to 27,409 control-group observations. In this procedure, a Twitter spike event defines a treatment event. We use a fixed-effects panel model to implement a differences-in-differences approach; fixed-effects are incorporated at the level of the firm and all of the stated time variable indicator units.

We used a lagged-model framework, such that the levels of Twitter activity are measured for the ten-minute increments immediately preceding the periods in which the dependent variable measures begin:

$$\Delta \text{TradingVolume}_{40min} = \text{Constant} + \beta_1 \text{TwitterSpike}_{i,t-1} + \beta_2 (\text{NasdaqAvgStockPrice}_{t-1} - \text{NasdaqAvgStockPrice}_{t-2})/\text{NasdaqAvgStockPrice}_{t-2} + \beta_3 (\text{NasdaqAvgVolume}_{t-1} - \text{NasdaqAvgVolume}_{t-2})/\text{NasdaqAvgVolume}_{t-2} + \varepsilon_{i,t} \quad (2)$$

Table 2: Procedure for Identifying Twitter Spike Events for a Firm
Treatment and control groups are defined at ten-minute increments during trading hours between June 21, 2012 and June 3, 2013.

Step	Description
1) Definition of training and determination period windows	Initial training period window is May 21, 2012 through June 20, 2012. This training window is used to determine Twitter, stock price, and trading volume spikes at 80^{th}, 85th, 90th and 95th percentiles in the initial determination period from June 21, 2012 through June 27, 2012. The training window subsequently expands one week at a time, for identification of spike events in the determination period in the following week; thus the training period window for the subsequent week of June 28, 2012 through July 3, 2012 expands to May 21, 2012 through June 27, 2012; and so on. (see Figure 4)
2) Definition of Twitter spike events	If the number of Twitter mentions in any ten-minute increment exceeds the 90^{th} percentile threshold for the firm in the current week of the study period, based on the training period that ends in the previous week, it is recorded as a Twitter spike event.
3) Control group pool specification	For each firm, we identified a set of historical ten-minute increments during trading hours in which Twitter mentions did not exceed the 80^{th} percentile threshold for that firm, based on the expanding training period window ending in the prior week.
4) Random assignment of treatment to control groups	We matched all of the ten-minute increments in which a Twitter spike event occurred to a ten-minute increment from the control group pool, randomly selecting from matches on the symbol, day of week, hour of day, and half-hour indicators.

Note that the subscript i represents a treatment-control group, and t represents a ten-minute interval. The above model in equation (2) enables us to interpret the estimate of the coefficient β_1 as the impact of a Twitter spike event in any ten-minute time increment on trading volume in the periods of 40, 30, 20, and 10 minutes beginning in the subsequent ten-minute time increment. The impact is stated as a difference over what normally be expected in the matched control time periods. In other words, the model measures the extent that a Twitter spike for a firm signals an oncoming surge in trading volume, over and above the level that is normally expected for the firm in the given day of week and time of day. A similar equation is used to model the absolute value of changes in subsequent stock price in periods of 10, 20, 30 and 40 minutes. We also control for movements in the stock price and trading volume in the prior ten-minute period, averaged over all Nasdaq 100 firms included in the sample.

While we do not posit any kind of causal relationship between chatter on Twitter and movements in the stock market, our model allows us to examine whether chatter on Twitter pertaining to a firm can serve as a signal to predict unusually high levels of trading volume or changes to the stock price in subsequent minutes within the next hour. Because trading activity for a given firm can be unusually high in regularly occurring intervals, for example in a time of the day for a particular day of the week (for example, Tuesdays between 10 and 10:30 am), our model accounts for this through the treatment-control group matching.

We also code the dependent variable as a binary indicator representing unusually high increases in trading volume or changes in stock market price, at the 80^{th} percentile based on an expanding training-period window. To identify trading volume or stock price events, we used the same expanding training-period window that was used to identify spikes in Twitter activity, again to ensure that only past data was used to define the thresholds that mark these events. We used firm-specific thresholds at the 80^{th} percentile to define trading volume and stock price change events. Here again we use a within-estimation empirical framework, this time implementing a differences-in-differences model through a logistic fixed-effects panel estimator appropriate for binary dependent variables. Equation (3) expresses this model:

$$\log(\textit{odds of trading volume surge})_{i,t} = \text{Constant} + \beta_1 \text{TwitterSpike}_{i,t-1} + \beta_2 (\text{NasdaqAvgStockPrice}_{t-1} - \text{NasdaqAvgStockPrice}_{t-2})/\text{NasdaqAvgStockPrice}_{t-2} + \beta_3 (\text{NasdaqAvgVolume}_{t-1} - \text{NasdaqAvgVolume}_{t-2})/\text{NasdaqAvgVolume}_{t-2} + \varepsilon_{i,t} \quad (3)$$

Dependent variable definitions are presented in Table 3, and control variables are presented in Table 4.

Table 3: Dependent Variables

Variable Name	Variable Construction/ Definition	Data Source
ΔTradingVolume [40min, 30 min, 20min, 10min)	Change in total number of shares traded of the stock from the beginning of the subsequent ten minute time increment, measured over periods of 10, 20, 30, and 40 minutes.	Yahoo!Finance
Abs(ΔStockPrice) [40min, 30 min, 20min, 10min)	Magnitude of the change in stock price from the beginning of the subsequent ten minute time increment, measured over periods of 10, 20, 30, and 40 minutes	Yahoo!Finance
Stock Price Spike	Binary variable indicating an excess of the 80^{th} percentile in Abs(ΔStockPrice); based on expanding training window ending in the prior week.	Yahoo!Finance
Volume Spike	Binary variable indicating an excess of the 80^{th} percentile in ΔTradingVolume; based on expanding training window ending in the prior week.	Yahoo!Finance

Table 4: Control Variables: Definitions and Data Sources

Variable Name	Variable Construction/ Definition	Data Source
Nasd100 avg. trading volume change (t-1)	Prior period percentage change in average Nasdaq 100 trading volume: [(NasdaqAvgVolume$_{t-1}$- NasdaqAvgVolume$_{t-2}$)/ NasdaqAvgVolume$_{t-2}$].	Yahoo!Finance
Nasd100 avg. stock price (t-1)	Prior period percentage change in average Nasdaq 100 stock price: [(NasdaqAvgStockPrice$_{t-1}$- NasdaqAvgStockPrice$_{t-2}$)/ NasdaqAvgStockPrice$_{t-2}$].	Yahoo!Finance
Fixed-effects units in treatment-control groups	Firm, day of week, hour of day, and half-hour	

5. Main Findings

Table 5 shows the results of paired t-tests comparing the magnitude in changes in volume and stock price within each treatment-control group pairing. The differences between Twitter spike events and their matched control group counterparts are statistically significant (at $p < 0.0001$). The paired t-test results suggest that a Twitter event is associated with an increase in trading volume of about 223,000 shares, on average. This is significant in magnitude, as the median in trading volume during the sample periods is 1,547,737 shares; thus the average effect represents about 14% of the median in total shares traded. The difference between treatment and control in changes of stock price, at approximately $0.01, may appear to be quite small in magnitude; but given the high statistical significance ($p < 0.0001$), the profits can become non-trivial over large volumes of trading. For example, trades on the order of millions of dollars can potentially bring additional profits on the order of tens of thousands of dollars. However, we would urge caution on interpreting the effects on stock price changes until further analysis is done to ensure that statistical significance is not an artifact of the large sample size. We took steps to address this concern, by random selection of sub-samples of between 500 and 600 treatment-control group pairs. We conducted the same t-tests several dozen times using different sub-samples each time, and one example is shown in the bottom two rows of Table 5. We found that the t-test statistics for trading volume remained strongly significant in every sub-sample instance. The small sub-sample t-test statistics for stock price were less reliably significant. Our random sub-sample results suggest that the average treatment effects (ATE) for trading volume are robust and insensitive to sample size, unlike the ATE for stock price.

We next consider the results of the fixed-effects panel implementations of the difference-in-differences, which we present in Tables 6 and 7. Because the panel units in the fixed-effects models are the same treatment-control groups based upon indicators of firm, day of week, and half-hour of day, they are basically the same as the paired t-tests, except that they incorporate additional controls for one-period lagged movements in overall average Nasdaq trading volume and stock price. These control variables are important because overall movements in the stock market can influence the trading of any particular stock. The results in Tables 6 through 8 also show side-by-side comparisons of the different effects for 40, 30, 20 and 10 minutes. Trading volume increases are greater following a Twitter spike event than they would otherwise be: About 50,000 shares greater over the course of ten minutes, to about 208,000 shares over forty minutes, as indicated in Table 6 results. As expected, these magnitudes are in line with the results of the paired t-test. Just as we did with the paired t-tests, we conducted the same panel regressions on a number of randomly selected sub-samples of between 500 and 600 treatment-group pairs. The observed Twitter spike effects on subsequent trading volume are invariably statistically significant at $p < 0.01$ in the dozens of randomly selected sub-samples that we have tested. One example of the small sub-sample results for trading volume is reported in Table 7. We find results consistent to those reported in Table 6 among many randomly selected sub-samples, alleviating concerns that the observed effect upon trading volume may be an artifact of the central limit theorem for large samples (Lin, Lucas, and Shmueli 2013).

Fixed-effects panel results for stock price changes show coefficient magnitudes similar to the results of the paired t-test; stock prices change by about 1 cent more than they otherwise would in the forty minutes following periods of a Twitter spike event. However, we suspect at this point that the statistical significance of the observed effects on stock price may be driven by the central limit theorem for large sample sizes, and further investigation is needed.

According to the fixed-effects panel logistic regression estimation results in Table 9, periods featuring a spike in Twitter mentions of a firm have a greater odds of being followed by an upsurge in trading volume by a factor of exp(0.472), representing a 60% increase in odds. Twitter spike events are also more likely to be followed by unusually large changes in stock price by a factor of exp(0.162) in odds, or an increase of 18%. Running the same models on numerous randomly selected sub-samples, we found the statistical significance of the results for trading volume to be robust to sample size. However, statistical significance disappears for stock price in the small-sample results.

Table 5: Comparison of changes in trading volume and stock price following Twitter Spike Events and Non-Event Control periods: Paired t-tests

	Twitter Spike Event (Treatment)		No Twitter Spike (Control)		Difference	Comparison test
	Mean (m_t)	Std. Err.	Mean (m_c)	Std. Err.	m_t – m_c (std. err.)	Paired t-test of Ha: |m_t – m_c| > 0
Main Sample: 31,041 treatment-group pairs						
$\Delta TradingVolume_{40min}$	781,267	11,910	557,970	6,215	223,297 (11,111)	t = 20.1; p < 0.0001
$Abs(\Delta StockPrice)_{40min}$	0.192	0.0023	0.181	0.0023	0.011 (0.002)	t = 4.6; p < 0.0001
Randomly selected sub-sample: 586 treatment-group pairs						
$\Delta TradingVolume_{40min}$	881,052	88,864	639,915	61,523	241,138 (76,239)	t = 3.16; p < 0.001
$Abs(\Delta StockPrice)_{40min}$	0.21	0.019	0.16	0.013	0.043 (0.016)	t=2.62; p< 0.01

Table 6 Effect of Twitter Event on Subsequent Trading Volume: Differences-in-differences with treatment-control group fixed-effects. Fixed-effects panel regressions within treatment-control groups, created by matching ten-minute time increments randomly by firm, day of week, hour of day, and half-hour units. Significant at *10%, **5%, and ***1% level for two-tailed t-tests.

	(1)	(2)	(3)	(4)
Subsequent time period:	40 minutes	30 minutes	20 minutes	10 minutes
Twitter event (t-1)	207,954***	159,310***	107,368***	49,539***
	(10,846)	(8,862)	(6,067)	(2,757)
Nasd100 avg. stock price change (t-1)	1.236e+08***	1.002e+08***	6.280e+07***	3.662e+07***
	(1.671e+07)	(1.366e+07)	(9.349e+06)	(4.248e+06)
Nasd100 avg. trading volume change (t-1)	27.13	-163.0	41.00	-307.3
	(738.8)	(603.7)	(413.3)	(187.8)
Observations	58,509	58,509	58,509	58,509
Treatment-control groups	27,409	27,409	27,409	27,409
F stat	144.9***	136.2***	118.6***	135.4***
F test	0	0	0	0

Table 7 Small-sample Effect of Twitter Event on Subsequent Trading Volume: Differences-in-differences with treatment-control group fixed-effects. Fixed-effects panel regressions within treatment-control groups, created by matching ten-minute time increments randomly by firm, day of week, hour of day, and half-hour units.
Significant at *10%, **5%, and ***1% level for two-tailed t-tests.

	(1)	(2)	(3)	(4)
Subsequent time period:	40 minutes	30 minutes	20 minutes	10 minutes
Twitter event (t-1)	205,052**	177,961***	109,155***	59,251***
	(87,399)	(62,575)	(41,949)	(21,733)
Nasd100 avg. stock price change (t-1)	-2.289e+08	-1.961e+08	-5.320e+07	2.034e+07
	(1.714e+08)	(1.227e+08)	(8.226e+07)	(4.261e+07)
Nasd100 avg. trading volume change (t-1)	-12,209	-10,536	-8,000	-2,282
	(11,541)	(8,263)	(5,539)	(2,870)
Observations	1,104	1,104	1,104	1,104
Treatment-control groups	518	518	518	518
F stat	2.97**	4.33***	3.23**	2.76**
F test	0.0313	0.00493	0.0222	0.0418

Table 8 Effect of Twitter Event on Subsequent Changes to Stock Price: Differences-in-differences with treatment-control group fixed-effects. Fixed-effects panel regressions within treatment-control groups, created by matching ten-minute time increments randomly by firm, day of week, hour of day, and half-hour units.
Significant at *10%, **5%, and ***1% level for two-tailed t-tests.

	(1)	(2)	(3)	(4)
Subsequent time period:	40 minutes	30 minutes	20 minutes	10 minutes
Twitter event (t-1)	0.0126***	0.00925***	0.00923***	0.00685***
	(0.00239)	(0.00214)	(0.00179)	(0.00137)
Nasd100 avg. stock price change (t-1)	42.67***	35.86***	26.18***	23.23***
	(3.678)	(3.303)	(2.751)	(2.106)
Nasd100 avg. trading volume change (t-1)	3.13e-05	-0.000101	-4.70e-05	6.88e-05
	(0.000163)	(0.000146)	(0.000122)	(9.31e-05)
Observations	58,509	58,509	58,509	58,509
Treatment-control groups	27,409	27,409	27,409	27,409
F stat	53.77***	45.19***	38.71***	48.93***
F test	0	0	0	0

Table 9 Effect of Twitter Event on Likelihood of Subsequent Spikes in Trading Volume and Stock Price: Logistic panel fixed-effects on main sample and randomly selected subsample. Differences-in-differences through logistic panel fixed-effects model within control groups assigned randomly by the firm, day of week, hour of day, and half-hour units.
Significant at *10%, **5%, and ***1% level for two-tailed t-tests.

	Main Sample Results		Small-sample results (through random selection)	
	(1) Trading Volume Event: 80th pctile	(2) Stock Price Event: 80th pctile	(1) Trading Volume Event: 80th pctile	(2) Stock Price Event: 80th pctile
Twitter event (t-1)	0.472***	0.162***	0.687***	0.153
	(0.0222)	(0.0200)	(0.173)	(0.147)
Nasd100 avg. stock price change (t-1)	130.6***	292.2***	-347.7	534.7**
	(33.14)	(33.42)	(352.5)	(272.2)
Nasd100 avg. trading volume change (t-1)	-0.000527	-0.00505**	-0.0130	-0.0154
	(0.00119)	(0.00202)	(0.0235)	(0.0158)
Observations	19,507	23,095	346	427
Treatment-control groups	8,831	10,433	156	195
Chi-sqr stat	480.7***	158.6***	17.9***	6.5*

6. Discussion

In practice, it is extremely difficult for individual investors to capitalize on newly released public information by trading in the stock markets. The speed of information exchange, ever-faster financial trading networks, and liquidity of financial markets, all ensure that market prices almost instantly absorb news as it is released to the public, denying arbitrage opportunities to all but exclusive groups of institutional traders (Amihud and Mendelson 1988). For all practical purposes as far as common investors are concerned, the efficient market hypothesis is robust in denying arbitrage opportunities in the stock market based upon newly released public information (Fama and French 1996).

Nevertheless, news can spread on Twitter much more quickly than it can be absorbed in the financial markets. For an individual person, the act of sending or relaying a message on Twitter is often fast and effortless, without immediate financial consequences. Executing a trade on the stock market based on the same information can be a much more cognitively and emotionally taxing process that requires more time. Thus, on Twitter we often observe a rapid spread of news, speculation, rumors, opinions, and ideas about specific firms. Information often spreads on this social network before the financial markets can process it.

Our results suggest that chatter on large-scale social networks such as Twitter about a firm may serve as a signal of impending movements in the stock market; to a reliable and non-trivial extent in trading volume, and to a modest extent in the stock prices of those firms. This not only presents an opportunity to harness valuable information for participants in financial markets; but also provides greater insight into the types of information that spread on large-scale social networks such as Twitter. We observe distinct speeds at which information diffuses in Twitter in comparison to the time it takes for financial traders to process and act upon that information. We observe effects that are statistically significant, transitory, and require finite amounts of time. As such, signals propagating in Twitter may be useful to traders seeking to exploit small delays in the diffusion of news and the relatively slow responsiveness of the markets. We believe that statistical models employing real-time data from large social networks can apply not just to the financial markets, but also to other areas of electronic commerce in which consumer sentiment can have measurable effects in real-time, such as in online auctions or online merchandising. Moreover, from the perspective of industry research, it is possible to reveal and quantify more clearly the dependencies between firms in an industry ecosystem.

References

Amihud, Y., & Mendelson, H. 1988. Liquidity, volatility, and exchange. *Journal of Accounting, Auditing, and Finance,* **3**(4), 369-395.

Bollen, J., H. Mao, and X.-J. Zeng. 2011. Twitter mood predicts the stock market. *Journal of Computational Science*. **2**(1) 1-8.

Chozick, A., and Perlroth, N. 2013. Twitter speaks, markets listen, and fears rise. *The New York Times.* (accessed April 29, 2013) www.nytimes.com/2013/04/29/business/media/social-medias-effects-on-markets-concern-regulators.html

Fama, Eugene F., & French K.R. 1996. Multifactor explanations of asset pricing anomalies. *The Journal of Finance,* **51**(1), 55-84.

Hristidis, V., E.J. Ruiz, C. Castillo, A. Gionis, & A. Jaimes. 2012. Correlating financial time series with micro-blogging activity. *Web Search and Data Mining*, Seattle, WA.

Larmer , M. 2012. Akamai, Limelight slump as Netflix pushes data delivery plans. *The Wall Street Journal Online.* (accessed June 20, 2012) http://webcache.googleusercontent.com/search?q=cache:wyXiY2GkMgsJ:online.wsj.com/article/BT-CO-20120605-708991.html

Lin, Mingfeng, Henry C. Lucas, and Galit Shmueli. "Too Big to Fail: Large Samples and the p-Value Problem." *Information Systems Research* (2013). Published online before print April 12, 2013, doi: 10.1287/isre.2013.0480.

Schmerken, I. 2012. Academic's Twitter model predicts stock returns. *Advanced Trading.* (accessed March 21, 2012) http://advancedtrading.com/blogs/232602958

Zack's Equity Research. 2012. Garmin Faces Heat from Apple Maps - Analyst Blog. *Nasdaq Community Blog.* (accessed June 19, 2012) http://community.nasdaq.com/News/2012-06/garmin-faces-heat-from-apple-maps-analyst-blog.aspx?storyid=148196