Expert Systems With Applications 77 (2017) 256–263

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

The hasty wisdom of the mob: How market sentiment predicts stock market behavior

M.S. Checkley[a,*], D. Añón Higón[b], H. Alles[c]

[a] The Business School, University of Gloucestershire, Park, Cheltenham, Gloucestershire, GL502RH. UK
[b] ERICES and Departamento de Estructura Económica. Universitat de València, Avda. de los Naranjos s/n, 46022 Valencia, Spain
[c] CEO Infolution. Schouwschuit 37 1613 CJ Grootebroek, The Netherlands

ARTICLE INFO

Article history:
Received 24 August 2016
Revised 10 October 2016
Accepted 24 January 2017
Available online 3 February 2017

Keywords:
Sentiment
Stock market
Social media
Forecasting
Micro-blogging
Analytics

ABSTRACT

We explore the ability of sentiment metrics, extracted from micro-blogging sites, to predict stock markets. We also address sentiments' predictive time-horizons. The data concern bloggers' feelings about five major stocks. Taking independent bullish and bearish sentiment metrics, granular to two minute intervals, we model their ability to forecast stock price direction, volatility, and traded volume. We find evidence of a causal link from sentiments to stock price returns, volatility and volume. The predictive time-horizon is minutes, rather than hours or days. We argue that diverse and high volume sentiment is more predictive of price volatility and traded volume than near-consensus is predictive of price direction. Causality is ephemeral. In this sense, the crowd is more a hasty mob than a source of wisdom.

1. Introduction

Sentiment metrics derived from social media are claimed to help predict financial market behavior (c.f. Bollen, Mao, & Zeng, 2011, Mittermayer & Knolmayer, 2006). By matching the wording employed, in a micro-blog, with a dictionary of mood-related words, millions of financial market-related messages can be sampled for the emotions shown by their authors. Mood can be extracted real-time from bloggers and, in a fashion, measured. These metrics are aggregated over assets and time to provide a dynamic sense of collective market mood. The resulting time series of sentiment metrics can be tested for its ability to predict market prices and returns, price volatility, and trading volume. In other words, sentiment is argued to predict market behavior.

While the research literature on this topic burgeons, there is divergence in scholars' views about the underlying predictive ability of social media sentiment metrics. Research findings range from sentiment analysis being consistently prophetic, to having modest or selective value, to having no forecasting value whatsoever, particularly for price direction (c.f. Nassirtoussi, Aghabozorgi, Ying Wah, & Ling Ngo, 2014, for a review of the research literature).

To help resolve the growing debate over the supposed value of financial market forecasting with sentiment metrics, we explore intraday sentiment and price data for each of five high-profile US stocks. Focusing on more granular data than prior authors, we employ intraday market data; sampled-every-two-minutes data on price direction, price volatility and trading volume. We match the market data to equally-granular sentiment metrics. These data facilitate the study of shorter time-horizons – relative to prior research - in the predictive ability of sentiments for prices. We find evidence of sentiments causing market behavior, albeit with selective ability to reduce forecast errors. Uniquely, we find strongest results over time horizons of minutes, rather than hours or days. In sum, sentiment has some short-term predictive value, but more-so when emotions are divergent across the market. Forecasts are better in the cases of predicting volatility and trading volume, than price direction. Bloggers collectively evince uncertainty more than clarity.

This study contributes to literature by adding to arguments in favour of emotion being salient to investment decision making, by finding unique evidence on the rate at which most valuable information diffuses between market participants, and by specifying related classes of trading strategy most apt to building on such a predictive foundation. The findings are further embedded in the work of Surowiecki (2004); we argue that the strongest causal links between sentiment and market behaviour are found in times of strident and discordant sentiment. Traders collectively form a mob

* Corresponding author.
E-mail addresses: mcheckley@glos.ac.uk (M.S. Checkley), m.dolores.anon@uv.es (D.A. Higón), ha@infolution.com (H. Alles).

http://dx.doi.org/10.1016/j.eswa.2017.01.029

more than a wise crowd, but elements of shared sentiment can be weakly-linked to predictable price direction.

This paper is organized as follows: the next section contains a literature review, followed by discussion of the study's data and methods. Then we address the results of the statistical tests. Finally, conclusions are drawn, along with discussion of the paper's limitations and opportunities for additional research.

2. Literature review

This section selectively addresses prior literature on two inter-related themes. First, how theory frameworks – Efficient Markets, (Fama, 1965), Behavioural Finance (c.f. Lowenstein & Lerner, 2003), Information Diffusion (c.f. Hong & Stein, 1999), and The Wisdom of the Crowds (Surowiecki, 2004) – attempt to account for the predictability of markets. Second, we discuss major findings in the literature specific to social media sentiment-based prediction of financial markets.

2.1. Theory of market prediction with sentiment metrics

To what extent stock market prices are predictable is a long-standing and high profile debate in the finance and economics literature. The Efficient Market Hypothesis (Fama, 1965) states that prices fluctuate randomly and hence the very act of forecasting is based on misapprehension. Put crudely, long-term supernormal trading profits (adjusted for risk) are implausible in efficient financial markets.

However, responding to critiques of the Efficient Market Hypothesis (c.f. Cutler, Poterba, & Summers, 1989), recent studies in Behavioural Finance suggest that emotion has a substantial role in investment decision making (Lowenstein & Lerner, 2003). With the speedy expansion of the Web and social media, the influence of investors' and web users' emotions has become increasingly noteworthy. Traditional news media have evolved into diverse forms of blogging and micro-blogging (e.g. Twitter and StockTwits), chat rooms and discussion boards. Information can be authored with ease, searched, an opinion formed and thereafter rapidly and widely shared (Oh, Agrawal, & Raghav Rao, 2013). Emotion about financial markets is somewhat contagious and can be diffused and pooled (Oliveira, Cortez, & Areal, 2013). Hence, most of the studies reviewed in this section share a supposition; en masse, social media displays measurable emotion and the derived sentiment metrics can plausibly predict buying and selling behaviour. This, in turn, affects market prices.

Consistent with the Gradual-Information-Diffusion model (c.f. Hong & Stein, 1999, Hong, Lim, & Stein, 2000, Matsubara, Yasushi, Prakash, Li, & Faloutsos, 2012), investors track market-salient news and opinion-flow, and gradually and incompletely incorporate such information into their trading decisions. The rate of information diffusion to and between market participants influences how market prices shift over time. Traded assets avidly-covered by mass and social media would be expected to respond to news flow faster than low-profile assets, for which mass media reportage would be sparse or non-existent (Matsubara et al., 2012).

A central issue, related to the diffusion of information, is whether such diverse and decentralized opinions on social media, when aggregated, reflect "the wisdom of crowds" (Surowiecki, 2004). Do the opinions of bloggers collectively form valuable predictions about markets? Perhaps financial markets are susceptible to "group think", wherein a trader's opinions are not wholly decentralized and independent from those of other traders. If so, this would violate the conditions necessary for aggregated opinion to outperform experts (Surowiecki, 2004). It has been argued, for example, that rather than being diverse and decentralized, opinions are diffused along lines strongly mediated by the social networks of influential bloggers (Armentano, Godoy, & Amandi, 2013).

The study of information diffusion within markets leads to the question of the rate at which information spreads. This, in turn, might affect the time horizon for a forecast based on such information. Prior literature, whether or not it pinpoints a price-causal role for social media sentiments, has tended to focus on forecast horizons of a day or more, with a few engaging with data granular to the hour (c.f. Nassirtoussi et al., 2014). Yet we know that market participants typically respond to salient market data within one hour (Chordia, Roll, & Subrahmanyam, 2005), and plausibly within seconds or thousandths of a second (c.f. Lewis, 2014). Meanwhile, social media tends to elicit responses to high-interest blogs within a timeframe of minutes, albeit with the potential for an ongoing stream of responses and re-tweets over hours or days (PsychSignal, 2014). Putting these insights together, and knowing that social media is already in trade-prompting use by a number of market participants (PsychSignal, 2014), it appears worthwhile to investigate causal relationships of shorter duration than those prevalent in the extant research literature. One might suppose that, for high-profile assets, the duration of any price-causal relationship - between social media sentiment metrics and market behaviour - would more typically last minutes, rather than hours or days.

2.2. Predicting markets with social media sentiment metrics

Having reviewed some central theoretical frameworks in the prior sub-section, we now consider research specific to predicting market behaviour (particularly asset price movements) with social media sentiment metrics.

The number of Twitter users has grown to several hundreds of millions, resulting in the composition of many hundreds or thousands of tweets in a typical trading day for any major stock (PsychSignal, 2014). Many of these messages respond to unfolding market events. They can be considered a real-time and near-continual evocation of mood throughout the trading day (Oliveira et al., 2013).

The idea of extracting sentiment from microblogs, in order to forecast stock markets, inspires a burgeoning research stream. Prior research argues that sentiment metrics are materially predictive of stock price returns (c.f. Mittermayer, 2004), volatility (c.f. Antweiler & Frank, 2004), and traded volume (c.f. Oliveira et al., 2013). In contrast, several researchers find that social media sentiment metrics proffer little predictive advantage, particularly for stock returns (for a literature review, see Nassirtoussi et al., 2014).

Bollen et al. (2011), sampled sentiments in six dimensions (Calm, Alert, Sure, Vital, Kind, and Happy) and composed, employing artificial neural networks, a model of improved forecasts of the Dow Jones Industrial Average. In contrast, much recent research (including this paper) takes emotion to have two dimensions: valence (positive or negative); and arousal or "level" (low or high) (Nassirtoussi et al., 2014).

With a novel method, Makrehchi, Shah, and Liao (2013), test the predictive ability of sentiments by retrospectively assessing what sentiments plausibly could have been able to predict just before large market movements. Using daily analysis, they show that sentiments have sufficient predictive value to enable supernormal trading profits. Similarly, Ruiz, Hristidis, Castillo, Gionis, and Jaimes (2012), explore links between the volume of microblog messages and social network properties of those blogging, and market price movement and traded volume. They find stronger correlations with traded volume than price movement, but nonetheless suggest their results offer promise as the basis of a trading signal.

While this stream of research provides novel findings, there remain reasons for doubt. Many studies use only short periods of sampled data (c.f. Yu, Duan, & Cao, 2013). Moreover, most test