Twitter Volume and First Day IPO Performance

Tammy Kwan

The Leonard N. Stern School of Business
Glucksman Institute for Research in Securities
Faculty Advisor: Jeffrey Wurgler
April 1, 2015

I. Introduction

Behavioral finance is a relatively new field that uses more realistic assumptions about human cognition to model financial decision-making. The field was largely born out of a set of anomalies that seemed contrary to the efficient market hypothesis. For example, the "January Effect" is a phenomenon in which the average monthly returns for small-cap stocks are consistently higher in January than in any other month of the year (Thaler, 1987). The "winner's curse" (Bajari, 2003) – the tendency for the winning bid to exceed the intrinsic value of the item purchased – is also at odds with the theory that investors are rational and will only pay accordingly.

Along these lines, one area of active research is on mood and financial returns. Past work has examined general mood variables such as weather, hours of sunlight, and results of major sporting events. Lucey (2005) found that negative moods lead to abnormally low returns and positive moods lead to abnormally high returns.

More recently, the rise in social media has opened new possibilities for tracking public sentiment in real time. Twitter is one platform that has been used in prior work. As of March 2015, Twitter had 288 million active monthly users, or 4% of the global population, who send 500 million 140-character-and-under Tweets per day.

As a result of this trend, there has been a surge of interest in analyzing the relationship between social media and finance and whether society's mood predicts economic indicators. Some of this research has yielded promising results. For example, Bollen et al. (2011) was able to predict whether the Dow Jones Industrial Average would close up or down with 87.6% accuracy based on the public mood on Twitter. Zhang et al. (2011) also found that Twitter use of the words "hope," "fear," and "worry" predicted the behavior of the Dow on the following day.

The focus of this paper is on predicting initial public offerings (IPOs), an interesting domain to explore in the context of behavioral finance. Companies that IPO typically have a limited history of earnings, suggesting that investors may rely more heavily on qualitative data in making their decisions. For instance, investor decisions may be especially prone to herding/cascade effects and influenced by general hype around a company's IPO. Some previous research has explored the effects of sentiment on IPO prices. For example, Cornelli, Goldreich and Ljungqvist (2006) found that over-optimism by small investors "can cause IPOs to trade at prices on the first day at 40.5% higher, on average, than they would have in the absence of sentiment demand." Loughran and McDonald (2013) found that IPOs with a high level of uncertain text in their S-1 document have higher first-day returns and absolute offer price revisions and subsequent volatility.

This paper investigates the relationship between Twitter data and a company's IPO. I test two hypotheses, one relating to how Twitter activity *reacts* to an IPO and one relating to how Twitter activity *predicts* an IPO. First, is there a higher volume of same-day Tweets for an IPO that performs well on the first day of trading? Second, does higher Twitter volume for days prior to an IPO predict higher first-day returns?

II. Data Selection

Though Twitter has a vast amount of data, getting the data and leveraging it for research poses some challenges. The public search API only returns data from the last seven days, and there is a rate limit of 180 queries for 15 minutes. This limitation makes it difficult to use the API to conduct research on Tweets that were posted over a week ago. Researchers have gotten around this limitation by collecting data on a week-to-week basis over a long period of time. The Twitter website did update their search functionality as of November 2014, so that it now indexes and returns all tweets. This is an excellent tool for one-off historical searching. However, this data is only accessible from Twitter's website, which makes it difficult for researchers to consume results at scale. Finally, historical data is available via Gnip, a data platform that was acquired by Twitter, but it is quite expensive, even for an academic license. Requests start at $1000—which would yield up to 40 consecutive days of coverage and less than 1 million Tweets retrieved. A year of data at the lowest volume of tweets (1 million) starts at $5200.

Based on these constraints, I had to think creatively on how to gather the appropriate Twitter data in order to evaluate my hypotheses. My initial strategy was to ask other academics who do research on Twitter data for their existing datasets. I first reached out to Zachary Steinert-Threlkhold, whose group has been downloading Twitter data with GPS coordinates on an hourly basis since September 1, 2013. Since only 2-3% of accounts have GPS coordinates (Leetaru, 2013), I decided to seek other datasets that are more representative of the entire Twitter user base. Bollen (2011) had a dataset of 9.85 million Tweets from 2.7 million users, but the time period was from February 28-December 19, 2008. The most recent dataset I could find was a 2009 Twitter dataset from Yang (2011). To supplement this, I also used the Twitter website to collect my own 2014 dataset.

A. 2009 Twitter Dataset

This is a collection of public tweets that were recorded from June 1, 2009 to December 31, 2009. The dataset has 476 million Tweets from 20 million users, consisting of about 20-30% of all public tweets published during that particular time frame. The dataset is courtesy of Yang et al. (2011). Each Tweet includes a timestamp, the username, and the 140-character-or-under text content.

B. 2014 Twitter Dataset

Since Twitter was still in its early form in 2009, I also analyzed Tweets from the January 1, 2014 - May 31, 2014 timeframe. I used Twitter's search functionality on their website to conduct these searches. For each of the 119 companies, I conducted two searches – one for day of Tweets and one for Tweets prior to trade date.

C. IPO Data

Data from IPOScoop.com was used to determine the set of US companies that had IPO's during the time range of the Twitter datasets (6/1/09-12/31/09) and (1/1/14-5/31/14). In total, there were a total of 53 companies that fit the 2009 criterion and 119 companies that fit the 2014 criterion. The spreadsheet also contained information on IPO offer price, opening price, first day close, and first day percent change for each company.

III. Methodology

For each company that had an IPO during the selected time frame, I compiled a count of the Tweets that were relevant to the company's IPO a week prior to the trade date. Tweets that were in the count had to match the following search criteria: first word of company OR ticker symbol and "IPO." For example, for the company Cypress Energy Partners that had its first trade date on 1/15/14, the search term was: "Cypress" or "CELP" and the word "IPO". Thus, any tweet must include the words "Cypress" and "IPO" or "CELP" and "IPO." The allowed range of dates was 1/7/14-1/14/14.

For the 2009 Twitter data, I was also able to compute a sentiment analysis on the relevant Tweets, determining the portion of the IPO's tweets were positive and negative.
Sentiment analysis was conducted by training a naïve Bayes classifier on a dataset of 100,000 Tweets that were labeled as either positive or negative (with no neutral category). The features of the classifier were whether or not a particular word was present in the Tweet. The most informative features from the training set included words like "sad," "canceled," and "congratulations" (Table 1).

IV. Data Overview

The first day returns for the 53 companies in the 2009 period and the 199 companies in the 2014 sample were on average positive for the set of companies that we analyzed (Table 2). This is consistent with the finding that firms that go public reward first-day investors with considerable underpricing (Ritter, 2002).

Twitter has experienced rapid growth as a platform. In 2009, companies that went public had a mean of 15 and median of 8 Tweets on the first trade day. In 2014, companies that went public had a mean of 181 and a median of 87 on the first trade day (Table 3). Because there was a wide distribution in Tweet size, I perform a natural log transformation of Tweet counts (Table 4). Correlations were then performed using the log Tweet counts.

To investigate whether Twitter users discuss one industry more than others, I also categorized the companies by industry. High-tech companies, on average, had higher Tweet counts and higher first day returns, for both prior and on the first trade day than pharma/biotech companies (Table 5).

Finally, sentiment analysis on the 2009 data demonstrated that the majority of Tweets regarding company IPO's are positive (Figure 1).

V. Results

A. First Day Returns versus Number of Tweets on Trade Day

I ran a correlation between first day returns and the natural log of the number of Tweets on trade day for data in 2009 (Figure A) and 2014 (Figure B). As hypothesized, higher Twitter volume on trade day is correlated with higher first day returns for both 2009 (r=0.35, p<.05) and 2014 (r=0.20, p<.05), suggesting that users of Twitter like to talk about top performing IPOs on their trade day or that, perhaps, continued communication during the course of the initial trade date may affect final closing price.

Figure A: 2009 - On Trade Date

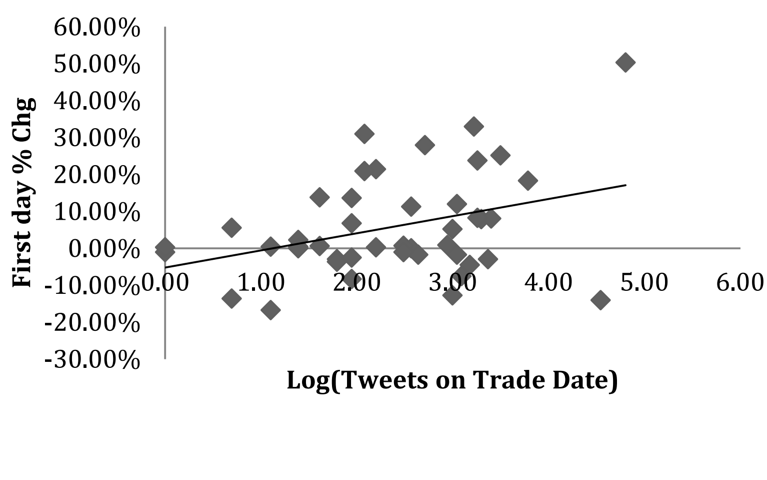

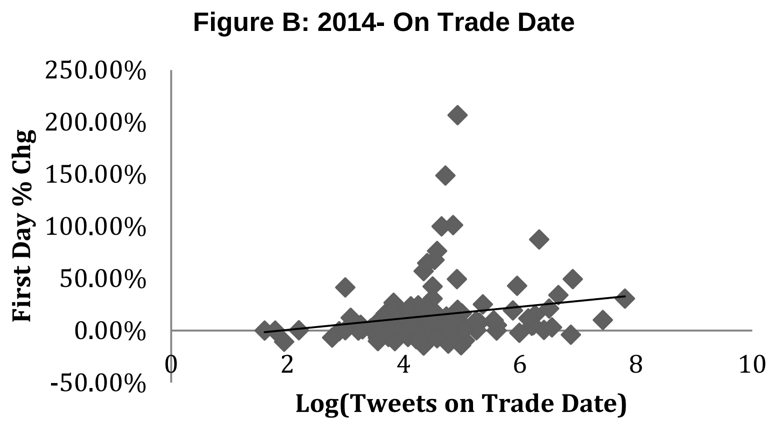

B. First Day Returns versus Number of Tweets Prior to Trade Day

An analysis of the 2009 dataset (Appendix Table 7) did not find a significant predictive relationship between Tweets and IPO performance, either for the total number of Tweets ($r=0.0028$, $p > 0.05$) as well as the number of Tweets with positive sentiment ($r=0.0105$, $p > 0.05$). See Figure 2 for the scatter plots. However, running the same analysis with the 2014 data (Figure C) yielded a weak but significant correlation between total Tweets and IPO performance ($r=0.1874$, $p < 0.05$).

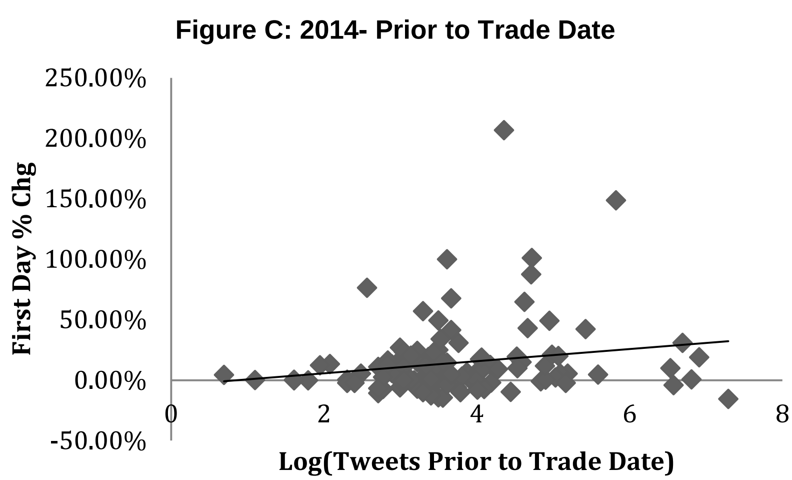

VI. Discussion

The results indicate that Twitter users tweet about IPOs that perform well on the first day. It is interesting to note that Twitter users talk more about first day winners than losers. I would have expected that IPOs that vastly underperformed might have also generated considerable chatter, but this was not the case.

The results also show that Twitter volume shortly prior to an IPO is weakly correlated with higher first day returns. It is interesting that only the 2014 dataset showed a significant result. Perhaps this was a data sparsity problem, and if I had more 2009 data points, the results would have been significant. It could also be the case that in 2009, Twitter was less influential and less representative of public sentiment around a company's IPO.

Although the results suggest that higher volume of Tweets does correlate with higher first day returns, there are other aspects of a company that one should consider as well. For example, the industry in which the IPO takes place in may play a role. King Digital Entertainment, maker of popular consumer game Candy Crush, had an impressive 1464 Tweets nine days prior to the IPO date, which was six standard deviations larger than the mean number of tweets regarding a company's IPO, and had over 3,000 tweets on the trade day. Despite all of this activity on Twitter, the King IPO performed poorly on the first day, with a first day loss of -15.56%. At the same time, Dicerna Pharmaceuticals, which had the largest first day returns at 206.67%, only had 78 tweets during the 9 days prior to the IPO and 138 tweets on trade day. Varonis, a B2B data protection company, also posted impressive 100% first day returns but only had 37 tweets 9 days prior and 105 tweets on trade day. These outliers suggest that the industry and company profile may matter- users may just like Tweeting about company IPO's for companies that they are familiar with.

VII. Future Direction

I assumed that day-of-IPO Tweets were attributed to people talking about IPOs that performed well. However, it is also possible that day-of-IPO Tweets regarding the company's IPO may also have a positive effect on closing price. Further research on breakdown of day-of-Tweets can help to determine which effects are occurring in the dataset.

Also, although the positive sentiment analysis did not yield significant results in the 2009 dataset, it would be interesting to run a similar test for the 2014 data to see if there is an effect, since the 2014 dataset is much larger. I would also like to run a similar analysis for another 6 months of data to validate that the significant correlations are replicable. Unfortunately, these additional tests are presently constrained by the availability of data.

Finally, this is a first attempt at using a sentiment classifier on IPO tweets. The sentiment classifier could be further refined to have a "neutral" category, along with a confidence interval for the sentiment categorization.

VIII. Conclusion

The analysis suggests that there is a significant relationship between Twitter data and first day IPO performance. Although this is a correlational study, pre-IPO companies that want to cover their tactical bases may consider increasing their social media presence. Investors that are providing funding pre-IPO may also want to monitor social media activity as another information source to leverage in their decision making regarding whether to invest.

Table 1: Most Informative Sentiment Words

Most Informative Features	Pos:Neg
vip	31.3:1.0
screwed	1.0:26.7
fml	1.0:26.1
sad	1.0:25.5
cramps	1.0:20.8
crying	1.0:19.7
dreading	1.0:18.1
canceled	1.0:18.1
dammit	1.0:17.5
congratulations	17.2:1.0

Table 2: Summary Statistics for IPO Companies

	First Day % Changes	
	2009 (n=53)	**2014 (n=119)**
Average	5.23%	14.33%
Median	0.43%	5.36%
Standard Dev.	13.95%	30.72%
Min	-16.75%	-15.56%
Max	50.30%	206.67%

Table 3: Summary Twitter Statistics for IPO Companies

Number Tweets for Company IPO				
	Prior to Trade Date		**On Trade Date**	
	2009 (n=53)	**2014 (n=119)**	**2009 (n=53)**	**2014 (n=117**)**
Average	8.32	96.86	14.91	181.42
Median	4.00	34.00	8.00	87.00
Standard Dev.	13.55	208.74	21.60	317.76
Min	0.00	2.00	0.00	5.00
Max	71.00	1464.00	122.00	2461.00

Note: For 2009, prior date count goes back 7 days. For 2014, prior date count goes back to 9 days.

Table 4: Summary Statistics for Transformed Tweet Counts

	LOG(Number Tweets)			
	Prior to Trade Date*		Trade Date	
	2009 (n=42)	2014 (n=119)	2009 (n=44)	2014 (n=117**)
Average	1.68	3.72	2.36	4.52
Median	1.61	3.53	2.48	4.47
Standard Dev.	1.15	1.14	1.06	1.09
Min	0	0.69	0	1.61
Max	4.26	7.29	4.80	7.81

Note: In the 2009 dataset, I excluded the companies where the number of Tweets was 0 (11 companies prior to trade date and 9 companies on trade date). For the 2014 data, I had to exclude two companies on trade date (King and Weibo) because the max count of Tweets was too large to count (>3,000).

Table 5: Breakdown of IPO and Twitter Data by Industry

High-Tech refers to any firm that produces technology that incorporates advanced computer electronics. In many cases, high-tech companies are consumer facing. Pharma refers to any company in the pharmaceutical or biotechnology industry that develops, produces and markets drugs.

Table 5a: 2009

			Number Tweets for Company IPO	
			7 days prior	Trade Date
Industry	**Number IPOs**	**First Day % Chg (Avg)**	**Mean**	**Mean**
High-Tech	8	5.66%	11.88	21.00
Pharma	5	-1.20%	16	8.00
Other	40	5.95%	6.65	14.55
Grand Total	**53**	**5.23%**	**8.32**	**14.91**

Table 5b: 2009 Transformed Data

		Number Tweets for Company IPO	
		7 days prior (N=42)	Trade Day (N=44)
Industry	**Number IPOs**	**LN**	**LN**
High-Tech	7	2.04	2.48
Pharma	4	1.86	1.84
Other	31 (33)	1.57	2.41
Grand Total	**42 (44)**	**8.32**	**2.36**

Table 5c: 2014

			Number Tweets for Company IPO			
			9 days prior**		Trade Date	
Industry	**Number IPOs**	**First Day % Chg (Avg)**	**Mean**	**LN(Mean)**	**Mean**	**LN(Mean)**
High-Tech	21	24.42%	272.95	4.87	519.47	5.71
Pharma	41	14.22%	38.83	3.44	74.00	4.19
Other	57	10.68%	73.72	3.49	146.00	4.36
Grand Total	**119**	**14.33%**	**96.86**	**3.72**	**181.42**	**4.52**

Figure 1: Sentiment Analysis

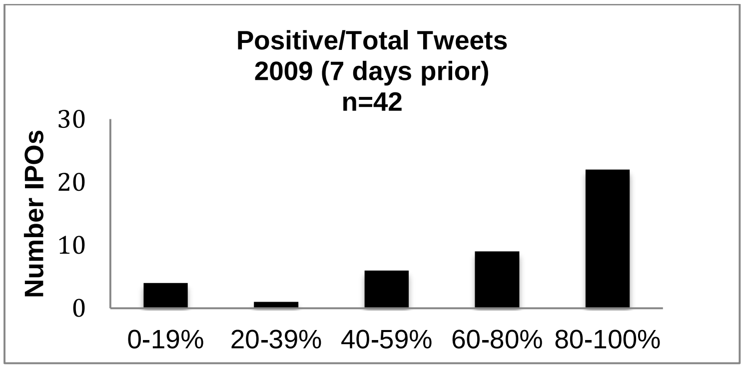

Figure 2: 2009 Correlations (p>.05)

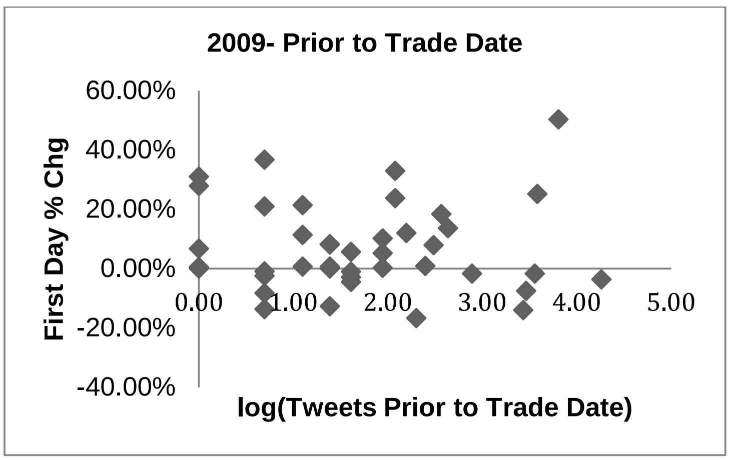

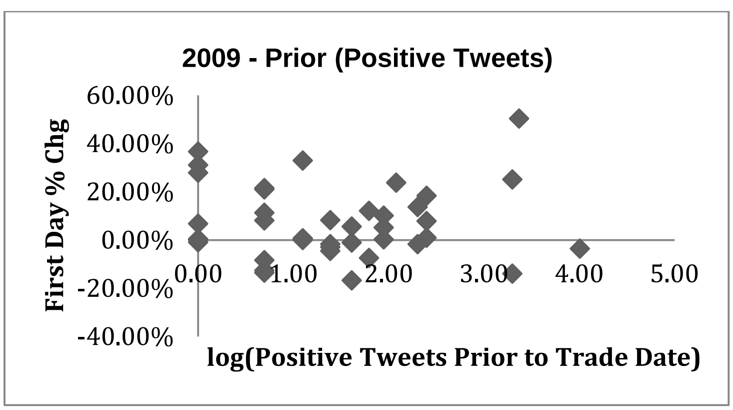

References

Bajari, Patrick, and Ali Hortacsu. "The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions." *RAND Journal of Economics* (2003): 329-355.

Bollen, Johan, Huina Mao, and Xiaojun Zeng. "Twitter mood predicts the stock market." *Journal of Computational Science* 2.1 (2011): 1-8.

Cornelli, Francesca, David Goldreich, and Alexander Ljungqvist. "Investor sentiment and pre-IPO markets." *The Journal of Finance* 61.3 (2006): 1187-1216.

Haugen, Robert A., and Josef Lakonishok. *The incredible January effect: The stock market's unsolved mystery*. Business One Irwin, 1988.

Leetaru, Kalev, et al. "Mapping the global Twitter heartbeat: The geography of Twitter." *First Monday* 18.5 (2013).

Loughran, Tim, and Bill McDonald. "IPO first-day returns, offer price revisions, volatility, and form S-1 language." *Journal of Financial Economics* 109.2 (2013): 307-326.

Lucey, Brian M., and Michael Dowling. "The Role of Feelings in Investor Decision-Making." *Journal of Economic Surveys* 19.2 (2005): 211-237.

Ritter, Jay R., and Ivo Welch. "A review of IPO activity, pricing, and allocations." *The Journal of Finance* 57.4 (2002): 1795-1828.

Thaler, Richard H. "Anomalies: The January Effect." *The Journal of Economic Perspectives* 1.1 (1987): 197-201.

Yang, Jaewon, and Jure Leskovec. "Patterns of temporal variation in online media." *Proceedings of the fourth ACM International Conference on Web Search and Data Mining*. ACM, 2011.

Zhang, Xue, Hauke Fuehres, and Peter A. Gloor. "Predicting stock market indicators through twitter “I hope it is not as bad as I fear”." *Procedia-Social and Behavioral Sciences* 26 (2011): 55-62.