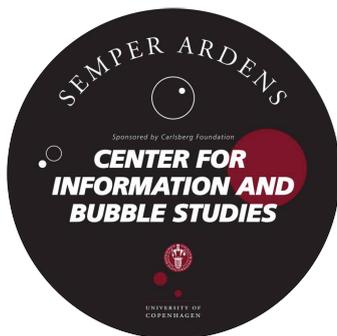


Department of Economics
FACULTY OF SOCIAL SCIENCES
UNIVERSITY OF COPENHAGEN



PhD Thesis

Laurs Randbøll Leth

Essays on High-Frequency Market Microstructure

Herding and Volume-Synchronized Probability of Informed Trading

Main advisor: Professor Peter Norman Sørensen

Co-advisor: Professor Vincent F. Hendricks

Date of Submission: May 27, 2019



This page intentionally left blank.

And, as I recall, I had to take some personal time, see my dad.

— Rust Cohle, *True Detective*

Acknowledgment

This dissertation is the product of years of struggle (sometimes referred to as doctoral studies) that mainly took place at Department of Economics of University of Copenhagen and Center for Information and Bubble Studies from January 2016 to May 2019. My journey to become Philosophiae Doctor would not have been possible without the support and encouragement from family, friends, and fellows. It is a great pleasure to have the opportunity to acknowledge them.

First and foremost, I would like to express my sincere gratitude to my advisors, Peter and Vincent. I am privileged to be motivated and supervised by two not only excellent researchers, but also kind and warm individuals. Their guidance and support have been significant for writing this thesis.

I owe Professor Antonio Guarino a special thanks for arranging my research stay at University College London in Spring 2017. During my visit, he spent hours meeting with me and discuss new ideas and results, which inspired and constituted parts of the first chapter in this dissertation.

I am also indebted to Professor David Easley and his staff for inviting me to Cornell University in Spring 2018 where I was enrolled in the Economics PhD program. I consider my stay in Ithaca, NY, as one of the best periods of my life. It was a unique opportunity to experience new cultures and at the same time studying market microstructure from recognized experts in the field. Insightful comments and invaluable guidance from Professor Maureen O'Hara and Professor Gideon Saar have contributed to writing the second chapter of this dissertation, and for that I am grateful.

Several other people have contributed to my doctoral studies. I have been fortunate to be one of the *guys* from the Center for Information and Bubble Studies, where I learned the importance of conducting research for the benefit of society and also made friends for life. The Center will always have a special place in my heart. I also wish to thank Patrick Thöni with whom I shared office at the Department of Economics. Patrick and I had many valuable conversations related to my research, and he made my life as PhD student even better.

Fortunately, academia does not exclusively reflect my person, and I am blessed with many friends who support and encourage me. To my close friends, Simon and Rasmus, I am truly grateful for having both of you in my life.

My doctoral studies have only been possible with funding from the Carlsberg Foundation. For that, I would like to express my special gratitude. Also, I deeply appreciate the financial support during my research stay at Cornell University from Danmark-Amerika Fondet, Augustinus Fonden, Louis-Hansen Fonden and Christian og Ottilia Brorsons Rejselegater.

It is impossible not to mention my wonderful girlfriend Majken and my beloved mom Birgitte. I will always rely on both of them, and their love and presence have been crucial through this entire process.

All those words become white noise when I think of my dad who taught me everything about life. He was a man of extraordinary intellect and talent, but I will remember him as my loving and caring father. He became sick in the first month of my doctoral studies and sadly passed few weeks before I was supposed to submit this thesis. Reflecting on my time as PhD fellow, it is impossible to abstract from father's disease and how it affected our family. I now regret not having completed my studies before he left this world. The least I can do is dedicating my dissertation to him.

Laurs Randbøll Leth
Copenhagen, May 2019

CONTENTS

Introduction	1
.....	
Chapter 1	
Rational Herding during a Stock Crash	7
Chapter 2	
Delta Hedging and the VPIN	69
Chapter 3	
Maximum Likelihood Estimation of VPIN: Toxic Order Flow and Warning Signals	133

This page intentionally left blank.

Introduction

This thesis consists of three single-authored chapters on high-frequency market microstructure. All of them are self-contained, and can be read independently. Chapter 1 examines the extent of herd behavior in a financial market from 2005 to 2008; in particular, whether herd selling increased during 2007 and 2008, consistent with a stock market crash. Chapter 2 investigates how portfolio hedging can take advantage of a microstructure measure (VPIN) of toxic order flow. Finally, Chapter 3 presents a study of parametric estimation of this measure. In the following, concepts and state of the art of market microstructure are reviewed, and the three chapters are introduced.

Market microstructure is a branch of financial economics, that studies trading mechanisms used for securities. Initial theoretical models of market microstructure date back to the mid-seventies, and as O'Hara (1995) points out, their relevance became clear after the crash of 1987, which revealed the true fragility of financial markets. Under explicit trading rules, the discipline seeks to explain the formation of prices. The standard econometrician settles for the supply-demand paradigm, whereas models of market microstructure provide an explanation for this emerged price.

Liquidity is a crucial concept with different interpretations. The liquidity of a security market reflects the economic concept of elasticity, whereas supply and demand for liquidity denote the passive and active side, respectively, of the market. A strict definition of a market's liquidity is essentially obtained with a transformation of the observed bid-ask spread. Importantly, models of market microstructure should feature a price-setting rule used by the passive side of the market, resulting in a positive spread. Additionally, the size of this spread should depend on the actions taken from the active side of the market.

The initial framework of theoretical analyses of security market microstructure began with the class of inventory models. Common to these models is the stochastic order flow, which may cause execution problems for traders and/or inventory risk for the market maker. For instance, Garmann (1976) considers a monopolistic market maker who quotes prices and clear trades. The focal point is to avoid failure, defined by running out

of inventory or cash. In this setting, the specialist is exposed to the risk stemming from random arrival of buy and sell orders¹. Solving the model leads to a positive bid-ask spread — the market maker sets the spread to protect herself against failure. Several inventory models were subsequently proposed but limitations of these eventually led to the rise of a new class of models.

Information-based trading models constitute the main pillar of the theoretical framework in this dissertation. The model by Glosten and Milgrom (1985) is considered as one of the two cornerstones of information-based models². Trading is seen as a game between a market maker and traders chosen to trade in a random sequence. Some traders only buy or sell the asset for exogenous reasons (e.g., liquidity or hedging), whereas the remaining group of traders are endowed with private information of the risky asset and want to maximize expected profit. Consequently, the learning market maker is exposed to adverse selection and sets a positive bid-ask spread for protection.

The first chapter of this thesis analyses herding in financial markets, which is related to mispricing and often blamed by media of instigating large stock market trends. Formally, the phenomenon occurs when privately informed agents trade against own information. Avery and Zemsky (1998) contribute to this area of the literature by presenting conditions in a quote-driven market under which herding is possible. In the spirit of Glosten and Milgrom (1985), they consider a model with event uncertainty and asymmetric information between traders and the price-setting market maker. These two ingredients form the basis of a herding-hospitable environment.

Easley et al. (1996, 1997) provide the framework for fitting financial transaction data to a fully parametric information-based model. Privately informed agents are only present in the market on days where the fundamental value of the risky asset moves. They receive either good or bad news about this change-in-value, which uniquely determines the direction of their market orders. The model fit enables inference of the level of informed and noise trading as well as the famous PIN measure, which has substantial importance in the market microstructure literature.

¹Arrival of these orders are modeled by independent Poisson processes.

²The other is the model of Kyle (1985).

Inspired by the work of Avery and Zemsky (1998) and Easley et al. (1997), Cipriani and Guarino (2014) developed a structural model of herding which may be fitted to financial transaction data. To illustrate their methodological innovation, they conduct an empirical study on a NYSE traded stock during 1995.

The first chapter of this dissertation contributes to the literature by extending the model by Cipriani and Guarino (2014), allowing for the distribution of information in good and bad news to be non-symmetric. This generalization makes it possible to separately identify the extent of herding on the buy and sell sides of the market. The main research question of this chapter is whether herd behavior of informed traders was more prominent during periods associated with extreme market turbulence compared to periods of tranquility. This question is answered by an empirical investigation of a NYSE-traded stock traded from 2005-2008, including the crash of 2008. Compared to trading in 2005 and 2006, it is statistically significant that the extent of herd selling increased during this crash, suggesting that rational herd behavior provides a part of the explanation of financial crises.

The rise of high-frequency traders in recent years has challenged empirical market microstructure. As O'Hara (2015) points out, technology has transformed markets, and trading frequencies have increased dramatically. Markets are now fragmented, and standard trade classification algorithms used to infer the active side of a trade are exposed to several numerical problems, e.g., linking trades with quotes, high order cancellation rates and quote volatility. Furthermore, high trading frequencies imply that variables in likelihood functions of information-based models will be raised in the power of millions and calibration becomes demanding.

Researchers and practitioners have therefore been forced to renew and rethink models of market microstructure. Easley et al. (2012) developed the volume-synchronized probability of informed trading (VPIN) to measure toxicity of order flow in high-frequency markets. In contrast to the PIN model, estimation of this new measure is nonparametric, and the metric is updated in stochastic time measured on the volume clock. Furthermore, Easley et al. (2012) present numerical evidence in favor of the VPIN as predictor of liquidity-induced return volatility. Strikingly, they argue that the VPIN anticipated the infamous Flash Crash on May 6, 2010. This work

led to an academic controversy with Andersen and Bondarenko (2014a,b, 2015), disputing the empirical findings in the original work. Several papers have subsequently either supported or contested the usefulness of the VPIN as a warning signal of market turbulence.

Chapter 2 contributes to this ongoing debate with an application of the VPIN in a new setting. From a portfolio manager's (PM's) point of view, the problem of hedging a short position in the European call is investigated. The PM rebalances her portfolio once at the end of each day using the risk-minimizing Δ -hedge in a Heston setting, and she is expected to collect the volatility risk premium during periods of tranquility. However, the hedging portfolio is subject to losses on days characterized by large intraday price movements and/or high volatility. The focal point in this chapter is whether the VPIN metric is able to signal future losses of the portfolio. The empirical study identifies red numbers in the hedging portfolio on days subsequent to the bankruptcy of Lehman Brothers in 2008. The VPIN is then computed using financial transaction data for the SPDR S&P 500 EFT (SPY), and empirical findings reveal that the metric in most cases signals losses for the hedging portfolio. Thus, the portfolio manager may use the VPIN to reduce losses, either by expanding her portfolio with VIX futures and/or adjust her daily risk exposure.

Performance evaluation of the VPIN in Chapter 2 is inspired by the computational research of Wu et al. (2013) where high VPIN readings are classified as either true positives or false positives. Assessment of the VPIN's ability to predict short-term return volatility is then determined by the corresponding false-discovery rate. One limitation in the moment estimation of VPIN is that information in volume time is neglected. A practitioner capable of including this component when computing the VPIN will presumably be rewarded with a lower false-discovery rate.

Lin and Ke (2017) provide a theoretical contribution to the literature by introducing maximum likelihood estimation of VPIN, capturing the information embedded in volume time. The parametric framework allows them to stress that VPIN and PIN are two different probability measures. Moreover, the mathematical analysis indicates that moment estimation of VPIN becomes unstable on small volume time intervals, whereas maximum likelihood estimation generates consistent estimates of the model parameter.

Eventually, Lin and Ke (2017) argue that their parametric estimation method will improve the VPIN's performance in terms of predicting return volatility, yet no studies in the literature have conducted this experiment. It is therefore pertinent to empirically apply the methodology by Lin and Ke (2017) to investigate their theoretical conjectures.

The final chapter of this dissertation contributes to the literature by conducting this study. Maximum likelihood estimation as well as the method of moment estimation are used to compute the VPIN of SPY traded from 2007-2015. Consistent with theory, empirical findings show that the VPIN does not approximately measure PIN, and that the moment estimation of VPIN becomes unstable on small volume intervals. Most importantly, the results reveal that maximum likelihood estimation of VPIN increases its predictive power of return volatility.

In conclusion, each of the three chapters in this dissertation has their own contribution to the literature of high-frequency market microstructure. Furthermore, they provide results and directions for future research. The common denominator is how information-based trading may cause large intraday price movements. The research presented in the last two chapters shares many common features. In particular, numerical findings in Chapter 3 can directly be used in the setting described in Chapter 2. A conceptual difference between Chapter 1 and chapters 2-3 is the perception of information and time in financial markets. The herding model assumes that information arrives over a fixed unit time, whereas the VPIN represents a new paradigm where information is measured on the volume clock. It is believed that extending the class of information-based models to this volume clock will improve their usefulness in markets affected by high-frequency trading.

References

- Andersen, T. G. and O. Bondarenko (2014a). Reflecting on the vpin dispute. *Journal of Financial Markets* 17, 53–64.
- Andersen, T. G. and O. Bondarenko (2014b). Vpin and the flash crash. *Journal of Financial Markets* 17, 1–46.
- Andersen, T. G. and O. Bondarenko (2015). Assessing measures of order flow toxicity and early warning signals for market turbulence*. *Review of Finance* 19(1), 1–54.
- Avery, C. and P. Zemsky (1998, September). Multidimensional Uncertainty and Herd Behavior in Financial Markets. *American Economic Review* 88(4), 724–748.
- Cipriani, M. and A. Guarino (2014, January). Estimating a structural model of herd behavior in financial markets. *American Economic Review* 104(1), 224–51.
- Easley, D., N. M. Kiefer, and M. O’Hara (1997). One day in the life of a very common stock. *Review of Financial Studies* 10(3), 805–835.
- Easley, D., N. M. Kiefer, M. O’Hara, and J. B. Paperman (1996). Liquidity, information, and infrequently traded stocks. *The Journal of Finance* 51(4), 1405–1436.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2012, 02). Flow toxicity and liquidity in a high frequency world. *Review of Financial Studies* 25(5), 1457–1493.
- Garman, M. B. (1976). Market microstructure. *Journal of Financial Economics* 3(3), 257–275.
- Glosten, L. R. and P. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14(1), 71–100.
- Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica* 53(6), 1315–35.
- Lin, H.-W. W. and W.-C. Ke (2017). An improved version of the volume-synchronized probability of informed trading. *Critical Finance Review* 6(2), 357–376.
- O’Hara, M. (1995). *Market Microstructure Theory*. Blackwell Publishers Cambridge, Mass.
- O’Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics* 116(2), 257–270.
- Wu, K., E. W. Bethel, M. Gu, D. Leinweber, and O. Ruebel (2013, 06). A big data approach to analyzing market volatility. *Algorithmic Finance* 2, 241–267.

CHAPTER 1

RATIONAL HERDING
DURING A STOCK CRASH

*I can calculate the motions of the heavenly
bodies, but not the madness of people*

— Isaac Newton

This page intentionally left blank.

Rational Herding During a Stock Crash^{*,**}

Laurs Randbøll Leth^{a,b,*}

^a*Department of Economics, University of Copenhagen*

^b*Center for Information and Bubble Studies, University of Copenhagen*

Abstract

The paper considers an asymmetric information model in market microstructure in which rational statistical herding occurs with positive probability. The model is fully parametric and may be fitted to financial tick data using maximum likelihood estimation. The empirical analysis is concentrated on a NYSE stock, Ashland Inc., traded from 2005-2008, including the Great Recession. Estimates show that herd selling behavior is more prevalent during the crisis, consistent with a stock market crash. In particular, compared to trading of 2005 and 2006, it is statistically significant that the proportion of herd sellers increased during this crash.

Keywords: rational statistical herding, market microstructure, observational learning, bounded private beliefs, adverse selection, stock market crash.

JEL: C58; D81; D82; D83; G12; G14.

***This paper has been invited to 'revise and resubmit' in Journal of Financial Markets.**

**The author thanks Peter Norman Sørensen, Antonio Guarino, Martin Gonzalez Eiras, and Paolo Galeazzi for constructive comments and valuable suggestions. The author also thanks Andreas Utheman for help with the numerical implementation. The author is grateful to an anonymous referee for providing insightful comments and directions for additional work. Finally, the author appreciates and acknowledges financial support from the Carlsberg Foundation grant to the Center for Information and Bubble Studies, University of Copenhagen.

*Corresponding author

Email address: bkt797@ku.dk (Laurs Randbøll Leth)

Conflict of Interest Disclosure Statement

I have nothing to disclose.

1. Introduction

Herding behavior in which investors tend to follow the crowd of others is often accused of instigating large stock market trends, such as asset price bubbles and financial crises. Thus, it is pertinent to analyze whether herd buying preceded the Great Recession, and whether herd selling contributed to the ensuing crash. This paper's focal point is to investigate whether *rational statistical herding* is affected by the macroeconomic environment, e.g., is herding more prominent during bubbles and crises compared to periods of tranquility?

Rational statistical herding occurs in financial markets when informed individuals trade against their private information in order to maximize expected utility. Private information is modeled as receiving a noisy signal about the unknown true value of the traded asset. The main problem of identifying and measuring herding is that private signals are unobservable, hence impossible to detect herd behavior with certainty. Following Cipriani and Guarino (2014), this problem is circumvented by presenting a probabilistic microstructure model in which the specific model structure allows for defining rational herd behavior in terms of quoted bid and ask prices. The model is fully parametric and classical maximum likelihood estimation may be carried out using tick data as input. Model estimates are used to draw inferences about latent private signals, allowing for measurements of intraday herding.

Even though this methodology is tractable, it still faces several empirical issues as new markets are transformed by technology and characterized by higher trading frequencies (O'Hara, 2015). Classification of the direction of trades (buy or sell) and estimation of model parameters using the number of trades is quite challenging, if not infeasible, in high frequency market microstructure (HFT)¹. To the best of my knowledge, the only empirical study to apply the methodology of Cipriani and Guarino (2014) is found in the very same paper based on tick data with significantly lower trading frequencies.

¹These drawbacks of the new markets led to the development of tools applicable for HFT research, such as *bulk classification* and the *volume-synchronized probability of informed trading* (VPIN) developed by Easley et al. (2013, 2012).

With this in mind, there is another question to address: Is measuring herding meaningful when facing higher trading frequencies for the stock? And if so, how do the empirical results relate to previous findings in the literature? The questions raised immediately lead to this paper's main contributions. From the empirical perspective, model estimation and analysis of herding is carried out in financial periods with significantly higher trading frequencies as well as different macroeconomic environments. The theoretical contribution is to generalize the structural herding model described by Cipriani and Guarino (2014), allowing for *asymmetries* in the distribution of informed traders' signals — this allows me to separately identify the extent of herding on the buy and sell sides of the market.

The research backbone relies on the theory of market microstructure. Over the past decades, theoretical market microstructure models have gained ground in the financial and economic literature. Some models focus on the market maker's learning problem determining ask and bid prices after observing the timing and sequence of trades. Two central cornerstones in theoretical market microstructure are the models developed by Kyle (1985) and Glosten and Milgrom (1985). The latter views trading as a game between traders and the learning market maker who is exposed to adverse selection caused by the presence of privately informed traders. The theoretical paper by Avery and Zemsky (1998) considers a model, which is a special case of the Glosten-Milgrom model. They show that herd behavior arises when there is event uncertainty as well as asymmetric information between traders and the market maker. Easley and O'Hara (1992) provide a more structural approach where trading is repeated over several days while Easley et al. (1997) demonstrate how to estimate this model using financial transaction data. The model considered in the latter article is very similar to the original PIN model proposed by Easley et al. (1996).

The concept of herding in financial markets is closely related to market microstructure and in particular to asymmetric information according to Bikhchandani and Sharma (2000)². A well-known and simple example of information-based herding in a non-financial setting is presented by

²See also some of the basic models of information-based herding and cascades by Bikhchandani et al. (1992) and Welch (1992).

Banerjee (1992), where privately informed agents — arriving in a random sequence — must choose between restaurant A and restaurant B after observing the predecessors' actions. It is shown how the first individual's dining choice, say restaurant A , may create an information cascade by Bayesian reasoning, i.e. all successive individuals mimic their peers and dine at restaurant A . In this case, herd behavior is rational and may cause an informational inefficiency since individuals trade against private information. Banerjee's scenario does not fully capture trading on financial markets where stock prices and trade imbalances move together, and prices will thus reflect the aggregated information in the direction of trades. However, if asymmetric information is allowed between the price-setting market maker and the traders, informed traders may decide to herd buy (sell) the asset, which in the short run will result in a higher (lower) and increasing (decreasing) price. Herding is associated with mis-pricing such that the traded price deviates from the value justified by fundamentals. This is significant when discussing and explaining reasons for financial bubbles and crashes characterized by the price rising above respectively falling beneath its fundamental value.

One issue of measuring herding relates to the fact that private information is modeled by a latent variable: What is observed is the aggressor's side of the trade but not the reason behind this action. The majority of the empirical herding literature is using statistical measures of trade clustering to detect herding, which is not in the spirit of Avery and Zemsky (1998), e.g., the reason for a high trade imbalance could be market stress³.

Fortunately, a new methodology to measure herding in financial markets has been developed in the recent paper by Cipriani and Guarino (2014), combining and renewing former literature on the theoretical and empirical work on herd behavior and microstructure modeling.

³Empirical evidence of *institutional* herding in financial markets where investors mimic each other has primarily been based on the LHV measure proposed by Lakonishok et al. (1992). Other studies are focusing on *market-wide* herding where individual stock returns are compared to overall market returns. Christie and Huang (1995) propose the CSSD-measure (cross-sectional standard deviation) to capture the extent of herding while Hwang and Salmon (2004) introduces the HS-measure in a similar setting. Finally, Chang et al. (2000) and Galariotis et al. (2015) look at the cross-sectional absolute deviation (CSAD) and overall market returns.

Their model is a continuation of the framework presented by Glosten and Milgrom (1985) and Avery and Zemsky (1998) where trading is repeated over days, and the daily asset value is a realization of a random variable depending on the arrival or absence of an information event prior to trading. Trades may come from traders privately informed of the event or from noise traders trading with exogenous motives (e.g., liquidity and hedging). Private information is modeled by the reception of a noisy signal of the unknown asset value. Assuming that the signal has conditional linear density functions with state-independent *precision*, allows Cipriani and Guarino (2014) to define the presence of herd buying (selling) behavior through a buying (selling) threshold depending on model parameters. In addition, the identification of the thresholds are used to derive accurate measures of herding. From the empirical perspective, the model described in Cipriani and Guarino (2014) is fitted to trade data with likelihood estimation using a technique similar to Easley et al. (1997). Estimation of model parameters is then used to estimate the buying and the selling threshold, which in turn are used to estimate intraday herding. The empirical study by Cipriani and Guarino (2014) focuses on the stock Ashland Inc. traded on NYSE in 1995, and their main findings are that on average herding occurs and causes informational inefficiencies in the market.

This paper contributes with an extension of the model by Cipriani and Guarino (2014), motivated by the presumption that a buy and a sell order convey different amounts of information. In contrast with Cipriani and Guarino (2014), it is not assumed that the distribution of the private signal is symmetric. Instead, it is only required that signals satisfy the reasonable assumption 2. Under this assumption, the definition of herd buying (selling) behavior is in equilibrium expressed through a buying (selling) threshold for the private signal. In addition, expressions for the two equilibrium thresholds are derived, which in general do not permit closed-form solutions. Using standard numerical techniques, one may still solve the two threshold-equations and thereby also estimate model parameters for a wider range of signal distributions using the methodology from Cipriani and Guarino (2014). Finally, it is demonstrated how model estimates may be used to measure intraday herding depending on model parameters, the distribution of the signal and the daily history of trades.

The concrete assumption that signals have conditional linear densities is appealing from a numerical point of view. Closed-form solutions for the threshold equations are obtained implying a more robust estimation procedure. This treatise specifies a model with conditional linear densities *and* a state-dependent precision. As a result, the ability of processing private information now depends on the unknown state of the world. It is then shown that the threshold equations still permit closed-form solutions. The likelihood ratio test may be used to determine if the conditional densities share the same precision since the model described by Cipriani and Guarino (2014) is nested in this model. Test statistics show that the precision parameter depends on the state — the signal distribution is not symmetric.

To rehearse, the main purpose of this paper is to investigate if rational statistical herding represent more or less of the total amount of trading activity during financial crises compared with tranquil periods. Indārs and Savin (2017) finds evidence of regular herding in the Moscow Exchange during days with negative market returns and periods of turmoil (e.g., the Financial Crisis of 2008), but their study is based on the CSAD-measure briefly mentioned in footnote 3⁴. To answer the main question, the empirical study in this paper is concentrated on one stock, Ashland Inc., traded on the New York Stock Exchange between 2005-2008, including a stock crash associated with the Great Recession⁵. The model is estimated for each trading year, and model fits are compared across years. The proportion of trading activity stemming from informed traders with correct signals (the PIN) is between 7.9% (2008) and 11.8% (2006). Estimates of model parameters are used to compute public as well as private beliefs. Two advantages come out of this: (1) Identification of trading periods within each day where herd behavior was present, and (2) compute the probability of herding for each trading period. The results show that herd buying and in particular herd selling increased during 2007 and 2008, consistent with a stock market crash. In particular, on event days during the stock crash, on average 4.9% of informed traders were herd selling, whereas 0.3% of

⁴Indārs and Savin (2017) also finds that extreme upward oil price movements is a factor associated with herding towards the Moscow Exchange.

⁵The stock Ashland Inc. is also used in the empirical studies by Easley et al. (1997) and Cipriani and Guarino (2014).

informed traders were herd selling before the crash. Finally, it is shown that herd selling behavior increased during high volatility periods in terms of the volatility index VIX.

The rest of this paper is organized as follows. Section 2 describes the model. Section 3 presents and solves the informed trader's decision problem. Section 4 analyzes the concept of rational statistical herding behavior. Measures to infer intraday herding are also derived. Section 5 considers the problem of estimating model parameters. Section 6 describes the data used for the empirical study while Section 7 presents the empirical findings. Section 8 concludes.

2. The Model

This paper considers a sequential trading model of asymmetric information. Trading of a risky asset is viewed as a game between liquidity providers and position takers repeated over several days. Let $d = 1, 2, \dots$ denote trading days and introduce the filtered probability space $\{\Omega, P, \mathcal{F}, (\mathcal{F}_d)_{d \geq 1}\}$ to capture the information flow where Ω represents all possible states of the economy, \mathcal{F} is the corresponding σ -algebra and P is the physical measure. Finally, $(\mathcal{F}_d)_{d \geq 1}$ is the σ -field generated by the value process of the risky asset.

Asset. The risky asset has a fundamental value on day d given by the realization of the random variable V_d depending on the state of the world. The asset value can only change between days: V_d is revealed at the end of trading day d and becomes common knowledge prior to trading the next day. Assume that $V_d \in \{V_d^H, V_d^L, v_{d-1}\}$ with $V_d^H > v_{d-1} > V_d^L$ where v_{d-1} denotes the realization of V_{d-1} . The σ -field $(\mathcal{F}_d)_{d \geq 1}$ is generated by the process (V_0, V_1, V_2, \dots) , and it is assumed that the closing price is a P -martingale with respect to (\mathcal{F}_d) , i.e., $E(V_d | \mathcal{F}_{d-1}) = V_{d-1}$ for $d = 1, 2, \dots$

Informational Events. If $V_d \neq v_{d-1}$, an information event has occurred, and the event is referred to as good (bad) if $V_d = V_d^H$ ($V_d = V_d^L$). Let $\omega \in \{H, L\}$ denote the state of the world during event days. Moreover, let $P(V_d \neq v_{d-1}) = \alpha$ and $P(V_d = V_d^H | V_d \neq v_{d-1}) = \delta$. Thus, prior to trading, α is the probability that an information event has occurred, and δ

is the probability that this event is good. Finally, the model assumes that information events are independently distributed across days.

Trading. Trading of the risky asset takes place over $D \geq 1$ days. For each day $d = 1, 2, \dots, D$, let $[0, T]$ denote the daily trading interval. Thus, time is continuous within days but discrete across days. Within each day d , a sequence of traders, $i = 1, 2, \dots, I_d$, enters the market. Individual i arrives at the random time $t_i \in [0, T]$ with a one-time opportunity to trade or refrain from trading at quoted prices⁶. Assume that if $i < j$ then $t_i < t_j$, and there is a bijection between trades and arrival times. Let $\{B, S, N\}$ be the *action space* (buy, sell or no-trade), and let $X_i^d \in \{B, S, N\}$ be the action taken by individual i observed at time $t_i \in [0, T]$ of day d . Figure 1 illustrates the structure of the observations.

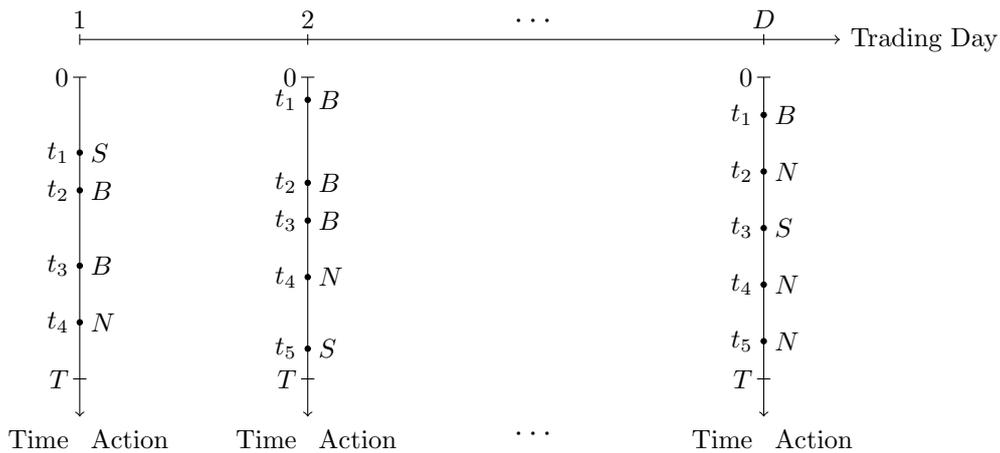


Figure 1: Example of the data structure.

Public History. The model structure is common knowledge among traders and the market maker. Let $h_t^d = \{X_1^d, \dots, X_t^d\}$ denote the publicly available history of trades before time $t + 1$ of day d . Moreover, let h^d denote the complete trading history of day d . Past transaction prices are assumed to be public.

⁶A trade is defined as buying or selling one unit of the stock.

Market Maker. The model assumes a risk-neutral and competitive price-setting market maker. The market maker is exposed to asymmetric information between him and privately informed traders (see **The Traders**). Consequently, the market maker protects himself against this adverse selection by setting a positive bid-ask spread. Let a_t^d (b_t^d) denote the ask (bid) price at time t of day d . The assumption that the market maker is competitive implies that

$$\begin{aligned} a_t^d &= \min \{a > 0 \mid a - E(V_d \mid h_{t-1}^d, X_t^d = B, a) = 0\} \\ b_t^d &= \max \{b > 0 \mid E(V_d \mid h_{t-1}^d, X_t^d = S, b) - b = 0\}. \end{aligned}$$

The expectation $E(V_d \mid h_{t-1}^d)$ is referred to as the *model price*.

Traders. Traders are placed in two groups: *informed* and *noise*. Informed traders are only present in the market on event days where they have private information of the unknown asset value. Noise traders are always present in the market, and they buy or sell with equal probability $\epsilon/2$ and refrain from trading with the remaining probability $1 - \epsilon^7$. On event days, the probability that an action stems from an informed trader is $\mu \in (0, 1)$.

Public Belief. The public belief of the asset value at time $t + 1$ of day d is given by $P(V_d = v \mid h_t^d)$ for $v \in \{V_d^H, V_d^L, v_{d-1}\}$. In addition, *the public expectation* of the asset value at time $t + 1$ of day d is given by $\sum_v v P(V_d = v \mid h_t^d)$. Let $\tilde{\alpha}_{t+1} := P(V_d \neq v_{d-1} \mid h_t^d)$ and $\tilde{\delta}_{t+1} := P(V_d = V_d^H \mid h_t^d, V_d \neq v_{d-1})$ denote, respectively, the posterior probability of being in an event day and the posterior probability that this event is good. Bayes' formula is used to update public beliefs, e.g.,

$$\tilde{\alpha}_{t+1} = \tilde{\alpha}_t \frac{P(X_t^d \mid h_{t-1}^d, V_d \neq v_{d-1})}{\sum_{v \in \{V_d^H, V_d^L, v_{d-1}\}} P(V_d = v \mid h_{t-1}^d) P(X_t^d \mid h_{t-1}^d, V_d = v)}.$$

⁷These traders act for exogenous motives such as liquidity reasons. The presence of noise traders ensures that the market does not break down (Avery and Zemsky, 1998).

Private Belief. On event days, the informed trader receives a private noisy signal about the unknown asset value. The signal is viewed as a realization of the random variable Z_t^d with support A . The conditional signal $Z_t^d | V_d^\omega$ has support A_ω and a distribution given by the measure ν^ω with density $f(\cdot | \tau_\omega, \omega)$, $\omega \in \{H, L\}$ ⁸. The *precision* parameter τ_ω may depend on the state ω and is interpreted as the trader's ability to process private information. The model also assumes that the conditional signals are IID. Finally, let $\theta = \{\alpha, \delta, \mu, \tau_H, \tau_L, \epsilon\} \in \Theta \subset \mathbb{R}^6$ denote the model parameter.

The signal measure $\nu = \delta\nu^H + (1 - \delta)\nu^L$ is the weighted average of the two state-dependent measures and defines the distribution of Z_t^d . It is natural to define the sequence of *posterior signals measures* $(\nu_t)_{t \geq 1}$ with $\nu_t = \tilde{\delta}_t\nu^H + (1 - \tilde{\delta}_t)\nu^L$. The signal realization ($Z_t^d = z$) is used to compute the informed trader's private belief of being in a good-event day, $p_t^d(z) := P(V_d = V_d^H | h_{t-1}^d, Z_t^d = z)$, as well as his private expectation of the fundamental asset value, $E(V_d | h_{t-1}^d, Z_t^d = z) = V_d^H p_t^d(z) + V_d^L (1 - p_t^d(z))$.

2.1. Modeling Private Signals

Private signals are by definition not observable and their distribution is of great importance for key features of the model (e.g., impact on prices and herding). The following example introduces the concept of *bounded* and *unbounded* private beliefs.

2.1.1. Example: Bounded Private Beliefs

Consider a signal realization $z \in A$ yielding the private belief $p(z) = P(V_d = V_d^H | Z_t^d = z) \in [0, 1]$. Assume first that $A_L = A_H$ (i.e., no signal realizations reveal the state of the world). As a result, the state-dependent measures ν^H and ν^L are mutually absolutely continuous with respect to each other, and the Radon-Nikodym derivative $g = \frac{d\nu^L}{d\nu^H} : A \rightarrow (0, \infty)$ exists and is unique. In particular, the private belief can be expressed by

⁸It is in general not assumed that the two state-dependent distributions share the same support, implying that some signals may reveal the true asset value. More technically, it is not required that ν^H and ν^L are equivalent measures, and the existence of the Radon-Nikodym derivative of ν^H with respect to ν^L (or ν^L with respect to ν^H) is not presumed.

$p(z) = 1/(f(z) + 1) \in (0, 1)$. Thus, neither of the two singletons $\{0\}$ and $\{1\}$ is an atom for the distribution of private beliefs, but it may be case that 0 and 1 is in A . Smith and Sørensen (2000) will in this setting say that beliefs are bounded if $\{0\}, \{1\} \notin A$ and unbounded if $\{0\}, \{1\} \in A$.

If $A_H \neq A_L$ and $P(Z_t^d \in A_H) > 0$ but $P(Z_t^d \in A_L) = 0$, there is positive probability attached to a signal realization z perfectly revealing state $\omega = H$, i.e., $p(z) = 1$. In a similar way, one may construct the corresponding case with $p(z) = 0$. Combining the two scenarios yields that both $\{0\}$ and $\{1\}$ are now atoms for the distribution of private beliefs, and beliefs are unbounded in the sense of Smith and Sørensen (2000). Consequently, it is natural to divide the different types of signals into three groups: bounded, unbounded and unbounded revealing.

With this in mind, it seems reasonable to expand the definition of bounded and unbounded signals depending on whether the two support sets coincide.

Definition 1. *Private beliefs are said to be bounded if $A_L = A_H$. Otherwise, beliefs are unbounded and*

$$\nu((A_H \cap A_L^c) \cup (A_H^c \cap A_L)) = \nu((A_H \cap A_L^c)) + \nu((A_H^c \cap A_L))$$

measures the level of unboundedness.

Figure 2 shows the trader's private belief as a function of the private signal in a setting with unbounded beliefs. Clearly, $\{0\}$ and $\{1\}$ are atoms for the distribution of the private belief, which is increasing in $z \in A_L \cap A_H$.

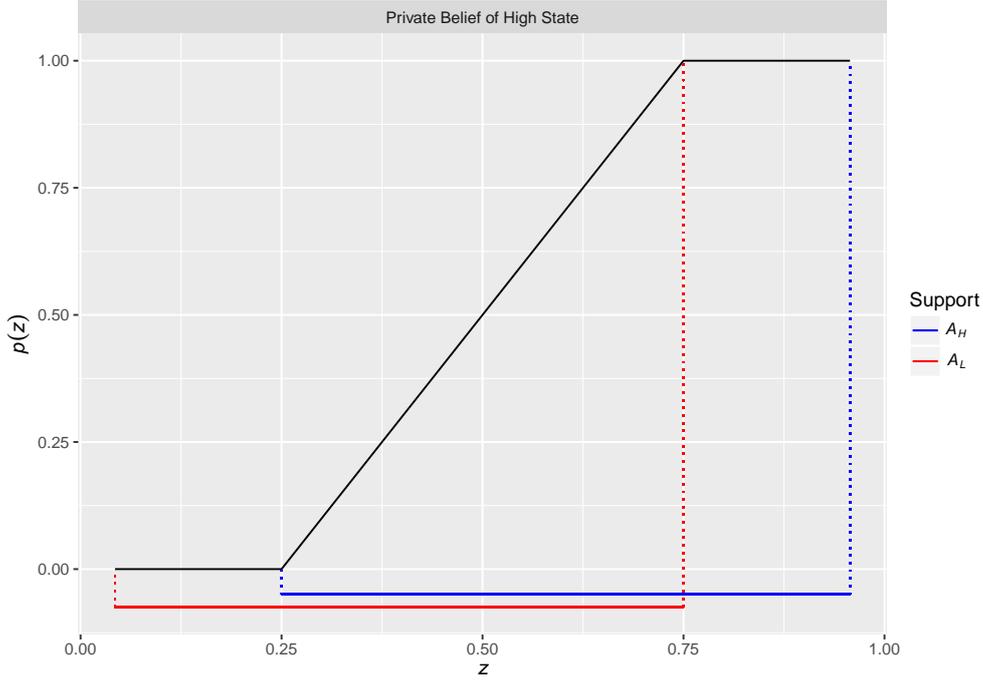


Figure 2: A signal distribution with conditional linear densities $f(z|H) = 4z - 1$ and $f(z|L) = 3 - 4z$, implying that $A_H = [1/4, 1/4 + \sqrt{1/2}]$ and $A_L = [3/4 - \sqrt{1/2}, 3/4]$.

Let $\underline{A} = A_H \cap A_L$ denote the common support of $Z_t^d|V_d^H$ and $Z_t^d|V_d^L$. The signal distribution will in the remainder of this paper be required to satisfy the following assumption.

Assumption 2. Let $z', z'' \in \underline{A}$ and assume

- i) that the interior of the common support of $Z_t^d|V_d^H$ and $Z_t^d|V_d^L$ is nonempty.
- ii) that the conditional signal distributions satisfy the monotone likelihood ratio property: For all $z', z'' \in \underline{A}$ with $z' > z''$ it must be the case that

$$\frac{f(z'|\tau_H, H)}{f(z'|\tau_L, L)} \geq \frac{f(z''|\tau_H, H)}{f(z''|\tau_L, L)}.$$

- iii) that there exists a signal realization $z^* \in \underline{A}$ satisfying $f(z^*|\tau_H, H) = f(z^*|\tau_L, L)$.

The signal z^* is called uninformative, since the informed trader's private expectation is unchanged after receiving z^* , i.e., $E(V_d | h_{t-1}^d, Z_t^d = z^*) = E(V_d | h_{t-1}^d, V_d \neq v_{d-1})$ for any trading history h_{t-1}^d . Similarly, the signal realization $z > z^*$ ($z < z^*$) will increase (decrease) the private expectation and is therefore interpreted as a good (bad) signal⁹. Clearly, the signal's credibility depends on the conditional distributions, which may exhibit important asymmetries.

Definition 3 (Symmetric Signal Distributions). *The signal distribution is symmetric if $f(z' | \tau_L, L) = f(z'' | \tau_H, H)$ for all $z', z'' \in \underline{A}$ with $z' = z^* - l$ and $z'' = z^* + l$ for $l > 0$. Otherwise, the signal distribution is asymmetric.*

If the signal distribution is symmetric, the probability of receiving a correct (or wrong) signal does not depend on the state ω . For the asymmetric case, the confidence in the private signal is state-dependent, which may affect the extent of herding on the buy and sell sides of the market¹⁰.

2.1.2. Example: Conditional Densities

Consider the case $f(z|H) = e^{-z}$ and $f(z|L) = 2e^{-2z}$ with $A_H = A_L = (0, \infty)$ (exponential distributions). Realize that $z^* = \log(2)$ and that the signal distribution is asymmetric. Assume for simplicity that $\delta = 1/2$, and notice that the Radon-Nikodym of ν^L with respect to ν^H exists and is given by $g(z) = 2e^{-z}$. As a result, the signal $z \in (0, \infty)$ yields a private belief of state $\omega = H$ given by

$$p(z) = \frac{1}{2e^{-z} + 1}$$

Furthermore, $p(z) \rightarrow 1/3$ for $z \rightarrow 0$, $p(z) \rightarrow 1$ for $z \rightarrow \infty$ and $p(z^*) = 1/2$.

⁹Use Bayes' formula to see that

$$\frac{P(V_d = V_d^H | h_{t-1}^d, z)}{P(V_d = V_d^L | h_{t-1}^d, z)} = \frac{f(z | \tau_H, H) P(V_d = V_d^H | h_{t-1}^d)}{f(z | \tau_L, L) P(V_d = V_d^L | h_{t-1}^d)},$$

and the interpretation of good and bad signals follows from the monotone likelihood ratio property.

¹⁰The informed trader is less likely to reject private information whenever he has great faith in his signal. Intuitively, greater confidence in signals during good-event days compared to bad-event days would imply more herd selling than herd buying in the market.

Signals are unbounded in the terminology of Smith and Sørensen (2000), because $\{1\}$ is included in the support for the distribution of private beliefs. However, signals are bounded in this setting by definition 1, since not even the most extreme signal realizations will reveal the true state of the world.

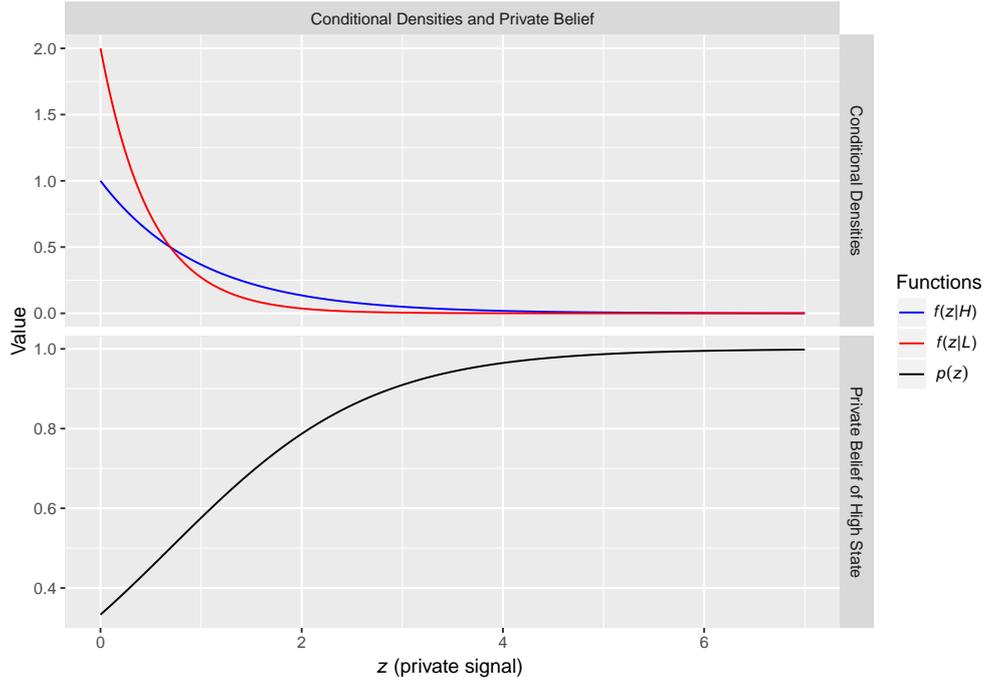


Figure 3: Top panel: Two densities satisfying the monotone likelihood ratio property. Lower panel: The updated private belief after observing z . The uninformative signal is given by the intersection of $f(z|H)$ and $f(z|L)$.

3. The Informed Trader's Decision Problem

The informed trader's decision problem (buy, sell or refrain from trading) is now formulated and then solved. Recall that noise traders buy and sell one unit of the asset with equal probability ($\epsilon/2$). Meanwhile, informed traders seek to maximize expected profit using private as well as public information. The trader's payoff depends on the realization of V_d at the end of day d , his action taken in day d and quoted prices when he was selected to trade. Thus, the utility function, $u : \{V_d^H, V_d^L\} \times \{B, S, N\} \times [V_d^L, V_d^H]^2$, is given by

$$u(V_d, X_t^d, a_t^d, b_t^d) = \begin{cases} V_d - a_t^d & \text{if } X_t^d = B, \\ b_t^d - V_d & \text{if } X_t^d = S, \\ 0 & \text{if } X_t^d = N \end{cases}.$$

Clearly, expected profit is positive whenever the trader's private expectation exceeds (is beneath) the ask (bid) price, and the trader will choose to buy (sell). That is, $X_t^d = B$ if $E(V_d | h_{t-1}^d, z_t^d) > a_t^d$, and $X_t^d = S$ if $E[V_d | h_{t-1}^d, z_t^d] < b_t^d$. Whenever the private valuation is inside the bid-ask spread, $E(V_d | h_{t-1}^d, z_t^d) \in [b_t^d, a_t^d]$, the trader will refrain from trading. Combining these observations with assumption 2 implies that the trader's decision problem may be expressed in terms of a buying threshold, β_t^d , and a selling threshold, σ_t^d , satisfying

$$\beta_t^d = \inf \{z \in \underline{A} \mid E(V_d | h_{t-1}^d, Z_t^d = z) > a_t^d\}, \quad (3.1)$$

$$\sigma_t^d = \sup \{z \in \underline{A} \mid E(V_d | h_{t-1}^d, Z_t^d = z) < b_t^d\}. \quad (3.2)$$

The convention $\beta_t^d = \sup \underline{A}$ is used if $\{z \in \underline{A} \mid E(V_d | h_{t-1}^d, z) > a_t^d\}$ is empty while $\sigma_t^d = \inf \underline{A}$ if $\{z \in \underline{A} \mid E(V_d | h_{t-1}^d, z) < b_t^d\}$ is empty¹¹. Summarizing, the solution to the informed trader's decision problem is given by

$$X_t^d = \begin{cases} B & \text{if } z_t^d \geq \beta_t^d \\ S & \text{if } z_t^d \leq \sigma_t^d \\ N & \text{if } z_t^d \in (\sigma_t^d, \beta_t^d) \end{cases}. \quad (3.3)$$

This suggests that the latent private signal (Z_t^d) may be inferred through the observable action X_t^d , since the probability of an action taken by the informed trader depends on β_t^d and σ_t^d ¹².

¹¹When ν^H and ν^L define continuous distributions, the inequalities in (3.1) and (3.2) become equalities, implying that $E(V_d | h_{t-1}^d, Z_t^d = \beta_t^d) = a_t^d$ and $E(V_d | h_{t-1}^d, Z_t^d = \sigma_t^d) = b_t^d$.

¹²E.g., the probability of observing $X_t^d = B$ from the informed trader is $\nu_t([\beta_t^d, \infty))$.

Formulas to compute the two thresholds established by equations (3.1) and (3.2), respectively, are given in the following proposition.

Proposition 4. *The threshold-process (β_t^d) is given recursively by*

$$\beta_t^d = \inf \left\{ z \in \underline{A} : 0 \leq \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t) \Delta_t^d(z, B) + (1 - \tilde{\alpha}_t) \frac{\epsilon}{2} \left(\tilde{\delta}_t (1 - \delta) f(z | \tau_H, H) - \delta (1 - \tilde{\delta}_t) f(z | \tau_L, L) \right) \right\} \quad (3.4)$$

while the threshold-process (σ_t^d) is given recursively by the equation

$$\sigma_t^d = \sup \left\{ z \in \underline{A} \mid 0 \geq \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t) \Delta_t^d(z, S) + (1 - \tilde{\alpha}_t) \frac{\epsilon}{2} \left(\tilde{\delta}_t (1 - \delta) f(z | \tau_H, H) - \delta (1 - \tilde{\delta}_t) f(z | \tau_L, L) \right) \right\} \quad (3.5)$$

with

$$\Delta_t^d(z, x) = f(z | \tau_H, H) P(X_t^d = x | h_{t-1}^d, V_d = V_d^L) - f(z | \tau_L, L) P(X_t^d = x | h_{t-1}^d, V_d = V_d^H)$$

for $z \in \underline{A}$ and $x \in \{B, S, N\}$. If the set in (3.4) ((3.5)) is empty, let $\beta_t^d = \sup \underline{A}$ ($\sigma_t^d = \inf \underline{A}$).

Proof.

See Appendix A. □

Prior to trading, posterior beliefs coincide with model parameters: $\tilde{\alpha}_1^d = \alpha$ and $\tilde{\delta}_1^d = \delta$. In particular, β_1^d and σ_1^d may be computed as solutions to equations (3.4) and (3.5), respectively. After observing the action X_1^d , $\tilde{\alpha}_2^d$ and $\tilde{\delta}_2^d$ are updated, and β_2^d and σ_2^d may be computed. The following scheme illustrates this iterative procedure.

Step 1 At time $t = 1$ equations (3.4) and (3.5) only depend on model parameters (θ) , and β_1^d and σ_1^d are both computed.

Step 2 At time $t = 2$, X_1^d is observed and the posterior probabilities $\tilde{\alpha}_2^d$ and $\tilde{\delta}_2^d$ are updated. Next, β_2^d and σ_2^d are computed using $P(X_2^d | h_1^d, V_d = V_d^\omega)$ for $\omega \in \{H, L\}$.

Step 3 Compute recursively the entire process of both β_t^d and σ_t^d : Repeat step 1 and step 2 for the full sequence of trades of day d .

When ν^H and ν^L are discrete, it is not guaranteed that there exists a signal realization β_t^d satisfying $E(V_d | h_{t-1}^d, Z_t^d = \beta_t^d) = a_t^d$ cf. equation (3.1) (and vice versa for σ_t^d), and the threshold is given by the smallest signal $z \in \underline{A}$ meeting the inequality $E(V_d | h_{t-1}^d, Z_t^d = z) > a_t^d$. When ν^H and ν^L are continuous distributions, the threshold β_t^d is the smallest solution to

$$0 = \tilde{\alpha}_t \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \Delta_t^d(z, B) + (1 - \tilde{\alpha}_t^d) \frac{\epsilon}{2} \left(\tilde{\delta}_t^d (1 - \delta) f(z | \tau_H, H) - \delta (1 - \tilde{\delta}_t^d) f(z | \tau_L, L) \right) \quad (3.6)$$

in $z \in \underline{A}$. Similarly, σ_t^d is the largest solution to

$$0 = \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \Delta_t^d(\sigma_t^d, S) + (1 - \tilde{\alpha}_t^d) \frac{\epsilon}{2} \left(\tilde{\delta}_t^d (1 - \delta) f(z | \tau_H, H) - \delta (1 - \tilde{\delta}_t^d) f(z | \tau_L, L) \right) \quad (3.7)$$

in $z \in \underline{A}$. The two thresholds are equivalent with, respectively, the smallest ask price and largest bid price resulting in zero expected profits for the market maker at time t of day d .

In addition to assumption 2, two extra criteria are introduced for the empirical study of herding: bounded private beliefs and tractability. Both criteria are met in a special case of the reputational cheap talk game described by Ottaviani and Sørensen (2006), where privately informed players have conditional linear densities. Let $f(z | \tau_x, x) = C_x + 2\tau_x z x$ where $x \in \{-1, 1\}$ is the state of the world (low or high), $\tau_x \in (0, 1)$ is the player's ability to process information, and $z \in [0, 1]$. The constant C_x is chosen such that $f(\cdot | \tau_x, x)$ becomes a probability density.

Model Specification I. *Assume that*

$$f(z | \tau_H, H) = 1 + \tau_H (2z - 1) \quad (3.8)$$

$$f(z | \tau_L, L) = 1 - \tau_L (2z - 1) \quad (3.9)$$

where $\tau_H, \tau_L \in (0, 1)$ and $z \in [0, 1] = A_L = A_H$.

Model Specification II. *Assume model specification I with $\tau_L = \tau_H = \tau$.*

Model specification I satisfies assumption 2. In accordance with definition 1, the restriction $\tau_\omega \in (0, 1)$ implies that beliefs are bounded while a state-dependent precision parameter, τ_ω , leads to an asymmetric signal distribution by definition 3. If the precision parameter is independent of the state ($\tau_L = \tau_H$), model specification I coincides with the model described by Cipriani and Guarino (2014). Finally, tractability is obtained since solutions to the threshold-equations in proposition 4 both become closed-form expressions.

Corollary 5. *Assume model specification I. Then, β_t^d and σ_t^d are given as the smallest solution to the quadratic equation (3.10) and the largest solution the quadratic equation (3.11), respectively*

$$0 = B_{\beta_t^d}(\theta, h_{t-1}^d)z^2 + C_{\beta_t^d}(\theta, h_{t-1}^d)z + D_{\beta_t^d}(\theta, h_{t-1}^d) \quad (3.10)$$

$$0 = B_{\sigma_t^d}(\theta, h_{t-1}^d)z^2 + C_{\sigma_t^d}(\theta, h_{t-1}^d)z + D_{\sigma_t^d}(\theta, h_{t-1}^d) \quad (3.11)$$

where the coefficients in the quadratic equations are given in the proof.

Proof.

See Appendix A. □

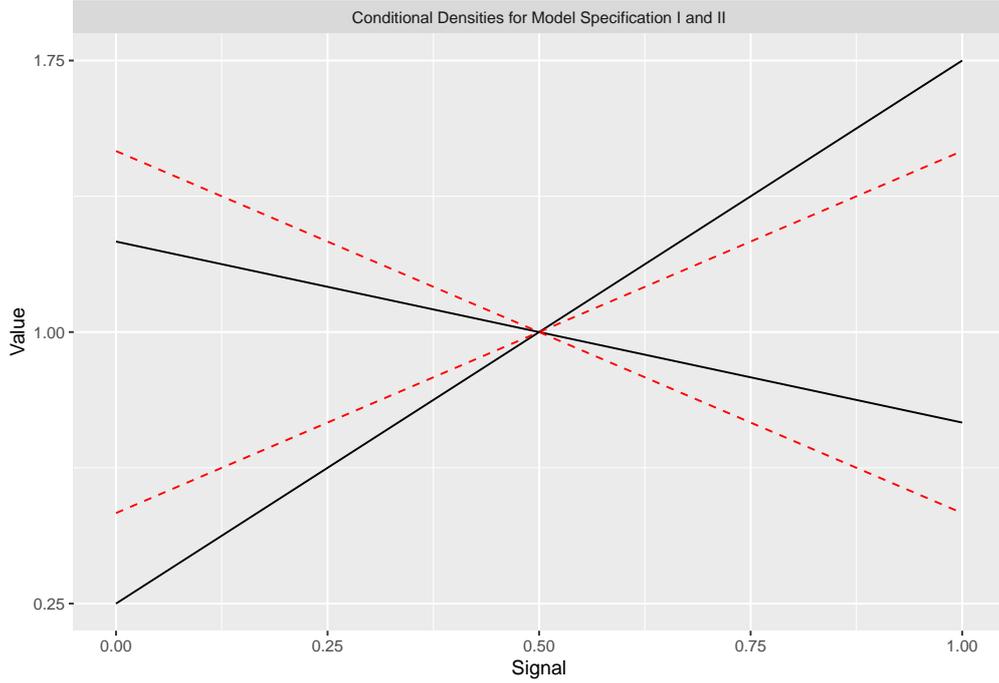


Figure 4: The black lines represent densities under model specification I with $\tau_H = 0.75 > 0.25 = \tau_L$. The red lines show the densities under model specification II ($\tau = 0.5$).

4. Herd Behavior

Combining the existence of the uninformative signal z^* with the buying (selling) threshold β_t^d (σ_t^d) in proposition 4 leads to a simple definition of herd buying (selling) behavior. As a starting point, rational herding is defined in accordance with Avery and Zemsky (1998).

Definition 6. *The informed trader is herd buying at time t of day d if*

$$E(V_d | h_{t-1}^d, z_t^d) > a_t^d \text{ for } z_t^d < z^* \quad (4.1)$$

$$E(V_d | h_{t-1}^d) > v_{d-1} \quad (4.2)$$

are both satisfied. The informed trader is herd selling at time t of day d if

$$E(V_d | h_{t-1}^d, z_t^d) < b_t^d \text{ for } z_t^d > z^* \quad (4.3)$$

$$E(V_d | h_{t-1}^d) < v_{d-1} \quad (4.4)$$

are both satisfied.

Thus, the informed trader exhibits herd buying (selling) behavior if 1) he receives a signal decreasing (increasing) his valuation of the asset prior to trading, and then 2) the public history of trades causes both the public and private valuation of the asset to rise (fall), and it is now profitable (in expectation) to buy (sell) the asset.

Lemma 7. *If $z_t^d \in [\beta_t^d, z^*)$ then $E(V_d | h_{t-1}^d) > v_{d-1}$. On the other hand, if $z_t^d \in (z^*, \sigma_t^d]$ then $E(V_d | h_{t-1}^d) < v_{d-1}$.*

Proof.

See Appendix A. □

Note that a signal realization $Z_t^d = z_t^d$ satisfying condition (4.1)—and thereby also condition (4.2)—occurs with positive probability for $\beta_t^d \in A \cap (-\infty, z^*)$. With this in mind, the definition of herd behavior can be rephrased in terms of the thresholds β_t^d and σ_t^d and the uninformative signal z^* .

Definition 8. *Herd behavior is present at time t of day d if the measure of engaging in herding (buying or selling) is positive, that is if*

$$\beta_t^d < z^* \text{ or } \sigma_t^d > z^*.$$

It is conjectured that herding occurs with probability 0 when the measure of perfectly informative signals tends toward 1, that herd buying occurs with positive probability when $A_H \cap A_L^C \neq \emptyset$ and that herd selling occurs with positive probability when $A_L \cap A_H^C \neq \emptyset$.

Assumption 9. *Assume that $1 - \epsilon > \nu_1((\sigma_1^d, \beta_1^d))$.*

The assumption simply says that prior to trading, noise traders are more likely to refrain from trading than informed traders are. The assumption is reasonable (verified by model estimates) and has the following implication: If $h_1^d = \{N\}$ then $X_2^d = N$ implies that $\nu_1((\sigma_1^d, \beta_1^d)) \geq \nu_2((\sigma_2^d, \beta_2^d))$ so $1 - \epsilon > \nu_2((\sigma_2^d, \beta_2^d))$ by assumption 9. Or more generally, if $h_{t-1}^d = \{N, \dots, N\}$ and $X_t^d = N$ then $1 - \epsilon > \nu_{t+1}((\sigma_{t+1}^d, \beta_{t+1}^d))$. This ensures

that each no-trade in a consecutive sequence of no-trades in the beginning of day d increases the probability of a buy or a sell for the informed trader in period $t+1$, which has the following interpretation: The learning market maker becomes more confident of being in a no-event day after observing the no-trade. Consequently, it is less likely that he is exposed to adverse selection due to asymmetric information, and he reduces the bid-ask spread, or, equivalently, the distance between the buying threshold and the selling threshold shrinks.

Lemma 10. *Consider the history of trades exclusively containing no-trades. For each no-trade, the market maker attaches lower probability to being in an event day. That is, $1 - \tilde{\alpha}_{t+1} > 1 - \tilde{\alpha}_t$ if $h_{t-1}^d = \{N, N, \dots, N\}$ and $X_t^d = N$.*

Proof.

See Appendix A. □

As in the models described by Avery and Zemsky (1998) and Cipriani and Guarino (2014), the combination of asymmetric information and event uncertainty leads to a positive probability of herd behavior. In addition, herding can be misdirected, meaning herd buying during a bad-event day or herd selling during a good-event day.

Proposition 11. *During a trading day, herding occurs with positive probability and can be misdirected.*

Proof.

See Appendix A. □

4.1. Measuring Rational Statistical Herding

Recall that the informed trader will herd at time t of day d if his private signal belongs to $[\beta_t^d, z^*)$ (buy) or $(z^*, \sigma_t^d]$ (sell). Thus, herding at time t on event day d can be represented by the binary variable $1_{[\beta_t^d, z^*)}(Z_t^d) + 1_{(z^*, \sigma_t^d]}(Z_t^d)$. This variable is — opposite to X_t^d — by definition not observed, and herd behavior cannot be detected with probability one. However, using the underlying model assumptions, accurate measures of herding can be derived (e.g., the probability that the informed trader herd buys is given by $P(Z_t^d \in [\beta_t^d, z^*) | h_{t-1}^d)$).

Thus, measuring herding on event days is directly related to inferring the distribution of the latent variable Z_t^d , which can be expressed by the state-dependent distributions, and the posterior probabilities of the bad-event and the good-event respectively:

$$\begin{aligned} P(Z_t^d | h_{t-1}^d) &= \sum_{v \in \{V_d^H, V_d^L\}} P(V_d = v | h_{t-1}^d, V_d \neq v_{d-1}) P(Z_t^d | V_d = v, h_{t-1}^d) \\ &= \sum_{v \in \{V_d^H, V_d^L\}} P(V_d = v | h_{t-1}^d, V_d \neq v_{d-1}) P(Z_t^d | V_d = v). \end{aligned} \quad (4.5)$$

From an econometric point of view, the informed trader's probability of statistical herding *when* he arrives at the market is not of interest. Instead, the entire daily history of trades, h^d , is used to infer the extent of herding during day d ¹³. At first, consider the probabilities $P(Z_t^d \in [\beta_t^d, z^*] | h^d)$ and $P(Z_t^d \in (z^*, \sigma_t^d] | h^d)$. The former (latter) probability is referred to as the measure of herd buying (selling). Now, let $\Psi_{\beta_t^d}$ ($\Psi_{\sigma_t^d}$) denote the ratio between the measure of herd buying (selling) and the *potential* measure of herd buying (selling), that is

$$\Psi_{\beta_t^d} := \frac{P(Z_t^d \in [\beta_t^d, z^*] | h^d)}{P(Z_t^d < z^* | h^d)} \quad \text{and} \quad \Psi_{\sigma_t^d} := \frac{P(Z_t^d \in (z^*, \sigma_t^d] | h^d)}{P(Z_t^d > z^* | h^d)} \quad (4.6)$$

with $\Psi_{\beta_t^d}, \Psi_{\sigma_t^d} \in [0, 1]$. Using the terminology of Cipriani and Guarino (2014), the two parameters are called the proportion of herd buyers and the proportion of herd sellers, respectively. Finally, let $\Psi_{\beta}^d := \frac{1}{I_d} \sum_{t=1}^{I_d} \Psi_{\beta_t^d}$ and $\Psi_{\sigma}^d := \frac{1}{I_d} \sum_{t=1}^{I_d} \Psi_{\sigma_t^d}$ denote the daily averages of the two ratios above. These two parameters will be used to estimate the extent of herd buying and herd selling behavior on the market.

¹³Public beliefs will converge toward 0 or 1 due to the high number of trading frequencies. Especially, $P(Z_t^d | h^d)$ will (approximately) equal either $P(Z_t^d | V_d = V_d^L)$ or $P(Z_t^d | V_d = V_d^H)$.

4.1.1. Example: The Measure of Herd Buying

Consider model specification II and assume that $V_d = V_d^H$. Note that $z^* = 1/2$ since $f(1/2|\tau_H, H) = f(1/2|\tau_L, L) = 1$. The measure of herd buying at time t is given by

$$\begin{aligned} P(Z_t^d \in [\beta_t^d, 1/2) | V_d = V_d^H) &= \int_{\beta_t^d}^{1/2} 1 + \tau(2z - 1) dz \\ &= \tau \left(\frac{1}{4} - (\beta_t^d)^2 \right) + (1 - \tau) \left(\frac{1}{2} - \beta_t^d \right) \end{aligned}$$

if $\beta_t^d < 1/2$ and 0 otherwise. The measure of herd selling can be computed in a similar way. Parameter values are specified such that herd buying behavior occurs with a positive probability. Let $\tilde{\alpha}_t^d = \delta = \mu = \epsilon = 0.5$ and $\tilde{\delta}_t^d = 0.6$. The parameter choice $\tilde{\delta}_t^d > \delta$ implies that the model price has increased during the trading day, which is a necessary condition for an informed trader to engage in herd buying according to definition 8. Figure 5 illustrates the measure of herd buying as well as the buying threshold β_t^d as a function of the precision τ .

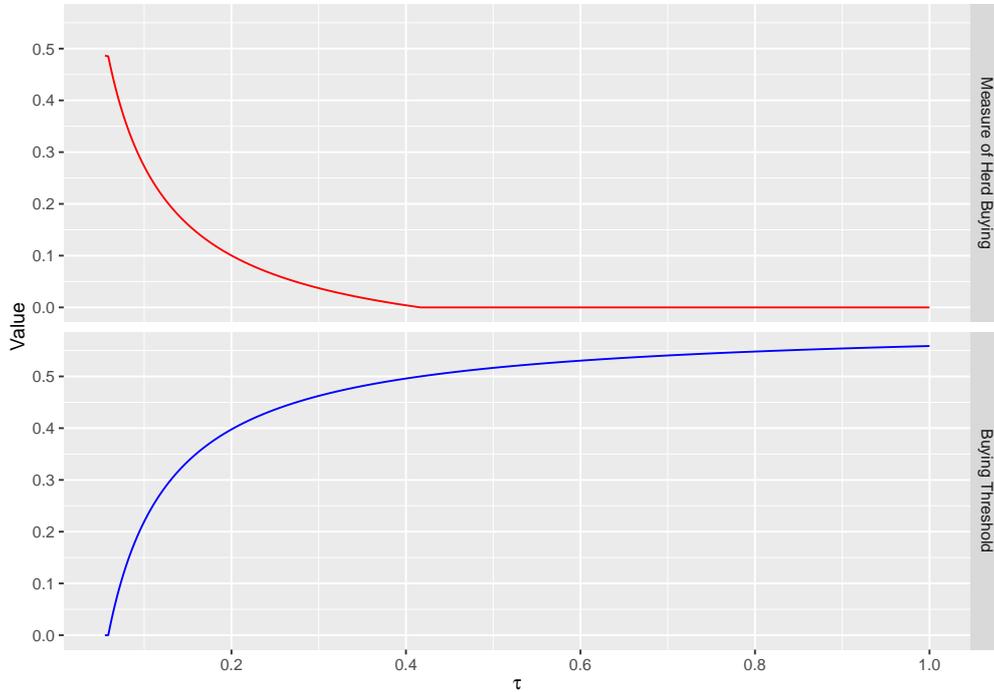


Figure 5: The buying threshold β_t^d (blue line) and the measure of herd buying behavior (red line) as a function of the precision τ .

Clearly, β_t^d is increasing in τ , whereas the measure of herd buying is decreasing. If $\tau \rightarrow 0$ the informed trader's signal becomes completely uninformative, and he chooses to buy the asset since $\tilde{\delta}_t^d > \delta$. This observation corresponds to $\beta_t^d \rightarrow 0$. Notice that the probability of receiving a bad (and a good) signal converges to $1/2$ as $\tau \rightarrow 0$, and the measure of herd buying also converges to $1/2$ as $\tau \rightarrow 0$. The proportion of herd buyers, $\Psi_{\beta_t^d}$, is obtained by multiplying $\gamma_{\beta_t^d}$ with $(0.5 - 0.25\tau)^{-1}$ and will as a function of τ have the same shape as the measure of herd buying.

When τ increases, the precision of private information increases, and the informed trader demands a higher signal realization in order to buy the asset. This is illustrated by the probability of receiving a bad signal conditioning on the high state

$$P(Z_t^d < 0.5 | V_d = V_d^H) = \frac{1}{2} - \frac{1}{4}\tau,$$

which is decreasing in τ . Consequently, if τ increases so does β_t^d , and the measure of signal realizations that result in herd buying behavior is decreasing. When β_t^d rises above 0.5, herd buying opportunities vanishes.

Contrary to the original PIN model by Easley et al. (1996) and the subsequent model by Easley et al. (1997), the private noisy signal does not reveal the true asset value. Consequently, the probability of information-based trading activity is independent of the precision of the signal in these models. If informed traders only receive correct signals in this setting ($\tau \rightarrow \infty$), the model boils down to the model described by Easley et al. (1997), and the PIN is equal to $\text{PIN}_{EKO} := \frac{\alpha\mu}{\alpha\mu + \epsilon(1 - \alpha\mu)}$ ¹⁴. When $\tau < \infty$, informed traders may receive incorrect signals or choose to herd, and therefore the former expression is adjusted by multiplying with the probability of receiving a correct signal:

$$\begin{aligned} \text{PIN} &:= \frac{\alpha\mu}{\alpha\mu + \epsilon(1 - \alpha\mu)} \cdot ((1 - \delta)P(Z_t^d < 0.5 | V_d = V_d^L) + \delta P(Z_t^d > 0.5 | V_d = V_d^H)) \\ &= \text{PIN}_{EKO} \cdot \left(\frac{1}{2} + \frac{1}{4}\tau \right). \end{aligned}$$

The PIN measures the proportion of trading activity from informed traders with a correct signal, and it is obviously increasing in the precision τ .

¹⁴If $\tau \geq 1$ beliefs are unbounded, and the support for the conditional signal shrinks.

5. The Likelihood

The model is fully parametric and may be fitted to data using classical maximum likelihood estimation. Recall that $\theta \in \Theta \subset \mathbb{R}^6$ is the parameter of interest where $\Theta = (0, 1)^6$. In this framework — and opposite to Easley et al. (1997), where only the daily total numbers of buys, sells and no-trades are used to infer information about model parameters — the specific order of trades will convey information. For any trading day d , the probability of the sequence of trades $h_t^d = \{X_1^d, \dots, X_t^d\}$ is given recursively by

$$P(h_t^d | \theta) = \prod_{s=1}^t P(X_s^d | h_{s-1}^d, \theta). \quad (5.1)$$

The value on the left-hand side in equation (5.1) is referred to as the action probability at time t of day d . Since informational events are independent across days, and V_{d-1} is revealed prior to trading of day d , the full model likelihood over D days is expressed by

$$L(\theta | \{h^d\}_{1 \leq d \leq D}) = P(\{h^d\}_{1 \leq d \leq D} | \theta) = \prod_{d=1}^D P(h^d | \theta). \quad (5.2)$$

In the numerical implementation it is convenient to rewrite the estimation problem in terms of the loglikelihood function, $l(\theta | \{h^d\}_{1 \leq d \leq D}) := -\log(L(\theta | \{h^d\}_{1 \leq d \leq D}))$, converting the problem of maximizing the product of action probabilities to minimizing the sum of the log-probabilities¹⁵. The product terms on the right-hand side of equation (5.1) are computed using the law of total probability:

¹⁵The complexity and non-linearity of the likelihood causes several empirical issues when estimating model parameters. The likelihood function is implemented in C++ due to the recursive structure as well as the high trading frequencies (need for speed). The estimation is carried out in the language R using the package `Rcpp` providing R-functions and C++-classes, which offers a seamless integration of R and C++. A global optimizer is required due the complexity of the likelihood function—solutions using general-purpose optimization functions (based on Nelder-Mead, quasi-Newton and conjugate-gradient algorithms) searching for a local minimum will depend on the initial parameter values. Mullen (2014) recommends using the optimizer from the package `rgenoud`. In addition, the R-package `DEoptim` provides a faster global optimizer (`DEoptim()`) based on the differential evolution algorithm, which also shows good performance (Mullen, 2014). The optimizer used will depend on the specific estimation problem.

$$\begin{aligned}
P(X_t^d | h_{t-1}^d) &= P(V_d = V_d^H | h_{t-1}^d)P(X_t^d | h_{t-1}^d, V_d = V_d^H) \\
&\quad + P(V_d = V_d^L | h_{t-1}^d)P(X_t^d | h_{t-1}^d, V_d = V_d^L) \\
&\quad + P(V_d = v_{d-1} | h_{t-1}^d)P(X_t^d | h_{t-1}^d, V_d = v_{d-1}), \quad (5.3)
\end{aligned}$$

and the problem of estimating θ is then reduced to compute all conditional probabilities on the right-hand side of equation (5.3).

The total number of trading days is used to infer the event probabilities α and δ , while intraday observations are used to draw inference of the remaining model parameters $(\mu, \epsilon, \tau_H, \tau_L)$. Therefore, model estimation is only sensible for a minimum of trading days, and standard deviations are expected to be highest for estimated event probabilities.

5.1. Computing Action Probabilities

Consider again the model likelihood (5.2) given by the action probabilities in (5.3). Computing all action probabilities requires computations of $P(X_t^d | V_d^\omega = v, h_{t-1}^d)$ and $P(V_d = v | h_{t-1}^d)$ for $\omega \in \{H, L\}$ and $v \in \{V_d^H, V_d^L, v_{d-1}\}$. If an information event has occurred on day d , the action $X_t^d \in \{B, S, N\}$ can either stem from the informed trader with probability μ , or it can come from the noise trader with probability $1 - \mu$. Recall the informed trader's optimal decision is given by equation (3.3) saying that the privately informed trader is only willing to buy or sell the asset for signal realizations above the threshold β_t^d or for signal realizations below σ_t^d , respectively. Summarizing,

$$P(X_t^d = B | V_d^\omega, h_{t-1}^d) = \mu \int_{\beta_t^d}^1 f(z | \tau_\omega, \omega) dz + (1 - \mu) \frac{\epsilon}{2} \quad (5.4)$$

$$P(X_t^d = S | V_d^\omega, h_{t-1}^d) = \mu \int_0^{\sigma_t^d} f(z | \tau_\omega, \omega) dz + (1 - \mu) \frac{\epsilon}{2} \quad (5.5)$$

for $\omega = H, L$, and $P(X_t^d = N | V_d^\omega, h_{t-1}^d)$ equal to the complement of the two probabilities. If no event occurs, trading is based solely on exogenous motives and $P(X_t^d = B | V_d = v_{d-1}, h_{t-1}^d) = P(X_t^d = S | V_d = v_{d-1}, h_{t-1}^d) = \epsilon/2$. Finally, the posterior event probabilities are given by Bayes' formula

$$P(V_d = v | h_t^d) = \frac{P(V_d = v | h_{t-1}^d) P(X_t^d | h_{t-1}^d, V_d = v)}{\sum_{j \in \{V_d^H, V_d^L, v_{d-1}\}} P(X_t^d | h_{t-1}^d, V_d = j) P(V_d = j | h_{t-1}^d)} \quad (5.6)$$

for $v \in \{V_d^H, V_d^L, v_{d-1}\}$. Thus, the action probability for any day d is obtained by computing the three probabilities (5.4), (5.5) and (5.6), all depending on the threshold-processes (β_t^d) and (σ_t^d) . Formulas for β_t^d and σ_t^d were given in proposition 4 followed by the iterative scheme to compute the two thresholds recursively.

6. Data

Financial intraday data are retrieved from the MTAQ (*monthly* trades and quotes) database. The dataset contains transactions and quotes for stocks traded on different US stock exchanges with time stamps within a given second. Data cleaning is essential for an accurate model calibration. This paper follows the cleaning procedures of the MTAQ-file described by Holden and Jacobsen (2013) and Barndorff-Nielsen et al. (2009)—see appendix B for a description.

The empirical analysis is concentrated on the stock Ashland Inc. traded on the New York Stock Exchange (NYSE) over the four-year period 2005-2008¹⁶¹⁷. Figure 6 shows closing prices for the stock as well as daily closing values for the VIX. Both time series are retrieved from <https://finance.yahoo.com/>¹⁸.

¹⁶As mentioned earlier, the NYSE stock Ashland Inc. is also used in the empirical studies by Easley et al. (1997) and Cipriani and Guarino (2014).

¹⁷Only transactions stemming from NYSE are considered. More crucially, the empirical study only uses quotes from NYSE based on an investigation of liquidity measurement problems for the given trading period. Angel et al. (2015) report that the NYSE's market share in NYSE-listed stocks has decreased from 80% in 2005 to under 40% in 2008. Comparing trade locations obtained from NYSE-quotes with trade locations obtained from the constructed NBBO-file (adjusted for withdrawn quotes) indicates that NYSE-quotes are sufficient for this study.

¹⁸Closing prices for the stock are adjusted for two splits in the given period: 10000/8212 on 07/01/2005 and 511/250 on 15/05/2017.

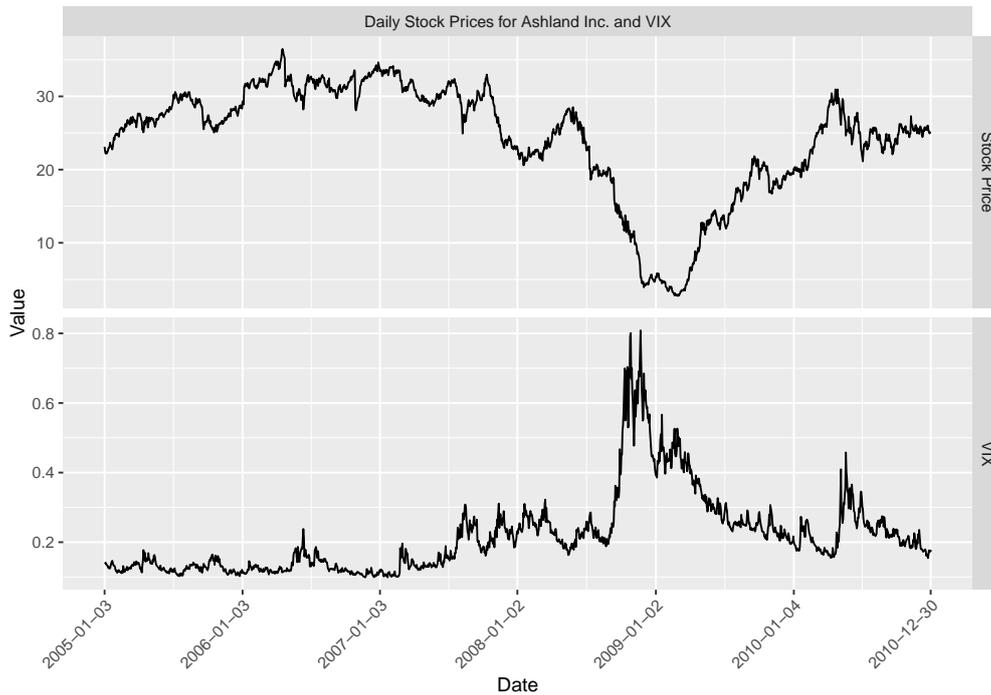


Figure 6: Daily closing prices for Ashland Inc. and the VIX from 2005 to 2010.

The trade classification algorithm proposed by Lee and Ready (1991), which requires Level 2 tick data (transactions and quotes), is used to determine the aggressor's side of each transaction^{19 20}. Table 1 presents the total numbers of trades as well as the allocation of buys and sells across the four trading years.

¹⁹Lee-Ready algorithm: If the transaction price is above (beneath) the midpoint for the quotes posted just before the trade, the trade is classified as a buy (sell). If the transaction price is equal to the midpoint, the transaction price is compared with the most recent price movement. If the price has increased (decreased), the trade is classified as a buy (sell).

²⁰Quotes are registered before trades and should therefore be moved ahead in time. The delay between quotes and transactions depends on the market. Cipriani and Guarino (2014) (sample from 1995) follow the suggestion by Lee and Ready (1991) and delay posted quotes 5 seconds. Vergote (2005) suggests a 2-second delay based on an analysis of stocks traded on NYSE after 2005. Comparing trade locations (at the quote, inside the spread and outside the spread) for different delays suggests a 1-second delay consistent with the majority of empirical studies, cf. Holden and Jacobsen (2013).

Table 1: SUMMARY OF TRANSACTION DATA

Trading Year	Trading Days	Trades	Buys	Sells
2005	251	490228	261227	228965
2006	251	513168	271121	242021
2007	251	446459	226624	219825
2008	253	463995	229658	234324

Description: Yearly allocations of buys and sells for the stock Ashland Inc. traded on NYSE between 2005 and 2008. Transaction data are retrieved from MTAQ (Monthly Trades and Quotes), and the classification of buys and sells are obtained using the Lee-Ready algorithm.

Table 1 shows that the order imbalance (difference between buy orders and sell orders) is positive for all trading years except 2008. Moreover, the order imbalance is decreasing throughout the entire sample period (on a yearly basis). The high number of trades emphasizes one important issue discussed by O'Hara (2015): the problem of estimating the likelihood function. Combining the speed of Bayesian learning with high trading frequencies yields serious econometric issues. Essentially, the likelihood function becomes rapidly flat in the parameter values when increasing the number of daily actions, and model fitting becomes very difficult if not infeasible.

In addition, there is the problem of classifying a no-trade action, which has been previously addressed by Easley et al. (1997) and Cipriani and Guarino (2014). In both model-settings, the no-trade rule is determined by intraday trade durations (e.g., if the time between two consecutive trades exceeds a specified limit, a no-trade action is inserted in between the two trades). Cipriani and Guarino (2014) propose that the no-trade interval should equal the ratio between total daily trading time (23400 seconds) and the average number of trades across days²¹.

The choice of the no-trade interval in this paper is also based on the distribution of intraday trade duration times across trading years; a summary

²¹A trade occurred approximately every 259 seconds on average for the 1995 sample. If the time between two consecutive trades was greater than 259 seconds and at the same time beneath $k \cdot 259$ seconds, $k = 1, 2, \dots$, k no-trades were inserted in between these two trades.

of duration times for the 2008 sample is presented in internet appendix I.A.1. For a given trading year, the no-trade interval is chosen as the average of trading duration times²²²³. As a result, the percentage of buys, sells and no-trades across years ranges from 29.6% (2007) to 30.2% (2006), 26.1% (2005) to 30.2% (2008) and 40.1% (2008) to 44.2% (2005).

7. Results

Model specifications I and II are fitted to the data described in the previous section. As a starting point, yearly updated estimates under model specification I are reported. Table 2 shows the results, while standard errors of the estimates as well as estimates under model specification II are reported in tables 10 and 11 in appendix C.

Table 2: MAXIMUM LIKELIHOOD ESTIMATES

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\tau}_H$	$\hat{\tau}_L$	$\hat{\epsilon}$	$l(\hat{\theta})$
2005	0.234	0.696	0.322	0.394	0.161	0.517	933698
2006	0.376	0.695	0.320	0.416	0.077	0.536	961528
2007	0.382	0.500	0.247	0.321	0.307	0.550	824148
2008	0.299	0.319	0.300	0.444	0.330	0.567	835564

Description: Yearly updated maximum likelihood estimates under model specification I. The last column reports log-likelihood values associated with the estimates.

The yearly updated parameter estimates show that the probability of the occurrence of an information event ranges between 23.4% (2005) and 38.2% (2007). The estimated probability of the good information event ($\hat{\delta}$) is decreasing across years, ranging between 69.6% (2005) and 30.0% (2008). Except for 2008, $\hat{\delta}$ is above or equal to 50%, indicating a non-negative drift of the stock value for those years. Furthermore, the probability that an action is information-based ($\hat{\mu}$) is between 24.7% (2007) and 32.2% (2005).

²²The no-trade intervals are 10 seconds for 2005 and 2006 and 12 seconds for 2007 and 2008.

²³As a robustness check, the empirical analysis in section 7 has been replicated for other choices of the no-trade interval. Even though the choice of no-trade interval has an impact on the model calibration, the main findings about herding are not affected.

A noise trader's probability of buying or selling is stable across years (between 51.7% and 56.7%) due to the choice of no-trade interval. Finally, estimates of τ_L and τ_H suggest that the private signal distribution exhibits important asymmetries. In particular, the *direction* of this asymmetry depends on the sample period: Recall that τ_L and τ_H are interpreted as the informed trader's ability to benefit from private information. For instance, in 2006 the estimated probability of receiving a bad signal ($Z_t^d < 0.5$) on a good-event day is $1/2 - 1/4 \cdot 0.416 = 0.396$, while the estimated probability of receiving a good signal during a bad-event day is 0.481^{24} . The private signal conditioning on the bad information event is almost completely uninformative in 2006.

According to definition 3, the signal distribution is said to be symmetric if $\tau_L = \tau_H$. The model under specification II is nested in the model under specification I, and the likelihood ratio test can be used to investigate the hypothesis of a symmetric signal distribution. Consider the parametric space $\Theta = (0, 1)^6 \subset \mathbb{R}^6$ and the subset $\Theta_0 = (0, 1)^5 \subset \Theta$. Moreover, consider the hypothesis $H : \theta = (\alpha, \delta, \mu, \tau_H, \tau_L, \epsilon) \in \Theta_0$. The corresponding size of the likelihood ratio test is given by

$$Q = \frac{\sup_{\theta \in \Theta_0} L(\theta; \{h^d\}_{1 \leq d \leq D})}{\sup_{\theta \in \Theta} L(\theta; \{h^d\}_{1 \leq d \leq D})} \in [0, 1],$$

which is critical for small values²⁵. Test sizes and corresponding P -values are reported in table 12 in appendix C. Except for 2007, the hypothesis is rejected for all samples at the 0.01-level, and it is statistically significant that the signal distribution is asymmetric²⁶. This contradicts the restriction by Cipriani and Guarino (2014), where the privately informed trader's ability to process private information is state-independent. As a result, estimates under model specification I are used for the further investigation.

²⁴Model assumptions imply that the probability of receiving incorrect signals is in $[1/4, 1/2]$.

²⁵Wilks's theorem states that the asymptotic distribution of the random variable $-2 \log Q$ follows a χ^2 -distribution with 1 degree of freedom.

²⁶Model estimates for the 2007 sample imply a test size equal to 2.85 with a corresponding P -value equal to 0.09.

Table 3 presents results of information-based trading activity where the PIN are reported across trading years²⁷.

Table 3: THE PROBABILITY OF INFORMATION-BASED TRADING

	2005	2006	2007	2008
PIN	0.079	0.118	0.092	0.087

Description: Yearly estimates of information-based trading activity under model specification I.

The PIN ranges between 7.9% (2005) and 11.8% (2006). Comparing with the PIN reported by Cipriani and Guarino (2014) (12% for the 1995 sample), the major increase in trading frequencies has a negative impact on information-based trading activity. This is mainly due to a lower fraction of informed traders (μ) with less precise private information (τ_L and τ_H).

Finally, estimates and the history of trades are used to compute the daily posterior beliefs ($\hat{\alpha}_t^d$ and $\hat{\delta}_t^d$). In addition, the two threshold-processes (β_t^d) and (σ_t^d) used to measure rational statistical herding are computed from posterior beliefs.

7.1. Herding

The empirical investigation of rational statistical herding is concentrated on event days. The high frequency of actions implies that public beliefs approximately converge to their true values, that is, $\min\{\tilde{\alpha}_{I_d}^d, 1 - \tilde{\alpha}_{I_d}^d\} \approx 0$ and $\min\{\tilde{\delta}_{I_d}^d, 1 - \tilde{\delta}_{I_d}^d\} \approx 0$. By construction, trading day d is classified as a good-event day if $\tilde{\delta}_{I_d}^d > K$ and as a bad-event day if $1 - \tilde{\delta}_{I_d}^d > K$ with the arbitrary choice $K = 0.99$ ²⁸. In particular, the number of event days across years are consistent with model estimates in table 2 and

²⁷The expression for the PIN is given by

$$\text{PIN} = \frac{\alpha\mu}{\alpha\mu + \epsilon(1 - \alpha\mu)} \left(0.5 + \delta \left(\frac{\tau_H - \tau_L}{4} \right) + \frac{\tau_L}{4} \right)$$

where the first term corresponds to the *standard* PIN introduced by Easley et al. (1997), and the second term is the probability that the informed trader receives a correct signal.

²⁸Non-reported results show that the main findings of herding are the same for $K = 0.9$ and $K = 0.999$.

total number of trading days reported in table 1²⁹. Intuitively, herding is observed when investors trade in the same direction over a period of time.

Figure 7 illustrates the price ($\tilde{\delta}_t^d$), the event probability ($\tilde{\alpha}_t^d$) and the trade imbalance for the first 300 actions during 23/06-2008.



Figure 7: The price (dashed black line), event probability (dashed red line) and trading imbalance (solid black line) during an event day. The upper plot shows the posterior beliefs, while the lower plot shows the trading imbalance (difference between buys and sells).

Figure 7 shows a positive dependence between the price (black dashed line) and the trading imbalance (black solid line). Prior to trading, the price is $\hat{\delta} = 0.299$, which immediately drops caused by a consecutive number of sells at the beginning of the day. The trading imbalance is constant around trading times 50 and 90, and the price becomes constant. After this quiet period, an increase in the trading imbalance (and thereby the price) is observed. The event probability fluctuates more than the price since updating $\tilde{\delta}_t^d$ also depends on the posterior belief $\tilde{\alpha}_t^d$. Clearly, the

²⁹The numbers of good-event and bad-event days are 64 and 29 for 2005, 53 and 34 for 2006, 39 and 40 for 2007 and finally 27 and 45 for 2008. Based on model estimates, it is, for instance, expected that the 2008 sample has $0.299 \cdot 0.319 \cdot 253 \approx 24$ good-event days and 51 bad-event days.

market maker's belief about the information event is almost 0 when the trading imbalance is constant, while it converges to 1 after observing a steady increase in the trading imbalance.

By definition 8, rational statistical herding is said to be present in data whenever $\hat{\beta}_t^d < 0.5$ (herd buying) or $\hat{\sigma}_t^d > 0.5$ (herd selling). Consequently, the threshold-processes are estimated using the parameter estimates reported in table 2.

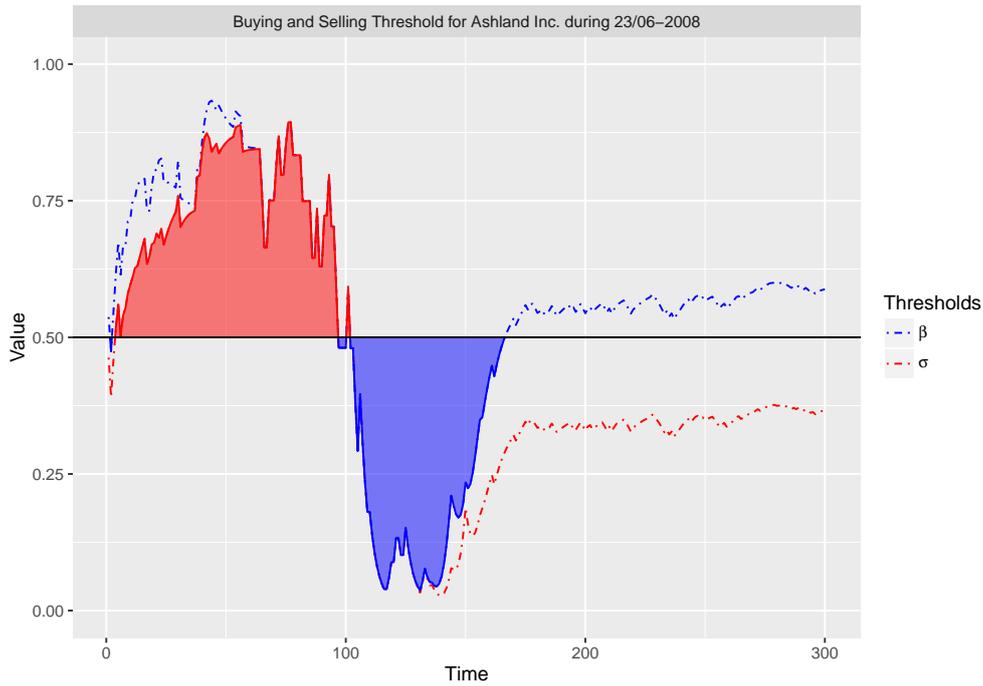


Figure 8: The blue (red) dashed line shows the buy (sell) threshold β_t^d (σ_t^d). The uninformative signal $z^* = 1/2$ is presented by the horizontal black line. The blue (red) filled regions indicate the presence (and the measure) of herd buying (selling) behavior.

Figure 8 shows that herd selling behavior is present during the first 90 trades approximately. This is in accordance with figure 7, where the price in this period is below the price prior to trading, and the trade imbalance has decreased. Between trading times 90 and 170, the price is above the initial level, and herd buying behavior is present. At some point, the market maker becomes almost certain that the good information event occurred ($\tilde{\alpha}_t^d \approx 1$), and the two thresholds boil down to

$$\beta_t^d = \frac{2\mu + (1 - \mu)\epsilon - \sqrt{(2\mu + (1 - \mu)\epsilon)^2 + 4\mu(-\mu - (1 - \mu)\frac{\epsilon}{2})}}{2\mu}$$

$$\sigma_t^d = \frac{-(1 - \mu)\epsilon + \sqrt{(1 - \mu)^2\epsilon^2 + 4\mu(1 - \mu)\frac{\epsilon}{2}}}{2\mu}.$$

In particular, estimated model parameters imply that $\beta_t^d > 0.5$ and $\sigma_t^d < 0.5$ —in the spirit of Avery and Zemsky (1998), herd behavior is no longer present in data due to the absence of event uncertainty.

From now on, the empirical investigation will be concentrated on the entire sample. Table 4 presents a summary of herding frequencies across days for all four samples.

Table 4: SUMMARY OF HERDING FREQUENCIES

	> 5%	> 10%	> 25%	Mean Time	Mean Length	Length SD	Max Length
2005							
Herd Buy	0.02	0.01	0.00	7.85	12.56	21.51	183
Herd Sell	0.02	0.01	0.00	9.12	13.02	18.70	139
2006							
Herd Buy	0.03	0.01	0.01	9.62	14.34	50.89	679
Herd Sell	0.03	0.00	0.00	10.53	11.63	19.83	183
2007							
Herd Buy	0.26	0.21	0.13	40.54	36.00	137.99	1783
Herd Sell	0.21	0.18	0.10	9.65	40.55	169.09	2104
2008							
Herd Buy	0.21	0.17	0.10	52.15	14.49	51.64	865
Herd Sell	0.25	0.21	0.11	17.13	23.59	80.48	1158

Description: The first three columns show the proportion of days in which the percentage of trading periods with herd behavior was higher than 5, 10 and 25. Column 4 reports the mean period of the day in which herd behavior first occurred, while the last three columns report mean, standard deviation and maximum for the number of consecutive trading periods in which herd behavior was present in data.

The left part of table 4 suggests that herding periods were most represented in data during the stock crash in 2007-2008. For instance, in 25% (21%) of all event days in 2008, the estimate of the threshold σ_t^d (β_t^d) was above (beneath) 0.5 in at least 5% of all trading periods, compared to 2% (2%) of event days in 2005. The numbers reported in the last four columns show that herding was also more pervasive during trading in 2007 and 2008.

The mean period of the day in which herding first occurred ranges between 7.85 (2005) and 51.15 (2007) for herd buying and between 9.12 (2005) and 17.13 (2009) for herd selling³⁰. The mean length of consecutive trading periods in which herd selling behavior was present during a trading day was between 11.63 (2006) and 40.55 (2007)—those numbers are quite variable according to standard deviations. Finally, a major increase in the maximum number of consecutive trading periods with herd selling behavior is detected after 2006.

The proportion of herd buyers at time t of day d , $\Psi_{\beta_t^d}$, was briefly discussed in subsection 4.1. This proportion was defined as the ratio between the measure of herd buying and the potential measure of herd buying

$$\begin{aligned}\Psi_{\beta_t^d} &= \frac{P(Z_t^d \in [\beta_t^d, 1/2] | h^d)}{P(Z_t^d < 1/2 | h^d)} \\ &= \frac{\left(\tilde{\delta}_t^d \tau_H - (1 - \tilde{\delta}_t^d) \tau_L\right) \left(\frac{1}{4} - (\beta_t^d)^2\right) + \left(\tilde{\delta}_t^d (1 - \tau_H) + (1 - \tilde{\delta}_t^d) (1 + \tau_L)\right) \left(\frac{1}{2} - \beta_t^d\right)}{\left(\tilde{\delta}_t^d \tau_H - (1 - \tilde{\delta}_t^d) \tau_L\right) \frac{1}{4} + \left(\tilde{\delta}_t^d (1 - \tau_H) + (1 - \tilde{\delta}_t^d) (1 + \tau_L)\right) \frac{1}{2}}\end{aligned}$$

if $\beta_t^d < 1/2$ and 0 otherwise³¹. The proportion of herd buyers of day d is then computed by the daily average of $\Psi_{\beta_t^d}$, $t = 1, 2, \dots$, using model estimates. Let $\Psi_{\beta}^{d,y}$ ($\Psi_{\sigma}^{d,y}$) denote the proportion of herd buyers (sellers) of day d in year $y \in \{2005, \dots, 2008\}$. Table 5 presents estimates (averages) of the proportion of herders for each trading year³².

³⁰For instance, based on the choice of the no-trade interval, herd selling behavior was on average present in data after one and a half minutes of trading during a trading day in 2007.

³¹The proportion of herd sellers at time t of day d is given by

$$\Psi_{\sigma_t^d} = \frac{\left(\tilde{\delta}_t^d \tau_H - (1 - \tilde{\delta}_t^d) \tau_L\right) \left((\sigma_t^d)^2 - \frac{1}{4}\right) + \left(\tilde{\delta}_t^d (1 - \tau_H) + (1 - \tilde{\delta}_t^d) (1 + \tau_L)\right) \left(\sigma_t^d - \frac{1}{2}\right)}{\left(\tilde{\delta}_t^d \tau_H - (1 - \tilde{\delta}_t^d) \tau_L\right) \frac{3}{4} + \left(\tilde{\delta}_t^d (1 - \tau_H) + (1 - \tilde{\delta}_t^d) (1 + \tau_L)\right) \frac{1}{2}}$$

if $\sigma_t^d > 1/2$ and 0 otherwise.

³²The average of the daily proportion of herders is not equal to the total average of proportion of herders due to the differing number of trades within days. However, non-reported results show that the two averages are almost identical.

Table 5: SUMMARY OF PROPORTION OF HERDERS

	Event Days			Good-Event Days	Bad-Event Days
	Average	SD	Max	Average	Average
2005					
Herd Buy	0.004	0.011	0.089	0.003	0.010
Herd Sell	0.003	0.010	0.081	0.002	0.008
2006					
Herd Buy	0.006	0.033	0.308	0.009	0.001
Herd Sell	0.002	0.006	0.054	0.002	0.002
2007					
Herd Buy	0.055	0.130	0.629	0.069	0.039
Herd Sell	0.053	0.140	0.649	0.012	0.098
2008					
Herd Buy	0.035	0.104	0.556	0.048	0.027
Herd Sell	0.046	0.111	0.536	0.044	0.048

Description: Yearly estimates of the proportion of herd buyers and the proportion of herd sellers. The left part of the table reports the average, standard deviation and maximum for all event days. The last two columns show the average during good-event days and bad-event days, respectively.

The left part of table 5 shows that on average, the proportion of herd buyers ranges between 0.4% (2004) and 5.5% (2007) across years, while the proportion of herd sellers is between 0.2% (2006) and 5.3% (2007). In particular, both the proportion of herd buyers and the proportion of herd sellers are highest during trading periods in 2007 and 2008, but vary considerably across days (e.g., for 2008 the standard deviation of the proportion of herd sellers was 14%, while the maximum was 64.9%). The two last columns reporting the proportion of herders on good-event days and bad-event days indicate that misdirected herding occurs. For instance, the proportion of herd sellers was 4.4% on good-event days in 2008.

For further comparison of the extent of herding between trading years, the log-transformation of the proportion of herders is considered, $\log(\Psi_\beta^{d,y})$ and $\log(\Psi_\sigma^{d,y})$, with $d = 1, 2, \dots, D_y$ and $y = 2005, \dots, 2008$ ³³. It is then

³³The study is then restricted to event days where the proportion of herd buyers and the proportion of herd sellers were positive. The log-samples for the proportion of herd buyers and the proportion of herd sellers then represent 89% and 82%, respectively, of the original samples. This also shows that the majority of event days are associated with herding.

assumed that the level of log-proportion of herders is the same within years but may differ across years, which is modeled with the linear factor model (model assumptions are discussed and accepted in internet appendix I.A.2). More specifically, for the log-proportion of herd sellers it is assumed that

$$\log(\Psi_{\sigma}^{d,y}) = \gamma_y + \epsilon_{d,y} \quad (7.1)$$

for $d = 1, \dots, D_y$, $y = 2005, \dots, 2008$ and $\epsilon_{d,y} \sim \mathcal{N}(0, \Sigma)$ where $\epsilon_{1,2005}, \dots, \epsilon_{D,2008}$ are independently distributed. This model is also assumed for herd buyers. Table 6 reports model fits with the 2007 sample used as reference.

Table 6: TEST STATISTICS FOR THE LOG-PROPORTION OF HERDERS (I)

(a) Herd buyers				
	Estimate	Std. Error	t value	Pr(> t)
2007 (Intercept)	-6.15	0.33	-18.74	< .01
2005	-0.67	0.44	-1.52	0.13
2006	-1.67	0.45	-3.69	< .01
2008	-0.89	0.48	-1.87	0.06
(b) Herd Sellers				
	Estimate	Std. Error	t value	Pr(> t)
2007 (Intercept)	-6.19	0.36	-17.39	< .01
2005	-0.99	0.48	-2.07	0.04
2006	-1.20	0.49	-2.46	0.01
2008	-0.20	0.49	-0.42	0.68

Description: Test statistics for the log-proportion of herd buyers (upper table) and herd sellers (lower table) with the "2007" group used as reference.

The first column reports estimates of the mean value across groups. For instance, the mean value for log-proportion of herd sellers is estimated at -6.19 for the 2007 sample and $-6.19 - 0.99 = -7.18$ for the 2005 sample, respectively. The exponential function of those estimates corresponds to estimates of the median of the proportion of herders for each sample³⁴.

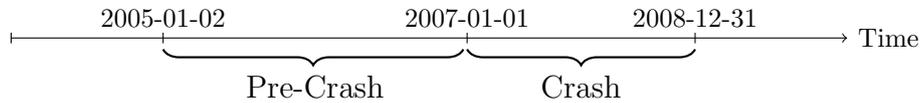
³⁴If $\log \Psi \sim \mathcal{N}(\bar{\Psi}, \Sigma)$, then $\bar{\Psi}$ is also the median of $\log \Psi$: $P(\Psi \leq e^{\bar{\Psi}}) = 0.5$.

Both subtables indicate that the proportion of herders was highest in 2007, consistent with the findings reported in table 4 and table 5 respectively. For herd buyers it is only statistically significant (at the 0.01-level) that the level of herding between 2006 and 2007 differs. However, estimates and P -values strongly suggest that herd buying behavior was most prominent in 2007. For herd sellers it is statistically significant (at the 0.05-level) that the level of herding in 2007 differs from the levels in 2005 and 2006. The estimate and P -value for the 2008 group indicate that the levels of herd selling in 2007 and 2008 were close to each other.

Summarizing, herding periods within trading days were present more frequently in 2007 and 2008 than they were in 2005 and 2006. The proportion of both herd buyers and herd sellers also increased in 2007 and 2008.

7.2. Herding during the Stock Crash

Previous results provide evidence that the extent of herd selling increased during and subsequent to 2007. To investigate this matter, the daily proportions of herders, $\Psi_{\beta}^{d,y}$ and $\Psi_{\sigma}^{d,y}$, $d = 1, \dots, D_y$ and $y = 2005, \dots, 2008$, are now placed in two groups. The first group associated with trading in 2005 and 2006 (182 event days) is labeled "Pre-Crash," while the second group associated with trading in 2007 and 2008 (150 event days) is labeled "Crash"³⁵. Formally, introduce the factor $p : \{2005, \dots, 2008\} \rightarrow \{\text{"Pre-Crash"}, \text{"Crash"}\}$ with $p(y) = \text{"Pre-Crash"}$ for $y \in \{2005, 2006\}$ and $p(y) = \text{"Crash"}$ for $y \in \{2007, 2008\}$ such that $\Psi_{\beta}^{d',y'}$ and $\Psi_{\beta}^{d'',y''}$ belongs to the same group if $p(y') = p(y'')$. The further investigation is then based on a pairwise comparison between the two periods.



³⁵In accordance with figure 6, the entire trading year 2008 is characterized by the stock crash, whereas 2007 contains periods associated with both the crash and less troubled times. Consequently, it is natural to refer to the 2005-2006 group as "Pre-Crash" and the 2007-2008 group as "Crash". The "Crash" period is of course arbitrary, but non-reported results with other (reasonable) choices of this period show that the main findings are the same.

Table 13 and table 14 in appendix C report, respectively, herding frequencies and a summary of the proportion of herders over the two periods. The findings suggest that the extent of herding was highest during the stock crash. The nonparametric Mann-Whitney U test is used to detect significant differences between two samples when the normal assumption is not satisfied. Let $\bar{\Psi}_p$ denote a randomly selected value from group $p \in \{PC, C\}$ (abbreviations for the two groups) and consider the null hypothesis $P(\bar{\Psi}_C > \bar{\Psi}_{PC}) = P(\bar{\Psi}_{PC} > \bar{\Psi}_C)$ with alternative hypothesis $P(\bar{\Psi}_C > \bar{\Psi}_{PC}) \neq P(\bar{\Psi}_{PC} > \bar{\Psi}_C)$. Test statistics are reported in table 7.

Table 7: SUMMARY OF MANN-WHITNEY U TEST STATISTICS

	Herd Buyers		Herd Sellers	
	w	Pr	w	Pr
$H : \Pr(\bar{\Psi}_C > \bar{\Psi}_{PC}) = \Pr(\bar{\Psi}_{PC} > \bar{\Psi}_C)$	14412.0	0.381	16151.0	< .01

Description: Test statistics for the proportion of herders. Test sizes (w) and corresponding P -values for the nonparametric Mann-Whitney U test are reported.

For the proportion of herd buyers, the null hypothesis is accepted. For the proportion of herd sellers, the hypothesis is rejected at the 0.01-level. This suggests that only the proportion of herd sellers changed during the stock crash.

Again, the linear model is assumed where the level of log-proportion of herders depends on the period

$$\log(\Psi_{\sigma}^{d,y}) = \gamma_{p(y)} + \epsilon_{d,y} \quad (7.2)$$

for $d = 1, \dots, D_y$ and $y = 2005, \dots, 2008$, where $\epsilon_{d,y} \sim \mathcal{N}(0, \Sigma)$ and $\epsilon_{1,2005}, \dots, \epsilon_{D_{2008},2008}$ are independently distributed. This model specification may be tested against the model specified by equation (7.1) using the statistical F -test, which strongly supports the presumption that only herd selling increased in 2007-2008³⁶. Consequently, only the model fit for the log-proportion of herd sellers is reported.

³⁶The F -test yields a test size equal to 0.198 (4.484), with corresponding P -value equal to 0.821 (0.012) for herd sellers (buyers).

Table 8: TEST STATISTICS FOR THE LOG-PROPORTION OF HERDERS (I)

(a) Herd Sellers				
	Estimate	Std. Error	t value	Pr(> t)
Pre-Crash (Intercept)	-7.28	0.23	-31.87	< .001
Crash	0.98	0.33	2.93	0.004

Description: Test statistics for the log-proportion of herd sellers in 2005-2006 and 2007-2008.

The first column reporting estimates of the mean value for the two periods ("Pre-Crash" as reference) indicates that herd selling increased during the stock crash. The two last columns report t -test statistics of the null hypothesis (the mean values of the two samples are equal). The hypothesis is rejected at the 0.01-level, and the alternative hypothesis is accepted: The log-proportion of herd sellers increased during 2007-2008, consistent with a stock market crash.

7.3. Herding and Market Uncertainty

Finally, one last factor is presumed to be associated with herding in the stock: market-wide uncertainty proxied by the VIX. Comparing empirical findings of herding in the stock with the daily VIX values illustrated in figure 6 suggests a positive dependence between herding in the stock and the market's expectation of volatility. To investigate this matter, the level of daily log-proportion of herders is assumed to depend on the VIX-level. The online Bloomberg article *What History Says About Low Volatility* by Brown (2017) states that a VIX value above 0.2 must be considered high, which is in accordance with the empirical distribution of VIX values in the sample period 2005 – 2008³⁷. Therefore, the daily log-proportion of herders is placed in the "High VIX" group if corresponding daily VIX values were above 0.2, and remaining values are placed in the "Low VIX" group³⁸. Table 9 presents statistics for the two model fits.

³⁷For the period 01/03/2005 – 12/31/2010, the median and the average of daily closing VIX values were 0.14 and 19.0, respectively.

³⁸Non-reported results show that the conclusion of the findings is the same when the thresholds are equal to 0.15, 0.175 and 0.225 respectively.

Table 9: TEST STATISTICS FOR THE LOG-PROPORTION OF HERDERS (III)

(a) Herd buyers				
	Estimate	Std. Error	t value	Pr(> t)
High VIX (Intercept)	-7.24	0.29	-24.85	< .01
Low VIX	0.38	0.35	1.08	0.28
(b) Herd Sellers				
	Estimate	Std. Error	t value	Pr(> t)
High VIX (Intercept)	-6.18	0.29	-21.26	< .01
Low VIX	-0.96	0.36	-2.69	0.01

Description: Test statistics for the log-proportion of herd buyers (upper table) and herd sellers (lower table) where the group "High VIX" is used as reference.

The VIX-level has an opposite effect on herd buying and herd selling respectively. During trading days with low market volatility, herd buying increases, while herd selling decreases. However, test statistics for herd buyers indicate that the level of herding is not affected by the VIX. This is not the case for the log-proportion of herd sellers: It is statistically significant (at the 0.05-level) that the log-proportion of herd sellers increased in periods with a high VIX, suggesting that market-wide uncertainty affects the extent of herd selling in stocks. The findings are consistent with some of the criticisms of herding, accusing investors of trading against their private information during tumultuous periods, which in the worst case can instigate stock crashes. Such a claim is somewhat supported by the empirical findings in table 5, reporting that the amount of misdirected herd selling behavior was highest in 2008 — a trading year associated with extreme uncertainty³⁹.

8. Conclusion

The paper described a structural microstructure model in which informed traders and noise traders sequentially trade one unit of a risky asset in a financial market repeatedly over many days.

³⁹The proportion of herd sellers was 4.4% during good-event days in 2008.

The fundamental asset value changes on a daily basis, and the informed trader's action depends on the realization of a noisy private signal of the unknown asset value. This results in asymmetric information between the learning market maker and privately informed traders.

The market maker also faces event uncertainty, resulting in a positive probability of intraday herding for privately informed traders. In particular, herd buying (selling) behavior may be expressed through a buying (selling) threshold given as a function of model parameters. Inspired by Ottaviani and Sørensen (2006), the noisy signal satisfies 1) bounded private beliefs and 2) conditional linear density functions where the trader's ability to gain from private information may depend on the state of the world. Consequently, the signal distribution exhibits asymmetries allowing for separately identifying the extent of herding on the buy and sell sides of the market.

Conditional linear densities are assumed to obtain tractability. The model specification implies closed-form solutions for the two thresholds and thereby a closed-form expression for the recursive likelihood function. The empirical study used financial tick data for the NYSE stock Ashland Inc. traded from 2005 to 2008, including tranquil as well as more tumultuous periods on the stock market. Model calibrations confirm that the signal distribution is asymmetric, and parameter estimates are used to compute public as well as private beliefs, which in turn are used to estimate rational statistical herding. The extent of herd behavior varies across years. In particular, herd selling behavior increased in 2007 and 2008, consistent with a stock market crash. These findings are supported by various test statistics. Finally, market-wide uncertainty in terms of the VIX-index is identified as a factor associated with herd selling behavior: Informed investors were more likely to ignore private beliefs during high volatility periods, and herd selling behavior increased.

For future research, it will be interesting to conduct the same experiment for a large number of stocks. Stock characteristics such as liquidity may have an impact on the extent of herding in the stock. Finally, it is appealing to apply the same methodology to cryptocurrency markets, which are associated with herding and extreme price movements.

Appendix A Proofs

Without loss of generality, write $V_d^H - v_{d-1} = \lambda_d^H > 0$ and $v_{d-1} - V_d^L = \lambda_d^L > 0$. As a result, $\lambda_d^L/\lambda_d^H = \delta/(1 - \delta)$ by the martingale assumption.

PROOF OF PROPOSITION 4. First, the threshold β_t^d is computed (i.e., equation (3.6) is derived). The market maker's ask price is expressed by equation (3.1). Assuming continuous state-contingent implies

$$\begin{aligned} E(V_d | h_{t-1}^d, Z_t^d = \beta_t^d) &= E(V_d | h_{t-1}^d, X_t^d = B) \Leftrightarrow \\ P(V_d = V_d^H | h_{t-1}^d, Z_t^d = \beta_t^d) - P(V_d = V_d^H | h_{t-1}^d, X_t^d = B) \\ &= \frac{\delta}{1 - \delta} \left(P(V_d = V_d^L | h_{t-1}^d, \beta_t^d) - P(V_d = V_d^L | h_{t-1}^d, X_t^d = B) \right), \end{aligned} \quad (\text{A.1})$$

since $\lambda_d^L/\lambda_d^H = \delta/(1 - \delta)$. Realize that equation (A.1) and equation (3.6) are equivalent: Apply Bayes' theorem to express the informed trader's private belief of the good-event by

$$P(V_d = V_d^H | h_{t-1}^d, Z_t^d = \beta_t^d) = \frac{f(\beta_t^d | \tau_H, H) \tilde{\delta}_t}{f(\beta_t^d | \tau_H, H) \tilde{\delta}_t + f(\beta_t^d | \tau_L, L)(1 - \tilde{\delta}_t)}, \quad (\text{A.2})$$

with $\tilde{\delta}_t = P(V_d = V_d^H | h_{t-1}^d, V_d \neq v_{d-1})$ and $1 - \tilde{\delta}_t = P(V_d = V_d^L | h_{t-1}^d, V_d \neq v_{d-1})$. Now, use Bayes' theorem to compute the probability of the good-event conditional on history of trades and a potential buy order at time t :

$$P(V_d = V_d^H | h_t^d) = \frac{P(X_t^d = B | h_{t-1}^d, V_d = V_d^H) \tilde{\alpha}_t \tilde{\delta}_t}{\sum_{j \in \{V^H, V^L, v\}} P(X_t = B | h_{t-1}^d, V = j) P(j | h_{t-1}^d)}, \quad (\text{A.3})$$

with $\tilde{\alpha}_t \tilde{\delta}_t = P(V^H | h_{t-1})$, $\tilde{\alpha}_t(1 - \tilde{\delta}_t) = P(V^L | h_{t-1})$ and $(1 - \tilde{\alpha}_t) = P(V_d = v_{d-1} | h_{t-1})$. In a similar way to (A.2) and (A.3), $P(V^L | h_{t-1}, \beta_t)$ and $P(V^L | h_{t-1}, X_t = B)$ are computed. To avoid notational clutter, omit the h_{t-1} notation and realize that equation (A.1) can be written as

$$\begin{aligned} (1 - \delta) (P(V_d = V_d^H | \beta_t^d) - P(V_d = V_d^H | X_t^d = B)) \\ = \delta (P(V_d = V_d^L | \beta_t^d) - P(V_d = V_d^L | X_t^d = B)) \end{aligned}$$

or equivalent

$$\begin{aligned} 0 &= \frac{f(\beta_t^d | \tau_H, H) \tilde{\delta}_t (1 - \delta) - f(\beta_t^d | \tau_L, L)(1 - \tilde{\delta}_t) \delta}{f(\beta_t^d | \tau_H, H) \tilde{\delta}_t + f(\beta_t^d | \tau_L, L)(1 - \tilde{\delta}_t)} \\ &\quad - \frac{P(X_t^d = B | V_d = V_d^H) \tilde{\alpha}_t \tilde{\delta}_t (1 - \delta) - P(X_t^d = B | V^L) \tilde{\alpha}_t (1 - \tilde{\delta}_t) \delta}{P(X_t^d = B | V_d = V_d^H) \tilde{\alpha}_t \tilde{\delta}_t + P(X_t^d = B | V_d = V_d^L) \tilde{\alpha}_t (1 - \tilde{\delta}_t) + \epsilon/2(1 - \tilde{\alpha}_t)}. \end{aligned}$$

In particular,

$$\begin{aligned}
0 &= \left(f(\beta_t | \tau_H, H) \tilde{\delta}_t (1 - \delta) - f(\beta_t^d | \tau_L, L) (1 - \tilde{\delta}_t) \delta \right) \times \\
&\quad \left\{ P(X_t^d = B | V_d = V_d^H) \tilde{\alpha}_t \tilde{\delta}_t + P(X_t^d = B | V_d = V_d^L) \tilde{\alpha}_t (1 - \tilde{\delta}_t) + \frac{\epsilon}{2} (1 - \tilde{\alpha}_t) \right\} \\
&\quad - \left(f(\beta_t | \tau_H, H) \tilde{\delta}_t + f(\beta_t | \tau_L, L) (1 - \tilde{\delta}_t) \right) \times \\
&\quad \left\{ P(X_t^d = B | V_d = V_d^H) \tilde{\alpha}_t \tilde{\delta}_t (1 - \delta) - P(X_t^d = B | V_d = V_d^L) \tilde{\alpha}_t (1 - \tilde{\delta}_t) \delta \right\}.
\end{aligned}$$

After some manipulations, the threshold-equation is a decomposition of the 10 terms,

$$0 = \tilde{\alpha}_t \tilde{\delta}_t^2 (1 - \delta) f(\beta_t^d | \tau_H, H) P(X_t^d = B | V_d = V_d^H) \quad (\text{A.4})$$

$$\begin{aligned}
&+ \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t (1 - \delta)) f(\beta_t^d | \tau_H, H) P(X_t^d = B | V_d = V_d^L) \\
&+ (1 - \tilde{\alpha}_t) \tilde{\delta}_t (1 - \delta) \frac{\epsilon}{2} f(\beta_t | \tau_H, H) \\
&- \tilde{\alpha}_t \tilde{\delta}_t \delta (1 - \tilde{\delta}_t) f(\beta_t^d | \tau_L, L) P(X_t^d = B | V_d = V_d^H) \\
&- \tilde{\alpha}_t \delta (1 - \tilde{\delta}_t)^2 f(\beta_t^d | \tau_L, L) P(X_t^d = B | V_d = V_d^L) \quad (\text{A.5})
\end{aligned}$$

$$\begin{aligned}
&- (1 - \tilde{\alpha}_t) (1 - \tilde{\delta}_t) \delta \frac{\epsilon}{2} f(\beta_t^d | \tau_L, L) \\
&- \tilde{\alpha}_t \tilde{\delta}_t^2 (1 - \delta) f(\beta_t^d | \tau_H, H) P(X_t^d = B | V_d = V_d^H) \quad (\text{A.6})
\end{aligned}$$

$$\begin{aligned}
&- \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t) (1 - \delta) f(\beta_t^d | \tau_L, L) P(X_t^d = B | V_d = V_d^H) \\
&+ \tilde{\alpha}_t \tilde{\delta}_t \delta (1 - \tilde{\delta}_t) f(\beta_t^d | \tau_H, H) P(X_t^d = B | V_d = V_d^L) \\
&+ \tilde{\alpha}_t \delta (1 - \tilde{\delta}_t)^2 f(\beta_t^d | \tau_L, L) P(X_t^d = B | V_d = V_d^L). \quad (\text{A.7})
\end{aligned}$$

Now, four terms cancel out: (A.4)–(A.6) and (A.5)–(A.7). In particular,

$$\begin{aligned}
0 &= \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t (1 - \delta)) f(\beta_t^d | \tau_H, H) P(X_t^d = B | h_{t-1}^d, V_d = V_d^L) \\
&+ (1 - \tilde{\alpha}_t) \tilde{\delta}_t (1 - \delta) \frac{\epsilon}{2} f(\beta_t^d | \tau_H, H) \\
&- \tilde{\alpha}_t \tilde{\delta}_t \delta (1 - \tilde{\delta}_t) f(\beta_t^d | \tau_L, L) P(X_t^d = B | h_{t-1}^d, V_d = V_d^H) \\
&- (1 - \tilde{\alpha}_t) (1 - \tilde{\delta}_t) \delta \frac{\epsilon}{2} f(\beta_t^d | \tau_L, L) \\
&- \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t) (1 - \delta) f(\beta_t^d | \tau_L, L) P(X_t^d = B | h_{t-1}^d, V_d = V_d^H) \\
&+ \tilde{\alpha}_t \tilde{\delta}_t \delta (1 - \tilde{\delta}_t) f(\beta_t^d | \tau_H, H) P(X_t^d = B | h_{t-1}^d, V_d = V_d^L) \\
&= \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t) \left\{ f(\beta_t^d | \tau_H, H) P(X_t^d = B | h_{t-1}^d, V_d = V_d^L) \right. \\
&\quad \left. - f(\beta_t^d | \tau_L, L) P(X_t^d = B | h_{t-1}^d, V_d = V_d^H) \right\} \\
&+ (1 - \tilde{\alpha}_t) \frac{\epsilon}{2} \left(\tilde{\delta}_t (1 - \delta) f(\beta_t^d | \tau_H, H) - \delta (1 - \tilde{\delta}_t) f(\beta_t^d | \tau_L, L) \right).
\end{aligned}$$

The first part of proposition 4 is proved. On the other hand, using the theoretical expression for the bid price as well as equation (3.2), equation (3.7) is obtained

$$\begin{aligned}
0 = & \tilde{\alpha}_t \tilde{\delta}_t (1 - \tilde{\delta}_t) \left\{ f(\sigma_t^d | \tau_H, H) P(X_t^d = S | V_d = V_d^L, h_{t-1}^d) \right. \\
& \left. - f(\sigma_t^d | \tau_L, L) P(X_t^d = S | V_d = V_d^H, h_{t-1}^d) \right\} \\
& + (1 - \tilde{\alpha}_t) \frac{\epsilon}{2} \left(\tilde{\delta}_t (1 - \delta) f(\sigma_t^d | \tau_H, H) - \delta (1 - \tilde{\delta}_t) f(\sigma_t^d | \tau_L, L) \right). \quad (\text{A.8})
\end{aligned}$$

The proposition is proved. \square

PROOF OF COROLLARY 5. The corollary is proved for the general case $\tau_L, \tau_H \in (0, \infty)$, allowing private beliefs to be unbounded. In this case, the support A_ω is chosen such that $f(\cdot | \tau_\omega, \omega)$ becomes a probability density for $\omega \in \{H, L\}$,

$$\begin{aligned}
A_H &= \begin{cases} [0, 1] & \text{if } \tau_H \leq 1 \\ \left[\frac{\tau_H - 1}{2\tau_H}, \frac{\tau_H - 1 + 2\sqrt{\tau_H}}{2\tau_H} \right] & \text{if } \tau_H > 1 \end{cases} \\
A_L &= \begin{cases} [0, 1] & \text{if } \tau_L \leq 1 \\ \left[\frac{\tau_L + 1 - 2\sqrt{\tau_L}}{2\tau_L}, \frac{\tau_L + 1}{2\tau_L} \right] & \text{if } \tau_L > 1 \end{cases}.
\end{aligned}$$

Let $a_\omega(\tau_\omega) = \inf_{\tau_\omega > 0} A_\omega$ and $b(\tau_\omega) = \sup_{\tau_\omega > 0} A_\omega$ for $\omega \in \{H, L\}$ and realize that

$$\begin{aligned}
& f(\beta_t^d | \tau_H, H) P(X_t^d = B | h_{t-1}^d, V_d = V_d^L) - f(\beta_t^d | \tau_L, L) P(X_t^d = B | h_{t-1}^d, V_d = V_d^H) \\
&= -(\tau_H + \tau_L) \mu (\beta_t^d)^2 + \left\{ 2\mu \tau_H \tau_L \left(- (b_L(\tau_L))^2 - b_H(\tau_H)^2 \right) + \left(\frac{b_L(\tau_L)}{\tau_L} + \frac{b_H(\tau_H)}{\tau_H} \right) \right. \\
&\quad \left. + (b_L(\tau_L) - b_H(\tau_H)) + (1 - \mu) \epsilon (\tau_H + \tau_L) \right\} \beta_t^d \\
&\quad + \left\{ \mu \left(- (\tau_L b_L(\tau_L))^2 + \tau_H b_H(\tau_H)^2 \right) + (b_L(\tau_L) - b_H(\tau_H)) \right. \\
&\quad \left. + (\tau_L b_L(\tau_L) + \tau_H b_H(\tau_H)) + \tau_H \tau_L (b_L(\tau_L)^2 - b_H(\tau_H)^2) \right. \\
&\quad \left. - (\tau_H b_L(\tau_L) + \tau_L b_H(\tau_H)) - \tau_H \tau_L (b_L(\tau_L) - b_H(\tau_H)) \right\} - (1 - \mu) \frac{\epsilon}{2} (\tau_H + \tau_L),
\end{aligned}$$

which is recognized as a quadratic equation in β_t^d . In addition, see that

$$\begin{aligned}
& \tilde{\delta}_t^d (1 - \delta) f(\beta_t^d | \tau_H, H) - \delta (1 - \tilde{\delta}_t^d) f(\beta_t^d | \tau_L, L) \\
&= (1 - \tilde{\alpha}_t^d) \frac{\epsilon}{2} \left(2(\delta \tau_L + \tilde{\delta}_t^d \tau_H - \delta \tilde{\delta}_t^d (\tau_H + \tau_L)) \beta_t^d + (\tilde{\delta}_t^d - \delta - \tilde{\delta}_t^d \tau_H - \delta \tau_L + \delta \tilde{\delta}_t^d (\tau_H + \tau_L)) \right),
\end{aligned}$$

which is a linear equation in β_t^d . Combining the above two expressions

implies that equation (3.6) can be written as

$$0 = B_{\beta_t^d}(\theta, h_{t-1}^d)(\beta_t^d)^2 + C_{\beta_t^d}(\theta, h_{t-1}^d)\beta_t^d + D_{\beta_t^d}(\theta, h_{t-1}^d),$$

where

$$\begin{aligned} B_{\beta_t^d}(\theta, h_{t-1}^d) &= -\tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \mu (\tau_H + \tau_L), \\ C_{\beta_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \left\{ 2\mu \tau_H \tau_L \left(- (b_L(\tau_L))^2 - b_H(\tau_H)^2 \right) + \left(\frac{b_L(\tau_L)}{\tau_L} + \frac{b_H(\tau_H)}{\tau_H} \right) \right. \\ &\quad \left. + (b_L(\tau_L) - b_H(\tau_H)) + (1 - \mu) \epsilon (\tau_H + \tau_L) \right\} \\ &\quad + (1 - \tilde{\alpha}_t^d) \epsilon \left(\delta \tau_L + \tilde{\delta}_t^d \tau_H - \delta \tilde{\delta}_t^d (\tau_H + \tau_L) \right), \\ D_{\beta_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \left\{ \mu \left(- (\tau_L b_L(\tau_L))^2 + \tau_H b_H(\tau_H)^2 \right) + (b_L(\tau_L) - b_H(\tau_H)) \right. \\ &\quad \left. + (\tau_L b_L(\tau_L) + \tau_H b_H(\tau_H)) + \tau_H \tau_L (b_L(\tau_L)^2 - b_H(\tau_H)^2) \right. \\ &\quad \left. - (\tau_H b_L(\tau_L) + \tau_L b_H(\tau_H)) - \tau_H \tau_L (b_L(\tau_L) - b_H(\tau_H)) \right) \\ &\quad \left. - (1 - \mu) \frac{\epsilon}{2} (\tau_H + \tau_L) \right\} \\ &\quad + (1 - \tilde{\alpha}_t^d) \frac{\epsilon}{2} \left(\tilde{\delta}_t^d - \delta - \tilde{\delta}_t^d \tau_H - \delta \tau_L + \delta \tilde{\delta}_t^d (\tau_H + \tau_L) \right). \end{aligned}$$

In a similar way, starting from equation (3.7), the quadratic equation for σ_t^d is derived

$$0 = B_{\sigma_t^d}(\theta, h_{t-1}^d)(\sigma_t^d)^2 + C_{\sigma_t^d}(\theta, h_{t-1}^d)\sigma_t^d + D_{\sigma_t^d}(\theta, h_{t-1}^d),$$

with

$$\begin{aligned}
B_{\sigma_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \mu (\tau_H + \tau_L), \\
C_{\sigma_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \left\{ 2\mu \left(\tau_H \tau_L (a_L(\tau_L)^2 - a_H(\tau_H)^2) - (\tau_H a_L(\tau_L) + \tau_L a_H(\tau_H)) \right. \right. \\
&\quad \left. \left. + \tau_H \tau_L (a_H(\tau_H) - a_L(\tau_L)) \right) + (\tau_H + \tau_L)(1 - \mu)\epsilon \right\} \\
&\quad + (1 - \tilde{\alpha}_t^d) \epsilon \left(\delta \tau_L + \tilde{\delta}_t^d \tau_H - \delta \tilde{\delta}_t^d (\tau_H + \tau_L) \right), \\
D_{\sigma_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \left\{ \mu \left((\tau_L a_L(\tau_L)^2 + \tau_H a_H(\tau_H)^2) + (a_H(\tau_H) - a_L(\tau_L)) \right. \right. \\
&\quad \left. \left. - (\tau_L a_L(\tau_L) + \tau_H a_H(\tau_H)) - \tau_H \tau_L (a_L(\tau_L)^2 - a_H(\tau_H)^2) \right. \right. \\
&\quad \left. \left. + (\tau_H a_L(\tau_L) + \tau_L a_H(\tau_H)) - \tau_H \tau_L (a_H(\tau_H) - a_L(\tau_L)) \right) \right. \\
&\quad \left. - (\tau_H + \tau_L)(1 - \mu) \frac{\epsilon}{2} \right\} \\
&\quad + (1 - \tilde{\alpha}_t^d) \frac{\epsilon}{2} \left(\tilde{\delta}_t^d - \delta - \tilde{\delta}_t^d \tau_H - \delta \tau_L + \delta \tilde{\delta}_t^d (\tau_H + \tau_L) \right).
\end{aligned}$$

In the special case $\tau_L, \tau_H \in (0, 1]$, the expressions boil down to

$$\begin{aligned}
B_{\beta_t^d}(\theta, h_{t-1}^d) &= -\tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \mu (\tau_L + \tau_H), \\
C_{\beta_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \{ 2\mu (\tau_L + \tau_H) + (1 - \mu)\epsilon (\tau_L + \tau_H) \} \\
&\quad + (1 - \tilde{\alpha}_t^d) \epsilon \left(\delta \tau_L + \tilde{\delta}_t^d \tau_H - \delta \tilde{\delta}_t^d (\tau_L + \tau_H) \right), \\
D_{\beta_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \left\{ -\mu (\tau_L + \tau_H) - (1 - \mu) \frac{\epsilon}{2} (\tau_L + \tau_H) \right\} \\
&\quad + (1 - \tilde{\alpha}_t^d) \frac{\epsilon}{2} \left(\tilde{\delta}_t^d - \delta - \tilde{\delta}_t^d \tau_H - \delta \tau_L + \delta \tilde{\delta}_t^d (\tau_H + \tau_L) \right), \\
B_{\sigma_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \mu (\tau_L + \tau_H), \\
C_{\sigma_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \{ (1 - \mu)\epsilon (\tau_L + \tau_H) \} \\
&\quad + (1 - \tilde{\alpha}_t^d) \epsilon \left(\delta \tau_L + \tilde{\delta}_t^d \tau_H - \delta \tilde{\delta}_t^d (\tau_L + \tau_H) \right), \\
D_{\sigma_t^d}(\theta, h_{t-1}^d) &= \tilde{\alpha}_t^d \tilde{\delta}_t^d (1 - \tilde{\delta}_t^d) \left\{ -(1 - \mu) \frac{\epsilon}{2} (\tau_L + \tau_H) \right\} \\
&\quad + (1 - \tilde{\alpha}_t^d) \frac{\epsilon}{2} \left(\tilde{\delta}_t^d - \delta - \tilde{\delta}_t^d \tau_H - \delta \tau_L + \delta \tilde{\delta}_t^d (\tau_L + \tau_H) \right).
\end{aligned}$$

The corollary is proved. \square

PROOF OF LEMMA 10. Consider an arbitrary history of trades h_t and use Bayes' theorem to realize that

$$P(V_d = v | h_t) = P(V_d = v | h_{t-1}) \frac{P(X_t^d | h_{t-1}, V_d = v)}{\sum_{v \in \{V_d^H, V_d^L, v_{d-1}\}} P(V_d = v | h_{t-1}) P(X_t^d | h_{t-1}, V_d = v)}.$$

For $X_t = N$, then $1 - \tilde{\alpha}_{t+1} > 1 - \tilde{\alpha}_t$ if and only if

$$\begin{aligned} 1 - \epsilon > (1 - \tilde{\alpha}_t)(1 - \epsilon) + \tilde{\alpha}_t \left\{ \tilde{\delta}_t \left(\mu \nu_t^H ((\sigma_t^d, \beta_t^d)) + (1 - \mu) \frac{\epsilon}{2} \right) \right. \\ \left. + (1 - \tilde{\delta}_t) \left(\mu \nu_t^L ((\sigma_t^d, \beta_t^d)) + (1 - \mu) \frac{\epsilon}{2} \right) \right\} \Leftrightarrow \\ 1 - \epsilon > \nu_t^d ((\sigma_t^d, \beta_t^d)), \end{aligned}$$

which is true under assumption 9. \square

PROOF OF PROPOSITION 11. The result is shown by contradiction for the case of herd buying. Assume that $V_d \neq v_{d-1}$ and consider an informed trader with the bad signal $Z_t^d \in [z^* - l, z^*)$ for some $l > 0$. Moreover, assume that herding occurs with probability 0 during the first K trades and consider the first $k_1 + k_2 < K$ periods of trades with $X_t^d = N$ for $t = 1, 2, \dots, k_1$ and $X_t^d = B$ for $t = k_1 + 1, k_1 + 2, \dots, k_1 + k_2$. The informed trader's belief is unchanged for the first k_1 periods. Choose k_2 and $l > 0$ such that

$$E(V_d | h_{k_1+k_2-1}^d, Z_{k_1+k_2}^d) > v_{d-1} + \varphi$$

for fixed $\varphi > 0$. By lemma 10, choose k_1 sufficiently large such that

$$(1 - \tilde{\alpha}_{k_1+1}) > 1 - \varphi \quad \Leftrightarrow \quad \varphi > \tilde{\alpha}_{k_1+1}.$$

Now, for k_1 and k_1/k_2 large and a fixed $\varphi > 0$, the following inequality is obtained

$$a_{k_1+k_2}^d < v_{d-1} + \varphi = E(V_d | h_{k_1+k_2-1}^d, Z_{k_1+k_2}^d),$$

and the trader will herd buy at time $t = k_1 + k_2 < B$. Since $P(h_{k_1+k_2}) > 0$ due to the presence of noise traders, the probability of herd buying is positive at time $k_1 + k_2 < K$, contradicting the assumption that herd buying occurred with probability 0 during the first K trades. \square

PROOF OF LEMMA 7. The lemma is proved for the case of herd buying behavior. The proof for herd selling behavior is analogous.

Consider a signal z_t^d such that $E(V_d | h_{t-1}^d, Z_t^d = z_t^d) > E(V_d | h_{t-1}^d, X_t = B)$ and $z_t^d < z^*$. Using assumption 2 to rewrite the inequality yields that $\tilde{\delta}_t^d > \delta$. Moreover, the inequality $E(V_d | h_{t-1}^d) > v_{d-1}$ can be rewritten as

$$\frac{\tilde{\delta}_t^d}{1 - \tilde{\delta}_t^d} > \frac{\lambda^L}{\lambda^H} = \frac{\delta}{1 - \delta},$$

which is trivially satisfied when $\tilde{\delta}_t^d > \delta$. The lemma is proved. \square

Appendix B Data Cleaning: MTAQ-File

Let a and b denote the ask price and the bid price, respectively. The MTAQ-file is treated in line with Holden and Jacobsen (2013)⁴⁰. Only trades and quotes with time stamps inside 9:30 am - 4:00 pm⁴¹ are considered.

For trades, entries are deleted if the transaction price is equal to zero. Corrected trades and trades with an abnormal sale condition are also removed from the dataset. Finally, only transactions originating from NYSE are used.

Only quotes with nonnormal quote conditions are considered. If the bid price b is greater than the ask price a for the same exchange, the entry is deleted if only both $a > 0$ and $b > 0$. If $b > 0$ and $a = 0$, the bid price is still considered valid since the ask is withdrawn. If the spread ($a - b$) is greater than 5 (dollars), the entry is deleted if only both $a > 0$ and $b > 0$. If $a > 0$ and $b = 0$, the ask price is valid and the bid has been withdrawn from the exchange. Finally, if the ask price or ask size equals zero or is missing, the ask is seen as withdrawn from the exchange. The same applies for the bid. If an exchange has two or more quotes with identical time stamps, the last quote is considered in-force for that exchange (order is preserved even though observations have the same time stamp). Using this approach, two datasets for quotes are constructed. First, the NBBO (National Best Bid and Offer) is constructed across exchanges, with one-second accuracy. The last dataset contains only quotes from NYSE.

Trades and quotes are then matched with a one-second delay (prior second rule), and transaction prices are compared to ask and bid prices. The proxies for trade location *at the ask/bid*, *inside the spread* and *outside the spread* are then used to assess the accuracy of the trades-and-quotes file. Empirical findings (at the ask/bid and outside the spread) suggest to use only quotes from the NYSE. Note: Empirical studies after 2008 should be based on the NBBO-file due to the major decrease in NYSE's market share in NYSE-listed stocks, reported by Angel et al. (2015).

⁴⁰A thorough description of the dataset is given in 'MONTHLY TAQ USER GUIDE': www.nyxdata.com/doc/2552.

⁴¹The time stamps for trades and quotes are measured in seconds after midnight.

Appendix C Empirical Findings

Table 10: STANDARD ERRORS OF ESTIMATES (MODEL SPECIFICATION I)

Year	s.e. $(\hat{\alpha})$	s.e. $(\hat{\delta})$	s.e. $(\hat{\mu})$	s.e. $(\hat{\tau}_H)$	s.e. $(\hat{\tau}_L)$	s.e. $(\hat{\epsilon})$
2005	1.52×10^{-7}	1.09×10^{-7}	2.36×10^{-7}	1.63×10^{-7}	3.28×10^{-7}	1.96×10^{-7}
2006	2.34×10^{-7}	6.46×10^{-7}	1.06×10^{-6}	1.45×10^{-6}	6.06×10^{-7}	1.39×10^{-7}
2007	6.56×10^{-7}	8.77×10^{-7}	7.48×10^{-7}	1.88×10^{-7}	8.07×10^{-7}	2.25×10^{-7}
2008	4.82×10^{-6}	1.32×10^{-5}	7.67×10^{-6}	7.56×10^{-6}	4.18×10^{-6}	2.40×10^{-7}

Description: Standard errors of the estimates reported in table 2. The numbers are computed using the BHHH-estimator.

Table 11: MAXIMUM LIKELIHOOD ESTIMATES (MODEL SPECIFICATION II)

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\tau}$	$\hat{\epsilon}$	$l(\hat{\theta})$
2005	0.208	0.729	0.320	0.335	0.516	933784
2006	0.412	0.743	0.318	0.283	0.533	961601
2007	0.395	0.445	0.247	0.299	0.550	824150
2008	0.289	0.344	0.281	0.350	0.568	835583

Description: Yearly maximum likelihood estimates under model specification II. The last column shows the log-likelihood values associated with the maximum likelihood estimates.

Table 12: LIKELIHOOD RATIO TESTS

	2005	2006	2007	2008
q	171.46	147.52	2.85	38.75
Pr(> q)	< .01	< .01	0.09	< .01

Description: The first row reports likelihood ratio test sizes for the hypothesis $H : \tau_H = \tau_L$ for each year. The second row reports corresponding P -values.

Table 13: SUMMARY OF HERDING FREQUENCIES

	> 5%	> 10%	> 25%	Mean Time	Mean Length	Length SD	Max Length
Pre-Crash							
Herd Buy	0.03	0.01	0.01	8.68	13.40	38.37	679
Herd Sell	0.03	0.01	0.00	9.78	12.32	19.27	183
Crash							
Herd Buy	0.23	0.19	0.11	46.03	21.86	91.52	1783
Herd Sell	0.23	0.19	0.11	13.62	30.00	122.01	2104

Description: The first three columns show the proportion of days in which the percentage of trading periods with herd behavior was higher than 5, 10 and 25. Column 4 reports the mean period of the day in which herd behavior first occurred, while the last three columns report mean, standard deviation and maximum for the number of consecutive trading periods in which herd behavior was present in data.

Table 14: SUMMARY OF DAILY PROPORTION OF HERDERS

	Event Days			Good-Event Days	Bad-Event Days
	Average	SD	Max	Average	Average
Pre-Crash					
Herd Buy	0.005	0.024	0.308	0.005	0.004
Herd Sell	0.003	0.008	0.081	0.002	0.004
Crash					
Herd Buy	0.045	0.118	0.629	0.061	0.032
Herd Sell	0.049	0.127	0.649	0.024	0.070

Description: Yearly estimates of the proportion of herd buyers and the proportion herd sellers. The left part of the table reports the results average, standard deviation and maximum for all event days. The last two columns show the average on good-event days and bad-event days, respectively.

Internet Appendix

I.A.1. Duration Times and the No-Trade Interval

Let $D_i^d := t_i^d - t_{i-1}^d$ denote the duration in seconds between trades i and $i - 1$ of day d . The no-trade interval is based on the empirical distribution of the D_i^d 's. Figure 9 shows the histogram of intratrade durations for the 2008 sample.

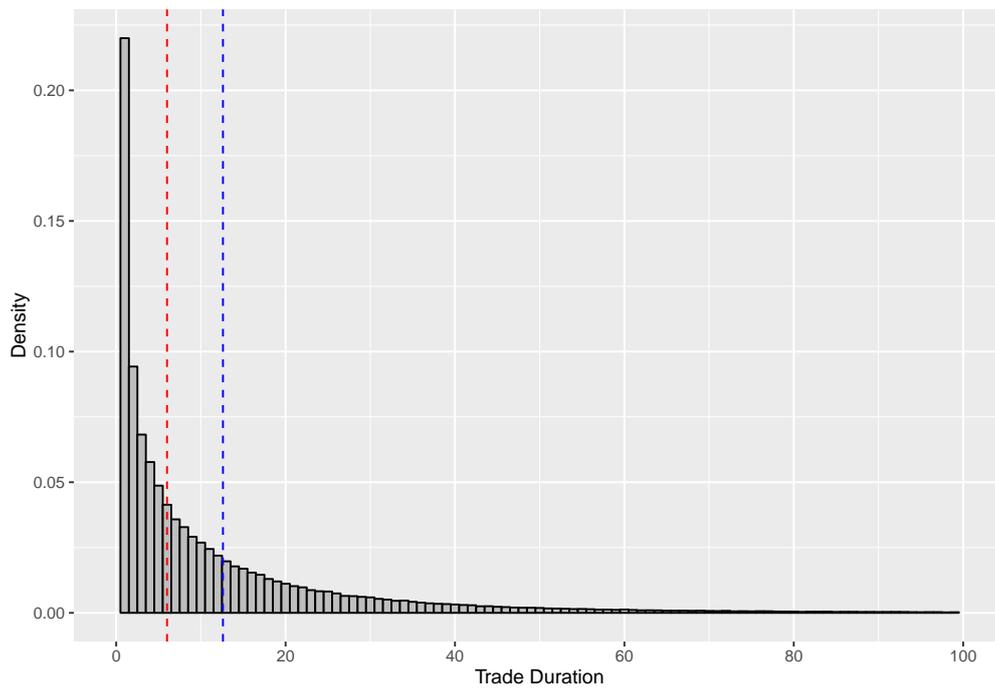


Figure 9: Histogram of intraday trade durations for Ashland Inc. during 2008. The vertical red line denotes the median of trade durations (6 seconds), whereas the blue line denotes the average of trade durations (12.67 seconds).

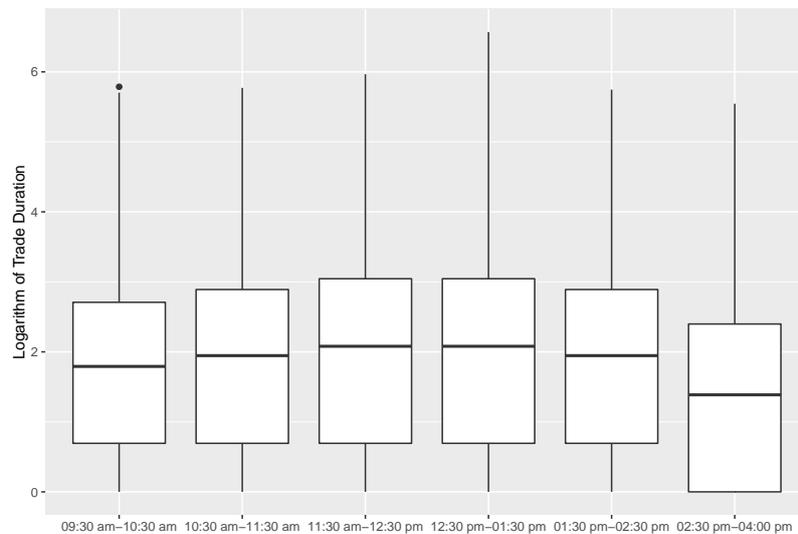
Obviously, the distribution for intraday trade durations is asymmetric. The median value is 6 seconds, and the mean value is 12.61 seconds. The geometric mean of duration times is 5.71, suggesting that duration times follow the log-normal distribution. The frequency of duration times may depend on the time of the day. Table 15 aggregates the frequencies of trades across daily trading intervals.

Table 15: SUMMARY OF TRADE DURATIONS

Trading Hour	Trades	Median	Mean	S.D.
09:30 am-10:30 am	74425	6	11.64	16.45
10:30 am-11:30 am	63778	7	14.19	19.96
11:30 am-12:30 pm	55049	8	16.39	23.39
12:30 pm-01:30 pm	53119	8	16.90	24.52
01:30 pm-02:30 pm	62933	7	14.19	19.84
02:30 pm-04:00 pm	154684	4	8.71	12.44

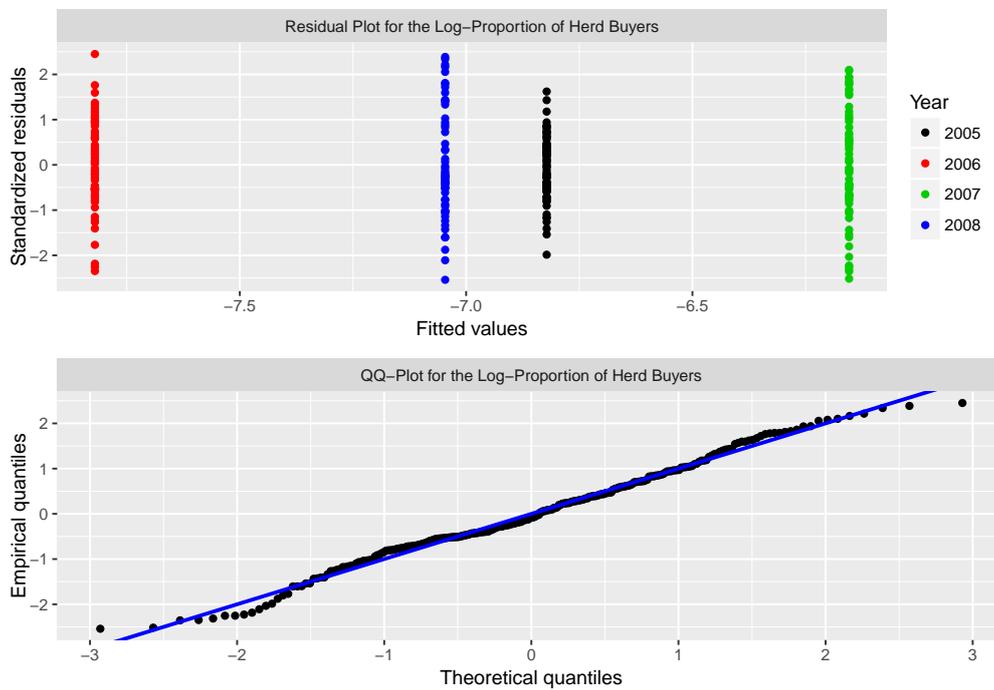
Description: Frequency table of trades across trading times. The last three columns report the median, the mean and the standard deviation of duration times, respectively.

Clearly, trading is more frequent in the beginning and in end of the day. Figure 10 presents a box plot of the logarithm of duration times across trading times.

**Figure 10:** Boxplot of intraday trade durations across trading hours during 2008.

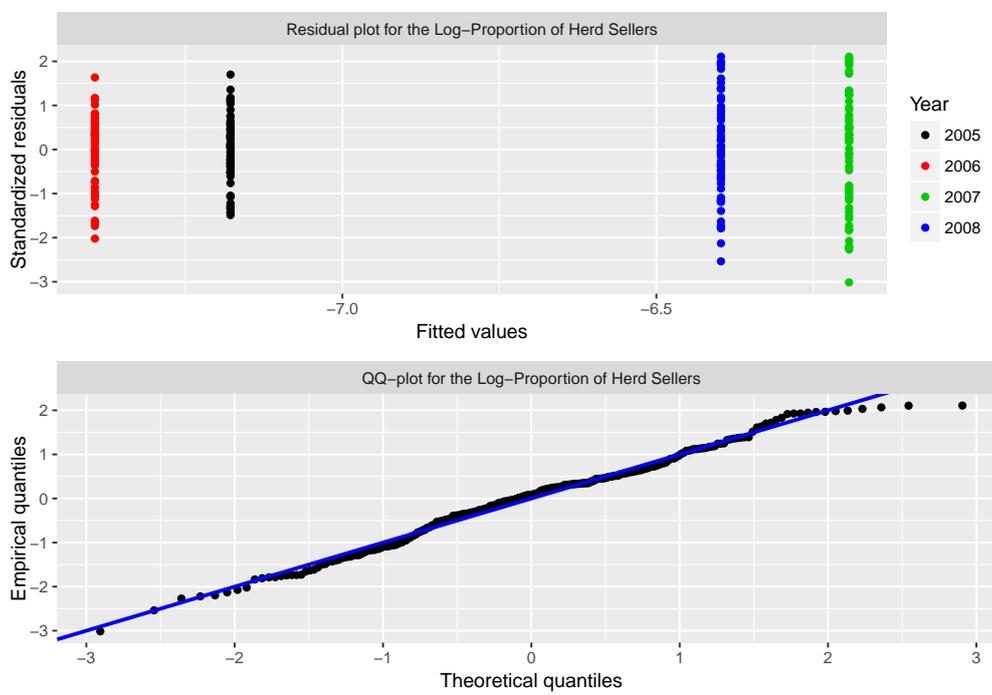
I.A.2. Model Diagnostics for the Log-Proportion of Herders

In section 7, it was assumed that the log-proportion of herders followed the linear model, with mean value depending on the trading year. Model assumptions are now discussed with relevant diagnostic plots.

Figure 11: MODEL DIAGNOSTICS FOR LOG-PROPORTION OF HERD BUYERS

Description: Residual plot (upper) and QQ-plot of residuals (lower) for the log-proportion of herd buyers.

Figure 11 shows the residual plot (standardized residuals vs. fitted values) as well as a QQ-plot for standardized residuals for the log-proportion of herd buyers. The residual plot indicates that the variance does not depend on mean level (variance homogeneity), while the lower plot is consistent with normality. Finally, figure 12 shows diagnostic plots for the log-proportion of herd sellers. The conclusion remains unchanged: model assumptions are accepted.

Figure 12: MODEL DIAGNOSTICS FOR LOG-PROPORTION OF HERD SELLERS

Description: Residual plot (upper) and QQ-plot of residuals (lower) for the log-proportion of herd sellers.

References

- Angel, J. J., L. E. Harris, and C. S. Spatt (2015). Equity trading in the 21st century: An update. *Quarterly Journal of Finance (QJF)* 05(01), 1–39.
- Avery, C. and P. Zemsky (1998, September). Multidimensional Uncertainty and Herd Behavior in Financial Markets. *American Economic Review* 88(4), 724–748.
- Banerjee, A. V. (1992). A simple model of herd behavior. *QUART. J. ECONOM* 107(3), 797–818.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2009). Realized kernels in practice: trades and quotes. *Econometrics Journal* 12(3), C1–C32.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy* 100(5), 992–1026.
- Bikhchandani, S. and S. Sharma (2000). Herd behavior in financial markets: A review. *IMF Staff Papers* 47(3), 279–310.
- Brown, A. (2017, July). What history says about low volatility. <https://www.bloomberg.com/view/articles/2017-07-03/what-history-says-about-low-volatility>.
- Chang, E. C., J. W. Cheng, and A. Khorana (2000). An examination of herd behavior in equity markets: An international perspective. *Journal of Banking & Finance* 24(10), 1651 – 1679.
- Christie, W. G. and R. D. Huang (1995). Following the pied piper: Do individual returns herd around the market? *Financial Analysts Journal* 51(4), 31–37.
- Cipriani, M. and A. Guarino (2014, January). Estimating a structural model of herd behavior in financial markets. *American Economic Review* 104(1), 224–51.
- Easley, D., N. M. Kiefer, and M. O’Hara (1997). One day in the life of a very common stock. *Review of Financial Studies* 10(3), 805–835.
- Easley, D., N. M. Kiefer, M. O’Hara, and J. B. Paperman (1996). Liquidity, information, and infrequently traded stocks. *The Journal of Finance* 51(4), 1405–1436.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2012, 02). Flow toxicity and liquidity in a high frequency world. *Review of Financial Studies* 25(5), 1457–1493.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2013, 03). Bulk classification of trading activity. *SSRN Electronic Journal*.

-
- Easley, D. and M. O'Hara (1992). Time and the process of security price adjustment. *Journal of Finance* 47(2), 576–605.
- Galariotis, E. C., W. Rong, and S. I. Spyrou (2015). Herding on fundamental information: A comparative study. *Journal of Banking & Finance* 50(Supplement C), 589 – 598.
- Glosten, L. R. and P. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14(1), 71–100.
- Holden, C. W. and S. Jacobsen (2013). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *The Journal of Finance* 69(4), 1747–1785.
- Hwang, S. and M. Salmon (2004). Market stress and herding. *Journal of Empirical Finance* 11(4), 585–616.
- Indārs, E. R. and A. Savin (2017, April). Herding behaviour in an emerging market: Evidence from moscow exchange. Bachelor thesis, Stockholm School of Economics in Riga.
- Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica* 53(6), 1315–35.
- Lakonishok, J., A. Shleifer, and R. W. Vishny (1992). The impact of institutional trading on stock prices. *Journal of Financial Economics* 32(1), 23 – 43.
- Lee, C. and M. Ready (1991). Inferring trade direction from intraday data. *Journal of Finance* 46(2), 733–46.
- Mullen, K. M. (2014). Continuous global optimization in r. *Journal of Statistical Software* 60(6).
- O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics* 116(2), 257–270.
- Ottaviani, M. and P. N. Sørensen (2006, January). Professional advice. *Journal of Economic Theory* 126(1), 120–142.
- Smith, L. and P. Sørensen (2000). Pathological outcomes of observational learning. *Econometrica* 68(2), 371–398. JEL Classification: D83.
- Vergote, O. (2005, March). How to Match Trades and Quotes for Nyse Stocks? Working Papers Department of Economics ces0510, KU Leuven, Faculty of Economics and Business, Department of Economics.
- Welch, I. (1992). Sequential sales, learning, and cascades. *The Journal of Finance* 47(2), 695–732.

This page intentionally left blank.

CHAPTER 2

DELTA HEDGING AND THE VPIN

*Which Free Lunch Would You Like Today, Sir?:
Delta Hedging, Volatility Arbitrage and Optimal
Portfolios*

— Riaz Ahmad & Paul Wilmott

This page intentionally left blank.

Delta Hedging and the VPIN^{*,**}

Laurs Randbøll Leth^{a,b,*}

^a*Department of Economics, University of Copenhagen*

^b*Center for Information and Bubble Studies, University of Copenhagen*

Abstract

This paper examines minimization of losses when hedging European options on a daily basis. In a Heston setting with stochastic volatility, a risk-minimizing strategy associated with the minimal-entropy martingale measure is considered. Options data on the S&P 500 index from 2004-2013 are used for the empirical study. Positive risk premiums are in most cases detected, but the hedging portfolio is exposed to significant daily losses during market turbulence. Applying the VPIN metric on Level 1 tick data of the SPDR S&P 500 ETF signals the impending losses. As a result, new strategies utilizing the VPIN are discussed with the purpose of eliminating risk when rebalancing the hedging portfolio on a daily basis.

Keywords: volatility arbitrage, stochastic volatility, European options, delta hedging, market microstructure, VPIN.

JEL: G11; G12; G13; G17; C58.

*The author appreciates and acknowledges the financial support of the Carlsberg Foundation through the grant to the Center for Information and Bubble Studies, University of Copenhagen.

**The author thanks Peter Norman Sørensen, Maureen O'Hara, Gideon Saar, and Patrick Thöni for their helpful comments and valuable suggestions.

*Corresponding author

Email address: bkt797@ku.dk (Laurs Randbøll Leth)

Conflict of Interest Disclosure Statement

I have nothing to disclose.

1. Introduction

This paper attempts to link two disjoint areas in finance, namely, high-frequency market microstructure (HFMM) and mathematical finance. Specifically, can risk tools from HFMM anticipate daily losses when hedging European options and thereby improve the hedging strategy?

Market makers supplying options are exposed to two sources of risk: adverse selection (volatility traders) and large intraday price movements (inventory risk). To bear this risk, the market-implied volatility is typically above the realized volatility, and trading European options results a in nonzero volatility risk premium.

The main problem is described from the portfolio manager's (PM's) point of view, with a short position in the European call. She will attempt to collect this risk premium by engaging in dynamic delta hedging of the underlying asset. Because of market frictions (such as transaction costs), the hedging portfolio is rebalanced only once at the end of each trading day. The PM is expected to reap the risk premium in periods of tranquility, whereas the hedging portfolio is subject to losses during market turbulence characterized by high volatility.

In mathematical finance, volatility arbitrage refers to a hedging strategy that almost surely results in profit. However, the existence of such strategies is contradicted by market data on options: The volatility is not constant as assumed in the famous model proposed by Black and Scholes (1973), but depends on the option's time-to-maturity and strike level¹. More likely, volatility itself is a random process. As a result, a model true to market data is incomplete, and it is impossible to perfectly hedge the European call. To account for stochastic volatility and at the same time maintain tractability, this paper assumes the option pricing model proposed by Heston (1993). The arbitrage-free price of the option is on semi-closed form, and derivatives (Greeks) used for hedging are computed. In comparison to the standard Black-Scholes Δ -hedge, the locally risk-minimizing Δ -hedge associated with the minimal-entropy martingale measure is introduced to minimize risk caused by the random nature of volatility.

¹These implied volatility patterns are illustrated by the *smile/skew*. See Cont and Tankov (2004) for more details.

Furthermore, the terminal hedging error generated by the model's risk-minimizing hedge is derived — a result Ellersgaard et al. (2017) refer to as the fundamental theorem of derivative trading. The terminal profit-and-loss of the hedging portfolio is random, and the option dealer is always exposed to losses when trading European options. Additionally, the discretization of time will add uncertainty to the terminal hedging error. With this consideration in mind, two research questions are addressed: 1) Is it possible to implement a hedging strategy that most likely will generate profit, and 2) Is the risk-minimizing strategy superior to the standard Black-Scholes hedge in a risk-minimizing context?

The empirical hedging experiment is conducted on 3-month at-the-money (ATM) European options written on the S&P 500 index from 2004-2013. Disjoint trading periods are considered, resulting in 37 hedging portfolios, each with a short position in the option. The locally risk-minimizing strategy is superior to the Black-Scholes Δ -hedge in terms of returns and standard deviations. In most of cases, the PM succeeds collecting the risk premium. However, significant losses between days are identified for one hedging portfolio subject to extreme market turbulence on September 15, 2008, and on the following trading days. These losses are caused by high (low) market exposure on days where the price of the underlying asset was falling (rising). Thus, a relevant question is whether losses could be reduced when using information from intraday patterns in the market.

This study considers the volume-synchronized probability of informed trading (VPIN) to explore certain stock characteristics (e.g., liquidity and volatility)². The VPIN metric is an indicator of toxic order flow and is best known as having sent a signal hours before the Flash Crash on May 6, 2010. Furthermore, Easley et al. (2012, 2010) present evidence that high VPIN values signal rises in return volatility, and in continuation of this, Easley et al. (2010, 2011) suggest that the VPIN may be used as an instrument for volatility arbitrage. In contrast to this, Andersen and Bondarenko (2014a,b, 2015) reject the VPIN as a good predictor of return volatility as well as its ability to signal ensuing market turbulence.

²Easley et al. (2012) developed the VPIN as a high-frequency approximation of the famous PIN measure introduced by Easley et al. (1996).

Several studies have since then both defended and disputed the applicability of the VPIN³. This study contributes to the literature, with the usefulness of the VPIN being assessed from the delta hedging PM's perspective.

The SPDR S&P 500 exchange-traded-fund (SPY) traded on NYSE Arca is used to track the index (SPX). Level 1 tick data of the ETF from 2004-2013 are used to compute the volume-synchronized probability of informed trading. The empirical study of the VPIN is concentrated on trading periods reflecting the 3-month ATM option's lifetime; that is, VPIN series are computed over quarterly periods. Eventually, the VPIN series for the trading period associated with the hedging portfolio whose P&L exhibited large fluctuations/losses are presented.

This paper contributes to the empirical market microstructure literature with results showing that high VPIN readings indirectly signaled impending daily losses for this hedging portfolio. The VPIN's ability to signal market turbulence is evaluated with the maximum intermediate return (MIR) of intraday prices. In particular, high MIR values with a wrong composition of market exposure may result in losses of the hedging portfolio. However, the VPIN is not a perfect risk tool; the metric was not able to signal all losses (type II errors), stressing a drawback of the methodology.

In conclusion, the portfolio manager is advised to rebalance with the risk-minimizing delta from Heston's model. This strategy is computationally tractable and superior to the standard Black-Scholes delta hedge. Furthermore, the VPIN signals daily losses in the hedging portfolio and therefore may be used to improve the hedge. Two strategies for which the PM can take advantage from the VPIN are discussed. The first strategy suggests expanding the portfolio with VIX futures immediately after the occurrence of VPIN events. This approach is supported by data, showing that on certain days, VIX increased dramatically after VPIN events were detected. The final strategy advises the portfolio manager to adjust the market exposure (delta) when VPIN events occur.

³For instance, see the empirical studies by Easley et al. (2012, 2010, 2011, 2014), Wu et al. (2013a,b) and Song et al. (2014) (supporting) or Andersen and Bondarenko (2014a,b, 2015), Pöppe et al. (2016) and Abad et al. (2018) (disputing).

2. Preliminary Concepts

The financial market consists of a risky asset and a risk-free asset, both traded continuously up to some fixed time horizon T . As usual, the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_s)_{0 \leq t \leq T}, P)$ is introduced to capture the information flow: Ω contains all possible states of the market, \mathcal{F} is the corresponding σ -algebra, P is the physical measure, and (\mathcal{F}_t) is the σ -field generated by the stochastic process $(S_t)_{0 \leq t \leq T}$, i.e., $\mathcal{F}_t = \sigma \{S_s \mid 0 \leq s \leq t\}$.

In the remainder of this paper, S_t denotes the value of the risky asset at time t , and is assumed to solve the stochastic differential equation

$$dS_t = \mu S_t dt + \sqrt{\nu_t} S_t dW_t$$

where $(W_t)_{t \geq 0}$ is a P -Brownian motion, and the volatility $(\sqrt{\nu_t})_{t \geq 0}$ is a positive process, possibly random. The risk-free asset's value is given by the process $(B_t)_{0 \leq t \leq T}$ with $dB_t = rB_t dt$ where r is constant⁴.

Let $\Phi(x) = (x - K)^+$ for $x > 0$ and some fixed constant $K > 0$. Then, $\Phi(S_T)$ denotes the terminal value of the European call with maturity T and strike price K . The arbitrage-free price at time $t < T$ of the European call is then given by

$$C(t, S_t, \nu_t) = E^Q \left(\frac{B_t}{B_T} \Phi(S_T) \mid \mathcal{F}_t \right)$$

under the pricing measure Q ⁵. It is assumed that the pricing function $C(\cdot, \cdot, \cdot)$ belongs to class of $\mathcal{C}^{1,2,2}$ -functions. To ease the notation, ∂_t and ∂_x denote the partial derivatives with respect to the time t and the direction x , respectively. Finally, $E_t(\cdot)$ denotes the conditional expectation given \mathcal{F}_t .

The remainder of this paper investigates the risk a portfolio manager encounters with a short position in the European call, especially how to improve the daily rebalancing in a risk-minimizing context. Suppose the portfolio manager sells one European call with expiry T at implied volatility $\sqrt{\nu^i}$ and wants to hedge this position using the delta of the option.

⁴The risk-free asset is interpreted as the money account with short rate of interest r .

⁵Assuming that there exists an equivalent martingale measure.

For simplicity, consider the Black-Scholes model where $(S_t)_{t \geq 0}$ follows a geometric Brownian motion with drift μ and constant volatility $\sqrt{\nu_t} = \sqrt{\nu}$, while the risk-free rate is fixed to 0. At time t , the portfolio consists of -1 European call and Δ_t units of the underlying stock with $\Delta_t \equiv \partial_s C(t, S_t)$. The money account B_t is chosen to satisfy the self-financing condition, i.e., the net value of this portfolio, Π_t^h , is simply

$$\Pi_t^h = -C^i(t, S_t) + B_t + \Delta_t^h S_t = 0,$$

where the superscript i implicitly refers to implied volatility (e.g., $C^i(t, S_t)$ is the market price of the option), whereas the superscript h emphasizes that $\sqrt{\nu^h}$ is used as hedging volatility. Combining the self-financing condition, Itô's formula on $C^h(t, S_t)$ and the Black-Scholes pricing PDE yields

$$d\Pi_t^h = dC^i(t, S_t) - dC^h(t, S_t) + \frac{1}{2} (\nu^h - \nu) S_t^2 \partial_{ss} C^h(t, S_t).$$

The terminal P&L holding this portfolio over $[0, T]$ is given by

$$P\&L_T^h \equiv \int_0^T d\Pi_t^h = C^i(0, S_0) - C^h(0, S_0) + \int_0^T \frac{1}{2} (\nu^h - \nu) S_t^2 \partial_{ss} C^h(t, S_t) dt$$

using the boundary conditions $C^h(T, S_T) = C^i(T, S_T) = (S_T - K)^+$. This result is called *the fundamental theorem of derivative trading* (FTODT) by Ellersgaard et al. (2017), and a more general version, where $(\sqrt{\nu_t})_{t \geq 0}$ itself is a diffusion is presented in the next section. Clearly, successful hedging in a misspecified model depends on the relationship between actual volatility and the hedging volatility chosen by the portfolio manager. The gamma of the option is positive, $\partial_{ss} C^h(t, S_t) > 0$, and the sign of the terminal hedging error equals the sign of $\nu^h - \nu$. As a result, if $\nu^h > \nu$, volatility arbitrage is obtained using either ν or ν^h as hedging volatility — in the first case, the P&L is deterministic, whereas it becomes random for the second case⁶. A numerical example of the FTODT originally produced by Wilmott and Ahmad (2005) is illustrated in figure 1.

⁶Selling the option when $\nu < \nu^h$ is in line with *buy low, sell high*.

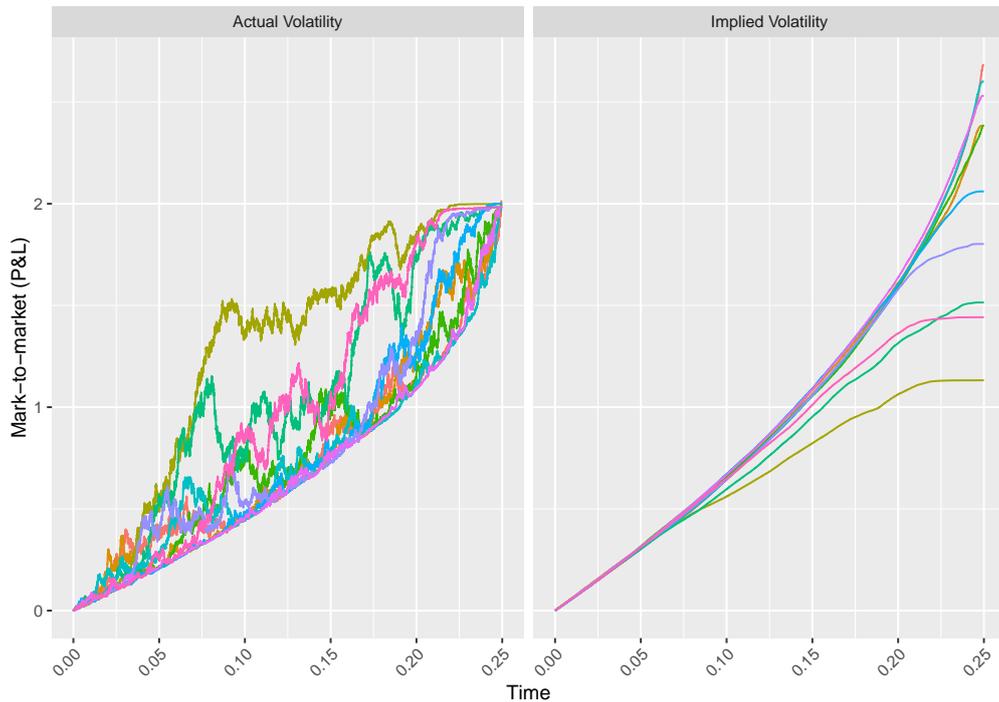


Figure 1: Ten sample paths of the geometric Brownian motion are simulated with parameter values $S_0 = 100$, $\mu = 0.1$ and $\sqrt{\nu} = 0.1$. Market parameters are $r = 0.02$, $\sqrt{\nu^i} = 0.2$, $K = 100$ and $T = 1/4$. The portfolio is rebalanced 5000 times for each combination of sample paths and hedging volatility. Left plot: Actual volatility is used for hedging. The terminal P&L is deterministic and equals $\Pi_T = e^{rT}(C^i(0, S_T) - C(0, S_0)) = 1.99$ using the Black-Scholes formula. Right plot: Implied volatility is used for hedging. The terminal P&L is random but positive.

At first sight, the above example clearly suggests an advantage obtained by Δ -hedging the option using either implied or actual volatility. However, the example implicitly stresses some of the weaknesses of the model: Volatility arbitrage is not present in the real world, which contradicts the outcomes of Wilmott's hedge example. One reason is that the actual volatility is unknown, and it is thereby impossible to determine whether the option is mispriced. Even if the option dealer were superior in estimating actual volatility, the terminal hedging error could still end up negative. The explanation is that volatility is not constant as prescribed in a Black-Scholes world. More likely, volatility itself is a diffusion process, and the terminal hedging error becomes random. Consequently, the option dealer is always exposed to losses when trading European call options.

3. The Portfolio Manager's Hedging Model

The stochastic volatility model proposed by Heston (1993) is assumed. That is, the bivariate process $(S_t, \nu_t)_{t \geq 0}$ satisfies the system of stochastic differential equations

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{\nu_t} S_t dW_t^1 \\ d\nu_t &= \kappa(\theta - \nu_t) dt + \sigma \sqrt{\nu_t} dW_t^2, \end{aligned}$$

where $(W_t^1)_{t \geq 0}$ and $(W_t^2)_{t \geq 0}$ are P -Brownian motions, possibly dependent with correlation coefficient $\rho \in [-1, 1]$. The variance process $(\nu_t)_{t \geq 0}$ (CIR) is continuous as well as non-negative under the assumption $\kappa, \theta, \sigma > 0$ ⁷. Finally, $(\nu_t)_{t \geq 0}$ is mean-reverting with long-term mean θ and rate of mean-reversion κ . The parameter σ is referred to as volatility-of-volatility.

3.1. Pricing Measures

A necessary condition for an equivalent martingale measure to exist is

$$\frac{\mu - r}{\sqrt{\nu_t}} = \sqrt{1 - \rho^2} \gamma_1(t) + \rho \gamma_2(t) \quad (3.1)$$

where $(\gamma(t))_{t \geq 0}$ is the bivariate risk premium process adapted (\mathcal{F}_t) and associated with the likelihood process (L_t) restricted to $0 \leq t \leq T$ ⁸. If $E[L_T] = 1$, Girsanov's theorem introduces an equivalent probability measure Q given by $L_T = dQ/dP$ on \mathcal{F}_T . In particular,

$$dW_t = -\gamma(t) dt + dW_t^Q$$

where $(W_t^Q)_{t \geq 0}$ is a 2-dimensional standard Brownian motion under Q .

⁷In addition, the Feller condition $\sigma^2 \leq \kappa\theta$ will imply that $\nu_t > 0$ for $t \in [0, T]$.

⁸The likelihood process is given by

$$L_t = e^{\int_0^t -(\gamma_1(s) dW_s^1 + \int_0^t \gamma_2(s) dW_s^2) - \frac{1}{2}(\int_0^t \gamma_1^2(s) ds + \int_0^t \gamma_2^2(s) ds)}.$$

The discounted process $(\hat{S}_t)_{t \geq 0}$ must be a local martingale under the pricing measure Q requiring that $(\gamma(t))_{t \geq 0}$ satisfies equation (3.1).

Condition (3.1) is satisfied for

$$\gamma^{\min}(t) = \frac{\mu - r}{\sqrt{\nu_t}} \cdot \begin{pmatrix} \sqrt{1 - \rho^2} \\ \rho \end{pmatrix}, \quad (3.2)$$

which Poulsen et al. (2009) show is the risk premium process associated with the minimal-entropy martingale measure. (Wong and Heyde, 2006) prove that the proposed choice of $\gamma_2^{\min}(t)$ implies that $E[L(T)] = 1$. Thus, the pricing measure Q exists, and the Q -dynamics of $(S_t, \nu_t)_{t \geq 0}$ are given by

$$\begin{aligned} dS_t &= rS_t dt + \sqrt{\nu_t} S_t \left(\sqrt{1 - \rho^2} dW_t^{1,Q} + \rho dW_t^{2,Q} \right) \\ d\nu_t &= \kappa^Q (\theta^Q - \nu_t) dt + \sigma \sqrt{\nu_t} dW_t^{1,Q} \end{aligned}$$

where $(W_t^{1,Q})_{t \geq 0}$ and $(W_t^{2,Q})_{t \geq 0}$ are independent standard Brownian motions under Q , $\kappa^Q = \kappa$ and $\theta^Q = \theta - \frac{(\mu-r)\rho\sigma}{\kappa}$; that is, the change-of-measure generates Q -dynamics with the same parametric form as under P .

3.2. Greeks and Hedging

Let $C(t, S_t, \nu_t)$ denote the time t price of the European call with maturity T and strike K . Using risk-neutral valuation, the fair price may be expressed by

$$\begin{aligned} C(t, S_t, \nu_t) &= e^{-r(T-t)} E_t^Q ((S_T - K)^+) \\ &= S_t E_t^{\tilde{Q}} (1_{S_T > K}) - e^{-r(T-t)} K E_t^Q (1_{S_T > K}) \end{aligned}$$

where \tilde{Q} is given by $d\tilde{Q}/dQ = (B_t/B_T)/(S_t/S_T)$. If $x_t = \log S_t$, the price can be expressed by

$$C(t, S_t, \nu_t) = S_t p_1(t, x_t, \nu_t) - e^{-r(T-t)} K p_2(t, x_t, \nu_t)$$

where $p_1(t, x_t, \nu_t)$ and $p_2(t, x_t, \nu_t)$ denote the conditional probability at time $t < T$ that the option expires in-the-money under \tilde{Q} and Q , respectively.

Theorem 1. *The expiry-in-money probabilities are given by*

$$p_1(t, x_t, \nu_t) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{e^{-i\phi \log K} \tilde{f}(\phi)}{i\phi} \right] d\phi \quad (3.3)$$

$$p_2(t, x_t, \nu_t) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{e^{-i\phi \log K} f(\phi)}{i\phi} \right] d\phi \quad (3.4)$$

where Re denotes the real part of a complex number, i is the imaginary unit, and $f(\phi)$ and $\tilde{f}(\phi) = f(\phi - i)/f(-i)$ are the conditional characteristic functions of x_T under the pricing measure Q , respectively, under the measure \tilde{Q} . Furthermore,

$$f(\phi) = e^{c(\phi) + d(\phi)\nu_t + i\phi x_t} \quad (3.5)$$

where the functions $d(\cdot)$ and $c(\cdot)$ are given in appendix B.

PROOF.

A version of the theorem is proved by Heston (1993). \square

Corollary 2 (Greeks). *The option's Δ is at time t given by*

$$\Delta_t = \partial_s C(t, S_t, \nu_t) = p_1(t, x_t, \nu_t) \quad (3.6)$$

where the conditional probability p_1 is derived in theorem 1. Moreover, the vega of the option is given by

$$\begin{aligned} \mathcal{V}_t &= \partial_\nu C(t, S_t, \nu_t) \\ &= \frac{1}{\pi} \left\{ S_t \int_0^\infty \operatorname{Re} \left[d(\phi - i) \frac{e^{-i\phi \log K} \tilde{f}(\phi)}{i\phi} \right] d\phi \right. \\ &\quad \left. - e^{-r(T-t)} K \int_0^\infty \operatorname{Re} \left[d(\phi) \frac{e^{-i\phi \log K} f(\phi)}{i\phi} \right] d\phi \right\} \quad (3.7) \end{aligned}$$

where $d(\cdot)$ is given in appendix B.

PROOF.

This well-known result is proved in appendix A. \square

The Δ -hedging option dealer is investing the \tilde{Q} -probability of the European call is in-the-money at expiry. However, this strategy does not account for volatility risk, and it is therefore not optimal in a risk-minimizing context.

3.2.1. The Risk-Minimizing Hedge

The risk-minimizing hedge refers to a non-self-financing strategy that replicates $(S_T - K)^+$ and at the same time minimizes the expected squared deviation between the portfolio value and the value of the option. Föllmer and Sondermann (1986) provide a formal definition of the risk-minimizing hedging strategy. However, such a strategy may not exist, but Schweizer (1991) overcomes this problem proposing the *locally* risk-minimizing hedge under the minimal-entropy martingale measure Q^{\min} .

In this model setting, the locally risk-minimizing hedge is given in the following proposition of Poulsen et al. (2009):

Proposition 3 (Locally risk-minimizing strategy). *At time t , the portfolio manager applying the locally risk-minimizing strategy invests*

$$\Delta_t^{\min} = \partial_s C(t, S_t, \nu_t) + \frac{\rho\sigma}{S_t} \partial_\nu C(t, S_t, \nu_t) \quad (3.8)$$

units in stock, and $C(t, S_t, \nu_t) - \Delta_t^{\min} S_t$ units in the risk-free asset. The price of the option is given by $C(t, S_t, \nu_t) = e^{-r(T-t)} E^{Q^{\min}} E_t((S_T - K)^+)$, where Q^{\min} is the pricing measure obtained from the process given by equation (3.2). Furthermore, the dynamics of $(S_t, \nu_t)_{t \geq 0}$ under Q^{\min} are given by

$$\begin{aligned} dS_t &= rS_t dt + \sqrt{\nu_t} S_t \left(\sqrt{1 - \rho^2} dW_t^{Q^{\min},1} + \rho dW_t^{Q^{\min},2} \right) \\ d\nu_t &= (\kappa(\theta - \nu_t) - \rho\sigma(\mu - r)) dt + \sigma\sqrt{\nu_t} dW_t^{Q^{\min},2} \end{aligned}$$

where $(W_t^{Q^{\min},1})_{t \geq 0}$ and $(W_t^{Q^{\min},2})_{t \geq 0}$ are Q^{\min} -Brownian motions.

PROOF. For convenience, a heuristic proof is given in appendix A. \square

The locally risk-minimizing hedge and the Δ -hedge coincide when $\rho = 0$. Moreover, the parameterization $\theta^{Q^{\min}} = \theta - \rho\theta \cdot (\mu - r)/\kappa$ yields Q^{\min} -dynamics of $(S_t, \nu_t)_{t \geq 0}$ with the same parametric form as under P . Finally, the locally risk-minimizing strategy depends on the option's vega, which was derived in corollary 2. Plugging this expression into Δ_t^{\min} yields

$$\Delta_t^{\min} = \Delta_t + \frac{\rho\sigma}{\pi} \left\{ \int_0^\infty \operatorname{Re} \left[D(\phi - i) \frac{e^{-i\phi \log K} \tilde{f}(\phi)}{i\phi} \right] d\phi - e^{-r(T-t)} \frac{K}{S_t} \int_0^\infty \operatorname{Re} \left[D(\phi) \frac{e^{-i\phi \log K} f(\phi)}{i\phi} \right] d\phi \right\}.$$

It is strongly hinted that in this setting, hedging the European call results in a nonzero hedging error almost surely. This is illustrated with an example similar to Wilmott's hedge example but in the Heston setting, where model and market parameters are specified in line with the literature (see table 1). The constant implied volatility is assumed greater than the initial (and long term) volatility. Thus, the option is overpriced prior to trading, and a short position in the European call is considered. The PM then applies the risk-minimizing Δ to hedge her position with either actual volatility or implied volatility.

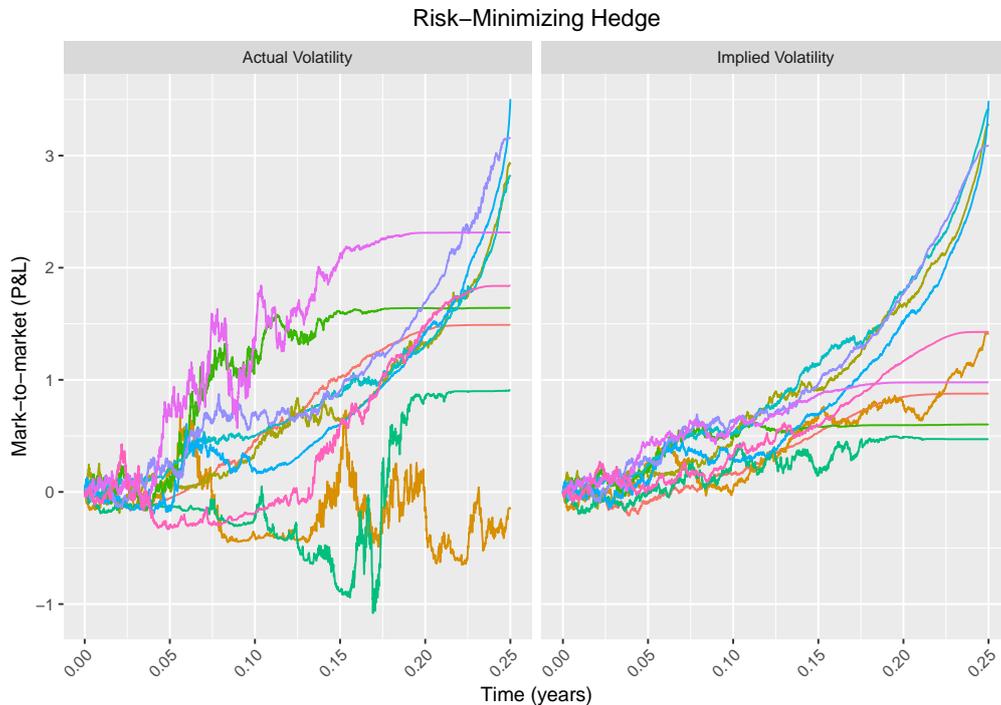


Figure 2: Ten sample paths of both the underlying and the variance process are simulated with Euler discretization. The portfolio is rebalanced 5000 times for each combination of sample paths and hedging volatility.

Figure 2 shows mark-to-market values of hedging portfolios obtained by the two strategies. In most cases profit is obtained, but a negative terminal hedging error also occurs; the PM is always exposed to volatility risk. Finally, hedging with implied volatility results in smoother P&L paths in line with Wilmott's example.

With simple delta hedging, Ellersgaard et al. (2017) derives an expression for the random terminal hedging error when $(S_t)_{t \geq 0}$ is a locally SDE, while Leth (2015) shows the result in the Heston setting. Clearly, their strategy neglects the source of randomness that arises from the volatility process; that is, risk associated with $(W_t^{2,Q})_{t \geq 0}$ is not reduced. By contrast, applying the risk-minimizing strategy will capture both components of risk, which is mathematically shown in theorem 4.

Theorem 4. *Let $C_t = C(t, S_t, \nu_t) \in \mathcal{C}^{1,2,2}([0, \infty) \times \mathbb{R} \times \mathbb{R})$ be the time- t price of European call. Assume that the option is sold at initiation for the market price $C(0, S_0, \nu_0^i)$, where $\sqrt{\nu_0^i}$ is the implied volatility. Applying the risk-minimizing strategy Δ^{\min} with hedging volatility $\sqrt{\nu_t^h}$ leads to a terminal P&L given by*

$$\begin{aligned} P\&L_T^{\min} = C(0, S_0, \nu_0^i) - C(0, S_0, \nu_0^h) \\ &+ \int_0^T e^{-rt} \left\{ (\nu_t^h - \nu_t) \left\{ \frac{1}{2} (S_t^2 \partial_{ss} C(t, S_t, \nu_t^h) + \sigma^2 \partial_{\nu\nu} C(t, S_t, \nu_t^h)) \right. \right. \\ &\quad \left. \left. + \sigma \rho S_t \partial_{s\nu} C(t, S_t, \nu_t^h) - \kappa \partial_\nu C(t, S_t, \nu_t^h) \right\} dt \right. \\ &\quad \left. + \sigma \sqrt{\nu_t} \left(\rho \partial_\nu C^h(t, S_t, \nu_t^h) dW_t^{1,Q} - dW_t^{2,Q} \right) \right\}. \quad (3.9) \end{aligned}$$

PROOF.

See appendix A. □

Unlike the Black-Scholes setting, volatility arbitrage is not present in the market. The present value of the $P\&L_T^{\min}$ depends on the second-order Greeks gamma, vomma, vanna and vega. In addition, the hedging error depends on two stochastic integrals. Even if the PM knew and preferred to hedge with actual volatility, a hedging error would still occur. In particular,

setting $\sqrt{\nu_t^h} = \sqrt{\nu_t}$ implies that the terminal P&L is decomposed by the upfront premium and the difference between the two stochastic integrals

$$\sigma \int_0^T e^{-rt} \sigma \sqrt{\nu_t} \left(\rho \partial_\nu C^h(t, S_t, \nu_t^h) dW_t^{1,Q} - dW_t^{2,Q} \right).$$

Thus, the sign of the hedging error will depend on the trajectories of two Brownian motions. As Leth (2015) shows, when hedging with Heston's Δ_t instead, the $dW_t^{1,Q}$ -term vanishes in equation (3.9). In comparison with the Δ -hedge, it is therefore expected that the risk-minimizing strategy will reduce the terminal hedging error due to $\text{Cov}(dW_t^{1,Q}, dW_t^{2,Q}) = \rho dt$.

3.2.2. Example: Model Risk

The risk-minimizing strategy introduced in proposition 3 was proposed to minimize risk induced by the random nature of volatility. In addition, model risk is an important consideration when choosing the hedging strategy. For instance, what is the outcome of using the standard Black-Scholes delta hedge in a Heston world?

Table 1: SPECIFICATION OF MARKET AND MODEL PARAMETERS

Parameters	S_0	K	r	T	ν_i	μ	ν_0	κ	θ	σ	ρ
Values	100	100	0.02	1/4	0.3 ²	0.1	$\sqrt{0.2}$	4.75	$\sqrt{0.2}$	0.5	-0.5

Using the model specification from table 1, 300 sample paths of the underlying and the variance process are simulated with Euler discretization⁹. Pricing and hedging are then conducted under two pricing measures: Q^{\min} and Q^{mg} . The former measure is the minimal martingale measure generated by $(\gamma_t^{\min})_{t \geq 0}$ from equation (3.2), while the latter measure is obtained by replacing the stock drift μ with the risk-free rate r ¹⁰. This simulation study is akin to that of Leth (2015), but with a different setting and objective¹¹. Finally, four different strategies are considered:

⁹Model parameters are consistent with empirical findings in the literature (e.g., Eraker (2004)).

¹⁰ $(\hat{S}_t)_{t \geq 0}$ is assumed to be a P -martingale, which corresponds to the standard change-of-measure in the Black-Scholes model.

¹¹The objective is to investigate drawbacks of using the Black-Scholes Δ -hedge in the Heston setting.

1. The risk-minimizing Δ -hedge under Q^{\min} .
2. The risk-minimizing Δ -hedge under Q^{mg} .
3. The Heston Δ -hedge under Q^{mg} .
4. The Black-Scholes Δ -hedge under Q^{mg} .

As usual, the portfolio consists of the short position in the option, long in the underlying, and the money account chosen to meet the self-financing condition. For each sample path (and each strategy), the portfolio is re-balanced on a daily basis (63 times). Terminal hedging errors under the different strategies are illustrated in figure 3.

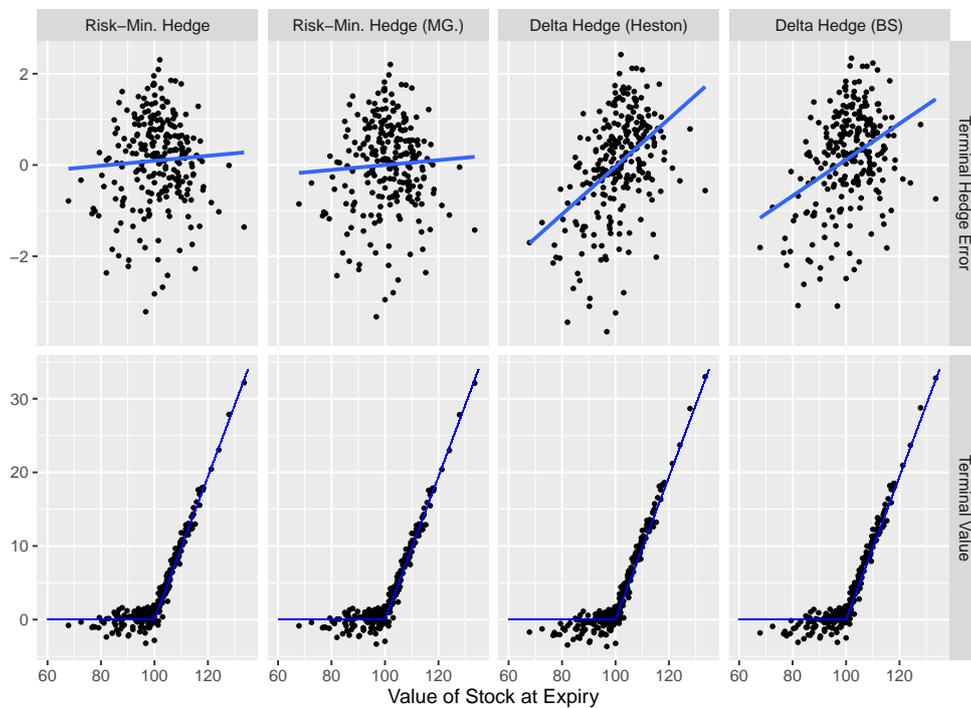


Figure 3: For each simulated path, the terminal hedge error is computed with the locally risk-minimizing hedge (under both Q^{\min} and Q^{mg}), Heston's Δ -hedge (Q^{mg}) and the Black-Scholes Δ (Q^{mg}). Top row: terminal hedging errors as a function of S_T . Bottom row: terminal values of the hedging portfolios as a function of S_T , where the blue lines represent $(S_T - K)^+$.

Terminal errors fluctuate most under the two Δ -strategies, where it appears that the out-the-money scenarios result in the largest hedging errors. The reason is that $\rho = -0.5$, and the Δ -hedging trader invests too heavily in the stock. Terminal hedging errors under the two risk-minimizing strategies do not differ significantly, and they appear to be 'symmetric' around $S_T = K$.

The estimate of the hedging error proposed by Poulsen et al. (2009) is used to assess the performance of the four strategies: the standard deviation of terminal hedge errors relative to the initial value of the option in percentages are computed¹². Table 2 reports the estimates.

Table 2: ESTIMATED HEDGING ERRORS

Strategy	Δ^{\min}	$\Delta^{\min} (Q^{\text{mg}})$	Heston- $\Delta (Q^{\text{mg}})$	BS- $\Delta (Q^{\text{mg}})$
Error	23.22	23.73	27.53	25.8

Description: Estimated hedging errors under the four strategies.

The risk-minimizing strategies are superior to both Δ -hedges in terms of hedging errors. For instance, the Heston Δ -hedging trader may reduce the hedging error by a factor of $\approx 15\%$ when switching to the risk-minimizing strategy. Comparing the two risk-minimizing strategies, the estimated hedging error is not significantly affected by the choice of pricing measure. This is convenient for the PM due to well-known numerical difficulties of estimating the stock drift μ . Finally, the results show that no advantage is obtained by using Heston's Δ instead of the standard B-S Δ .

4. The VPIN

Easley et al. (2012) developed the volume-synchronized probability of informed trading (VPIN) used as an indicator of toxic order flow. They argue that the VPIN has incremental predictive power for models of short-term return volatility and may be used to signal imminent market turmoil. Strikingly, Easley et al. (2012, 2010) conclude that the VPIN metric signaled the Flash Crash on May 6, 2010. Those findings led to a controversial debate with Andersen and Bondarenko (2014a,b, 2015) disputing the VPIN's capacity to predict return volatility as well as signaling ensuing market turbulence. Thus, in isolation, an empirical study of the VPIN has relevance, and if the metric signals intraday price movements, it may be used to reduce losses of the hedging portfolio, which is rebalanced on a daily basis.

¹²Algebraically, this corresponds to $100 \cdot \sqrt{\text{Var}^P(\Pi_T) / (e^{-T} E^Q((S_T - K)^+))}$.

Moment estimation of VPIN is nonparametric and utilizes intraday trade data to estimate the probability of informed trading given a fixed-sized trading volume¹³. Let V_t^B and V_t^S be buy and sell volume of trading period t such that $OI_t = |V_t^B - V_t^S|$ and $TT_t = V_t^B + V_t^S$ denote the absolute order imbalance and total trades, respectively. The VPIN metric is then the moment estimator of $E(OI_t)/VBS$ with $V_t^B + V_t^S = VBS$ and is derived as follows:

1. **Trade classification:** Buy and sell volume are computed. The bulk Volume Classification (BVS) introduced by (Easley et al., 2013) is used to determine the direction of order flow:

- (a) Sequential trades are aggregated using time or volume bars with bar price equal to the last transaction¹⁴. Let P_i and V_i denote the price and volume, respectively, of bar i .
- (b) Assuming $\Delta P_i := P_i - P_{i-1} \sim \mathcal{N}(0, \sigma_{\Delta P})$ where $\sigma_{\Delta P}$ is the standard deviation of price changes between all bars, the classification algorithm then assigns buy and sell volume of bar i by

$$\begin{aligned} \text{barBuy}_i &= V_i \cdot N\left(\frac{\Delta P_i}{\sigma_{\Delta P}}\right) \\ \text{barSell}_i &= V_i - \text{barBuy}_i, \end{aligned}$$

where $N(\cdot)$ is the CDF of the standard normal distribution.

2. **Volume bucketing:** Sequential bars are grouped into equal volume buckets of an exogenously defined size VBS , resulting in N buckets. Moreover, let OI_τ and TT_τ denote absolute order imbalance and total trades, respectively, of bucket τ .
3. **Estimation:** By construction, $TT_\tau = VBS$ for all τ . If the mapping b assigns bar i to bucket $b(i)$, order imbalance for bucket τ , $OI_\tau = V_\tau^B - V_\tau^S$, is determined by

¹³This is not equivalent to the PIN measure, which measures the probability of informed trading for a given fixed time interval.

¹⁴Using the last price is an arbitrary choice. One could also use the average of transaction prices within the bar.

$$V_\tau^B = \sum_{i: b(i)=\tau} V_i Z\left(\frac{\Delta P_i}{\sigma_{\Delta P}}\right)$$

$$V_\tau^S = \sum_{i: b(i)=\tau} V_i \left(1 - Z\left(\frac{\Delta P_i}{\sigma_{\Delta P}}\right)\right) = VBS - V_j^B.$$

Eventually, $E(OI_\tau | VBS)$ is estimated by the moment estimator

$$E(OI_\tau | VBS) \approx \frac{1}{n} \sum_{j=\tau-(n-1)}^{\tau} OI_j,$$

where n denotes the sample length.

Combining the three above steps leads to the VPIN metric

$$VPIN_\tau = \frac{1}{n \times VBS} \sum_{j=\tau-(n-1)}^{\tau} OI_j \quad (4.1)$$

for $\tau > n - 1$. Summarizing, computation of the VPIN requires 1) determining order flow, 2) choosing the equal-volume bucket size VBS and 3) the sample length n . It is important to emphasize that step 1 is solely used to measure order flow in a high-frequency market microstructure, and other trade classification algorithms may be used¹⁵¹⁶.

Given a trade classification rule, the equal-volume bucket size VBS and the sample length n are still left to be specified. One specification in the literature is setting VBS equal to one-fiftieth of average daily volume. If $n = 50$, the VPIN then reflects order imbalances of buckets across a trading day of average volume. Lin and Ke (2017) contribute by mathematically showing that the VPIN becomes unstable for *small* volume buckets, which is empirically supported in a study by Leth (2019)¹⁷.

¹⁵For instance, the trade classification algorithm proposed by Lee and Ready (1991) that requires Level 2 tick data (transactions *and* quotes).

¹⁶Easley et al. (2013) provides numerical evidence that standard classification algorithms depending on Level 2 tick data do not offer greater accuracy, whereas the BVS-procedure is obviously superior in terms of computational cost.

¹⁷Additionally, Lin and Ke (2017) propose an improved estimator of the VPIN (in terms of predicting order flow toxicity) that captures the information in volume time.

The parameter setting $(n, VBS) = (100, 1/100)$ is used for the empirical investigation, where $VBS = 1/k$ is shorthand notation for $1/k$ 'th of average daily volume for a given trading period. Appendix C reports findings under different combinations of n and VBS , including $(n, VBS) = (50, 1/50)$.

4.1. Performance Evaluation of the VPIN

High VPIN readings proxies toxic order flow, which will affect liquidity when market makers start demanding liquidity. This effect may imply large price returns. The hypothesis is that the VPIN serves as warning signal of impending market turbulence.

Inspired by the work of Song et al. (2014) and Leth (2019), a *daily* VPIN event is defined as the first time the CDF of the VPIN hits an exogenously given threshold c from below. For notational convenience, trading day d contains n_d subsequent trades, $(1, \dots, n_d)$, and p_i denotes the price of trade i .

Definition 5. *On trading day d , the VPIN event is specified by the hitting time*

$$i_d^* = \inf\{i \in \{1, \dots, n_d\} : CDF(VPIN)_i > c \ \& \ CDF(VPIN)_{i-1} \leq c\}$$

if $\max_i CDF(VPIN)_i > 0$ for $i \in \{1, \dots, n_d\}$.

The daily rebalanced hedging portfolio is subject to losses if volatility is high and/or large price movements occur. To capture both scenarios, the *maximum intermediate return* is used as proxy for intraday return volatility. On trading day d , let $R_{j,k} := p_k/p_j - 1$ be the return between p_j and p_k for $1 \leq j < k \leq n_d$. The daily maximum intermediate return is then given by

$$\text{MIR} = R_{j^*, k^*} \tag{4.2}$$

where the pair of trades (j^*, k^*) are called sentinels and maximize the intermediate returns over all combinations of trades between 1 and n_d , i.e.,

$$(j^*, k^*) = \arg \max_{0 \leq j < k \leq n_d} |R_{j,k}|. \tag{4.3}$$

The daily MIR value is either negative or positive¹⁸. For instance, a negative maximum intermediate return indicates that the largest price movement was downward. The VPIN's ability to signal market turbulence is then measured with respect to MIR. For a given period (e.g., the lifetime of the European option), daily MIR values are computed, and positive and negative values are separated from each other. Eventually, the average of negative outcomes, \underline{M} , and the average of positive outcomes, \overline{M} , are both computed. Finally, binary VPIN events are classified as true positives or false positives by comparing the MIR value between the VPIN event and the final trade of the day with the two daily averages.

Definition 6. *On trading day d , a VPIN event with hitting time i_d^* is labeled*

- *True positive (TP) if the $MIR_{i_d^*,n_d} > \overline{M}$ or $MIR_{i_d^*,n_d} < \underline{M}$.*
- *False positive (FP) if $MIR_{i_d^*,n_d} \in [\underline{M}, \overline{M}]$.*

If no VPIN event was detected but $MIR_{0,n_d} \notin [\underline{M}, \overline{M}]$, the day is labeled false negative (FN).

False negatives (type II errors) refer to trading days characterized by high volatility but with normal VPIN readings. Additionally, days with small MIR values and no VPIN events are classified as true negatives.

5. Data

The empirical experiment is concentrated on periods of trading from 2004/04/01-2013/07/03 (2331 trading days). The initial trading date is chosen such that hedging portfolios will reflect quarterly periods. The experiment requires market data on both the underlying index (index spot, zero rate and dividend yield) and on options (daily market prices expressed through implied volatilities for various combinations of strike prices and time-to-maturities relative to the spot index). In addition, time series for realized and implied volatilities are needed. Finally, financial tick data for a stock selected to track the underlying index are required in order to compute reliable VPIN series.

¹⁸In principle, the daily MIR may equal 0 if all prices are identical.

5.1. Options Data

Market data on options are similar to data used by Ellersgaard et al. (2017) and Leth (2015) and may be downloaded from <http://web.math.ku.dk/~rolf/Svend/>¹⁹. On each trading day, the index spot, Black-Scholes implied volatilities depending on strike prices and time-to-maturities, zero rates and dividend yields are all observed. Strike prices and time-to-maturities are relative to the spot index, where the moneyness (K/S) ranges between 1.3 and 0.7, and time-to-maturity is between 1 month and 36 months (182 combinations/prices/implied volatilities). Table 6 in appendix C shows market data on options on a given trading day.

A 3-month European option on the S&P 500 index is considered. Initially, the option is purchased at-the-money, and the position is then hedged on a daily basis prior to expiry. After expiry, a new 3-month ATM European option is purchased, and the procedure is repeated. This design results in a total of 37 portfolios over disjoint trading intervals where the first option is traded on 2004/04/01, and the last option expires on 2013/07/03.

5.2. Volatility Data

Time series of daily volatilities are used for two purposes:

1. As input in the daily Δ -hedge formula.
2. As valuation of the option (implied volatility).

5.2.1. Implied Volatility

The raw market data on options reports implied volatilities for various combinations of maturities and strike prices relative to the spot index. This is illustrated by figure 11 in appendix C, showing implied volatilities from 2008/07/02 - 2008/09/30 for ATM options with maturities of 1-month, 2-months and 3-months. However, valuation of the hedging portfolio requires daily updates of the implied volatility for the 3-month ATM option purchased at initiation of the trading period.

Cubic splines are used to recreate the series of implied volatilities. For instance, a 3-month option purchased on 2008/07/02 expires on 2008/09/30,

¹⁹Data from 2004-2009 are provided by a major commercial bank, while more recent data are retrieved from the database OptionsMetrics.

and daily updated implied volatilities for that specific option may be reconstructed with 2-dimensional interpolation and extrapolation utilizing all implied volatilities observed from 2008/07/02 - 2008/09/30²⁰.

5.2.2. Realized Volatility

For the S&P 500 index, (true) daily spot volatilities are proxied by high frequency volatility estimates from the Oxford-Man Institute of Quantitative Finance Realized Library²¹. The series of daily realized volatilities are retrieved for all trading days from 2004/04/01 - 2013/07/03. Figure 4 shows the spot index as well as realized and daily updated implied volatilities for each trading period.

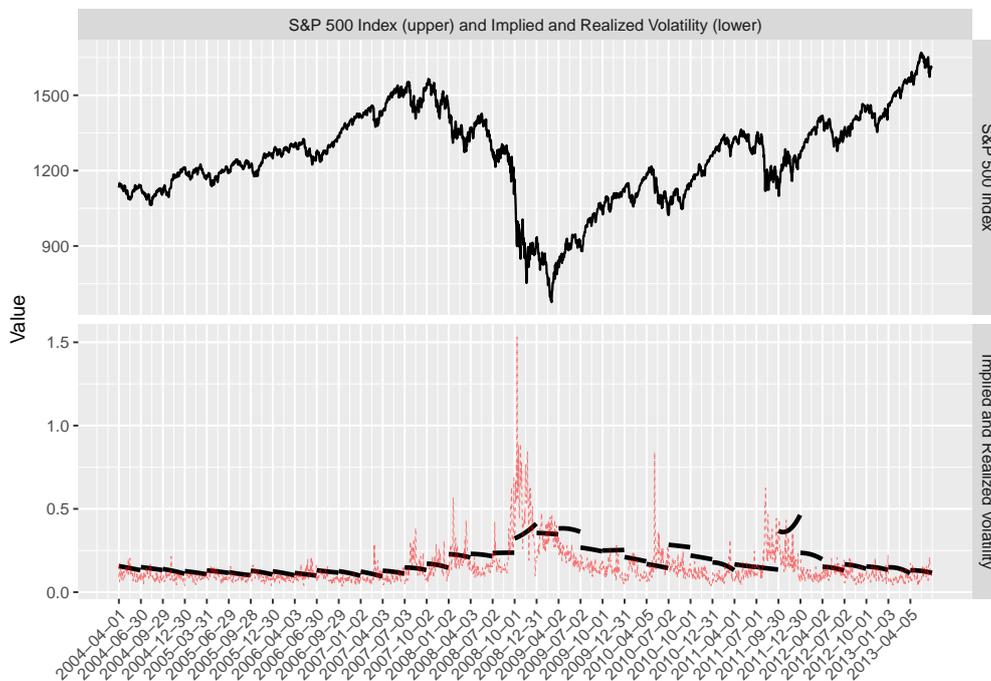


Figure 4: The upper plot shows the S&P 500 index. The lower plot shows realized volatilities (red dashed line) as well as the daily implied volatility (solid black lines) for a 3-month call option, initially purchased at-the-money.

²⁰At initiation, both the index spot and implied volatilities for ATM options with 1-month, 2-month and 3-month maturities are observed. This scenario repeats itself the next trading day. Now, interpolate first in the strike-direction and then in the maturity-direction to obtain a daily updated estimate of the implied volatility for the 3-month ATM option purchased at initiation.

²¹Daily realized variance estimates may be downloaded from: <https://realized.oxford-man.ox.ac.uk/data/download>.

5.3. Intraday Stock Data

The SPDR S&P 500 EFT traded on NYSE Arca with ticker SPY is used to track the S&P 500 index. Among ETFs, the SPY is one of the most actively traded and is widely considered as the best to track the S&P 500 index. Level 1 tick data from 2004-2013 are retrieved from the TAQ (trades and quotes) database, and data cleaning follows the procedure described by Holden and Jacobsen (2013). Additionally, only transactions during the core trading session (9:30 a.m. to 4:15 p.m. ET) are considered²². Table 7 in appendix C reports summary statistics (#trades, total volume, daily volume, shares per trade and price) of the data for each of the 37 trading periods. After 2004/09/29, the average daily volume was never below 10 million, peaking around and subsequent to the Great Recession in 2008 with almost 100 million shares per day (in average).

5.3.1. Data Aggregation: Time Bars

As explained in the previous section, the first step of computing the volume-synchronized probability of informed trading is to determine buy and sell volume. This paper uses the bulk volume classification, grouping sequential trades into time bars. Buy and sell volume is then computed as a fraction of total volume. The buy fraction for bar i is simply given by $N(\Delta P_i / \sigma_{\Delta P})$, where $N(\cdot)$ is the CDF of the standard normal distribution, $\Delta P_i = P_i - P_{i-1}$ is the price change between bar i and bar $i - 1$, while $\sigma_{\Delta P}$ is the standard deviation of price changes between bars. A limitation of the data is that trading is closed between days opposite to the empirical study by Easley et al. (2013) that uses futures traded around the clock. In particular, the price change between two consecutive bars overnight is expected to be higher than the price change between two consecutive bars of the same trading day. This problem is circumvented by omitting price changes overnight²³.

²²Corrected trades and trades with an abnormal sale condition are also removed from the dataset. Only trades from NYSE Arca are considered.

²³In the beginning of day d , the first trade is not included in a time bar but instead used to compute the price change for the initial time bar. Bulk classification is then applied on all transactions of day d . The procedure is then repeated for the next day.

6. Numerical Results

The hedging experiment is constructed in the spirit of Ellersgaard et al. (2017) and Leth (2015). In particular, partial results on the daily hedging outcome are related to results reported by Leth (2015)²⁴.

At initiation of trading period j , the European call is shorted, the hedging portfolio contains Δ_1 units long in the underlying, and the money account chosen in accordance with the self-financing condition, i.e.,

$$\Pi_1^h = -C_1 + \Delta_1 S_1 + B_1 = 0.$$

This procedure is repeated for the remaining trading days $k = 2, \dots, n_j$ of period j . For each $k \in \{2, \dots, n_j\}$, the daily mark-to-market value of the hedging portfolio $\tilde{\Pi}_k$ and the gain/loss between today and yesterday $dP\&L_k$ are computed, and finally the portfolio is rebalanced to meet the self-financing condition, that is,

$$\begin{aligned} \tilde{\Pi}_k &= -C_k + \Delta_{k-1} S_k e^{q_{k-1}(t_k - t_{k-1})} + B_{k-1} e^{r_{k-1}(t_k - t_{k-1})} \\ dP\&L_k &= \tilde{\Pi}_k - \Pi_{k-1}, \end{aligned}$$

and B_k solves $\Pi_k = -C_k + \Delta_k S_k + B_k = \tilde{\Pi}_k$. At maturity n_j , the terminal hedge error $\tilde{\Pi}_{n_j}$ and an estimate of the portfolio's quadratic variation over the option's lifetime, $QV_j = \sum_{k=1}^{n_j} |dP\&L_k|^2$, are both reported. The experiment is conducted for all 37 trading periods, hedging with both realized and implied volatility used as input in the hedging formula²⁵.

6.1. Profit-and-Losses

As a starting point, the upper plot in figure 5 shows the P&L on a day-to-day basis for all 37 hedging portfolios using the risk-minimizing strategy hedging with implied volatility.

²⁴Leth (2015) examines the terminal P&L under different hedging models where the objective is to maximize the Sharpe ratio. The purpose of this hedging experiment is to illustrate that the optimal risk-minimizing strategy can not prevent significant daily losses of the hedging portfolio. Moreover, the deterministic findings differ since both the hedging periods and the volatility series are different.

²⁵Also, model estimates reported by Eraker (2004) are used in the hedging formula.

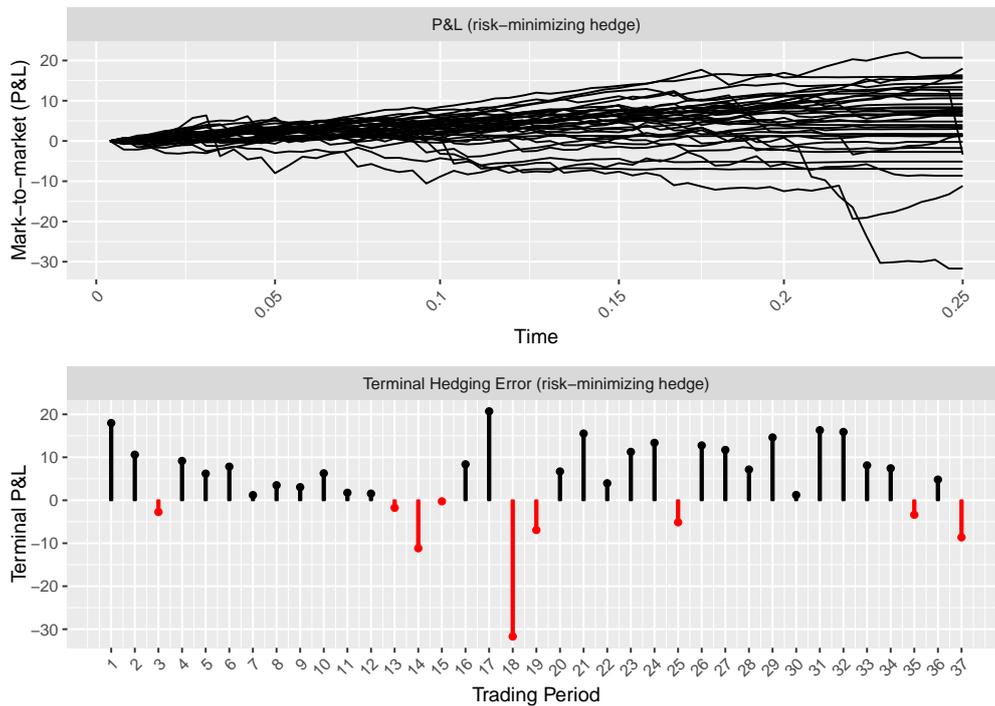


Figure 5: Top: P&Ls on a day-to-day basis for each of the 37 hedging portfolios. The locally risk-minimizing Δ -hedge is used with implied volatility. Bottom: Corresponding terminal hedging errors for each of the 37 portfolios.

In most of the cases, the terminal hedge error is positive (28), and the option dealer succeeded picking up the risk premium. However, 9 hedging portfolios end up on the *wrong* side. In particular, one trading period is associated with a major loss as well as significantly fluctuations in the value difference between consecutive trading days. The lower plot reveals that this is period 18 (2008-07-02 - 2008-09-30). Finally, panels in figure 5 indicate that hedging with implied volatility results in smoother P&L paths, which is consistent with the numbers reported in the last two columns in table 8 in appendix C. Consequently, hedging with implied volatility will produce the smallest quadratic variation (on average). This outcome is supported by table 3, reporting summary statistics of terminal hedge errors and quadratic variations for combinations of hedging strategies and hedging volatilities.

Table 3: SUMMARY OF HEDGE ERRORS AND QUADRATIC VARIATIONS

(a) Risk-Minimizing Hedge						
Hedging Volatility	$\overline{P\&L}$	s.e. $(\overline{P\&L})$	t value	Pr($> t $)	\overline{QV}	s.e. (\overline{QV})
Realized	5.17	1.85	2.80	0.008	2.04	0.82
Implied	4.80	1.61	2.97	0.005	0.71	0.16

(b) Δ -Hedge (Black-Scholes)						
Hedging Volatility	$\overline{P\&L}$	s.e. $(\overline{P\&L})$	t value	Pr($> t $)	\overline{QV}	s.e. (\overline{QV})
Realized	5.13	2.26	2.27	0.029	3.37	1.52
Implied	4.18	2.09	2.00	0.053	0.55	0.16

Description: Summary statistics of terminal hedge errors and quadratic variations. Upper table: Estimates obtained using the locally risk-minimizing hedge. Bottom table: Estimates obtained using the standard Δ -hedge.

The risk-minimizing hedge is preferred over the Δ -hedge in terms of standard deviations of terminal hedge errors. Clearly, a skilled forecaster may be rewarded by hedging with actual volatility due to a higher P&L. However, this comes with a cost in terms of a more volatile hedging portfolio measured by the quadratic variation as well as a higher standard deviation of the P&L. For instance, consider the hypothesis $H_0 : P\&L_T^h = 0$ with alternative hypothesis $P\&L_T^h \neq 0$. Under the risk-minimizing hedge, the two-sided t -test leads to p -values of 0.008 and 0.005 for actual volatility and implied volatility, respectively — profit is obtained at the 0.01 level.

In the remainder of this paper, only findings under the risk-minimizing strategy hedging with implied volatility are considered. This strategy was optimal in a risk-minimizing setting but still could not prevent significant losses from 2008/07/02 - 2008/09/30²⁶.

6.2. VPIN

The VPIN series are computed across all 37 periods, where bulk volume classification is used to determine the order flow. Tables 12–17 in appendix C report summary statistics of the volume buckets and the VPIN series under different parameter settings and bar sizes across trading periods.

²⁶All hedging strategies realized major losses in that period cf. table 8-10.

It appears that

- The average, median and standard deviation of the VPIN are increasing in the bar size
- The VPIN is decreasing in the volume bucket size consistent with the findings by Abad-Diaz and Yagüe (2012)²⁷.
- The average and the median of the VPIN are approximately for all combinations of (n, VBS) (see figure 6).

In the remainder of this study, only empirical findings under the specification $(n, VBS) = (100, 1/100)$ are presented. The results and conclusions obtained with the common specification $(n, VBS) = (50, 1/50)$ are very similar and reported in appendix C.

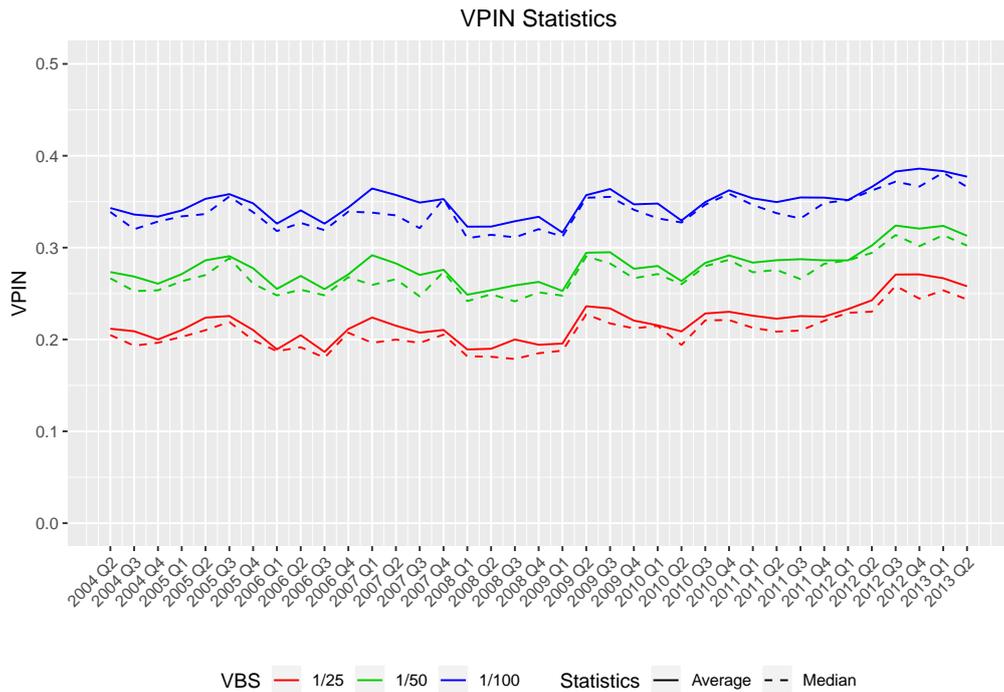


Figure 6: Average and median of the VPIN series across trading periods under the specifications $(25, 1/25)$, $(50, 1/50)$, and $(100, 1/100)$, respectively.

²⁷They argue that "VPIN specifications employing lower VBS to compute order imbalance may incorporate transitory as well as permanent information, whereas VPIN specifications that consider higher VBS to compute order imbalance may mainly incorporate fundamental information about stocks (as in PIN model).".

6.2.1. Signaling Losses (2008/07/02 – 2008/09/30)

The top plot in figure 7 shows daily gains/losses for the hedging portfolio of period 2008/07/02 – 2008/09/30 (trading period 18). The red points indicate daily losses greater than 10% relative to the initial valuation of the European option (see table 8). In particular, the hedging portfolio was subject to significant losses from 2008/09/15 – 2008/09/22.

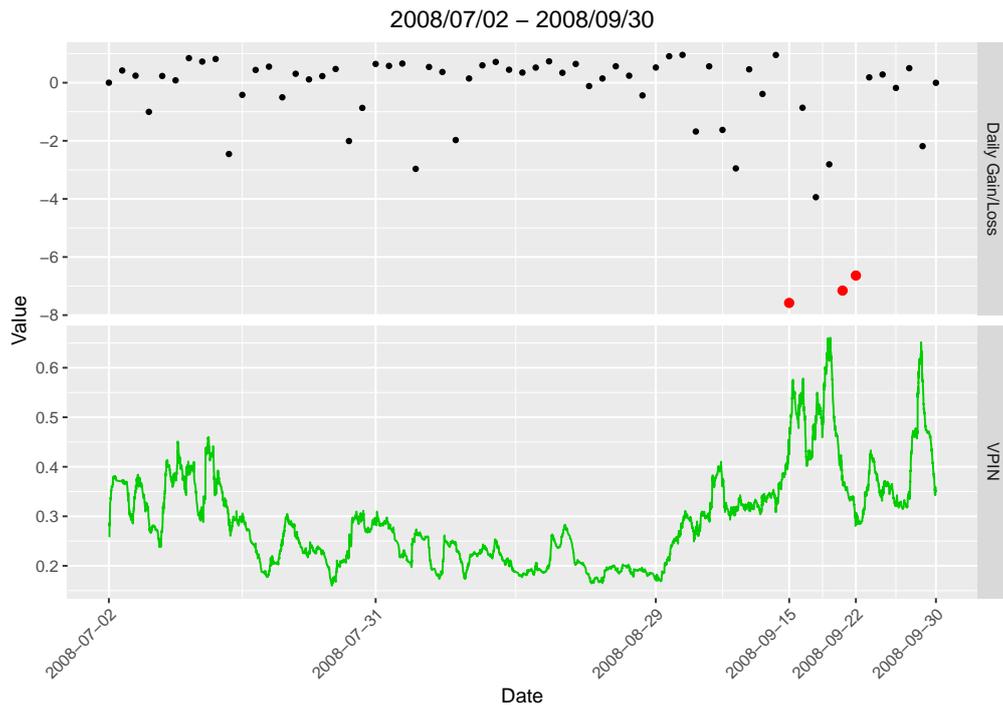


Figure 7: Top plot: Daily gains/losses for hedging portfolio 18 (2008/07/02 – 2008/09/30). Bottom plot: The VPIN during the same period.

To investigate if losses during bad days were anticipated, the VPIN series for the relevant period are considered. The bottom plot in figure 7 indicates a dependence between significant losses in the hedging portfolio and high readings of the VPIN. The performance evaluation of the VPIN described in section 4.1 enables determining whether the VPIN signaled or was (implicitly) impacted by losses in the hedging portfolio.

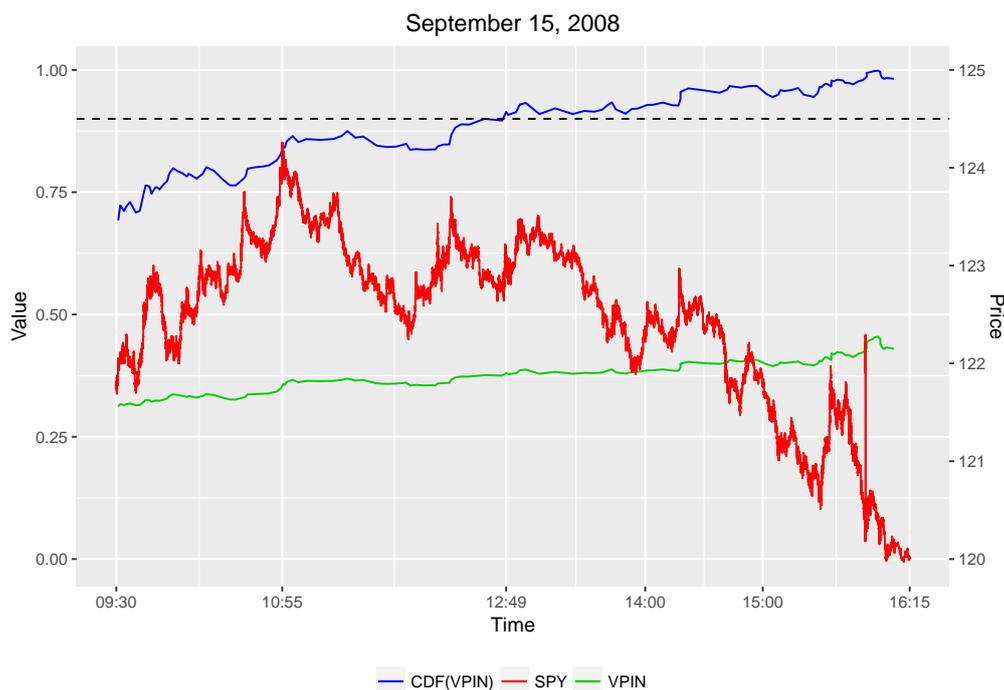


Figure 8: CDF(VPIN), VPIN, and price of the SPY on September 15, 2008. The left axis shows the values of CDF(VPIN) and the VPIN, while the right axis shows the value of the SPY. The dashed line represents the threshold 0.9.

Figure 8 shows the VPIN, the CDF(VPIN)²⁸ and the price of SPY on September 15, 2008. The VPIN is steadily increasing through the day, whereas the SPY hits its maximum at 10:55 and exhibits extreme fluctuations from 14:18 and until the end of the day. Compared to this, CDF(VPIN) reaches the threshold 0.9 at 12:49 and stays above 0.9 in the remainder of the day. A VPIN event was detected before extreme intraday price movements. This observation is supported by numerical evidence. The daily average of negative (positive) MIR readings of the option's lifetime was $\underline{M} = -0.022$ ($\overline{M} = 0.027$), whereas the MIR between the VPIN event and the final trade of the day was -0.029 , i.e., a true positive.

Moreover, figure 12 in appendix C shows that the period from 2008/09/15-2008/09/22 was characterized by a VPIN event until 12:10 at 2008/09/19. It appears that order flow became and stayed toxic in the vast majority of this period except for trading day 2008/09/22. This observation is

²⁸The CDF is estimated from the out-of-sample empirical distribution.

supported by numerical evidence in table 4, reporting daily values of P&L, valuation of the option, standardized moneyness²⁹, risk-minimizing Δ , realized volatility relative to implied volatility, MIR and classification of events from 2008/09/12-2008/09/22.

Table 4: P&L AND MIR

	$dP\&L$	C	$m(\nu^i, T - t)$	Δ^{\min}	$\sqrt{\nu^r}$	MIR	VPIN-event
2008-09-12	0.95	21.46	-0.151	0.425	0.239	0.017	TN
2008-09-15	-7.58	3.93	-1.134	0.098	0.462	-0.029	TP
2008-09-16	-0.86	6.84	-0.821	0.179	0.512	0.044	TP
2008-09-17	-3.94	0.54	-1.944	0.012	0.492	-0.036	TP
2008-09-18	-2.81	3.95	-1.057	0.120	0.630	0.070	TP
2008-09-19	-7.16	16.92	-0.130	0.438	0.525	-0.090	TP
2008-09-22	-6.64	2.52	-1.207	0.092	0.268	-0.035	FN

Description: Summary statistics of the hedging portfolio, MIR values and VPIN events from 2008/09/12-2008/09/22.

September 12 is described by quiet market conditions. The daily MIR was low and no VPIN event was detected. Additionally, realized volatility was approximately equal to the option's implied volatility. The standardized moneyness close to 0 reflects the size of Δ^{\min} . The following five days are characterized by extreme volatility ($\sqrt{\nu^r}$) and losses in the hedging portfolio. Prices rose when Δ^{\min} was low and dropped when Δ^{\min} was high. For all five days, high VPIN readings led to corresponding high MIR values (absolute terms), and events were classified as true positives. On 2008-09-19, CDF(VPIN) fell and stayed below the threshold 0.9 (see figure 9). However, figure 13 in appendix C reveals that large price movements also occurred on 2008-09-22, consistent with the huge loss (-6.64). This outcome is only captured by a high MIR value — neither the VPIN nor realized volatility indicates market turbulence. In this case, large price movements were not caused by toxic order flow but market uncertainty³⁰. Importantly, using the VPIN as a warning signal may imply type II errors.

²⁹The standardized moneyness is defined by $m(\nu^i, T - t) = \log(S/K)/(\sqrt{\nu^i(T - t)})$, i.e, the number of standard deviations the current spot price is above the strike.

³⁰Prices were constantly decreasing during the day due to uncertainty about the government's \$700 billion bailout plan after the bankruptcy of Lehman Brothers.

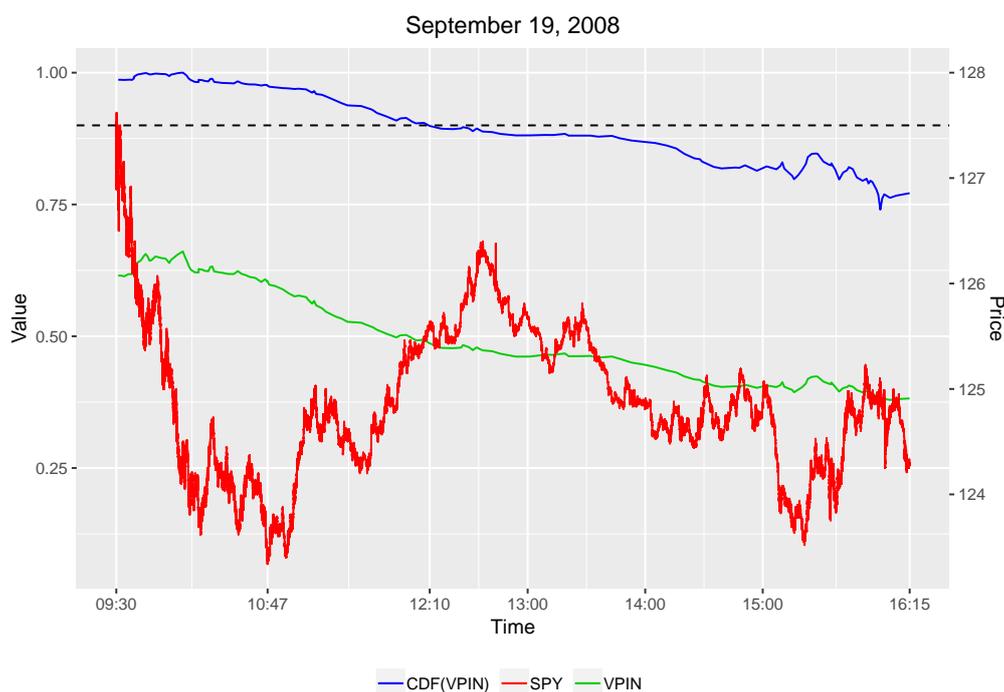


Figure 9: CDF(VPIN), VPIN, and price of the SPY on September 19, 2008. The left axis shows the values of CDF(VPIN) and the VPIN, while the right axis shows the value of the SPY. The dashed line represents the threshold 0.9.

6.2.2. Signaling Losses: May 6, 2010 (*The Flash Crash*)

A final study of the Flash Crash on May 6, 2010, is conducted. This date has been the root of the disagreement between Easley et al. (2012, 2010) and Andersen and Bondarenko (2014a,b, 2015). Did the VPIN anticipate the temporary crash on May 6? Previous studies are based on the E-mini S&P 500 futures contract traded around the clock, and the trading period — affecting both the BVS-scheme and *VBS* — is chosen in order to analyze the Flash Crash. In contrast, this paper relies on the ETF SPY, and the choice of trading period simply reflects the option's lifetime (quarterly periods). Additionally, this study illustrates the impact of the Flash Crash on the hedging portfolio.

Figure 10 demonstrates the VPIN, CDF(VPIN) and the SPY price on May 6, 2010. The SPY was subject to significant losses and crashed at 14:46, whereas the VPIN was steadily increasing through the day, and CDF(VPIN) reached the threshold 0.9 at 12:03. From 14:27 and some time after the crash, CDF(VPIN) was close to 1. Between the VPIN event

at 12:03 and the final trade of the day, MIR was equal to -0.096 . In comparison, the daily average of negative MIR values during the hedging period was -0.025 . Table 5 reveals that the hedging portfolio was subject to a loss of 1.50 (4.1% of the initial price of the European call) on May 6, 2010. This loss was the second largest during the option's lifetime and caused by large intraday price movements foreseen by the VPIN event (true positive). Additionally, the fast recovery of the crash prevented an even greater loss.

Table 5: P&L AND MIR (THE FLASH CRASH)

	$dP\&L$	C	$m(\nu^i, T - t)$	Δ^{\min}	$\sqrt{\nu^r}$	MIR	VPIN-event
2010-05-05	0.26	19.43	-0.285	0.384	0.199	0.016	TN
2010-05-06	-3.14	8.12	-0.810	0.208	0.839	-0.096	TP

Description: Summary statistics of the hedging portfolio, MIR values and VPIN events from 2010/05/05-2010/05/06.

The findings in figure 10 also support the study by Easley et al. (2012) arguing that extreme values of the VPIN (given by the $CDF(VPIN)$) anticipated the crash. Andersen and Bondarenko (2014b) dispute this since the VPIN kept rising after the crash. However, the metric is updated in real time and is based on order imbalances. If volume falls but imbalance remains high, the VPIN will increase. In contrast to the moment estimation of VPIN, maximum likelihood estimation of the metric proposed by Lin and Ke (2017) is capable of capturing this phenomenon³¹. Applying their methodology, Leth (2019) shows that the toxicity of order flow was falling ex-post the crash³².

³¹Maximum likelihood estimation of the VPIN also captures the information in volume time.

³²In particular, the VPIN peaked ex ante the crash, and the level of informed trading was decreasing after the crash.



Figure 10: CDF(VPIN), VPIN, and price of the SPY on May 6, 2010. The left axis shows the values of CDF(VPIN) and the VPIN, while the right axis shows the value of the SPY. The dashed line represents the threshold 0.9.

6.3. Forces at Work: Improving the Hedging Strategy with VPIN

The empirical findings suggest that volume-synchronized probability of informed trading anticipated daily losses in the hedging portfolio. In this setting, it is assumed that the PM rebalances her portfolio once at the end of day. Relaxing this assumption, enables the PM to reduce the losses on days where VPIN events are detected.

Expanding the Hedging Portfolio with VIX

Tables 4—5 showed that losses in the hedging portfolio occurred on days with extreme levels of realized volatility. In addition, on these days, high VPIN readings led in all cases to significant large maximum intermediate returns; that is, the VPIN signaled high daily volatility. This observation suggests that expanding the portfolio with long VIX futures when VPIN events are detected will reduce downside risk in the hedging portfolio³³. The benefit of this strategy is illustrated in figures 16-17 in appendix C:

³³Easley et al. (2011) propose a similar strategy used for volatility arbitrage.

On September 15, 2008, and May 6, 2010, respectively, the volatility index increased dramatically after the VPIN event occurred; profit could have been made by buying VIX.

Adjusting Δ

In this setting, a portfolio manager applies the risk-minimizing Δ -hedge to rebalance her portfolio one time at the end of day. As theorem 4 states, the hedging portfolio is always exposed to losses due the random nature of volatility. However, the discretization of time will also impact the daily P&L. The numbers reported in table 4 and 5 show that extreme intraday return patterns led to significant losses. The explanation is that the portfolio manager does not adjust her Δ in continuous time, and the option's convexity will favor itself. Two undesirable scenarios are observed from table 4:

- Intraday falling prices can amplify losses of the hedging portfolio when the PM invests heavily in the underlying; that is, when Δ^{\min} is sufficiently large.
- Intraday rising prices can amplify losses of the hedging portfolio when the PM is too precautionary; that is, when Δ^{\min} is sufficiently small.

The two cases above may be prevented by immediately readjusting the Δ^{\min} when observing a VPIN event at time t . Table 4 showed that high VPIN readings in most cases signaled extreme MIR values that reflect large intraday price movements (and/or high realized volatility). Furthermore, this situation will affect the Δ^{\min} , with impact determined by the second-order Greek gamma, $\partial_S \Delta^{\min} > 0$. Thus, in the case of low risk exposure, $\Delta_t^{\min} < k$, and a VPIN event at time t , the PM may consider to immediately rebalance her portfolio with the modified delta

$$\Delta_t^{\text{VPIN}} := \Delta_t^{\min} \cdot (1 + 1_{[0,k]}(\Delta_t^{\min}) \cdot g(\text{CDF}(\text{VPIN}_t))) \quad (6.1)$$

where $g : (0, 1] \rightarrow [0, (1 - \Delta_t^{\min})/\Delta_t^{\min}]$ is an increasing function, and k is sufficiently small.³⁴

³⁴Clearly, a similar expression for a modified delta can be derived in the case of high risk exposure.

The expression (6.1) can be perceived as a forecast/expectation of the future value of Δ_t^{\min} , and the economic reasoning is that the PM increases her exposure when risk premiums are high.

It is important to emphasize that the performance of the proposal above depends on the direction of subsequent price movements. However, moment estimation of the VPIN proposed by Easley et al. (2012) ignores the signed order imbalance when updating the metric in real time. This problem may be circumvented when using a more sophisticated estimation method of the VPIN. Lin and Ke (2017) provide a parametric framework for estimating the VPIN. Their method captures the information from volume time (time taken to fill a bucket), and generates consistent estimates of the model parameter for the underlying model, including the bucket-updated probability of good-news, which reflects the direction in which the market is moving. Combining intraday estimates of this parameter with a potential VPIN event enables the PM to benefit from the above strategy. Such an experiment is beyond the scope of this paper, but Leth (2019) conducts a large empirical study on maximum likelihood estimation of the VPIN for SPY from 2007-2015. Using findings from that study will form the basis of a new experiment, investigating the numerical outcome from adjusting the risk-minimizing delta when the VPIN signals market turbulence.

7. Conclusion

This paper investigates how to eliminate downside risk when delta hedging European options, particularly considering whether risk tools from high-frequency market microstructure can improve sophisticated hedging strategies from mathematical finance.

The problem is presented from the portfolio manager's (PM's) point of view, who takes a short position in the European call. By engaging in dynamic delta hedging, she may take advantage of the excess spread between implied and realized volatility and collect this volatility risk premium. Furthermore, the PM is restricted to rebalance her portfolio only once at the end of each the trading day. The market is incomplete, and therefore, the terminal hedging error becomes random. In addition, the discretization of time adds extra uncertainty to this hedging error.

First, the paper seeks to answer whether a risk-minimizing strategy from the Heston setting outperforms the standard Black-Scholes delta hedge, assessed by the return and standard deviation of the hedging portfolio; second, the paper aims to determine whether the volume-synchronized probability of informed trading (VPIN), used as predictor of short-term return volatility, can signal impending losses of the hedging portfolio.

Market data on 3-month at-the-money European call options written on the SPX from 2004-2013 are used in an empirical hedging experiment. The PM takes a short position in the option, and creates her delta hedged portfolio. The terminal P&L will in most cases end up positive, and profit is obtained at the 0.01 level. In addition, the risk-minimizing strategy from the Heston setting is superior to the Black-Scholes Δ -hedge in a risk-minimizing context. However, the risk-minimizing hedge is not able to prevent significant losses in the hedging portfolio during extreme market turbulence on trading days of September 2008.

This paper contributes with an empirical investigation of the VPIN in a new setting. The study is conducted using the ETF SPY for quarterly periods, reflecting the option's lifetime. The findings show that high VPIN readings signaled large intraday price movements on days, where the portfolio was subject to losses. These losses occurred when the PM either had high market exposure and intraday prices were falling, or low market exposure and intraday prices were rising. In particular, if the portfolio manager had reacted to warning signals from the VPIN, losses could have been prevented.

Eventually, two suggestions for reducing losses in the hedging portfolio using the VPIN are discussed. The first strategy is to expand the hedging portfolio with VIX futures immediately after a VPIN is detected. The second strategy is to scale market exposure in the underlying asset, conditional on the VPIN event. Arguably, this strategy must rely on the parametric estimation method of the VPIN developed by Lin and Ke (2017) that enables the portfolio manager to infer the direction in which the market is moving.

Trivially, inventory risk is of great importance for investors, supplying options to the market. In future research, it is therefore pertinent to conduct a study that investigates the performance of the two above strategies.

Appendix A Proofs

PROOF OF COROLLARY 2. Consider a European call with strike αK written on the underlying $(\alpha S_t)_{t \geq 0}$, with α and K being positive constants. Risk-neutral valuation yields

$$C(t, \alpha S_t, \nu_t) = \alpha S_t p_1(t, x_t, \nu_t) - \alpha e^{-r(T-t)} K p_2(t, x_t, \nu_t) = \alpha C(t, S_t, \nu_t),$$

and the price function is homogeneous of degree one with respect to S_t and K . Thus, Euler's homogeneous function theorem may be applied to obtain

$$\begin{aligned} C(t, S_t, \nu_t) &= S_t \partial_s C(t, S_t, \nu_t) + K \partial_K C(t, S_t, \nu_t) && \Leftrightarrow \\ \partial_s C(t, S_t, \nu_t) &= \frac{C(t, S_t, \nu_t) - K \partial_K C(t, S_t, \nu_t)}{S_t} \end{aligned}$$

Now, Leibniz's integral rule yields

$$\begin{aligned} \partial_K C(t, S_t, \nu_t) &= \partial_K e^{-r(T-t)} E^Q((S_T - K)^+) \\ &= e^{-r(T-t)} \int_K^\infty \partial_K(y - K) q(y) dy \\ &= -e^{-r(T-t)} p_2(t, x_t, \nu_t). \end{aligned}$$

Combining these expressions, the Δ of the option is given by

$$\Delta_t = \partial_s C(t, S_t, \nu_t) = p_1(t, x_t, \nu_t). \quad (\text{A.1})$$

For deriving vega, realize that

$$\partial_\nu C(t, S_t, \nu_t) = S_t \partial_\nu p_1(t, x_t, \nu_t) - e^{-r(T-t)} K \partial_\nu p_2(t, x_t, \nu_t).$$

Combining Leibniz's integral rule with Cauchy-Riemann equations results in semi-closed expressions for the two derivatives:

$$\begin{aligned} \partial_\nu p_1(t, x_t, \nu_t) &= \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[d(\phi - i) \frac{e^{-i\phi \log K} \tilde{f}(\phi)}{i\phi} \right] d\phi \\ \partial_\nu p_2(t, x_t, \nu_t) &= \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[d(\phi) \frac{e^{-i\phi \log K} f(\phi)}{i\phi} \right] d\phi. \end{aligned}$$

Summarizing, the vega of the option is given by

$$\mathcal{V}_t = \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[d(\phi - i) \frac{e^{-i\phi \log K} \tilde{f}(\phi)}{i\phi} \right] - e^{-r(T-t)} \frac{K}{S_t} \operatorname{Re} \left[d(\phi) \frac{e^{-i\phi \log K} f(\phi)}{i\phi} \right] d\phi.$$

The corollary is proved. \square

PROOF OF PROPOSITION 3 (HEURISTIC). El Karoui et al. (1997) show how to determine the locally risk-minimizing hedge. Otherwise, consider the self-financing hedging portfolio with the short position in the European call, and long Δ_t in the underlying. The value process of this portfolio, Π_t , has dynamics given by

$$d\Pi_t = (\dots)dt + (\Delta_t - \partial_s C(t, S_t, \nu_t))dS_t + \partial_\nu C(t, S_t, \nu_t)d\nu_t.$$

Taking the conditional variance of this expression yields

$$\begin{aligned} \text{Var}_t(d\Pi_t) = & \left\{ (\Delta_t - \partial_s C(t, S_t, \nu_t))^2 S_t^2 \nu_t + (\partial_\nu C(t, S_t, \nu_t))^2 \sigma^2 \nu_t \right. \\ & \left. + 2(\Delta_t - \partial_s C(t, S_t, \nu_t))\partial_\nu C(t, S_t, \nu_t)S_t \sigma \rho \nu_t \right\} dt, \end{aligned}$$

which is convex in $\Delta(t)$, and thereby minimized locally in time by

$$\Delta_t^{\min} = \partial_s C(t, S_t, \nu_t) + \frac{\rho\sigma}{S_t} \partial_\nu C(t, S_t, \nu_t)$$

as promised. □

PROOF OF THEOREM 4. Consider the self-financing hedging portfolio with the short position in the European call, and long Δ_t^{\min} units in the underlying, where

$$\Delta_t^{\min} = \partial_s C(t, S_t, \nu_t^h) + (\rho\sigma/S_t) \cdot \partial_\nu C(t, S_t, \nu_t^h),$$

and $\sqrt{\nu^h}$ denotes the hedging volatility. The desired result is obtained by combining dynamics of the self-financing condition, dynamics of the pricing function, and Heston's PDE, respectively.

First, the money account B satisfies by construction the self-financing condition

$$\Pi_t^h = -C(t, S_t, \nu_t^i) + (\partial_s C(t, S_t, \nu_t^h) + (\rho\sigma/S_t) \cdot \partial_\nu C(t, S_t, \nu_t^h)) S_t + B_t = 0.$$

Applying Itô's formula on the above expression yields

$$\begin{aligned} d\Pi_t^h &= -dC^i(t, S_t, \nu_t) + (\partial_s C(t, S_t, \nu_t^h) + (\rho\sigma/S_t) \cdot \partial_\nu C(t, S_t, \nu_t^h)) dS_t + dB_t \\ &= -dC^i(t, S_t, \nu_t) + rC^i(t, S_t, \nu_t)dt \\ &\quad + (\partial_s C(t, S_t, \nu_t^h) + (\rho\sigma/S_t) \cdot \partial_\nu C(t, S_t, \nu_t^h)) (dS_t - rS_t dt). \quad (\text{A.2}) \end{aligned}$$

Second, using the bivariate version of Itô's formula on $C(t, S_t, \nu_t^h)$ implies that

$$\begin{aligned}
dC(t, S_t, \nu_t^h) &= \partial_t C(t, S_t, \nu_t^h) dt + \partial_s C(t, S_t, \nu_t^h) dS_t + \partial_\nu C(t, S_t, \nu_t^h) d\nu_t \\
&\quad + \frac{1}{2} (\partial_{ss} (dS_t)^2 + \partial_{\nu\nu} C(t, S_t, \nu_t^h) (d\nu_t)^2) + \partial_{s\nu} C(t, S_t, \nu_t^h) dS_t d\nu_t \\
&= \partial_t C(t, S_t, \nu_t^h) dt + \partial_s C(t, S_t, \nu_t^h) dS_t + \kappa\theta \partial_\nu C(t, S_t, \nu_t^h) dt \\
&\quad + \nu_t^r \left\{ \frac{1}{2} (S_t^2 \partial_{ss} C(t, S_t, \nu_t^h) + \sigma^2 \partial_{\nu\nu} C(t, S_t, \nu_t^h)) \right. \\
&\quad \left. + \sigma \rho S_t \partial_{s\nu} C(t, S_t, \nu_t^h) - \kappa \partial_\nu C(t, S_t, \nu_t^h) \right\} dt + \sigma \sqrt{\nu_t^r} dW_t^{2,Q}.
\end{aligned} \tag{A.3}$$

Third, the price of the European call satisfies Heston's partial differential equation; thus, for $\nu_t = \nu_t^h$ the following equality holds almost surely,

$$\begin{aligned}
rC(t, S_t, \nu_t^h) &= \partial_t C(t, S_t, \nu_t^h) + rS_t \partial_s C(t, S_t, \nu_t^h) + \kappa\theta \partial_\nu C(t, S_t, \nu_t^h) \\
&\quad + \nu_t^h \left\{ \frac{1}{2} (S_t^2 \partial_{ss} C(t, S_t, \nu_t^h) + \sigma^2 \partial_{\nu\nu} C(t, S_t, \nu_t^h)) \right. \\
&\quad \left. + \rho \sigma S_t \partial_{s\nu} C(t, S_t, \nu_t^h) - \kappa \partial_\nu C(t, S_t, \nu_t^h) \right\}.
\end{aligned} \tag{A.4}$$

Substituting expression (A.4) into expression (A.3) yields

$$\begin{aligned}
0 &= -dC(t, S_t, \nu_t^h) + rC(t, S_t, \nu_t^h) dt - \partial_s C(t, S_t, \nu_t^h) (rS_t dt - dS_t) \\
&\quad + (\nu_t - \nu_t^h) \left\{ \frac{1}{2} (S_t^2 \partial_{ss} C(t, S_t, \nu_t^h) + \sigma^2 \partial_{\nu\nu} C(t, S_t, \nu_t^h)) \right. \\
&\quad \left. + \sigma \rho S_t \partial_{s\nu} C(t, S_t, \nu_t^h) - \kappa \partial_\nu C(t, S_t, \nu_t^h) \right\} dt + \sigma \sqrt{\nu_t} dW_t^{2,Q}.
\end{aligned} \tag{A.5}$$

Finally, subtract equation (A.5) from the dynamics of the hedging portfolio (A.2) to obtain dynamics of Π_t^h given by

$$\begin{aligned}
d\Pi_t^h &= dC(t, S_t, \nu_t^h) - dC(t, S_t, \nu_t^i) - r(C(t, S_t, \nu_t^h) - C(t, S_t, \nu_t^i)) dt \\
&\quad + (\nu_t^h - \nu_t) \left\{ \frac{1}{2} (S_t^2 \partial_{ss} C(t, S_t, \nu_t^h) + \sigma^2 \partial_{\nu\nu} C(t, S_t, \nu_t^h) \sigma^2) \right. \\
&\quad \left. + \sigma \rho S_t \partial_{s\nu} C(t, S_t, \nu_t^h) - \partial_\nu C(t, S_t, \nu_t^h) \kappa \right\} dt \\
&\quad + \sqrt{\nu_t} \sigma \left(\rho \partial_\nu C(t, S_t, \nu_t^h) dW_t^{1,Q} - dW_t^{2,Q} \right) \\
&= e^{rt} d \left(e^{-rt} (C(t, S_t, \nu_t^h) - C(t, S_t, \nu_t^i)) \right) \\
&\quad + (\nu_t^h - \nu_t) \left\{ \frac{1}{2} (S_t^2 \partial_{ss} C(t, S_t, \nu_t^h) + \sigma^2 \partial_{\nu\nu} C(t, S_t, \nu_t^h)) \right. \\
&\quad \left. + \sigma \rho S_t \partial_{s\nu} C(t, S_t, \nu_t^h) - \kappa \partial_\nu C(t, S_t, \nu_t^h) \right\} dt \\
&\quad + \sqrt{\nu_t} \sigma \left(\rho \partial_\nu C(t, S_t, \nu_t^h) dW_t^{1,Q} - dW_t^{2,Q} \right)
\end{aligned}$$

At last, the present value of the portfolio equals $P\&L_T^{\min} = \int_0^T e^{-rt} d\Pi_t^h$ and is given by

$$\begin{aligned}
P\&L_T^{\min} &= C(0, S_0, \nu_0^i) - C(0, S_0, \nu_0^h) \\
&\quad + \int_0^T e^{-rt} \left\{ (\nu_t^h - \nu_t) \left\{ \frac{1}{2} (S_t^2 \partial_{ss} C(t, S_t, \nu_t^h) + \sigma^2 \partial_{\nu\nu} C(t, S_t, \nu_t^h)) \right. \right. \\
&\quad \left. \left. + \sigma \rho S_t \partial_{s\nu} C(t, S_t, \nu_t^h) - \kappa \partial_\nu C(t, S_t, \nu_t^h) \right\} dt \right. \\
&\quad \left. + \sqrt{\nu_t} \sigma \left(\rho \partial_\nu C(t, S_t, \nu_t^h) dW_t^{1,Q} - dW_t^{2,Q} \right) \right\}
\end{aligned}$$

using the boundary condition $C(T, S_T, \nu_T^i) - C(T, S_T, \nu_T^h) = (S_T - K)^+$. The theorem is proved. \square

Appendix B Formulas for $c(\cdot)$ and $d(\cdot)$

The two functions $c(\cdot)$ and $d(\cdot)$ in theorem 1 are given by

$$c(\phi) = r i \phi (T - t) + \frac{\kappa \theta}{\sigma^2} \left((\kappa - \rho \sigma i \phi + h)(T - t) - 2 \log \left(\frac{1 - g e^{h(T-t)}}{1 - g} \right) \right)$$

$$d(\phi) = \frac{\kappa - \rho \sigma i \phi + h}{\sigma^2} \left(\frac{1 - e^{h(T-t)}}{1 - g e^{h(T-t)}} \right)$$

with

$$h = \sqrt{(\rho \sigma i \phi - \kappa)^2 - \sigma^2(-\phi^2 - i \phi)} \quad \text{and} \quad g = \frac{\kappa - \rho \sigma i \phi + h}{\kappa - \rho \sigma i \phi - h}.$$

Appendix C Empirical Findings: Figures and Tables

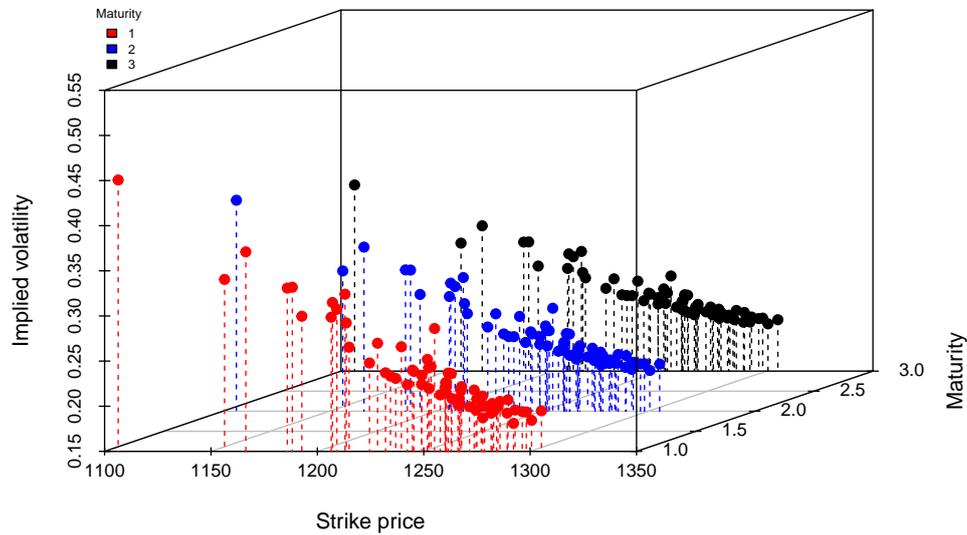


Figure 11: Daily implied volatilities for European options with maturity 1-month, 2-months and 3-months, respectively. Implied volatilities are observed from 2008/07/02 to 2008/09/30 (63 trading days and 189 observations).

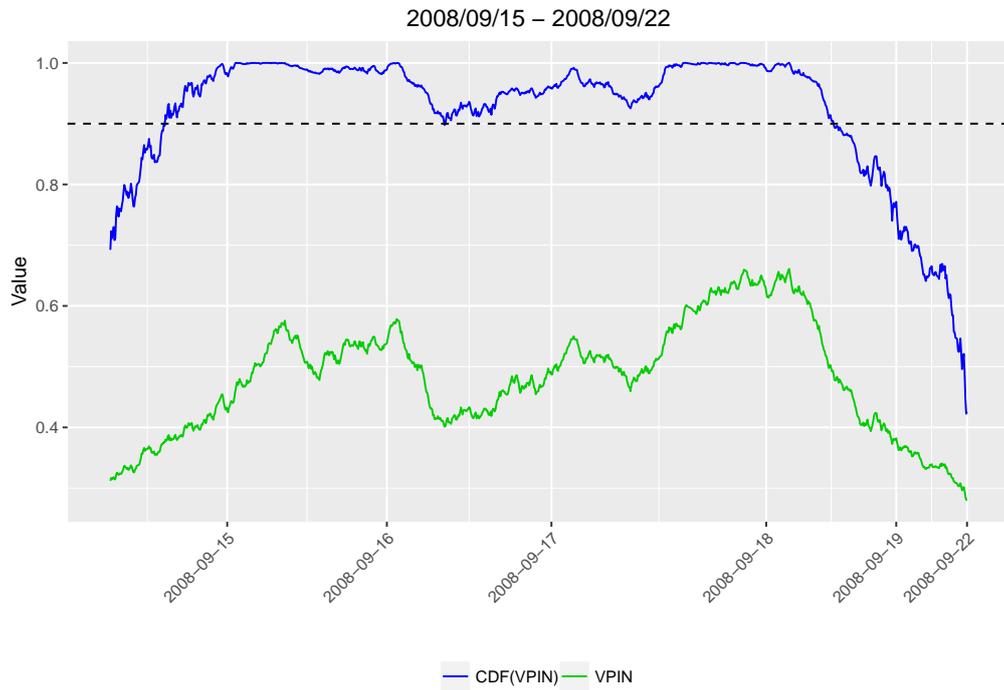


Figure 12: CDF(VPIN) and the VPIN from September 15, 2008, to September 22, 2008.

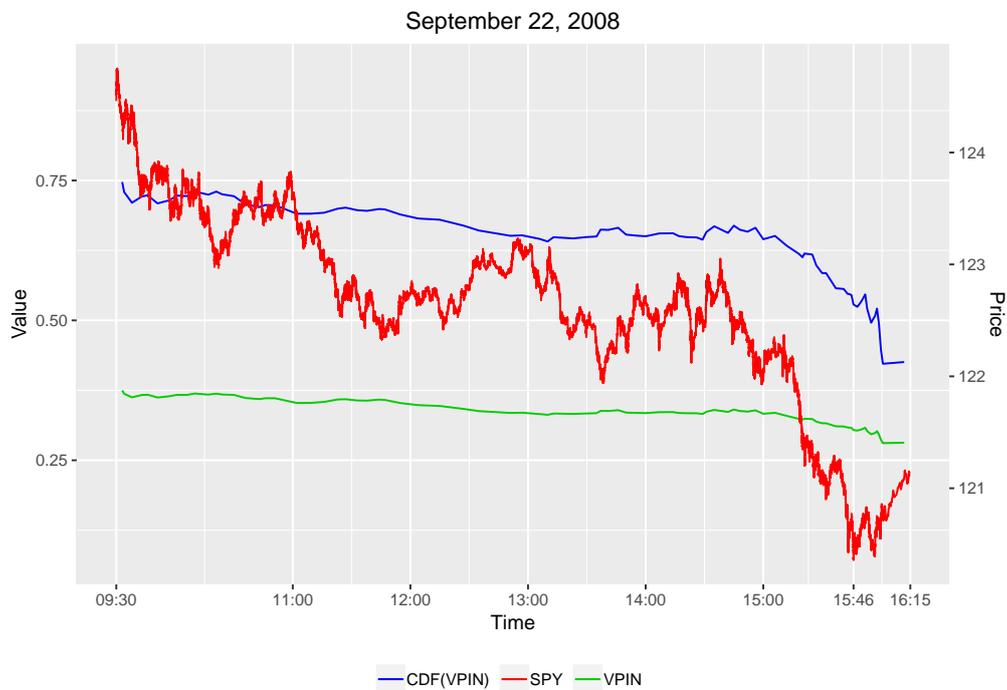


Figure 13: CDF(VPIN), the VPIN, and the price of the SPY on September 22, 2008.

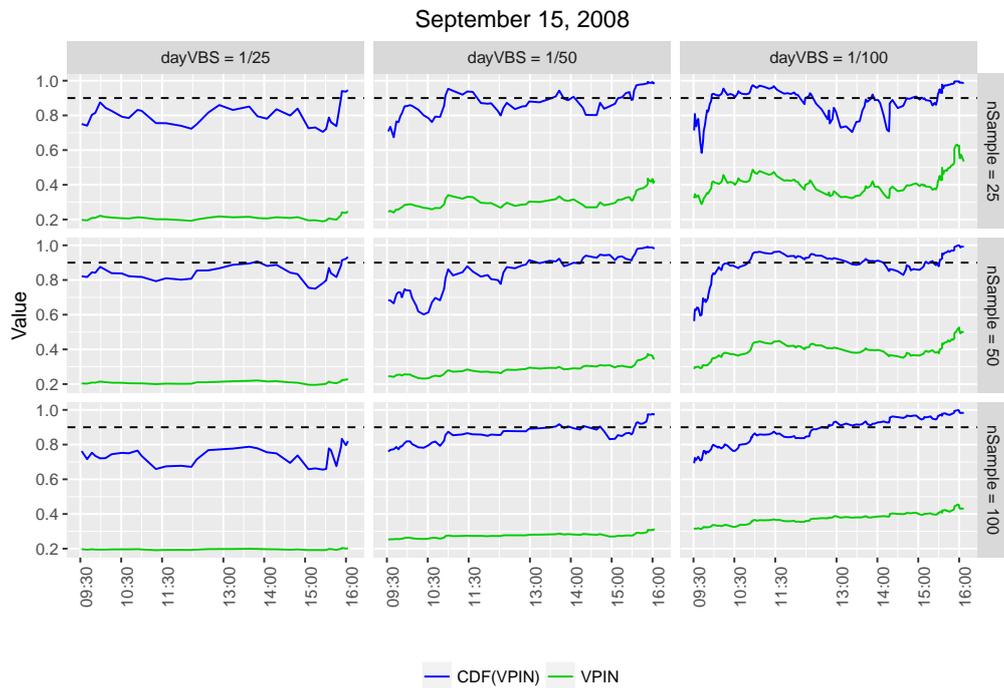


Figure 14: CDF(VPIN) and the VPIN on September 15, 2008, under various VPIN parameter specifications.

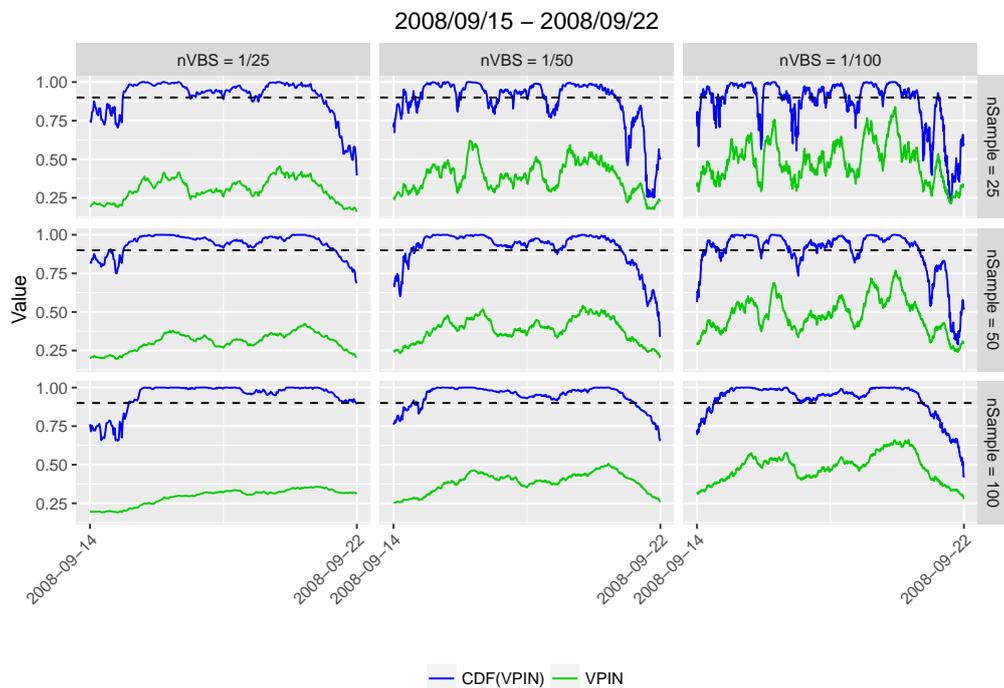


Figure 15: CDF(VPIN) and the VPIN from September 15, 2008, to September 22, 2008, under various VPIN parameter specifications.

VIX on September 15, 2008

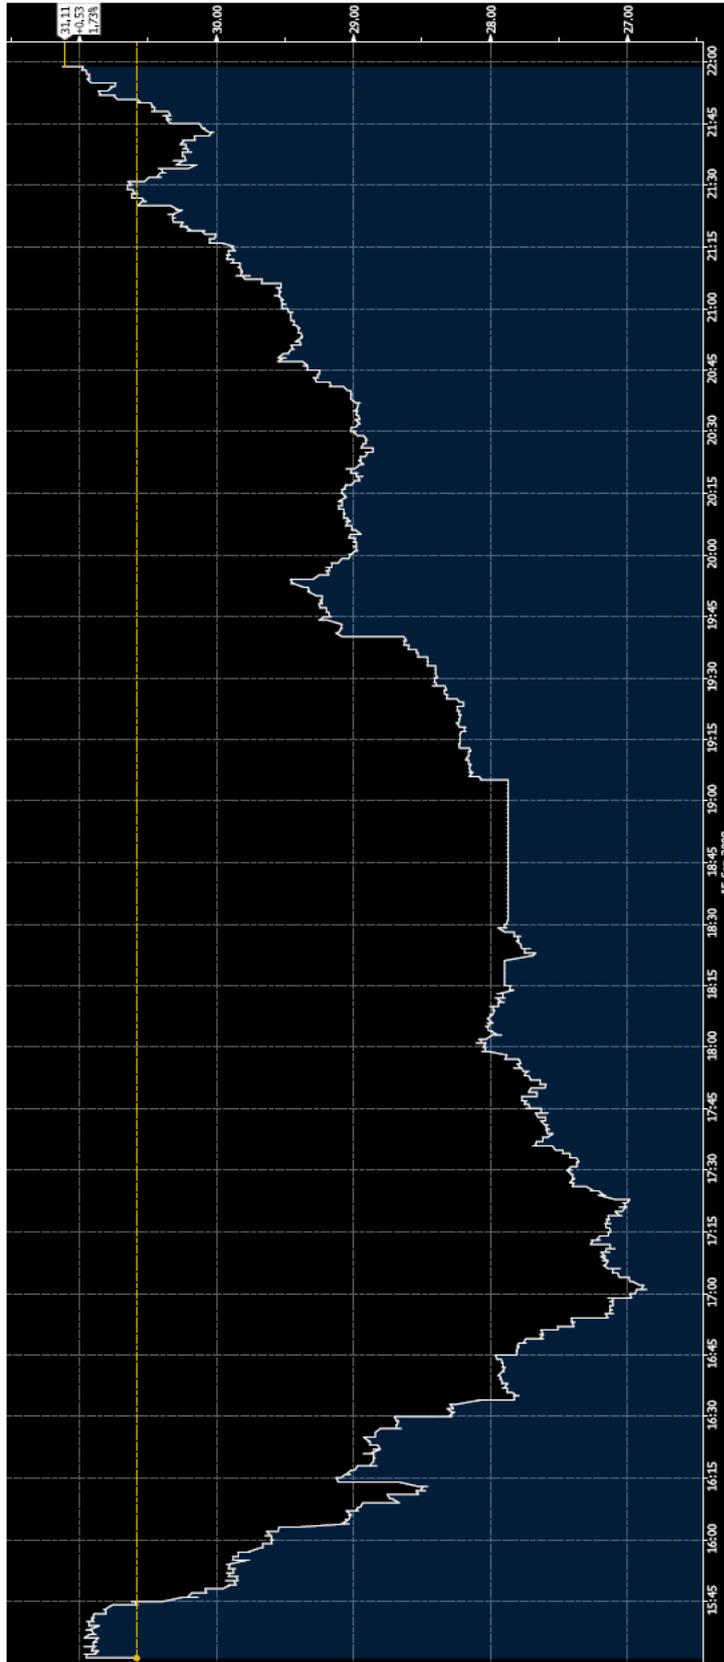


Figure 16: The VIX volatility index on September 15, 2008. The time zone is CEST. Source: Bloomberg.

VIX on May 6, 2010

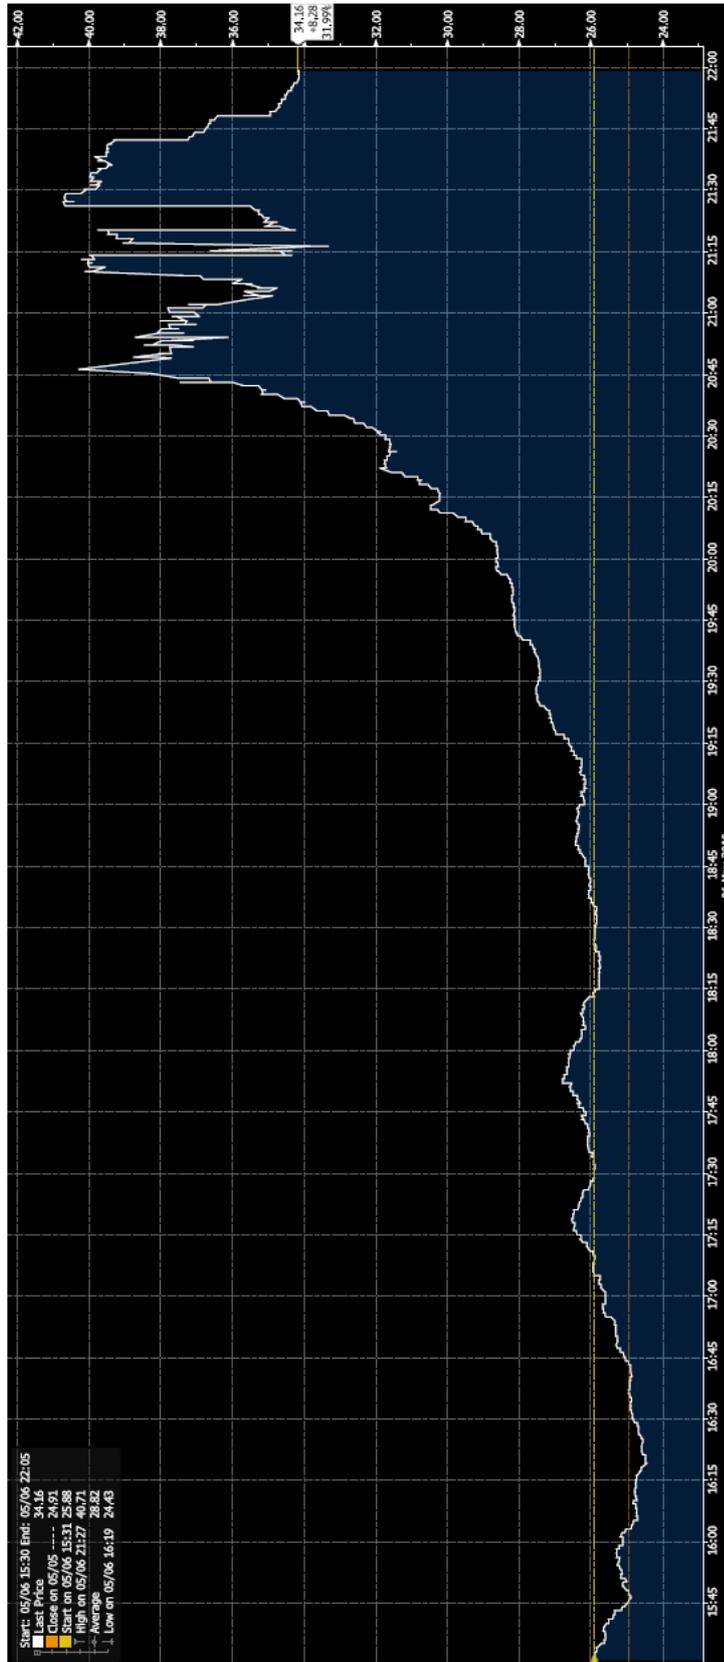


Figure 17: The VIX volatility index on May 6, 2010. The time zone is CEST. Source: Bloomberg.

Table 6: IMPLIED VOLATILITIES FOR A GIVEN TRADING DAY

Date:		31-07-2009														
Index Spot:	987.48	Strike, cash	Expiry/Strike rel. to spot													Div.-Yield
Tenor	Maturity	1283.72	1234.35	1184.98	1135.60	1086.23	1036.85	987.48	938.11	888.73	839.36	789.98	740.61	691.24	ZeroRate	Div.-Yield
		1.3	1.25	1.2	1.15	1.1	1.05	1	0.95	0.9	0.85	0.8	0.75	0.7		
1M	2009-08-31	0.14	0.14	0.15	0.17	0.18	0.21	0.23	0.26	0.30	0.34	0.38	0.44	0.50	0.00	0.03
2M	2009-09-30	0.15	0.16	0.17	0.19	0.20	0.22	0.24	0.27	0.29	0.32	0.35	0.39	0.43	0.00	0.03
3M	2009-11-02	0.17	0.18	0.19	0.20	0.21	0.23	0.25	0.27	0.29	0.32	0.34	0.37	0.40	0.00	0.02
6M	2010-02-01	0.19	0.19	0.20	0.21	0.23	0.24	0.26	0.27	0.29	0.31	0.33	0.35	0.37	0.00	0.02
9M	2010-04-30	0.19	0.20	0.21	0.22	0.23	0.24	0.26	0.27	0.29	0.30	0.32	0.33	0.35	0.01	0.02
12M	2010-08-02	0.20	0.21	0.21	0.22	0.23	0.25	0.26	0.27	0.28	0.30	0.31	0.33	0.34	0.01	0.02
15M	2010-11-01	0.20	0.21	0.22	0.23	0.24	0.25	0.26	0.27	0.28	0.30	0.31	0.32	0.33	0.01	0.02
18M	2011-01-31	0.21	0.21	0.22	0.23	0.24	0.25	0.26	0.27	0.28	0.29	0.31	0.32	0.33	0.01	0.02
21M	2011-05-02	0.21	0.22	0.23	0.23	0.24	0.25	0.26	0.27	0.28	0.29	0.30	0.32	0.33	0.01	0.02
24M	2011-08-01	0.22	0.22	0.23	0.24	0.25	0.25	0.26	0.27	0.28	0.29	0.30	0.31	0.33	0.01	0.02
27M	2011-10-31	0.22	0.23	0.23	0.24	0.25	0.26	0.26	0.27	0.28	0.29	0.30	0.31	0.32	0.02	0.02
30M	2012-01-31	0.22	0.23	0.24	0.24	0.25	0.26	0.27	0.28	0.28	0.29	0.30	0.31	0.32	0.02	0.02
33M	2012-04-30	0.23	0.23	0.24	0.25	0.25	0.26	0.27	0.28	0.29	0.29	0.30	0.31	0.32	0.02	0.02
36M	2012-07-31	0.23	0.24	0.24	0.25	0.26	0.26	0.27	0.28	0.29	0.29	0.30	0.31	0.32	0.02	0.02

Description: Black-Scholes implied volatilities for European option traded on a given trading day. For each combination of strike price and maturity ($13 \times 14 = 182$ combinations) relative to the spot index, the price of the European option (given by the implied volatility) is observed. The two last columns report the zero-rate and the dividend yield. More digits are available.

Table 7: SUMMARY OF SPY TICK DATA

Option	Period	#Trades	Total Volume	Daily Volume	Shares per Trade	Price
1	2004-04-01 - 2004-06-29	1.74×10^6	5.53×10^8	9.07×10^6	317	112.47
2	2004-06-30 - 2004-09-28	1.65×10^6	5.68×10^8	9.01×10^6	344	110.81
3	2004-09-29 - 2004-12-29	2.58×10^6	6.62×10^8	1.03×10^7	256	116.68
4	2004-12-30 - 2005-03-30	3.41×10^6	8.03×10^8	1.30×10^7	235	119.27
5	2005-03-31 - 2005-06-28	2.25×10^6	8.89×10^8	1.41×10^7	395	117.90
6	2005-06-29 - 2005-09-27	2.49×10^6	9.27×10^8	1.50×10^7	372	122.37
7	2005-09-28 - 2005-12-29	2.52×10^6	1.14×10^9	1.75×10^7	452	122.56
8	2005-12-30 - 2006-03-31	2.20×10^6	1.08×10^9	1.72×10^7	492	128.13
9	2006-04-03 - 2006-06-29	2.10×10^6	1.47×10^9	2.37×10^7	699	127.74
10	2006-06-30 - 2006-09-28	1.83×10^6	1.37×10^9	2.18×10^7	749	128.53
11	2006-09-29 - 2006-12-29	1.76×10^6	1.36×10^9	2.12×10^7	769	138.82
12	2007-01-03 - 2007-04-02	2.20×10^6	1.91×10^9	3.08×10^7	867	141.94
13	2007-04-03 - 2007-07-02	3.17×10^6	2.48×10^9	3.94×10^7	781	150.28
14	2007-07-03 - 2007-10-01	5.82×10^6	4.39×10^9	6.96×10^7	753	148.03
15	2007-10-02 - 2007-12-31	4.73×10^6	3.33×10^9	5.28×10^7	702	149.02
16	2008-01-02 - 2008-04-02	6.06×10^6	4.42×10^9	7.01×10^7	728	134.70
17	2008-04-03 - 2008-07-01	5.05×10^6	3.01×10^9	4.78×10^7	595	136.69
18	2008-07-02 - 2008-09-30	8.78×10^6	4.39×10^9	6.97×10^7	500	124.20
19	2008-10-01 - 2008-12-30	1.53×10^7	6.43×10^9	1.02×10^8	419	90.75
20	2008-12-31 - 2009-04-01	1.40×10^7	5.76×10^9	9.15×10^7	411	80.34
21	2009-04-02 - 2009-07-01	8.68×10^6	3.83×10^9	6.08×10^7	441	89.07
22	2009-07-02 - 2009-09-30	5.29×10^6	2.52×10^9	3.99×10^7	475	99.22
23	2009-10-01 - 2009-12-30	4.96×10^6	1.96×10^9	3.11×10^7	394	108.50
24	2009-12-31 - 2010-04-01	6.27×10^6	2.46×10^9	3.90×10^7	391	111.97
25	2010-04-05 - 2010-07-01	9.71×10^6	3.52×10^9	5.59×10^7	362	112.31
26	2010-07-02 - 2010-09-30	6.57×10^6	2.59×10^9	4.11×10^7	394	109.66
27	2010-10-01 - 2010-12-30	5.37×10^6	2.05×10^9	3.26×10^7	382	119.70
28	2010-12-31 - 2011-03-31	4.98×10^6	1.99×10^9	3.15×10^7	398	130.13
29	2011-04-01 - 2011-06-30	5.15×10^6	2.03×10^9	3.22×10^7	393	131.53
30	2011-07-01 - 2011-09-29	8.40×10^6	3.53×10^9	5.60×10^7	419	120.78
31	2011-09-30 - 2011-12-29	6.08×10^6	2.72×10^9	4.31×10^7	446	121.59
32	2011-12-30 - 2012-03-30	3.24×10^6	1.60×10^9	2.54×10^7	494	134.40
33	2012-04-02 - 2012-06-29	3.81×10^6	1.92×10^9	3.05×10^7	505	134.80
34	2012-07-02 - 2012-09-28	2.69×10^6	1.40×10^9	2.22×10^7	521	139.74
35	2012-10-01 - 2013-01-02	3.20×10^6	1.81×10^9	2.87×10^7	565	141.80
36	2013-01-03 - 2013-04-04	2.61×10^6	1.62×10^9	2.58×10^7	621	151.80
37	2013-04-05 - 2013-07-03	3.84×10^6	1.96×10^9	3.10×10^7	509	161.27

Description: Summary statistics of SPY trade data across the 38 trading periods. Numbers reported in the last three columns are averages.

Table 8: RISK-MINIMIZING HEDGE: SUMMARY OF P&L

Option	Period	Initial Price	$\sigma_1^{\text{imp.}}$	$\sigma_1^{\text{real.}}$	$P\&L_T^{\text{imp.}}$	$P\&L_T^{\text{real.}}$	$QV^{\text{imp.}}$	$QV^{\text{real.}}$
1	2004-04-01 - 2004-06-29	34.06	0.156	0.087	17.98	15.61	0.370	0.429
2	2004-06-30 - 2004-09-28	33.25	0.150	0.089	10.61	11.40	0.136	0.443
3	2004-09-29 - 2004-12-29	30.37	0.139	0.070	-2.72	-3.11	0.451	0.401
4	2004-12-30 - 2005-03-30	31.10	0.126	0.037	9.16	9.94	0.113	0.202
5	2005-03-31 - 2005-06-28	31.88	0.131	0.083	6.20	8.06	0.229	0.196
6	2005-06-29 - 2005-09-27	30.91	0.120	0.066	7.83	6.56	0.448	0.355
7	2005-09-28 - 2005-12-29	33.61	0.129	0.092	1.21	1.05	0.288	0.286
8	2005-12-30 - 2006-03-31	34.56	0.125	0.084	3.51	3.35	0.485	0.398
9	2006-04-03 - 2006-06-29	33.96	0.115	0.077	3.05	2.75	0.313	0.322
10	2006-06-30 - 2006-09-28	38.50	0.132	0.076	6.28	7.45	0.360	0.313
11	2006-09-29 - 2006-12-29	38.52	0.126	0.050	1.75	1.88	0.353	0.298
12	2007-01-02 - 2007-04-02	40.52	0.124	0.094	1.56	3.00	1.041	1.237
13	2007-04-03 - 2007-07-02	42.48	0.129	0.092	-1.76	0.19	0.564	0.529
14	2007-07-03 - 2007-10-01	50.74	0.148	0.051	-11.16	-1.46	2.024	1.775
15	2007-10-02 - 2007-12-31	58.21	0.171	0.065	-0.24	-3.11	0.635	0.875
16	2008-01-02 - 2008-04-02	69.65	0.228	0.168	8.38	15.95	0.551	1.515
17	2008-04-03 - 2008-07-01	63.40	0.230	0.138	20.70	17.53	0.640	1.266
18	2008-07-02 - 2008-09-30	59.04	0.235	0.170	-31.70	-35.33	3.692	7.953
19	2008-10-01 - 2008-12-30	76.01	0.323	0.336	-6.92	9.93	2.930	27.985
20	2008-12-31 - 2009-04-01	63.89	0.355	0.214	6.71	13.00	0.697	1.135
21	2009-04-02 - 2009-07-01	60.74	0.382	0.341	15.53	12.09	0.278	0.857
22	2009-07-02 - 2009-09-30	44.87	0.268	0.225	3.96	8.39	0.152	0.311
23	2009-10-01 - 2009-12-30	48.43	0.250	0.176	11.27	11.39	0.183	0.375
24	2009-12-31 - 2010-04-01	43.92	0.211	0.072	13.37	13.28	0.251	0.447
25	2010-04-05 - 2010-07-01	37.03	0.169	0.072	-5.13	-8.01	0.841	1.672
26	2010-07-02 - 2010-09-30	55.26	0.284	0.193	12.76	14.49	0.285	0.542
27	2010-10-01 - 2010-12-30	47.08	0.220	0.151	11.72	11.91	0.279	0.240
28	2010-12-31 - 2011-03-31	41.73	0.179	0.054	7.16	6.93	0.402	0.325
29	2011-04-01 - 2011-06-30	41.21	0.167	0.080	14.62	13.20	0.224	0.496
30	2011-07-01 - 2011-09-29	37.05	0.151	0.107	1.24	-24.23	0.155	12.899
31	2011-09-30 - 2011-12-29	79.71	0.368	0.323	16.31	15.51	0.533	3.149
32	2011-12-30 - 2012-03-30	55.53	0.235	0.061	15.90	20.78	0.187	0.280
33	2012-04-02 - 2012-06-29	39.62	0.153	0.079	8.13	10.70	0.155	0.358
34	2012-07-02 - 2012-09-28	41.78	0.167	0.116	7.44	10.33	0.602	0.484
35	2012-10-01 - 2013-01-02	40.19	0.154	0.136	-3.37	-6.22	4.287	4.179
36	2013-01-03 - 2013-04-04	39.20	0.149	0.080	4.81	4.63	0.280	0.178
37	2013-04-05 - 2013-07-03	36.90	0.132	0.174	-8.63	-8.56	0.834	0.959

Description: Summary of terminal hedge errors and quadratic variations across trading periods using Heston's risk-minimizing hedge. The highlighted rows indicate the hedging portfolios associated with biggest losses. The superscript 'impl' emphasizes that implied volatility is used as hedging volatility.

Table 9: HESTON DELTA HEDGE: SUMMARY OF P&L

Option	Period	Initial Price	$\sigma_1^{\text{imp.}}$	$\sigma_1^{\text{real.}}$	$P\&L_T^{\text{imp.}}$	$P\&L_T^{\text{real.}}$	$QV^{\text{imp.}}$	$QV^{\text{real.}}$
1	2004-04-01 - 2004-06-29	34.06	0.156	0.087	17.77	14.57	0.405	0.427
2	2004-06-30 - 2004-09-28	33.25	0.150	0.089	5.84	4.77	0.213	0.160
3	2004-09-29 - 2004-12-29	30.37	0.139	0.070	4.43	4.07	0.126	0.108
4	2004-12-30 - 2005-03-30	31.10	0.126	0.037	5.57	5.36	0.332	0.262
5	2005-03-31 - 2005-06-28	31.88	0.131	0.083	8.17	9.83	0.364	0.372
6	2005-06-29 - 2005-09-27	30.91	0.120	0.066	11.13	9.82	0.101	0.093
7	2005-09-28 - 2005-12-29	33.61	0.129	0.092	5.30	5.56	0.285	0.282
8	2005-12-30 - 2006-03-31	34.56	0.125	0.084	8.47	8.47	0.083	0.091
9	2006-04-03 - 2006-06-29	33.96	0.115	0.077	-3.99	-4.73	0.630	0.744
10	2006-06-30 - 2006-09-28	38.50	0.132	0.076	10.09	10.96	0.150	0.172
11	2006-09-29 - 2006-12-29	38.52	0.126	0.050	8.30	8.89	0.076	0.074
12	2007-01-02 - 2007-04-02	40.52	0.124	0.094	0.48	2.54	1.998	2.456
13	2007-04-03 - 2007-07-02	42.48	0.129	0.092	6.32	8.84	0.122	0.136
14	2007-07-03 - 2007-10-01	50.74	0.148	0.051	-20.03	-5.33	2.991	3.898
15	2007-10-02 - 2007-12-31	58.21	0.171	0.065	-12.56	-17.34	1.792	2.143
16	2008-01-02 - 2008-04-02	69.65	0.228	0.168	-3.78	3.43	0.651	1.463
17	2008-04-03 - 2008-07-01	63.40	0.230	0.138	18.49	14.75	0.961	1.939
18	2008-07-02 - 2008-09-30	59.04	0.235	0.170	-37.34	-37.01	5.554	10.713
19	2008-10-01 - 2008-12-30	76.01	0.323	0.336	-24.49	-3.70	2.537	34.312
20	2008-12-31 - 2009-04-01	63.89	0.355	0.214	-2.10	2.77	0.366	0.513
21	2009-04-02 - 2009-07-01	60.74	0.382	0.341	20.69	20.14	0.377	1.648
22	2009-07-02 - 2009-09-30	44.87	0.268	0.225	11.67	17.56	0.110	0.392
23	2009-10-01 - 2009-12-30	48.43	0.250	0.176	17.48	19.50	0.304	1.060
24	2009-12-31 - 2010-04-01	43.92	0.211	0.072	16.42	17.69	0.533	0.576
25	2010-04-05 - 2010-07-01	37.03	0.169	0.072	-15.35	-18.78	1.260	3.767
26	2010-07-02 - 2010-09-30	55.26	0.284	0.193	19.43	22.64	0.345	1.354
27	2010-10-01 - 2010-12-30	47.08	0.220	0.151	18.60	20.17	0.207	0.540
28	2010-12-31 - 2011-03-31	41.73	0.179	0.054	13.84	15.20	0.320	0.528
29	2011-04-01 - 2011-06-30	41.21	0.167	0.080	10.04	8.45	0.490	0.644
30	2011-07-01 - 2011-09-29	37.05	0.151	0.107	-10.93	-35.88	0.875	17.674
31	2011-09-30 - 2011-12-29	79.71	0.368	0.323	23.80	26.97	0.740	5.848
32	2011-12-30 - 2012-03-30	55.53	0.235	0.061	24.29	31.47	0.235	0.608
33	2012-04-02 - 2012-06-29	39.62	0.153	0.079	-1.21	0.27	0.280	0.199
34	2012-07-02 - 2012-09-28	41.78	0.167	0.116	14.60	17.36	0.232	0.264
35	2012-10-01 - 2013-01-02	40.19	0.154	0.136	-7.07	-11.13	4.436	3.815
36	2013-01-03 - 2013-04-04	39.20	0.149	0.080	13.82	14.67	0.157	0.305
37	2013-04-05 - 2013-07-03	36.90	0.132	0.174	1.99	2.74	0.464	0.634

Description: Summary of terminal hedge errors and quadratic variations across trading periods using the Heston's delta hedge. The highlighted rows indicate the hedging portfolios associated with biggest losses. The superscript 'impl' emphasizes that implied volatility is used as hedging volatility.

Table 10: BLACK-SCHOLES DELTA HEDGE: SUMMARY OF P&L

Option	Period	Initial Price	$\sigma_1^{\text{imp.}}$	$\sigma_1^{\text{real.}}$	$P\&L_T^{\text{imp.}}$	$P\&L_T^{\text{real.}}$	$QV^{\text{imp.}}$	$QV^{\text{real.}}$
1	2004-04-01 - 2004-06-29	34.06	0.156	0.087	17.88	11.91	0.258	0.309
2	2004-06-30 - 2004-09-28	33.25	0.150	0.089	9.59	11.11	0.101	0.407
3	2004-09-29 - 2004-12-29	30.37	0.139	0.070	2.92	2.35	0.058	0.215
4	2004-12-30 - 2005-03-30	31.10	0.126	0.037	6.57	5.65	0.071	0.152
5	2005-03-31 - 2005-06-28	31.88	0.131	0.083	6.46	8.29	0.112	0.170
6	2005-06-29 - 2005-09-27	30.91	0.120	0.066	7.59	4.58	0.067	0.248
7	2005-09-28 - 2005-12-29	33.61	0.129	0.092	3.76	4.32	0.093	0.265
8	2005-12-30 - 2006-03-31	34.56	0.125	0.084	5.85	3.26	0.086	0.389
9	2006-04-03 - 2006-06-29	33.96	0.115	0.077	-2.39	-4.02	0.263	0.393
10	2006-06-30 - 2006-09-28	38.50	0.132	0.076	8.23	9.21	0.108	0.240
11	2006-09-29 - 2006-12-29	38.52	0.126	0.050	7.67	10.29	0.044	0.227
12	2007-01-02 - 2007-04-02	40.52	0.124	0.094	-2.27	-1.87	1.886	3.311
13	2007-04-03 - 2007-07-02	42.48	0.129	0.092	4.94	12.50	0.052	0.246
14	2007-07-03 - 2007-10-01	50.74	0.148	0.051	-13.23	0.64	1.687	2.860
15	2007-10-02 - 2007-12-31	58.21	0.171	0.065	-5.87	-10.44	0.749	1.287
16	2008-01-02 - 2008-04-02	69.65	0.228	0.168	1.68	11.19	0.289	1.493
17	2008-04-03 - 2008-07-01	63.40	0.230	0.138	19.14	10.24	0.637	2.029
18	2008-07-02 - 2008-09-30	59.04	0.235	0.170	-32.74	-38.94	4.199	9.792
19	2008-10-01 - 2008-12-30	76.01	0.323	0.336	-35.18	-9.46	1.835	53.100
20	2008-12-31 - 2009-04-01	63.89	0.355	0.214	-2.22	3.67	0.267	0.327
21	2009-04-02 - 2009-07-01	60.74	0.382	0.341	17.97	14.49	0.217	1.255
22	2009-07-02 - 2009-09-30	44.87	0.268	0.225	7.53	16.30	0.074	0.558
23	2009-10-01 - 2009-12-30	48.43	0.250	0.176	13.84	15.91	0.130	1.195
24	2009-12-31 - 2010-04-01	43.92	0.211	0.072	16.20	16.80	0.228	0.543
25	2010-04-05 - 2010-07-01	37.03	0.169	0.072	-11.83	-18.64	0.677	3.733
26	2010-07-02 - 2010-09-30	55.26	0.284	0.193	15.44	16.90	0.203	1.485
27	2010-10-01 - 2010-12-30	47.08	0.220	0.151	15.18	14.41	0.145	0.786
28	2010-12-31 - 2011-03-31	41.73	0.179	0.054	8.81	9.63	0.225	1.291
29	2011-04-01 - 2011-06-30	41.21	0.167	0.080	11.99	7.64	0.219	0.558
30	2011-07-01 - 2011-09-29	37.05	0.151	0.107	-6.19	-30.56	0.404	21.727
31	2011-09-30 - 2011-12-29	79.71	0.368	0.323	19.03	18.01	0.440	6.490
32	2011-12-30 - 2012-03-30	55.53	0.235	0.061	20.20	32.83	0.154	1.074
33	2012-04-02 - 2012-06-29	39.62	0.153	0.079	2.95	7.86	0.063	0.398
34	2012-07-02 - 2012-09-28	41.78	0.167	0.116	12.24	18.11	0.204	0.663
35	2012-10-01 - 2013-01-02	40.19	0.154	0.136	-3.93	-10.05	3.930	3.816
36	2013-01-03 - 2013-04-04	39.20	0.149	0.080	10.71	13.03	0.097	0.854
37	2013-04-05 - 2013-07-03	36.90	0.132	0.174	-3.97	2.67	0.212	0.935

Description: Summary of terminal hedge errors and quadratic variations across trading periods using the Black-Scholes delta hedge. The highlighted rows indicate the hedging portfolios associated with biggest losses. The superscript 'impl' emphasizes that implied volatility is used as hedging volatility.

Table 11: SUMMARY OF VPIN ESTIMATES I

Option	Period	barSize	s.d.(#Buckets)	Min(#Buckets)	Max(#Buckets)	Average(VPIN)	Median(VPIN)	s.d.(VPIN)
1	2004-04-01 - 2004-06-29	0.5	14.55	27	95	0.236	0.236	0.059
2	2004-06-30 - 2004-09-28	0.5	14.83	28	93	0.228	0.216	0.050
3	2004-09-29 - 2004-12-29	0.5	12.00	17	89	0.220	0.215	0.037
4	2004-12-30 - 2005-03-30	0.5	12.82	22	83	0.231	0.227	0.045
5	2005-03-31 - 2005-06-28	0.5	15.95	21	90	0.242	0.233	0.053
6	2005-06-29 - 2005-09-27	0.5	15.05	28	92	0.250	0.251	0.047
7	2005-09-28 - 2005-12-29	0.5	18.56	10	115	0.234	0.226	0.055
8	2005-12-30 - 2006-03-31	0.5	13.58	27	103	0.217	0.214	0.034
9	2006-04-03 - 2006-06-29	0.5	22.15	15	119	0.222	0.216	0.056
10	2006-06-30 - 2006-09-28	0.5	13.48	20	89	0.209	0.204	0.041
11	2006-09-29 - 2006-12-29	0.5	14.17	21	94	0.225	0.221	0.042
12	2007-01-03 - 2007-04-02	0.5	26.04	21	143	0.243	0.211	0.098
13	2007-04-03 - 2007-07-02	0.5	21.91	9	105	0.238	0.219	0.069
14	2007-07-03 - 2007-10-01	0.5	26.82	15	134	0.217	0.202	0.071
15	2007-10-02 - 2007-12-31	0.5	21.67	11	108	0.223	0.223	0.054
16	2008-01-02 - 2008-04-02	0.5	15.98	21	101	0.196	0.188	0.048
17	2008-04-03 - 2008-07-01	0.5	12.88	26	92	0.208	0.205	0.040
18	2008-07-02 - 2008-09-30	0.5	22.11	23	130	0.210	0.191	0.074
19	2008-10-01 - 2008-12-30	0.5	20.19	6	118	0.211	0.201	0.055
20	2008-12-31 - 2009-04-01	0.5	10.80	23	71	0.218	0.210	0.054
21	2009-04-02 - 2009-07-01	0.5	11.96	27	96	0.244	0.231	0.061
22	2009-07-02 - 2009-09-30	0.5	13.12	26	90	0.258	0.242	0.068
23	2009-10-01 - 2009-12-30	0.5	17.18	8	93	0.244	0.230	0.061
24	2009-12-31 - 2010-04-01	0.5	19.07	20	119	0.233	0.225	0.054
25	2010-04-05 - 2010-07-01	0.5	20.05	21	126	0.223	0.216	0.073
26	2010-07-02 - 2010-09-30	0.5	11.09	31	71	0.236	0.231	0.052
27	2010-10-01 - 2010-12-30	0.5	18.46	15	103	0.254	0.247	0.057
28	2010-12-31 - 2011-03-31	0.5	19.68	22	131	0.247	0.243	0.068
29	2011-04-01 - 2011-06-30	0.5	15.72	19	101	0.244	0.235	0.050
30	2011-07-01 - 2011-09-29	0.5	23.04	17	124	0.237	0.218	0.080
31	2011-09-30 - 2011-12-29	0.5	16.96	16	103	0.240	0.236	0.054
32	2011-12-30 - 2012-03-30	0.5	10.72	32	72	0.264	0.259	0.057
33	2012-04-02 - 2012-06-29	0.5	10.74	32	85	0.260	0.247	0.065
34	2012-07-02 - 2012-09-28	0.5	15.48	25	99	0.286	0.280	0.068
35	2012-10-01 - 2013-01-02	0.5	16.41	24	115	0.292	0.270	0.090
36	2013-01-03 - 2013-04-04	0.5	13.27	25	89	0.291	0.281	0.075
37	2013-04-05 - 2013-07-03	0.5	17.61	27	111	0.269	0.254	0.078

Description: Summary of the VPIN series across trading periods. Time bars of 30-seconds are used, and VBS equals one-fiftieth of the average daily volume.

Table 12: SUMMARY OF VPIN ESTIMATES II

Option	Period	barSize	s.d.(#Buckets)	Min(#Buckets)	Max(#Buckets)	Average(VPIN)	Median(VPIN)	s.d.(VPIN)
1	2004-04-01 - 2004-06-29	1	14.55	27	95	0.273	0.267	0.067
2	2004-06-30 - 2004-09-28	1	14.83	28	93	0.268	0.253	0.067
3	2004-09-29 - 2004-12-29	1	12.00	17	89	0.261	0.253	0.048
4	2004-12-30 - 2005-03-30	1	12.82	22	83	0.271	0.263	0.052
5	2005-03-31 - 2005-06-28	1	15.95	21	90	0.286	0.270	0.066
6	2005-06-29 - 2005-09-27	1	15.05	28	92	0.291	0.288	0.060
7	2005-09-28 - 2005-12-29	1	18.56	10	115	0.277	0.261	0.069
8	2005-12-30 - 2006-03-31	1	13.58	27	103	0.255	0.248	0.044
9	2006-04-03 - 2006-06-29	1	22.15	15	119	0.269	0.254	0.074
10	2006-06-30 - 2006-09-28	1	13.48	20	89	0.255	0.248	0.052
11	2006-09-29 - 2006-12-29	1	14.17	21	94	0.271	0.268	0.052
12	2007-01-03 - 2007-04-02	1	26.04	21	143	0.292	0.259	0.115
13	2007-04-03 - 2007-07-02	1	21.91	9	105	0.283	0.266	0.077
14	2007-07-03 - 2007-10-01	1	26.82	15	134	0.270	0.247	0.093
15	2007-10-02 - 2007-12-31	1	21.67	11	108	0.276	0.273	0.073
16	2008-01-02 - 2008-04-02	1	15.98	21	101	0.249	0.242	0.061
17	2008-04-03 - 2008-07-01	1	12.88	26	92	0.254	0.249	0.050
18	2008-07-02 - 2008-09-30	1	22.11	23	130	0.259	0.242	0.092
19	2008-10-01 - 2008-12-30	1	20.19	6	118	0.263	0.251	0.072
20	2008-12-31 - 2009-04-01	1	10.80	23	71	0.253	0.248	0.056
21	2009-04-02 - 2009-07-01	1	11.96	27	96	0.294	0.290	0.056
22	2009-07-02 - 2009-09-30	1	13.12	26	90	0.295	0.283	0.073
23	2009-10-01 - 2009-12-30	1	17.18	8	93	0.277	0.267	0.065
24	2009-12-31 - 2010-04-01	1	19.07	20	119	0.280	0.271	0.063
25	2010-04-05 - 2010-07-01	1	20.05	21	126	0.264	0.260	0.091
26	2010-07-02 - 2010-09-30	1	11.09	31	71	0.283	0.280	0.059
27	2010-10-01 - 2010-12-30	1	18.46	15	103	0.291	0.287	0.073
28	2010-12-31 - 2011-03-31	1	19.68	22	131	0.284	0.273	0.083
29	2011-04-01 - 2011-06-30	1	15.72	19	101	0.286	0.276	0.063
30	2011-07-01 - 2011-09-29	1	23.04	17	124	0.287	0.266	0.096
31	2011-09-30 - 2011-12-29	1	16.96	16	103	0.286	0.282	0.068
32	2011-12-30 - 2012-03-30	1	10.72	32	72	0.286	0.286	0.055
33	2012-04-02 - 2012-06-29	1	10.74	32	85	0.302	0.294	0.071
34	2012-07-02 - 2012-09-28	1	15.48	25	99	0.324	0.314	0.076
35	2012-10-01 - 2013-01-02	1	16.41	24	115	0.321	0.302	0.100
36	2013-01-03 - 2013-04-04	1	13.27	25	89	0.324	0.314	0.080
37	2013-04-05 - 2013-07-03	1	17.61	27	111	0.313	0.302	0.091

Description: Summary of the VPIN series across trading periods. One-minute time bars are used, and VBS equals one-fiftieth of the average daily volume.

Table 13: SUMMARY OF VPIN ESTIMATES III

Option	Period	barSize	s.d.(#Buckets)	Min(#Buckets)	Max(#Buckets)	Average(VPIN)	Median(VPIN)	s.d.(VPIN)
1	2004-04-01 - 2004-06-29	2	14.55	27	95	0.320	0.311	0.076
2	2004-06-30 - 2004-09-28	2	14.83	28	93	0.326	0.313	0.079
3	2004-09-29 - 2004-12-29	2	12.00	17	89	0.312	0.307	0.061
4	2004-12-30 - 2005-03-30	2	12.82	22	83	0.329	0.322	0.060
5	2005-03-31 - 2005-06-28	2	15.95	21	90	0.339	0.325	0.076
6	2005-06-29 - 2005-09-27	2	15.05	28	92	0.344	0.338	0.071
7	2005-09-28 - 2005-12-29	2	18.56	10	115	0.329	0.309	0.083
8	2005-12-30 - 2006-03-31	2	13.58	27	103	0.311	0.302	0.053
9	2006-04-03 - 2006-06-29	2	22.15	15	119	0.332	0.318	0.089
10	2006-06-30 - 2006-09-28	2	13.48	20	89	0.314	0.303	0.071
11	2006-09-29 - 2006-12-29	2	14.17	21	94	0.324	0.320	0.061
12	2007-01-03 - 2007-04-02	2	26.04	21	143	0.343	0.314	0.132
13	2007-04-03 - 2007-07-02	2	21.91	9	105	0.346	0.320	0.100
14	2007-07-03 - 2007-10-01	2	26.82	15	134	0.339	0.318	0.122
15	2007-10-02 - 2007-12-31	2	21.67	11	108	0.340	0.342	0.091
16	2008-01-02 - 2008-04-02	2	15.98	21	101	0.322	0.307	0.086
17	2008-04-03 - 2008-07-01	2	12.88	26	92	0.309	0.305	0.062
18	2008-07-02 - 2008-09-30	2	22.11	23	130	0.322	0.300	0.113
19	2008-10-01 - 2008-12-30	2	20.19	6	118	0.328	0.314	0.093
20	2008-12-31 - 2009-04-01	2	10.80	23	71	0.314	0.308	0.056
21	2009-04-02 - 2009-07-01	2	11.96	27	96	0.341	0.334	0.061
22	2009-07-02 - 2009-09-30	2	13.12	26	90	0.341	0.334	0.075
23	2009-10-01 - 2009-12-30	2	17.18	8	93	0.335	0.321	0.079
24	2009-12-31 - 2010-04-01	2	19.07	20	119	0.332	0.318	0.080
25	2010-04-05 - 2010-07-01	2	20.05	21	126	0.312	0.310	0.110
26	2010-07-02 - 2010-09-30	2	11.09	31	71	0.336	0.334	0.071
27	2010-10-01 - 2010-12-30	2	18.46	15	103	0.346	0.341	0.081
28	2010-12-31 - 2011-03-31	2	19.68	22	131	0.335	0.323	0.094
29	2011-04-01 - 2011-06-30	2	15.72	19	101	0.343	0.335	0.071
30	2011-07-01 - 2011-09-29	2	23.04	17	124	0.341	0.314	0.126
31	2011-09-30 - 2011-12-29	2	16.96	16	103	0.339	0.329	0.077
32	2011-12-30 - 2012-03-30	2	10.72	32	72	0.339	0.341	0.065
33	2012-04-02 - 2012-06-29	2	10.74	32	85	0.359	0.355	0.069
34	2012-07-02 - 2012-09-28	2	15.48	25	99	0.370	0.361	0.084
35	2012-10-01 - 2013-01-02	2	16.41	24	115	0.379	0.364	0.097
36	2013-01-03 - 2013-04-04	2	13.27	25	89	0.375	0.359	0.095
37	2013-04-05 - 2013-07-03	2	17.61	27	111	0.368	0.358	0.102

Description: Summary of the VPIN series across trading periods. Two-minute time bars are used.

Table 14: SUMMARY OF VPIN ESTIMATES IV

Option	Period	barSize	s.d.(#Buckets)	Min(#Buckets)	Max(#Buckets)	Average(VPIN)	Median(VPIN)	s.d.(VPIN)
1	2004-04-01 - 2004-06-29	3	14.55	27	95	0.361	0.354	0.088
2	2004-06-30 - 2004-09-28	3	14.83	28	93	0.366	0.353	0.086
3	2004-09-29 - 2004-12-29	3	12.00	17	89	0.354	0.351	0.065
4	2004-12-30 - 2005-03-30	3	12.82	22	83	0.363	0.358	0.067
5	2005-03-31 - 2005-06-28	3	15.95	21	90	0.377	0.360	0.086
6	2005-06-29 - 2005-09-27	3	15.05	28	92	0.378	0.368	0.074
7	2005-09-28 - 2005-12-29	3	18.56	10	115	0.368	0.355	0.093
8	2005-12-30 - 2006-03-31	3	13.58	27	103	0.349	0.341	0.063
9	2006-04-03 - 2006-06-29	3	22.15	15	119	0.367	0.356	0.099
10	2006-06-30 - 2006-09-28	3	13.48	20	89	0.347	0.337	0.069
11	2006-09-29 - 2006-12-29	3	14.17	21	94	0.356	0.352	0.069
12	2007-01-03 - 2007-04-02	3	26.04	21	143	0.384	0.351	0.145
13	2007-04-03 - 2007-07-02	3	21.91	9	105	0.384	0.354	0.108
14	2007-07-03 - 2007-10-01	3	26.82	15	134	0.386	0.365	0.135
15	2007-10-02 - 2007-12-31	3	21.67	11	108	0.384	0.388	0.103
16	2008-01-02 - 2008-04-02	3	15.98	21	101	0.356	0.342	0.086
17	2008-04-03 - 2008-07-01	3	12.88	26	92	0.351	0.346	0.069
18	2008-07-02 - 2008-09-30	3	22.11	23	130	0.361	0.344	0.120
19	2008-10-01 - 2008-12-30	3	20.19	6	118	0.368	0.355	0.103
20	2008-12-31 - 2009-04-01	3	10.80	23	71	0.351	0.350	0.056
21	2009-04-02 - 2009-07-01	3	11.96	27	96	0.381	0.379	0.068
22	2009-07-02 - 2009-09-30	3	13.12	26	90	0.372	0.359	0.075
23	2009-10-01 - 2009-12-30	3	17.18	8	93	0.369	0.364	0.083
24	2009-12-31 - 2010-04-01	3	19.07	20	119	0.371	0.359	0.092
25	2010-04-05 - 2010-07-01	3	20.05	21	126	0.357	0.351	0.123
26	2010-07-02 - 2010-09-30	3	11.09	31	71	0.372	0.374	0.076
27	2010-10-01 - 2010-12-30	3	18.46	15	103	0.379	0.377	0.085
28	2010-12-31 - 2011-03-31	3	19.68	22	131	0.383	0.373	0.109
29	2011-04-01 - 2011-06-30	3	15.72	19	101	0.368	0.364	0.077
30	2011-07-01 - 2011-09-29	3	23.04	17	124	0.374	0.348	0.139
31	2011-09-30 - 2011-12-29	3	16.96	16	103	0.378	0.373	0.087
32	2011-12-30 - 2012-03-30	3	10.72	32	72	0.378	0.377	0.066
33	2012-04-02 - 2012-06-29	3	10.74	32	85	0.387	0.382	0.074
34	2012-07-02 - 2012-09-28	3	15.48	25	99	0.403	0.397	0.084
35	2012-10-01 - 2013-01-02	3	16.41	24	115	0.401	0.387	0.097
36	2013-01-03 - 2013-04-04	3	13.27	25	89	0.403	0.396	0.100
37	2013-04-05 - 2013-07-03	3	17.61	27	111	0.395	0.377	0.107

Description: Summary of the VPIN series across trading periods. Three-minute time bars are used, and VBS equals one-fiftieth of the average daily volume.

Table 15: SUMMARY OF VPIN ESTIMATES V

Option	Period	VBS	Average(VPIN)	Median(VPIN)	s.d.(VPIN)
1	2004-04-01 - 2004-06-29	1/25	0.213	0.209	0.050
2	2004-06-30 - 2004-09-28	1/25	0.209	0.196	0.052
3	2004-09-29 - 2004-12-29	1/25	0.201	0.201	0.033
4	2004-12-30 - 2005-03-30	1/25	0.210	0.207	0.041
5	2005-03-31 - 2005-06-28	1/25	0.224	0.217	0.044
6	2005-06-29 - 2005-09-27	1/25	0.225	0.217	0.042
7	2005-09-28 - 2005-12-29	1/25	0.211	0.201	0.055
8	2005-12-30 - 2006-03-31	1/25	0.189	0.190	0.029
9	2006-04-03 - 2006-06-29	1/25	0.205	0.201	0.053
10	2006-06-30 - 2006-09-28	1/25	0.187	0.182	0.033
11	2006-09-29 - 2006-12-29	1/25	0.212	0.210	0.036
12	2007-01-03 - 2007-04-02	1/25	0.224	0.201	0.087
13	2007-04-03 - 2007-07-02	1/25	0.215	0.203	0.053
14	2007-07-03 - 2007-10-01	1/25	0.209	0.197	0.064
15	2007-10-02 - 2007-12-31	1/25	0.212	0.209	0.047
16	2008-01-02 - 2008-04-02	1/25	0.190	0.184	0.045
17	2008-04-03 - 2008-07-01	1/25	0.189	0.182	0.034
18	2008-07-02 - 2008-09-30	1/25	0.199	0.186	0.072
19	2008-10-01 - 2008-12-30	1/25	0.195	0.188	0.049
20	2008-12-31 - 2009-04-01	1/25	0.195	0.188	0.043
21	2009-04-02 - 2009-07-01	1/25	0.235	0.234	0.040
22	2009-07-02 - 2009-09-30	1/25	0.234	0.226	0.052
23	2009-10-01 - 2009-12-30	1/25	0.221	0.213	0.041
24	2009-12-31 - 2010-04-01	1/25	0.216	0.215	0.047
25	2010-04-05 - 2010-07-01	1/25	0.209	0.202	0.070
26	2010-07-02 - 2010-09-30	1/25	0.227	0.223	0.039
27	2010-10-01 - 2010-12-30	1/25	0.231	0.229	0.048
28	2010-12-31 - 2011-03-31	1/25	0.226	0.217	0.063
29	2011-04-01 - 2011-06-30	1/25	0.223	0.212	0.041
30	2011-07-01 - 2011-09-29	1/25	0.227	0.211	0.067
31	2011-09-30 - 2011-12-29	1/25	0.225	0.219	0.045
32	2011-12-30 - 2012-03-30	1/25	0.232	0.229	0.034
33	2012-04-02 - 2012-06-29	1/25	0.243	0.237	0.053
34	2012-07-02 - 2012-09-28	1/25	0.269	0.266	0.055
35	2012-10-01 - 2013-01-02	1/25	0.271	0.259	0.077
36	2013-01-03 - 2013-04-04	1/25	0.266	0.260	0.054
37	2013-04-05 - 2013-07-03	1/25	0.258	0.251	0.062

Description: Summary of the VPIN series across trading periods. One-minute time bars are used, and *VBS* equals one-twenty-fifth of the average daily volume.

Table 16: SUMMARY OF VPIN ESTIMATES VI

Option	Period	VBS	Average(VPIN)	Median(VPIN)	s.d.(VPIN)
1	2004-04-01 - 2004-06-29	1/50	0.273	0.267	0.067
2	2004-06-30 - 2004-09-28	1/50	0.268	0.253	0.067
3	2004-09-29 - 2004-12-29	1/50	0.261	0.253	0.048
4	2004-12-30 - 2005-03-30	1/50	0.271	0.263	0.052
5	2005-03-31 - 2005-06-28	1/50	0.286	0.270	0.066
6	2005-06-29 - 2005-09-27	1/50	0.291	0.288	0.060
7	2005-09-28 - 2005-12-29	1/50	0.277	0.261	0.069
8	2005-12-30 - 2006-03-31	1/50	0.255	0.248	0.044
9	2006-04-03 - 2006-06-29	1/50	0.269	0.254	0.074
10	2006-06-30 - 2006-09-28	1/50	0.255	0.248	0.052
11	2006-09-29 - 2006-12-29	1/50	0.271	0.268	0.052
12	2007-01-03 - 2007-04-02	1/50	0.292	0.259	0.115
13	2007-04-03 - 2007-07-02	1/50	0.283	0.266	0.077
14	2007-07-03 - 2007-10-01	1/50	0.270	0.247	0.093
15	2007-10-02 - 2007-12-31	1/50	0.276	0.273	0.073
16	2008-01-02 - 2008-04-02	1/50	0.249	0.242	0.061
17	2008-04-03 - 2008-07-01	1/50	0.254	0.249	0.050
18	2008-07-02 - 2008-09-30	1/50	0.259	0.242	0.092
19	2008-10-01 - 2008-12-30	1/50	0.263	0.251	0.072
20	2008-12-31 - 2009-04-01	1/50	0.253	0.248	0.056
21	2009-04-02 - 2009-07-01	1/50	0.294	0.290	0.056
22	2009-07-02 - 2009-09-30	1/50	0.295	0.283	0.073
23	2009-10-01 - 2009-12-30	1/50	0.277	0.267	0.065
24	2009-12-31 - 2010-04-01	1/50	0.280	0.271	0.063
25	2010-04-05 - 2010-07-01	1/50	0.264	0.260	0.091
26	2010-07-02 - 2010-09-30	1/50	0.283	0.280	0.059
27	2010-10-01 - 2010-12-30	1/50	0.291	0.287	0.073
28	2010-12-31 - 2011-03-31	1/50	0.284	0.273	0.083
29	2011-04-01 - 2011-06-30	1/50	0.286	0.276	0.063
30	2011-07-01 - 2011-09-29	1/50	0.287	0.266	0.096
31	2011-09-30 - 2011-12-29	1/50	0.286	0.282	0.068
32	2011-12-30 - 2012-03-30	1/50	0.286	0.286	0.055
33	2012-04-02 - 2012-06-29	1/50	0.302	0.294	0.071
34	2012-07-02 - 2012-09-28	1/50	0.324	0.314	0.076
35	2012-10-01 - 2013-01-02	1/50	0.321	0.302	0.100
36	2013-01-03 - 2013-04-04	1/50	0.324	0.314	0.080
37	2013-04-05 - 2013-07-03	1/50	0.313	0.302	0.091

Description: Summary of the VPIN series across trading periods. One-minute time bars are used, and *VBS* equals one-fiftieth of the average daily volume.

Table 17: SUMMARY OF VPIN ESTIMATES VII

Option	Period	VBS	Average(VPIN)	Median(VPIN)	s.d.(VPIN)
1	2004-04-01 - 2004-06-29	1/100	0.343	0.339	0.074
2	2004-06-30 - 2004-09-28	1/100	0.336	0.320	0.067
3	2004-09-29 - 2004-12-29	1/100	0.334	0.328	0.051
4	2004-12-30 - 2005-03-30	1/100	0.340	0.334	0.056
5	2005-03-31 - 2005-06-28	1/100	0.353	0.337	0.076
6	2005-06-29 - 2005-09-27	1/100	0.358	0.356	0.062
7	2005-09-28 - 2005-12-29	1/100	0.348	0.339	0.076
8	2005-12-30 - 2006-03-31	1/100	0.326	0.318	0.051
9	2006-04-03 - 2006-06-29	1/100	0.341	0.327	0.082
10	2006-06-30 - 2006-09-28	1/100	0.326	0.319	0.060
11	2006-09-29 - 2006-12-29	1/100	0.344	0.339	0.055
12	2007-01-03 - 2007-04-02	1/100	0.364	0.338	0.126
13	2007-04-03 - 2007-07-02	1/100	0.357	0.335	0.091
14	2007-07-03 - 2007-10-01	1/100	0.349	0.321	0.112
15	2007-10-02 - 2007-12-31	1/100	0.353	0.353	0.086
16	2008-01-02 - 2008-04-02	1/100	0.323	0.310	0.074
17	2008-04-03 - 2008-07-01	1/100	0.323	0.314	0.056
18	2008-07-02 - 2008-09-30	1/100	0.329	0.311	0.112
19	2008-10-01 - 2008-12-30	1/100	0.334	0.320	0.087
20	2008-12-31 - 2009-04-01	1/100	0.316	0.312	0.056
21	2009-04-02 - 2009-07-01	1/100	0.357	0.354	0.054
22	2009-07-02 - 2009-09-30	1/100	0.364	0.355	0.072
23	2009-10-01 - 2009-12-30	1/100	0.347	0.341	0.072
24	2009-12-31 - 2010-04-01	1/100	0.348	0.332	0.082
25	2010-04-05 - 2010-07-01	1/100	0.329	0.327	0.106
26	2010-07-02 - 2010-09-30	1/100	0.350	0.347	0.062
27	2010-10-01 - 2010-12-30	1/100	0.362	0.359	0.078
28	2010-12-31 - 2011-03-31	1/100	0.354	0.346	0.097
29	2011-04-01 - 2011-06-30	1/100	0.350	0.338	0.066
30	2011-07-01 - 2011-09-29	1/100	0.355	0.332	0.112
31	2011-09-30 - 2011-12-29	1/100	0.354	0.349	0.079
32	2011-12-30 - 2012-03-30	1/100	0.352	0.352	0.056
33	2012-04-02 - 2012-06-29	1/100	0.366	0.362	0.073
34	2012-07-02 - 2012-09-28	1/100	0.383	0.372	0.076
35	2012-10-01 - 2013-01-02	1/100	0.386	0.366	0.101
36	2013-01-03 - 2013-04-04	1/100	0.383	0.382	0.080
37	2013-04-05 - 2013-07-03	1/100	0.377	0.366	0.103

Description: Summary of the VPIN series across trading periods. One-minute time bars are used, and *VBS* equals one percentage of the average daily volume.

References

- Abad, D., M. Massot, and R. Pascual (2018). Evaluating VPIN as a trigger for single-stock circuit breakers. *Journal of Banking & Finance* 86(C), 21–36.
- Abad-Diaz, D. and J. Yagiüe (2012, 07). From pin to vpin: An introduction to order flow toxicity. *The Spanish Review of Financial Economics* 10, 74–83.
- Andersen, T. G. and O. Bondarenko (2014a). Reflecting on the vpin dispute. *Journal of Financial Markets* 17, 53–64.
- Andersen, T. G. and O. Bondarenko (2014b). Vpin and the flash crash. *Journal of Financial Markets* 17, 1–46.
- Andersen, T. G. and O. Bondarenko (2015). Assessing measures of order flow toxicity and early warning signals for market turbulence*. *Review of Finance* 19(1), 1–54.
- Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy* 81, 637–654.
- Cont, R. and P. Tankov (2004). *Financial Modelling With Jump Processes* (1 ed.). Chapman and Hall/CRC Financial Mathematics Series.
- Easley, D., N. M. Kiefer, M. O’Hara, and J. B. Paperman (1996). Liquidity, information, and infrequently traded stocks. *The Journal of Finance* 51(4), 1405–1436.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2012, 02). Flow toxicity and liquidity in a high frequency world. *Review of Financial Studies* 25(5), 1457–1493.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2013, 03). Bulk classification of trading activity. Working Paper.
- Easley, D., M. M. López de Prado, and M. O’Hara (2010, 11). The microstructure of the "flash crash": Flow toxicity, liquidity crashes and the probability of informed trading. *The Journal of Portfolio Management* 37, 118–128.
- Easley, D., M. M. López de Prado, and M. O’Hara (2011). The exchange of flow toxicity. *The Journal of Trading* 6(2), 8–13.
- Easley, D., M. M. López de Prado, and M. O’Hara (2014). Vpin and the flash crash: A rejoinder. *Journal of Financial Markets* 17(C), 47–52.
- El Karoui, N., S. Peng, and M. C. Quenez (1997). Backward stochastic differential equations in finance. *Mathematical Finance* 7, 1–71.
- Ellersgaard, S., M. Jönsson, and R. Poulsen (2017). The fundamental theorem of derivative trading - exposition, extensions and experiments. *Quantitative Finance* 17(4), 515–529.

- Eraker, B. (2004). Do stock prices and volatility jump? reconciling evidence from spot and option prices. *The Journal of Finance* 3, 1367–1404.
- Föllmer, H. and D. Sondermann (1986). Hedging of non-refundant contingent claims. In W. Hildenbrand and A. Mas-Colell (Eds.), *Contributions to Mathematical Economics*, Volume 1, Chapter 12, pp. 204–223. North-Holland: Elsevier Science Publishers.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies* 6(2), 327–343.
- Holden, C. W. and S. Jacobsen (2013). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *The Journal of Finance* 69(4), 1747–1785.
- Lee, C. and M. Ready (1991). Inferring trade direction from intraday data. *Journal of Finance* 46(2), 733–46.
- Leth, L. R. (2015). The fundamental theorem of derivative trading. Master’s thesis, Department of Mathematical Sciences, University of Copenhagen.
- Leth, L. R. (2019). Maximum likelihood estimation of vpin: Toxic order flow and warning signals. Working paper.
- Lin, H.-W. W. and W.-C. Ke (2017). An improved version of the volume-synchronized probability of informed trading. *Critical Finance Review* 6(2), 357–376.
- Pöppe, T., S. Moos, and D. Schiereck (2016). The sensitivity of vpin to the choice of trade classification algorithm. *Journal of Banking & Finance* 73(C), 165–181.
- Poulsen, R., K. R. Schenk-Hoppè, and C. O. Ewald (2009). Risk minimization in stochastic volatility models: Model risk and empirical performanc. *Quantative Finance* 9(6), 693–704.
- Schweizer, M. (1991). Option hedging for semimartingales. *Stochastic Processes and their Applications* 37, 339–363.
- Song, J. H., K. Wu, and H. Simon (2014). Parameter analysis of the vpin (volume synchronized probability of informed trading) metric. In C. Zopounidis (Ed.), *Quantitative Financial Risk Management: Theory and Practice*. Wiley.
- Wilmott, P. and R. Ahmad (2005). Which free lunch would you like today, sir?: Delta hedging, volatility arbitrage and optimal portfolios.
- Wong, B. and C. C. Heyde (2006). On changes of measure in stochastic volatility models. *Journal of Applied Mathematics and Stochastic Analysis* 2006, 1–13.

Wu, K., E. W. Bethel, M. Gu, D. Leinweber, and O. Ruebel (2013a, 06). A big data approach to analyzing market volatility. *Algorithmic Finance 2*, 241–267.

Wu, K., E. W. Bethel, M. Gu, D. Leinweber, and O. Ruebel (2013b, 01). Testing vpin on big data response to 'reflecting on the vpin dispute'. *SSRN Electronic Journal*.

This page intentionally left blank.

CHAPTER 3

MAXIMUM LIKELIHOOD
ESTIMATION OF VPIN:
TOXIC ORDER FLOW AND
WARNING SIGNALS

All models are wrong, but some are useful

— George E. P. Box

This page intentionally left blank.

Maximum Likelihood Estimation of VPIN: Toxic Order Flow and Warning Signals^{*,**}

Laurs Randbøll Leth^{a,b,*}

^a*Department of Economics, University of Copenhagen*

^b*Center for Information and Bubble Studies, University of Copenhagen*

Abstract

This paper examines the performance of an improved metric of volume-synchronized probability of informed trading (VPIN) in providing a warning for imminent price movements. In contrast to the original (moment) estimation method of VPIN, the improved measure is computed using maximum likelihood estimation utilizing information in volume time. The empirical study considers Level 1 tick data of the SPDR S&P 500 ETF (SPY) from 2007 to 2015. The VPIN's performance is assessed by its capacity to predict short-term return volatility. Numerical results show that MLE improves the VPIN's predictability for movement in intraday price changes.

Keywords: VPIN, warning signals, high-frequency market microstructure, market turbulence.

JEL: G01; G12; G14; G17; C52; C58; C13.

*The author appreciates and acknowledges the financial support of the Carlsberg Foundation grant to the Center for Information and Bubble Studies, University of Copenhagen.

**The author thanks Peter Norman Sørensen and Marcos Lopez de Prado for their helpful comments and valuable suggestions.

*Corresponding author

Email address: bkt797@ku.dk (Laurs Randbøll Leth)

Conflict of Interest Disclosure Statement

I have nothing to disclose.

1. Introduction

This paper investigates whether maximum likelihood estimation (MLE) of the volume-synchronized probability of informed trading (VPIN) can improve its predictive power for short-term volatility. VPIN was developed by Easley et al. (2012) to measure the toxic order flow in markets characterized by high-frequency trading. The VPIN metric originated from the market microstructure, particularly the well-known probability of informed trading (PIN) model proposed by Easley et al. (1996). Easley et al. (2008) considered a modified PIN model focusing on trading volume. They could also derive the expected order imbalance and total traded volume, respectively, for a fixed unit of time. Their computations reveal that the ratio of the two expectations approximately measures the PIN.

Inspired by this result, Easley et al. (2012) proposed moment estimation to measure the PIN in high-frequency markets. The estimator (for method of moment estimation: MME) was labeled VPIN and derived from the fundamental assumption that the arrival of information is modeled on a volume clock, and not calendar time as for the PIN model. However, Lin and Ke (2017) show mathematically that the moment estimator approximates a probability measure for a given fixed-size trading volume labeled VPIN, whereas PIN is a probability measure for a given fixed time interval. Theoretically, the two measures do not coincide. The MLE utilizing information in volume enabled Lin and Ke (2017) to fully estimate both PIN and VPIN parametrically.

Presumably, moment estimation and full parametric estimation of the VPIN will lead to different outcomes; thus, it is relevant to compare the important characteristics of the two estimators. Moreover, it is pertinent to investigate whether the moment estimator can numerically measure PIN.

Abstracting from the disagreement between the two probability measures, (through moment estimation) VPIN has played a central role in the empirical market microstructure literature in recent years. Easley et al. (2012) find that VPIN is a good predictor of short-term return volatility and may be used as a warning signal for imminent market turmoil. The economic interpretation is that high VPIN values are related to toxic order flow, which negatively affects the market makers (high-frequency traders).

Over time, the market makers will withdraw from their positions and start demanding liquidity. This impact on market liquidity may affect the intraday price returns in what might be referred to as liquidity-induced volatility. Importantly, Easley et al. (2012, 2010) conclude that VPIN anticipated the Flash Crash of May 6, 2010, emphasizing that the metric is a powerful tool for signaling market turbulence.

This empirical work began a debate with Andersen and Bondarenko (2014a, 2015) on the opposite side. These studies claim that VPIN has no incremental predictive power for models of short-term volatility¹. As regards the Flash Crash of May 6, 2010, Andersen and Bondarenko (2014b) observe that VPIN kept on rising after the crash. They question whether VPIN was impacted by the crash or was predicting the crash. Their main conclusion is that VPIN is an inferior warning signal for market turmoil.

Numerous empirical studies have contributed to this debate, the majority of them based on in-sample performance². Wu et al. (2013a) advocate the usefulness of the VPIN on the basis of an enormous study of around 100 of the most liquid futures contracts over a period of five and a half years involving about 3 billion trades. They propose to assess the performance of the VPIN using a binary outcome where the events related to high VPIN values are labeled either true or false positives depending on maximum intermediate returns (MIR) for a given event horizon. Considering over 16,000 VPIN parameter combinations (including the event horizon), Wu et al. (2013a) identify the optimal parameter settings that lead to false discovery rates of around 7%. As Abad et al. (2018) correctly pointed out, this in-sample optimization approach reminds one of data snooping.

This work contributes an out-of-sample study of the VPIN's predictive power for short-term volatility. VPIN is computed using both the moment estimator and the recently developed MLE method capturing information from volume time. The MLE is also used to compute the PIN measure.

¹Especially, Andersen and Bondarenko (2015) states that VPIN predicts short-term volatility solely because ex-ante volatility distorts the trade classification algorithm; this is used in the estimation of VPIN.

²See the empirical work by Abad-Diaz and Yagüe (2012), Easley et al. (2012, 2010, 2011, 2014), Wu et al. (2013a,b), and Song et al. (2014) (supporting), or Andersen and Bondarenko (2014a,b, 2015), Pöppe et al. (2016), and Abad et al. (2018) (disputing).

This paper is strongly believed to be the first to conduct a comprehensive empirical study on the maximum likelihood estimation of VPIN.

Assessment of VPIN's (PIN's) performance is similar to the binary evaluation proposed by Wu et al. (2013a), preferring *small* false discovery rates. Inspired by the in-sample optimal VPIN parameter settings identified by Wu et al. (2013a) and Song et al. (2014), only a few volume bucket size and support window combinations are considered. The remaining parameters are fixed, with the daily average trading volume from the previous year used to compute the relevant volume bucket size. This setting is appropriate for answering the research questions of this paper: Does VPIN (PIN) have incremental predictive power for short-term return volatility? Can the MLE of VPIN improve this predictive power? Does VPIN empirically measure PIN?

This empirical study focuses on the SPRD S&P500 ETF (SPY) traded between 2007-2015. The findings reveal that VPIN does not approximately measure PIN. Thus, only MLE may be used to estimate PIN. Moreover, the moment estimation of VPIN produces false discovery rates ranging from 15% to 23%. In comparison, the MLE of VPIN leads to false discovery rates ranging from 10% – 15%, indicating a clear improvement. Strikingly, PIN is an inferior predictor of short-term return volatility, with false discovery rates of approximately 50%. This result supports the theoretical disagreement between VPIN and PIN, and further suggests that the toxic order flow in high-frequency markets should be modeled on the volume clock.

Empirical findings are supplemented with a minor case study, illustrating a significant shortcoming of the methodology proposed by Wu et al. (2013a): complete neglect of type II errors. The false discovery rate is likely to move opposite to the number of type II errors, and performance evaluation of VPIN must be based on the practitioner's specific problem. The F_1 -score, the harmonic mean of precision and sensitivity, is used to capture this performance trade-off. From the results, the lowest false discovery rate does not coincide with the optimal F_1 -score. This indirectly criticizes the empirical findings of Wu et al. (2013a). Finally, the MLE of VPIN yields the best outcomes in terms of larger F_1 -scores. This result supports the findings on false discovery rates: the MLE of VPIN improves its predictive power for short-term volatility.

Further, this paper presents an example on the Flash Crash of May 6, 2010. VPIN anticipated this crash, and advantages of the MLE of VPIN are demonstrated. For instance, the method shows the selling volume of the informed traders dramatically increasing (falling) prior (subsequent) to the crash. In addition, the toxic order flow affected the selling volume in accordance with theory. The VPIN metric stopped increasing after the crash, in contrast to the findings from the moment estimation of VPIN. If volume falls but imbalance remains high, the moment estimator of VPIN might increase and the market makers may incorrectly identify this as toxic order flow. In contrast, the MLE method captures the information in volume time, which increases (decreases) after (before) the crash. Now, the market makers would realize that the sell orders after the crash were actually submitted by noise traders.

The rest of this paper is organized as follows. Section 2 presents an overview of the PIN model. Section 3 presents the VPIN model, including the moment estimation and maximum likelihood estimation of the VPIN metric. Section 4 describes the trade data used for empirical investigation. Section 5 discusses how to assess the performance of the VPIN. The empirical design is described in section 6, and the numerical findings are presented in section 7. Section 8 concludes the paper.

2. The PIN Model (Easley et al., 1996)

The original PIN model views the trading of a risky asset as a game between liquidity providers (the market maker) and liquidity consumers (the traders). Trading takes place over D days, and is continuous within days where the market maker quotes the bid and ask prices at any time in $[0, T]$, while the traders submit their market orders of one unit of the risky asset. These market orders are driven by exogenous causes (e.g., demand for liquidity) or private information. Consequently, the bid-ask spread occurs from information-based trading, and is related to adverse selection.

Private information is modeled by an information event occurring prior to trading, affecting the daily fundamental value of the asset. Events occur on a daily basis with probability $\alpha \in (0, 1)$, and are assumed to be independently distributed. Additionally, events are either good news with probability $\delta \in (0, 1)$ or bad news with probability $1 - \delta$.

Formally, assume that $\mathbb{S} = \{\text{good-news}, \text{bad-news}, \text{no-news}\}$. Prior to trading, nature draws an $\omega \in \mathbb{S}$, and only privately informed traders have access to this information. When trading is closed, ω is revealed to all market participants.

For a unit trading period (e.g., one day), (Easley et al., 1996) assume that the arrival of buy and sell orders from noise traders is modeled by two independent Poisson processes with arrival rates ϵ_B and ϵ_S , respectively. Furthermore, the arrival of informed market orders is modeled by a Poisson process with rate μ . Summarizing, $\theta = (\alpha, \delta, \mu, \epsilon_B, \epsilon_S)$ is the model parameter, with the likelihood of observing any daily trading sequence of orders containing b_d buys and s_d sells, $h_d := (b_d, s_d)$, given by

$$\begin{aligned} L(\theta; h_d) &= \sum_{\omega \in \mathbb{S}} P(\omega) P(h_d | \omega) \\ &= \alpha \delta \cdot \text{dpois}(b_d; \mu + \epsilon_B) \cdot \text{dpois}(s_d; \epsilon_S) \\ &\quad + \alpha(1 - \delta) \cdot \text{dpois}(b_d; \epsilon_B) \cdot \text{dpois}(s_d; \mu + \epsilon_S) \\ &\quad + (1 - \alpha) \cdot \text{dpois}(b_d; \epsilon_B) \cdot \text{dpois}(s_d; \epsilon_S), \end{aligned}$$

where $\text{dpois}(x; \lambda)$ denotes the density for the Poisson distribution with arrival rate $\lambda > 0$ ³. Thus, the daily likelihood is a mixture of three Poisson probabilities. Since the information events are assumed to be independently distributed across days, the likelihood of observing the full history $\{h_d\}_{1 \leq d \leq D}$ is given by

$$L(\theta; \{h_d\}_{1 \leq d \leq D}) = \prod_{d=1}^D L(\theta; h_d), \quad (2.1)$$

and the maximum likelihood estimator $\hat{\theta}$ is obtained by maximizing equation (2.1).

³The density function is given by

$$\text{dpois}(x; \lambda) = \frac{\lambda^x e^{-\lambda}}{x!},$$

where $x \in \mathbb{N} \cup \{0\}$ and $\lambda > 0$.

Assuming $\delta = 1/2$, the opening quotes' spread is given by

$$\text{Ask} - \text{Bid} = \frac{\alpha\mu}{\epsilon_S + \epsilon_B + \alpha\mu}(\bar{V} - \underline{V}),$$

where \bar{V} and \underline{V} denote the expectation of V conditional on good news and bad news, respectively. The first term, $\alpha\mu/(\epsilon_S + \epsilon_B + \alpha\mu)$, denotes the probability that the first trade that occurs is information-based; this is called the probability of information-based trading (PIN). The relationship between the PIN and certain stock characteristics (e.g., liquidity and asset returns) has been an important topic in the market microstructure literature⁴. However, empirical PIN estimation studies are very demanding in high-frequency markets characterized by (among other factors) a tremendously large number of daily trades. For instance, the variables in the likelihood function of equation (2.1) will be raised to the power of millions when considering the most traded stocks, and estimation becomes problematic, if not infeasible.

3. Volume-Synchronized Probability of Informed Trading

The VPIN model proposed by Easley et al. (2012) is a generalization of the PIN model in a high-frequency setting. The PIN model is estimated for a given fixed time interval and depends on the total number of buy and sell orders. In contrast, the VPIN provides the probability of informed trading given a fixed-size trading volume, and relies on the buy and sell volume used to compute the order imbalances. On a given trading day, the arrival of the buy (sell) volume, $\{v_i^B\}_{i \geq 1}$ ($\{v_i^S\}_{i \geq 1}$), is modeled by a Poisson process, with the rate depending on the current state of the world⁵. Sequential trades are placed in equal-volume buckets of exogenous size vbs , where each bucket is treated as equivalent to a period for information arrival in the PIN model; that is, an information event affects the volume bucket with probability α ⁶. Thus, the model parameter is still θ with similar interpretation, but the

⁴For instance, see Easley et al. (2002) and Easley et al. (2010).

⁵The random variable V_i^B is Poisson distributed with parameter $\epsilon_B + \mu$ conditional on good news. In the event of bad news or no news, only liquidity traders would buy the asset, implying that v_i^B is Poisson distributed with parameter ϵ_B .

⁶Similar to the PIN model, the event is good news with probability δ .

arrival of news is now measured on a volume clock instead of equal time units (e.g., one trading day). The objective is to approximate the original PIN measure; this new high-frequency measure, labeled VPIN, denotes the probability of informed trading, given a fixed-size trading volume (\mathbf{vbs}).

Both models need to classify the direction of the order flow, which turns out to be extremely difficult in high-frequency market microstructure⁷. A solution to this problem is discussed in section 4, but for now, both buy and sell volumes are assumed given.

The remainder of this section presents two VPIN estimation methods. The first approach originates from Easley et al. (2012), and is based on simple moment estimation of order imbalance and total volume measured on the volume clock. The second method proposed by Lin and Ke (2017) uses parametric maximum likelihood estimation to calibrate the VPIN model. In contrast to the moment estimation of VPIN, this method captures the information in volume time **and** provides the estimates used to determine the level of (un)informed trading.

3.1. VPIN: Moment Estimation

Given a trade classification algorithm to determine the order flow direction, let v_t^B and v_t^S denote, respectively, the total buy and sell volume of trading period t . Also, let $\mathbf{OI}_t = v_t^B - v_t^S$ and $\mathbf{TT}_t = v_t^B + v_t^S$ denote, respectively, the order imbalance and total trade of period t . Assuming the PIN model, where the buy and sell orders are modeled by two independent Poisson processes with constraint $\epsilon_B = \epsilon_S = \epsilon$, Easley et al. (2008) show that $E|\mathbf{OI}_t| \approx \alpha\mu$ and $E(\mathbf{TT}_t) = 2\epsilon + \alpha\mu$. Consequently, they conclude that

$$\frac{E(|\mathbf{OI}_t|)}{E(\mathbf{TT}_t)} \in (0, 1] \quad (3.1)$$

serves as a good approximation of PIN ($\alpha\mu/(\epsilon_B + \epsilon_S + \alpha\mu)$) when $\epsilon_B = \epsilon_S$.

⁷Trade classification algorithms utilizing Level 2 data encounter several problems in high-frequency market microstructure, such as size and memory of databases, timing and merging of quotes and trades, high-order cancellation rates, and quote volatility. Additionally, informed traders may use limit orders to trade, and discerning information from trade data could be problematic.

The expression (3.1) is then estimated by the averages of $E|\text{OI}_t|$ and $E(\text{TT}_t)$, respectively, based on a volume time scale. Formally, this moment estimator, labeled VPIN, is computed in the following two steps:

1. **Volume bucketing:** Let \mathbf{vbs} be an exogenously defined volume size. Sequential trades (bars) are then grouped into equal-volume buckets of size \mathbf{vbs} . The buy and sell volumes of bucket τ are given by

$$V_\tau^B = \sum_{i: b(i)=\tau} v_i^B$$

$$V_\tau^S = \sum_{i: b(i)=\tau} v_i^S = \mathbf{vbs} - V_\tau^B,$$

where b is the surjective mapping assigning trade i to bucket $b(i)$. Thus, the order imbalance of bucket τ is given by $|\text{OI}_\tau| = |V_\tau^B - V_\tau^S|$, while total trades is equal to \mathbf{vbs} by construction.

2. **Estimation:** At bucket level, the expected order imbalance is estimated by the moment estimator

$$E|\text{OI}_\tau| \approx \frac{1}{n} \sum_{j=\tau-(n-1)}^{\tau} \text{OI}_j,$$

where n is a user-defined sample length.

Summarizing, the procedure above results in the VPIN metric given by

$$\text{VPIN}_\tau^{\text{mme}} = \frac{1}{n \times \mathbf{vbs}} \sum_{j=\tau-(n-1)}^{\tau} |\text{OI}_j| \quad (3.2)$$

for $\tau > n - 1$. However, this metric is not a moment estimator of the expression in equation (3.1) due to the binding constraint $V_\tau^B + V_\tau^S = \mathbf{vbs}$, but rather the moment estimator of

$$\begin{aligned}
\text{VPIN} &\equiv \frac{E(|V_\tau^B - V_\tau^S| | \mathbf{vbs})}{E(V_\tau^B + V_\tau^S | \mathbf{vbs})} \\
&= \frac{E(|V_\tau^B - V_\tau^S| | \mathbf{vbs})}{\mathbf{vbs}} \\
&\approx \frac{\alpha\mu}{2\epsilon + \mu} \\
&\neq \frac{\alpha\mu}{2\epsilon + \alpha\mu}
\end{aligned}$$

for $\alpha \in (0, 1)$ and $\epsilon_B = \epsilon_S$. Consequently, the estimator VPIN^{mme} measures approximately the VPIN, and not PIN. The difference between the two measures is emphasized in definition 1.

Definition 1. *The original PIN metric measures the probability of informed trading for a fixed (unit) time interval, and is given by*

$$\text{PIN} \equiv \frac{\alpha\mu}{\epsilon_B + \epsilon_S + \alpha\mu}. \quad (3.3)$$

The VPIN metric measures the probability of informed trading for a fixed-size trading volume \mathbf{vbs} , and is given by

$$\text{VPIN} \equiv \frac{\alpha\mu}{\epsilon_B + \epsilon_S + \mu}. \quad (3.4)$$

Summarizing, VPIN_τ^{mme} from equation (3.2) is the VPIN moment estimator in definition 1 when $\epsilon_B = \epsilon_S$, and is computed in three steps: (1) determine the order flow in terms of buy and sell volume, (2) choose an equal-volume bucket size \mathbf{vbs} , and (3) set the sample length n^8 .

3.2. VPIN: Maximum Likelihood Estimation

The moment estimator of VPIN neither estimates PIN nor approximately measures VPIN when $\epsilon_B \neq \epsilon_S$. Furthermore, the estimator does not capture the information in the duration time of equal-volume buckets.

⁸Step (1) can be independent of the VPIN computation.

At the bucket level, V_τ^B and V_τ^S are observed and satisfy condition $V_\tau^B + V_\tau^S = \mathbf{vbs}$ for $\tau = 1, \dots, N$, where \mathbf{vbs} is the fixed-size trading volume. Moreover, the duration time of bucket τ , t_τ , is available; that is, the time taken to fill bucket τ is observed. Formally, this stopping time is recursively given by

$$t_\tau = \inf \left\{ t \geq t_{\tau-1} : \sum_{i:b(i)=\tau} V_i^B + V_i^S = \mathbf{vbs} \right\} - t_{\tau-1},$$

with the convention $t_0 := 0$. Thus, the joint distribution of $(V_\tau^B, V_\tau^S, t_\tau)_{1 \leq \tau \leq n}$ may be used to estimate the unobserved model parameter $\theta \in \Theta$. Let g denote the density of $(V_\tau^B, V_\tau^S, t_\tau)$, conditional on $V_\tau^B + V_\tau^S = \mathbf{vbs}$, which may be expressed by the law of total probability,

$$\begin{aligned} g(b, s, t; \Theta) &= \sum_{\omega \in \mathbb{S}} P(\omega) g(b, s, t | \omega; \Theta) \\ &= \sum_{\omega \in \mathbb{S}} P(\omega) g_{(b,s)}(b, s | \omega, t; \theta) g_t(t | \omega; \theta), \end{aligned} \quad (3.5)$$

where $\mathbb{S} = \{\text{good-news}, \text{bad-news}, \text{no-news}\}$. Here, $g_{(b,s)}(\cdot; \theta)$ denotes the density of (V_τ^B, V_τ^S) , while $g_t(\cdot; \theta)$ is the density of the random variable t_τ .

Let $\omega \in \mathbb{S}$ and $j \in \{B, S\}$ be given, and realize that

$$V_\tau^j | t_\tau = \sum_{i:b(i)=\tau} V_i^j | t_\tau \sim \text{Pois}(t_\tau(\epsilon_j + \mu(j, \omega)))$$

with

$$\mu(j, \omega) = \begin{cases} \mu & \text{if } ((\omega = \text{good-news}) \cap (j = B)) \cup ((\omega = \text{bad-news}) \cap (j = S)) \\ 0 & \text{otherwise.} \end{cases}$$

Now,

$$V_\tau^j | (\omega, t_\tau, V_\tau^B + V_\tau^S = \mathbf{vbs}) \sim \text{Bin} \left(\mathbf{vbs}, \frac{\epsilon_j + \mu(j, \omega)}{\epsilon_S + \epsilon_B + \mu(j, \omega)} \right)$$

for $\omega \in \mathbb{S}$ and $j \in \{B, S\}$, where $\text{Bin}(n, p)$ denotes the Binomial distribution of $n \in \{0\} \cup \mathbb{N}$ trials, each with success probability $p \in (0, 1)$. Furthermore, condition $V_\tau^B + V_\tau^S = \mathbf{vbs}$ implies

$$\begin{aligned}
g_{(b,s)}(b, s|t_\tau, \omega; \theta) &= g_b(b|t_\tau, \omega; \theta) \\
&= g_s(s|t_\tau, \omega; \theta) \\
&= \text{dbin} \left(s; \mathbf{vbs}, \frac{\epsilon_S + \mu(S, \omega)}{\epsilon_S + \epsilon_B + \mu(S, \omega)} \right) \\
&= \text{dbin} \left(b; \mathbf{vbs}, \frac{\epsilon_B + \mu(B, \omega)}{\epsilon_S + \epsilon_B + \mu(B, \omega)} \right),
\end{aligned}$$

where $\text{dbin}(\cdot; n, p)$ denotes the density of the Binomial distribution with n trials, each with success probability p .

Finally, t_τ is, by definition, the waiting time for occurrence of the \mathbf{vbs} 'th event for the Poisson distributed variable $(V_i^B + V_i^S)$, implying that

$$t_\tau | (\omega, V_\tau^B + V_\tau^S = \mathbf{vbs}) \sim \Gamma(\mathbf{vbs}, (\epsilon_S + \epsilon_B + \mu(B, \omega) + \mu(S, \omega))),$$

where $\Gamma(n, \lambda)$ denotes the Gamma distribution with rate $\lambda > 0$ and shape $n > 0$. Summarizing,

$$\begin{aligned}
g(b, s, t; \theta) &= \alpha \delta \cdot \text{dbpin} \left(s; \mathbf{vbs}, \frac{\epsilon_S}{\epsilon_S + \epsilon_B + \mu} \right) \cdot \text{dgam}(t; \mathbf{vbs}, \epsilon_S + \epsilon_B + \mu) \\
&+ \alpha(1 - \delta) \cdot \text{dbin} \left(s; \mathbf{vbs}, \frac{\epsilon_S + \mu}{\epsilon_S + \epsilon_B + \mu} \right) \cdot \text{dgam}(t; \mathbf{vbs}, \epsilon_S + \epsilon_B + \mu) \\
&+ (1 - \alpha) \cdot \text{dbin} \left(s; \mathbf{vbs}, \frac{\epsilon_S}{\epsilon_S + \epsilon_B} \right) \cdot \text{dgam}(t; \mathbf{vbs}, \epsilon_S + \epsilon_B),
\end{aligned}$$

where $\text{dgam}(\cdot; n, \lambda)$ is the density of the Gamma distribution with rate $\lambda > 0$ and shape n^9 . For bucket τ , the (proportional) likelihood of observing $h_\tau := (b_\tau, s_\tau, t_\tau)$ is then given by

⁹The density is given by

$$\text{dgam}(t; n, \lambda) = \frac{1}{\Gamma(n)} \lambda e^{-\lambda t} (\lambda t)^{n-1}$$

for $t > 0$, where $\Gamma(n) := \int_0^\infty x^{n-1} e^{-x} dx$ is the Gamma function.

$$\begin{aligned}
L(\theta; h_\tau) &= \alpha \delta (\epsilon_B + \mu)^{b_\tau} \epsilon_S^{s_\tau} e^{-t_\tau(\epsilon_B + \epsilon_S + \mu)} \\
&\quad + \alpha(1 - \delta) (\epsilon_S + \mu)^{s_\tau} \epsilon_B^{b_\tau} e^{-t_\tau(\epsilon_B + \epsilon_S + \mu)} \\
&\quad + (1 - \alpha) \epsilon_B^{b_\tau} \epsilon_S^{s_\tau} e^{-t_\tau(\epsilon_B + \epsilon_S)}. \tag{3.6}
\end{aligned}$$

Assuming the observations are independently and identically distributed across buckets, the complete history of bucket data $\{h_\tau\}_{1 \leq \tau \leq n}$ results in the full model log-likelihood¹⁰

$$l(\theta; \{h_\tau\}_{1 \leq \tau \leq n}) := - \sum_{\tau=1}^n \log L(\theta; h_\tau). \tag{3.7}$$

Finally, the maximum likelihood estimator is given by

$$\hat{\theta} = \arg \min_{\theta \in \Theta} l(\theta; \{h_\tau\}_{1 \leq \tau \leq n}).$$

Using MLE, the estimators of PIN and VPIN for bucket τ are given respectively by

$$\text{PIN}_\tau^{\text{mle}}(\hat{\theta}) = \frac{\hat{\alpha}\hat{\mu}}{\hat{\epsilon}_S + \hat{\epsilon}_B + \hat{\alpha}\hat{\mu}} \quad \text{and} \quad \text{VPIN}_\tau^{\text{mle}}(\hat{\theta}) = \frac{\hat{\alpha}\hat{\mu}}{\hat{\epsilon}_S + \hat{\epsilon}_B + \hat{\mu}}, \tag{3.8}$$

both expressions implicitly depending on the sample length n and bucket size vbs .

The MLE of the model parameter θ allows for the practitioner to compute both PIN and VPIN, whereas the moment estimation method (algebraic) is restricted to the VPIN metric under $\epsilon_B = \epsilon_S$. Also, MLE enables a more precise interpretation of informed trading. If the volume falls but imbalance remains high, the VPIN will rise, and one may incorrectly conclude that this is due to informed trading. However, if $\epsilon_B \neq \epsilon_S$, the higher imbalance may be from liquidity trading in a no news volume period¹¹.

¹⁰The global optimizer used in this paper minimizes the object function. Thus, the log-likelihood is scaled with -1 .

¹¹On May 6, 2010, the moment estimation of VPIN generated values that kept on rising after the ‘‘Flash Crash’’.

3.2.1. Numerical Stability

Following Lin and Ke (2009), the marginal likelihood obtained from equation (3.6) is rewritten. The reformulated likelihood expression is based on two principles:

- If e^{x+y} , the expression $e^{x+y} (\text{sign}(x)e^{\log|x|+y})$ is more stable than $e^x e^y$ and (xe^y) when computing e^{x+y} ¹².
- The absolute computing error of a function f will increase with $|f'(x)|$ according to Eldén and Wittmeyer-Koch (1990). Especially, $\log(e^{x+y-m} + e^{z-m}) + m$ with $m = \max(x + y, z)$ is more accurate than $\log(e^{x+y} + e^z)$.

From the two principles, the reformulated likelihood is given by

$$L(\theta; h_\tau) = \log \left(\sum_{i=1}^3 e^{m_i(\tau) - m_{\max}(\tau)} \right) + m_{\max}(\tau), \quad (3.9)$$

where

$$\begin{aligned} m_1(\tau) &= \log(\alpha\delta) + b_\tau \log(\epsilon_B + \mu) + s_\tau \log(\epsilon_S) - t(\epsilon_B + \epsilon_S + \mu) \\ m_2(\tau) &= \log(\alpha(1 - \delta)) + b_\tau \log(\epsilon_B) + s_\tau \log(\epsilon_S + \mu) - t(\epsilon_B + \epsilon_S + \mu) \\ m_3(\tau) &= \log(1 - \alpha) + b_\tau \log(\epsilon_B) + s_\tau \log(\epsilon_S) - t(\epsilon_B + \epsilon_S), \end{aligned}$$

and $m_{\max}(\tau) = \max(c_1(\tau), c_2(\tau), c_3(\tau))$. Importantly, Lin and Ke (2017) introduce a typo in their reformulated likelihood function, which thus differs from the correct expression in equation (3.9).

The likelihood is implemented in C++ and estimated in R using the package `Rcpp`, providing R-functions as well as a C++-library facilitating the interface between the two languages. Because of the non-convex programming problem, the estimation procedure follows the suggestion of Mullen (2014), and is carried out in two steps:

1. The global optimizer `DEoptim()` in the R-package `DEoptim()` generates the initial model parameter estimates.

¹²For instance, the language R has $\approx e^{709.78}$ as upper benchmark. If $x = 1000$ and $y = -500$, the expression $e^x e^y$ results in an overflow, whereas e^{x+y} is computable.

2. The initial solution is used as start guess in R-function `optim()` utilizing the gradient-based Nelder-Mead method.

As with the moment estimation of VPIN, a rolling window of n buckets is considered for the estimation of θ at time τ . For each sample, $\hat{\theta}_\tau$ is the maximum likelihood estimator of $\theta \in \Theta$ computed applying the two steps above. The MLE of (V)PIN is given by equation (3.8). The practitioner is rewarded with an improved estimator of (V)PIN capturing the information in volume time and clearly interpreting the degree of informed trading and liquidity trading, respectively, but the reward comes with a price. Compared to the moment estimation of VPIN, the MLE approach has a significant larger computational cost. Moreover, the method is relevant only for markets in which the parameters can be accurately estimated.

4. Data

Level 1 trading data of the SPRD S&P 500 ETF (SPY) between 2007-2015 from the TAQ (trades and quotes) database are used for the empirical investigation. The data are treated in line with Holden and Jacobsen (2013), and the cleaning procedure is described in details in appendix A. Note that only the transactions with positive price and normal sale condition reported in the core trading session are considered; they represent approximately 96% of all trades. The summary statistics of trades and volume across years are reported in tables 12 and 13 in appendix B.

A significant shortcoming of the TAQ datafile is that the 'aggressor' flag is not reported; that is, it does not appear whether the buyer or seller initiated the trade. Classifying the order flow in a high-frequency world is non-trivial, and needs a suitable trade classification algorithm. Additionally, the model setting in section 3 requires that the chosen trade classification samples with some fixed (unit) time (e.g., 10 seconds). The bulk volume classification (BVC) algorithm meets all the criteria.

4.1. Bulk Volume Classification

The BVC algorithm introduced by Easley et al. (2013) generates bars by aggregating sequential trades over short time or volume intervals. Thus, a bar is a collection of trades assigned the total volume of all associated

transactions as well as the nominal bar price depending on the prices included in the bar. By construction, the volume is random for time bars, whereas the time becomes random when aggregating over volume intervals.

Formally, let v_i and p_i denote the volume and price, respectively, of bar i with $\Delta p_i := p_i - p_{i-1}$. BVC utilizes the standardized price change, $\Delta p_i / \sigma_{\Delta p}$, where $\sigma_{\Delta p}$ is the standard deviation of price changes across bars, to classify the fraction of buy and sell volume in probabilistic terms. In particular, it is assumed that $\Delta p_i / \sigma_{\Delta p}$ has CDF Z , and buy and sell volume are computed as

$$\begin{aligned} v_i^B &:= v_i \cdot Z\left(\frac{\Delta p_i}{\sigma_{\Delta p}}\right) \\ v_i^S &:= v_i - v_i^B. \end{aligned}$$

Other trade classification algorithms utilizing individual trades and quotes may be considered¹³. Clearly, the BVS algorithm is superior in terms of lower computational cost. Furthermore, BVC is not subject to problems such as quote volatility and the timing between trades and quotes. Easley et al. (2013) present numerical evidence that standard classification algorithms do not result in greater accuracy, whereas Easley et al. (2016) show that bulk volume classifications are better linked to information-based trading proxies.

As suggested by Song et al. (2014), the nominal bar price is by construction equal to the average of prices in the bar. The specification of nominal bar price will, along with the bar type (time or volume), alter the distribution of $\Delta p_i / \sigma_{\Delta p}$. Easley et al. (2012) assume the standard normal distribution, whereas Easley et al. (2016) consider the t -distribution to account for fat tails present in the data. Finally, Wu et al. (2013a) find the t -distribution preferred in some cases when optimizing the performance of the VPIN.

¹³For instance, the famous trade classification algorithm proposed by Lee and Ready (1991) depending on Level 2 tick data (transactions *and* quotes).

4.1.1. Time Bars

This subsection considers three different choices of time bars: 10-second bars, 30-second bars, and one-minute bars. For instance, sampling with 10-second bars will result in 2340 observations per day, which on average compresses the original data size by more than 99%. If the volume bucket is equal to one-fiftieth of the daily average volume, 47 time bars are on average used to fill a bucket.

Figure 1 illustrates the distribution of $\Delta p_i / \sigma_{\Delta p}$ for the different lengths of time bars, where the normal distribution appears to be a reasonable approximation¹⁴.

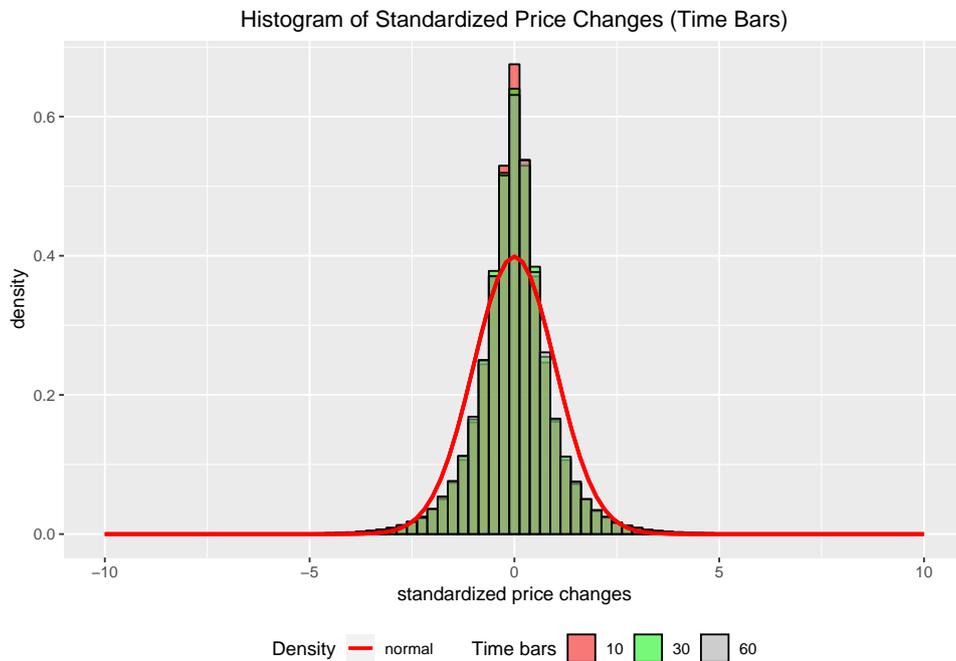


Figure 1: Histograms of standardized price changes of time bars in 2015. The red line illustrates the density of the (standard) normal distribution.

In contrast to SPY, most of the futures used in a vast part of the published literature are traded around the clock. Consequently, the overnight market is expected to distort the distribution of price changes for SPY. This problem is avoided by only calculating the time bars when the market is open and thereby excluding the overnight price changes¹⁵.

¹⁴The distribution is symmetric around 0, and fat tails are not present in the data.

¹⁵At the beginning of day d , the opening trade is excluded from time bar 1, and is instead used to compute the price change of the initial bar. The BVC algorithm is then applied on all the remaining transactions of day d , and the procedure is repeated for the next day.

5. Assessing the Performance of VPIN

Evaluating the ability of VPIN to signal impending market turbulence requires strict definitions of warning signals/VPIN events. The intuition behind warning signals is that high VPIN values proxy toxic order flows, impacting liquidity; this effect can result in large price returns. Following Song et al. (2014), a VPIN event is defined by the CDF of the VPIN hitting an exogenously given threshold c from below. The time when $CDF(VPIN)$ stays above c is called the event time.

Definition 2. *VPIN events are recursively specified by*

$$\begin{aligned}\underline{i}_j &= \inf\{i' \geq \underline{i}_{j-1} : CDF(VPIN)_{i'} > c \ \& \ CDF(VPIN)_{i'-1} \leq c\} \\ \bar{i}_j &= \inf\{i' \geq \bar{i}_{j-1} : CDF(VPIN)_{i'} \leq c \ \& \ CDF(VPIN)_{i'-1} > c\}\end{aligned}$$

with $\bar{i}_0 = \underline{i}_0 = 0$. That is, \underline{i}_j is the j th event measured in volume time (bucket number). The event time is given by $\eta_j := \bar{i}_j - \underline{i}_j$.

The definition of event and event time is a prerequisite for assessing the ability of VPIN to signal impending market stress. Intuitively, the event succeeds in signaling market turbulence when the price volatility during the event time is larger than expected. This is formalized in the following definition.

Definition 3. *Assuming M to be a given volatility measure, consider a VPIN event with event time η . The event is labeled*

- *true positive (TP) if the corresponding M value is greater than the average value of M over randomly selected time intervals of equal duration η , and*
- *false positive (FP) if the corresponding M value is less than the average value of M over randomly selected time intervals of equal duration η .*

From this event classification, Wu et al. (2013a) consider the false discovery rate

$$\text{FDR} = \frac{\#\text{TP}}{\#\text{FP} + \#\text{TP}} \quad (5.1)$$

for the performance evaluation of VPIN. The FDR is directly linked to the *precision* (or positive predictive value) by $\text{PPV} = 1 - \text{FDR}$; that is, a low FDR indicates that VPIN events most likely lead to a true positive.

Standard volatility measures based on the realized return over the period η are inappropriate for classifying true and false positives. The VPIN is updated in volume time, and the average return between buckets may be small, but a huge intermediate return could have occurred. To capture scenarios like this, the volatility measure M in definition 3 is modeled by the *maximum intermediate return* (MIR).

5.1. Maximum Intermediate Return

Consider a list of N prices, (p_0, \dots, p_{N-1}) , and let $R_{j,k} := p_k/p_j - 1$ denote the return between prices j and k for $0 \leq j < k < N$. The MIR is then given by

$$\text{MIR} = R_{j^*,k^*}, \quad (5.2)$$

where the pair of trades (j^*, k^*) are the sentinels and maximize the intermediate returns over all trade combinations; that is,

$$(j^*, k^*) = \arg \max_{0 \leq j < k < N} |R_{j,k}|. \quad (5.3)$$

The MIR may be computed by brute force with a double-nested loop to search for sentinels among all feasible combinations in equation (5.3). The computational cost of this approach is undesirable since the algorithm will read the N prices $N(N+1)/2$ times¹⁶. This computational problem is circumvented by using the algorithm introduced by Wu et al. (2013a).

¹⁶Such algorithms are in computational complexity labeled N^2 -algorithms. If $N = 2 \cdot 10^6$ (see table 12), the double-nested loop will read the N prices $2 \cdot 10^{12} + 10^6$ times!

5.1.1. MIR Algorithm

If \bar{R} and \underline{R} denote the largest and smallest intermediate returns, respectively, where \bar{R} (\underline{R}) is referred to as the maximum intermediate gain (loss), then MIR is either equal to \bar{R} or \underline{R} . Let \bar{j} be the position of the last appearance of the maximum price and \underline{j} be the first occurrence of the minimum price. MIR is then computed with the following algorithm:

1. If $\underline{j} < \bar{j}$, the maximum intermediate return is given by $\text{MIR} = \bar{R}$.
2. If $\underline{j} = \bar{j}$, the maximum intermediate return is equal to 0.
3. If $\underline{j} > \bar{j}$, the intermediate return $p_{\underline{j}}/p_{\bar{j}} - 1$ may be equal to the maximum intermediate loss \underline{R} , but there may exist other returns with larger absolute values. Therefore, the N prices are placed in three groups depending on their position relative to \underline{j} and \bar{j} : (L) $j \leq \bar{j}$, (M) $\bar{j} < j < \underline{j}$, and (R) $j > \underline{j}$ for $j = 0, \dots, N - 1$. Each group then contains its own maximum intermediate return, to serve as candidate for the global maximum intermediate return:

(L) The minimum price in this group is used to compute the maximum intermediate gain \bar{R}_L .

(R) The maximum price in this group is used to compute the maximum intermediate gain \bar{R}_R .

(M) The recursion to create three new groups is used to compute the middle group's maximum intermediate return \bar{R}_M .

Finally, the MIR is equal to either \underline{R} , \bar{R}_L , \bar{R}_M , or \bar{R}_R .

To further decrease the computational cost, time bars are considered instead of actual trades. If each bar carries the minimum and maximum trade prices in the bar along with their indexes, the algorithm above will in general generate MIR values identical to the values obtained from individual trades¹⁷.

¹⁷If a bar includes both the minimum and maximum price, with the former entering last, the maximum intermediate loss can be computed, but it may be impossible to determine whether an intermediate gain with larger absolute value exists.

6. Empirical Design

Following the notation used by Wu et al. (2013a), let β denote the number of buckets per day (on average), σ denote the support window (in days), and η denote the event horizon (also in days). Parameters σ and η represent a fraction of buckets per day; for example, fifty buckets are used to compute VPIN when $\beta = 100$ and $\sigma = 1/2$. Table 1 reports the VPIN parameter specifications used for the empirical design; the values are in accordance with the optimal parameter settings proposed by Wu et al. (2013a) and Song et al. (2014).

Table 1: SPECIFICATION OF VPIN PARAMETERS

	Description	Values
β	Buckets per day (average)	{100, 200, 500, 100}
σ	Support window (days)	{1/4, 1/2, 1}
η	Event duration (days)	{1/10, 1/4, 1/2}
c	VPIN threshold	{0.9, 0.99}
BVC	Bulk volume classification (time bars)	10 seconds

Description: Specification of the VPIN parameters.

First, the VPIN is computed across trading years under different β and σ combinations. Both moment estimation and maximum likelihood estimation are considered. The latter method is also used to compute the PIN. The characteristics of the three series are then presented; it is possible to determine whether VPIN approximately measures PIN. In contrast to the majority of empirical studies in the literature, the computations are out-of-sample. The daily average of volume from the previous sample (year) is used to compute the volume bucket size for the current sample, and the (V)PIN is updated in real time. From the practitioner's perspective, a (V)PIN metric performing well in out-of-sample will reinforce the argument that (V)PIN is a good tool for signaling market turbulence.

Second, (V)PIN events are detected from the computed (V)PIN series under the threshold c . From figure 5 in appendix B, the log-normal distribution is a reasonable approximation of V(PIN), and the corresponding CDF is used to identify extremely toxic order flow. The events are then classified as true or false positives under η . For each event, the associated MIR

value is computed and compared with the average MIR of 10000 randomly selected time intervals. If the positive (negative) MIR is above (below) the average MIR, the event is classified as a true positive; otherwise, it is a false positive. Finally, the false discovery rates for a given η between 2008-2015 for all the β and σ combinations are computed and reported. The optimal specification of (β, σ) is the pair yielding the smallest false discovery rate. This experiment allows for comparing the predictive power for short-term return volatility between the three series. For instance, is (V)PIN a good predictor of short-term return volatility? Is the maximum likelihood estimation of (V)PIN superior to the moment estimation of VPIN in terms of false discovery rates?

Eventually, two less extensive studies are conducted. The first highlights a significant drawback when the false discovery rate is used as performance evaluation for VPIN; that is, type II errors are neglected. Raising the CDF-threshold will likely increase the likelihood of true positives at the expense of more type II errors. This performance trade-off is completely suppressed in the research by Wu et al. (2013a) and Song et al. (2014), but may be captured by using other methods, as illustrated in a simple but useful example.

The final study focuses on the Flash Crash of May 6, 2010, where the relative advantages of the MLE of VPIN in a stressed market are illustrated, such as capturing information from the volume time and inferring the level of informed and noise trading prior to the crash.

7. Results

As a starting point, the (V)PIN series are computed across all trading years, with the relevant (β, σ) combinations listed in table 1. The fixed-size volume for a trading year will clearly depend on β as well as the average daily traded volume from the preceding year reported in table 13¹⁸.

Tables 6a–6i in appendix B present the summary statistics for the maximum likelihood estimator of θ , which is used to compute VPIN^{mle} , VPIN^{mle} , and PIN^{mle} , respectively.

¹⁸For instance, when $\beta = 100$, the volume bucket size in 2010 was approximately equal to 2.16×10^6 .

The estimates vary with the specific choice of (β, σ) but are, in general, quite stable. The left-hand side of tables 7a–7i shows the average and standard error (only for the two parametric measures) of the three (V)PIN series across the trading years¹⁹.

7.1. Does VPIN Measure PIN?

Easley et al. (2012) claim that the moment estimator of VPIN overcomes numerical difficulties when estimating PIN models in highly active markets. This assertion implicitly assumes that VPIN approximates PIN, which may be tested using the maximum likelihood estimator of PIN (and VPIN). Two loss functions are used to assess whether the moment estimator of VPIN approximately measures VPIN and/or PIN. First, consider the quadratic distance between $(\text{VPIN}_\tau^{mle})_{\tau \geq 1}$ and x ,

$$L_1(x) = \frac{1}{N} \sum_{\tau=1}^N (\text{VPIN}_\tau^{mme} - x_\tau)^2,$$

for x equal to either VPIN^{mle} or PIN^{mle} . This approach yields a direct comparison between the two series. The second loss function computes the ratios between the series x and the reference VPIN^{mme} :

$$L_2(x) = \frac{1}{N} \sum_{\tau=1}^N \frac{x_\tau}{\text{VPIN}_\tau^{mme}}$$

The right-hand side of tables 7a–7i in appendix B reports the findings across trading years for all the nine β and σ combinations. For simplicity, only the findings for $(\beta, \sigma) = (100, 1)$ and $(\beta, \sigma) = (500, 1)$ reported in the following table are discussed in this section; however, the conclusion remains unchanged.

¹⁹Maximum likelihood estimation of the underlying model parameter is also used to estimate the Hessian matrix, which is used to compute standard errors. For VPIN^{mme} , the standard deviation is reported.

Table 2: COMPARISON OF VPIN ESTIMATORS

(a) $(\beta, \sigma) = (100, 1)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.258 (0.083)	0.243 (0.028)	0.344 (0.026)	0.005	0.014	0.997	1.417
2009	0.215 (0.041)	0.214 (0.029)	0.359 (0.029)	0.004	0.040	1.027	1.721
2010	0.228 (0.058)	0.248 (0.03)	0.372 (0.028)	0.003	0.026	1.126	1.707
2011	0.243 (0.074)	0.24 (0.03)	0.372 (0.029)	0.005	0.024	1.041	1.624
2012	0.213 (0.032)	0.216 (0.03)	0.398 (0.032)	0.004	0.052	1.032	1.888
2013	0.237 (0.05)	0.201 (0.033)	0.478 (0.039)	0.008	0.087	0.867	2.077
2014	0.216 (0.061)	0.242 (0.03)	0.354 (0.028)	0.004	0.024	1.185	1.743
2015	0.241 (0.071)	0.241 (0.03)	0.368 (0.029)	0.004	0.023	1.049	1.609

(b) $(\beta, \sigma) = (500, 1)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.419 (0.12)	0.349 (0.015)	0.514 (0.011)	0.009	0.018	0.861	1.302
2009	0.377 (0.065)	0.31 (0.015)	0.501 (0.012)	0.009	0.025	0.840	1.373
2010	0.386 (0.085)	0.332 (0.016)	0.533 (0.012)	0.005	0.027	0.876	1.431
2011	0.398 (0.11)	0.325 (0.016)	0.525 (0.012)	0.010	0.024	0.840	1.394
2012	0.354 (0.046)	0.284 (0.016)	0.524 (0.014)	0.008	0.036	0.807	1.499
2013	0.383 (0.073)	0.293 (0.017)	0.577 (0.014)	0.012	0.050	0.773	1.554
2014	0.367 (0.094)	0.325 (0.015)	0.493 (0.012)	0.005	0.022	0.915	1.408
2015	0.39 (0.1)	0.325 (0.016)	0.512 (0.012)	0.009	0.023	0.854	1.376

Description: The left-hand side of table a (b) reports the summary statistics for the three (V)PIN series across the trading years under model specification $(\beta, \sigma) = (100, 1)$ ($(\beta, \sigma) = (500, 1)$). Standard errors are computed using the BHHH-estimator. The right-hand side presents the values computed by the two loss functions $L_1(\cdot)$ and $L_2(\cdot)$ for comparison of the (V)PIN series.

The left-hand side of the two subtables shows the averages and standard errors of the VPIN series for all trading years²⁰. The last four columns report the values for the two loss functions. Consistent with theory, the values for $L_1(\cdot)$ indicate that the moment estimator of VPIN is most precise for measuring VPIN. This is supported by the ratios reported in the last two columns: on average, PIN^{mle} is 72% (42%) higher than PIN^{mle} in the upper (lower) table, whereas VPIN^{mle} is 4% (15%) higher (lower). Finally, the numbers reveal that the VPIN^{mme} metric becomes unstable for small volume buckets (equivalently for large β) when measuring VPIN. This is consistent with the theoretical conjecture by Lin and Ke (2017).

Summarizing, moment estimation of VPIN does not measure PIN, nor is it a compelling measure of the VPIN for large β . It is therefore pertinent to compare all three metrics' ability to predict short-term volatility.

²⁰For VPIN^{mme}, the standard deviation is reported.

7.2. False Discovery Rates

The main research question is whether maximum likelihood estimation would improve the predictability of VPIN for changes in short-term volatility. To answer this question, the false discovery rates for both methods are computed and compared. Finally, the PIN metric's ability to predict flow toxicity is investigated.

MME of VPIN

Tables 8–9 in appendix B report the false discovery rates under moment estimation of VPIN for the two thresholds $c = .9$ and $c = .99$, respectively. The VPIN's ability to predict short-term volatility is significantly stronger considering the .99-threshold. For instance, if the event horizon is $\eta = 1/4$, the false discovery rates are between 0.297 and 0.381 when the VPIN events are detected with a CDF threshold of 0.9. At best, the VPIN succeeded only in anticipating short-term volatility in two-thirds of all events. When $c = 0.99$, the false discovery rates shrink to the 0.149–0.372 range²¹. This result is consistent with the empirical findings of Wu et al. (2013a) and Song et al. (2014) showing the importance of the CDF threshold. The case of $c = 0.9$ is not considered in the remainder of this analysis.

Table 3 identifies the best pair of (β, σ) resulting in the lowest false discovery rates for $\eta \in \{1/10, 1/4, 1/2\}$. In this setting, the optimal choice of buckets per day and support window coincides for all the three event horizons; that is, $(\beta, \sigma) = (500, 1)$. The resulting false discovery rates are 0.2299, 0.1494, and 0.1602, respectively²².

Table 3: MME OF VPIN: OPTIMAL FALSE DISCOVERY RATES

β	σ	η	False discovery rate
500	1	1/10	0.2299
500	1	1/4	0.1494
500	1	1/2	0.1602

Description: Optimal false discovery rates under moment estimation of VPIN given the threshold $c = 0.99$.

²¹The same conclusion is obtained for $\eta = 1/10$ or $\eta = 1/2$.

²²For instance, 84% of the times a VPIN event is detected, the price of SPY would rise or fall within 3.25 hours more than expected.

Finally, the false discovery rates for $\beta = 1000$ are also reported in the appendix; the predictability of short-term volatility is not monotonically decreasing in the fixed-size volume.

MLE OF VPIN

The false discovery rates computed from the VPIN^{mme} metric are presented in table 10. This metric is superior to the MME for predicting flow toxicity: 22 of the 27 scenarios show lower false discovery rates. Furthermore, the optimal choice for VPIN^{mme} , $(\beta, \sigma) = (500, 1)$ still leads to lower false discovery rates under MLE for all the three event horizons. The two metrics agree only on the optimal bucket size ($\beta = 500$), as illustrated in table 4, which reports the lowest false discovery rates for each η .

Table 4: MLE OF VPIN: OPTIMAL FALSE DISCOVERY RATES

β	σ	η	False discovery rate
500	1/2	1/10	0.1538
500	1	1/4	0.1007
500	1/4	1/2	0.1476

Description: Optimal false discovery rates under maximum likelihood estimation of VPIN given the threshold $c = 0.99$.

A direct comparison between the two VPIN estimators shows strong evidence that the MLE method indeed improves the predictability for short-term volatility. This improved performance comes with a price in terms of computational cost in parametric estimation using high-frequency trade data. However, both the estimators agree on the optimal bucket size ($\beta = 500$), and the (fast) method of moment estimation may be used to identify parameter β before the maximum likelihood estimation of VPIN.

MLE of PIN

Section 3 (definition 1) stated that VPIN and PIN are two different measures. Moreover, numerical evidence has shown that the moment estimation of VPIN does not approximately measure PIN. However, the MLE of the model parameter $(\alpha, \delta, \mu, \epsilon_B, \epsilon_S)$ also may be used to estimate PIN, thus raising the question whether PIN is a good predictor of short-term volatility.

This matter has never been investigated in published the high-frequency market microstructure literature, since all previous studies have relied on the method of moment estimation of VPIN, incorrectly perceived as a proxy for PIN.

The false discovery rates for the PIN^{mle} -metric are reported in table 11. Strikingly, this method yields extremely high false discovery rates ranging from 38.2% to 55.8%. The majority of rates are around 50%, indicating a significantly worse outcome than similar values obtained from the VPIN metric. These findings reinforce the argument that toxic order flow in a high-frequency market microstructure landscape is more accurately modeled on the volume clock rather than a fixed time interval.

7.3. Exhibit I: Type II Errors

A significant drawback of concentrating solely on minimizing the false discovery rate is that type II errors/false negatives (FN) as well as true negatives (TN) are neglected. For instance, in an empirical study, Leth (2019) shows that VPIN failed to predict large intraday price movements on September 22, 2008, under the $c = 0.99$ threshold. A high threshold (such as $c = 0.99$) may result in a low false positive rate, $FPR = \#FP/(\#FP + \#TN)$, at the expense of a low true positive rate (sensitivity), $TPR = \#TP/(\#TP + \#FN)$, as graphically illustrated by the ROC (receiver operating characteristic) curve. This performance trade-off is not captured by the false discovery rate and therefore completely ignored in empirical studies by Wu et al. (2013a) and Song et al. (2014). However, consider the F_1 -score,

$$F_1(c) = \frac{2 \times \#TP}{2 \times \#TP + \#FP + \#FN},$$

interpreted as the harmonic mean of precision and sensitivity under the threshold c . The combined metric takes values in the unit interval with large values preferred, and computation of the score requires classification of type II errors. However, the methodology used to compute false discovery rates is unsuitable for counting false negatives (and true negatives) owing to the overwhelmingly large number of trade sequence combinations for a given event horizon.

Thus, a novel but simple approach is proposed instead. The MIR is computed on a daily basis for all the trades resulting in about 252 MIR values. The trading days with MIR value below (above) the 0.025-quantile (0.975-quantile) of all MIR values are labeled high volatility. On trading days with VPIN events, the MIR is computed between the first hitting time and final trade, thus allowing for classifying the true and false positives. High-volatility days with no VPIN events are classified as false negatives, whereas the days with neither event type are true negatives. This empirical study is conducted with the specification $(\beta, \sigma) = (500, 1)$ under the two thresholds $c_1 = 0.9$ and $c_2 = 0.99$. Table 5 reports the relative performance of the F_1 score for relevant cases.

Table 5: F_1 -SCORES

$F_1^{mle}(0.99)/F_1^{mle}(0.9)$	$F_1^{mme}(0.99)/F_1^{mme}(0.9)$	$F_1^{mle}(0.9)/F_1^{mme}(0.9)$	$F_1^{mle}(0.99)/F_1^{mme}(0.99)$
0.770	0.763	1.113	1.123

Description: The first two columns present the relative performance of thresholds $c = 0.99$ and $c = 0.9$ for the given VPIN estimator. The last two columns show the relative performance of the two VPIN estimators for the given threshold c .

The F_1 -scores ratio is used as a measure for relative performance. The left-hand side of table 5 shows the threshold $c = 0.99$ relative to $c = 0.9$ reducing the F_1 -score with 23.0% (23.7%) under maximum likelihood (moment) estimation of VPIN. Thus, the CDF threshold $c = 0.9$ yields the best performance for VPIN as predictor of volatility events in terms of larger F_1 -score from fewer type II errors. This highlights a potential shortcoming from solely focusing on the false discovery rate, and performance evaluation of VPIN must be based on the practitioner's specific problem²³.

The right-hand side of table 5 reports the relative performance of the two VPIN estimators. Compared with the moment estimation of VPIN, MLE will improve the F_1 -score by 11.3% (12.3%) under the threshold $c = 0.9$ ($c = 0.99$). In line with the results of false discovery rates, the VPIN^{mme} -metric is a superior predictor of short-term volatility.

²³For instance, a portfolio manager investing with buy-and-hold strategies is prone to change the composition of his portfolio if volatility events are likely to occur. On the other hand, an option dealer rebalancing her Δ -hedged portfolio on a daily basis will prefer fewer type II errors at the expense of lower precision.

7.4. Exhibit II: The Flash Crash

Finally, the VPIN^{mle} -metric's performance as a warning signal for extreme market turbulence is illustrated. The analysis is concentrated on the Flash Crash of May 6, 2010, an event that has drawn much attention in the academic literature in recent years.

The VPIN^{mle} utilizes the volume time when estimating the model parameter θ . When $\beta = 500$, the average duration time to fill a volume bucket in 2010 was 55.7 seconds, with the median duration time of 38.5 seconds. On the other hand, the average and median of the May 6 duration times were 17.9 seconds and 12.9 seconds, respectively. Figure 2 indicates that the duration decreased dramatically around 14:00, with the falling prices accelerating right around 14:32. The price reached its minimum at 14:46, and by 15:00, the SPY had recovered from the crash²⁴.

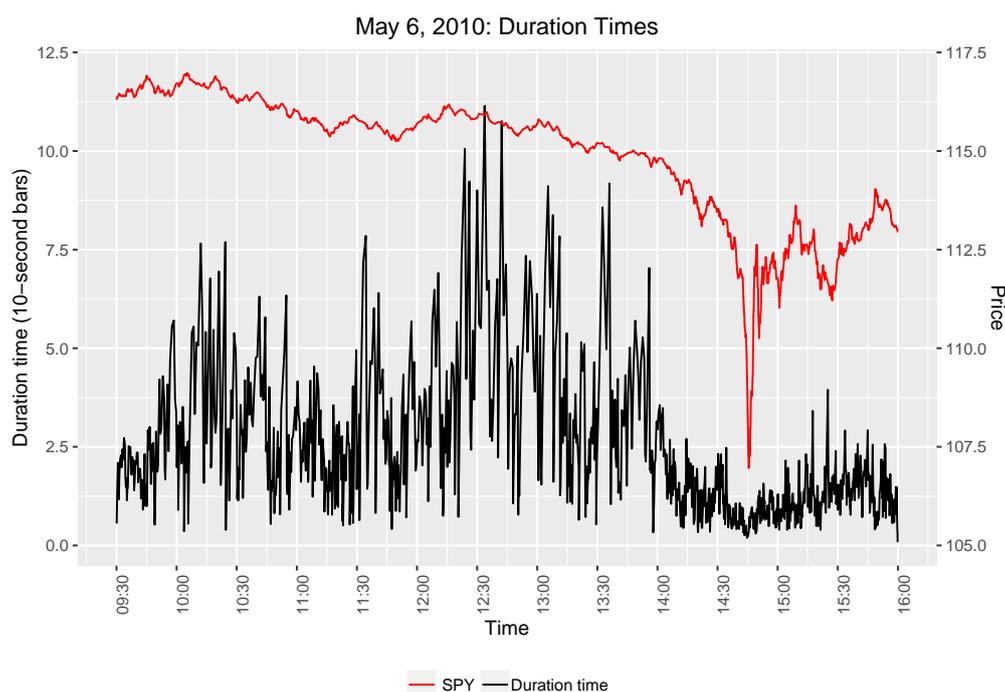


Figure 2: The black line illustrates the time elapsed to fill a volume bucket (size) measures with 10-second bars. The red line represents the price of SPY.

²⁴The SEC/CFTC report states that the crash began around 14:32 and ended by 15:00.

According to Easley et al. (2012), the order flow became increasingly toxic before the crash, contributing to liquidity providers (high-frequency traders) withdrawing from the market and instead consumed liquidity. This observation is consistent with extremely low duration times almost one hour before the crash. Figure 3 illustrates the $VPIN^{mle}$ -metric along with its CDF. The VPIN starts increasing just before 14:00. At 14:26, the CDF reaches the user-determined threshold 0.99. The VPIN rapidly increased prior to the crash, followed by a similar decline in value, abstracting from the sudden spike around 14:52 due to extreme market condition (see figure 4).

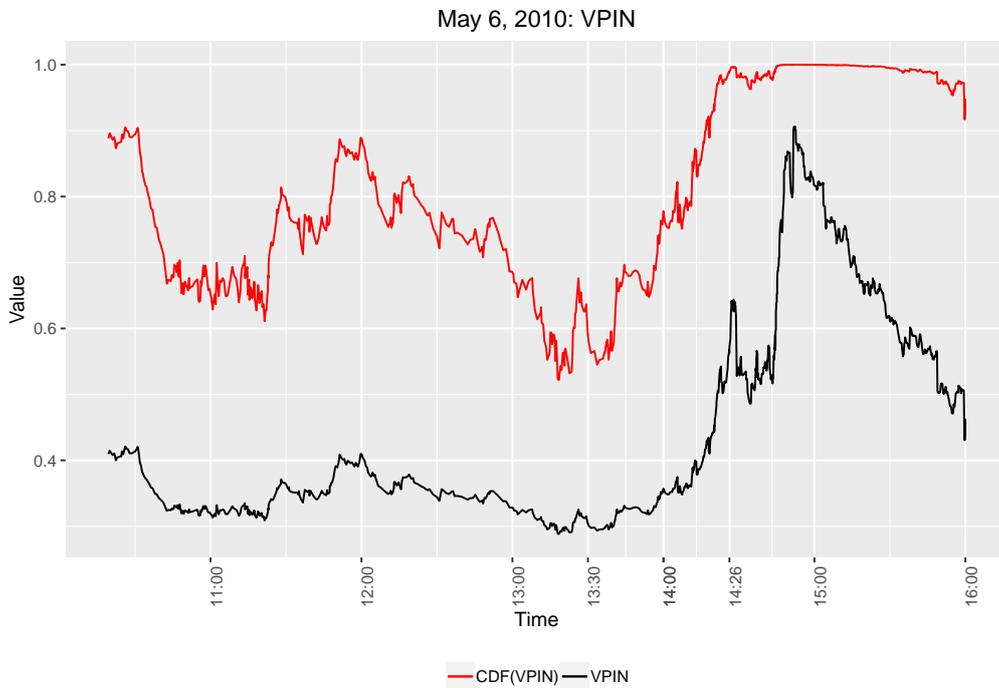


Figure 3: The black line shows the VPIN on May 6, 2010, obtained with maximum likelihood estimation. The VPIN parameters are set to $(\beta, \sigma) = (500, 1/4)$. The red line illustrates the corresponding CDF of VPIN.

The method of moment estimation associates extreme VPIN values with high order imbalances. In contrast, the parametric framework identifies the forces at work. Four new measures are illustrated in figure 4: $EOI = |2\alpha\delta\mu + \epsilon_B - \alpha\mu - \epsilon_B|$ (expected order imbalance), $IV = \mu$ (informed volume), $NV = \epsilon_B + \epsilon_S$ (noise volume), and $NOI = |\epsilon_B - \epsilon_S|$ (noise order imbalance).

During the crash, the order imbalance is extremely high, which is caused by both informed and noise trading. Informed volume increases from 14:00, and the order flow becomes toxic. However, the noise order imbalance is close to 0 at 14:26, and thereafter starts to rise; this constitutes the beginning of the crash. Finally, informed trading starts declining before the end of the crash, whereas the VPIN reaches its maximum subsequent to the crash after a temporary spike. The reason for this is an unforeseen drop in sell volume from uninformed traders. Eventually, the noise traders' selling volume accelerates again.

Andersen and Bondarenko (2014b) argue that the original VPIN metric is useless as a warning signal for imminent market turbulence, since this measure keeps on rising after the crash. However, the moment estimator is derived under the assumption $\epsilon_B = \epsilon_S$, clearly in contrast to the panel "NOISE ORDER IMBALANCE". If the volume falls but order balance remains high, the VPIN^{mme} metric will keep on rising. By capturing the information in volume time, the maximum likelihood estimator will enable the market makers to perceive the sell orders as in fact stemming from noise traders.

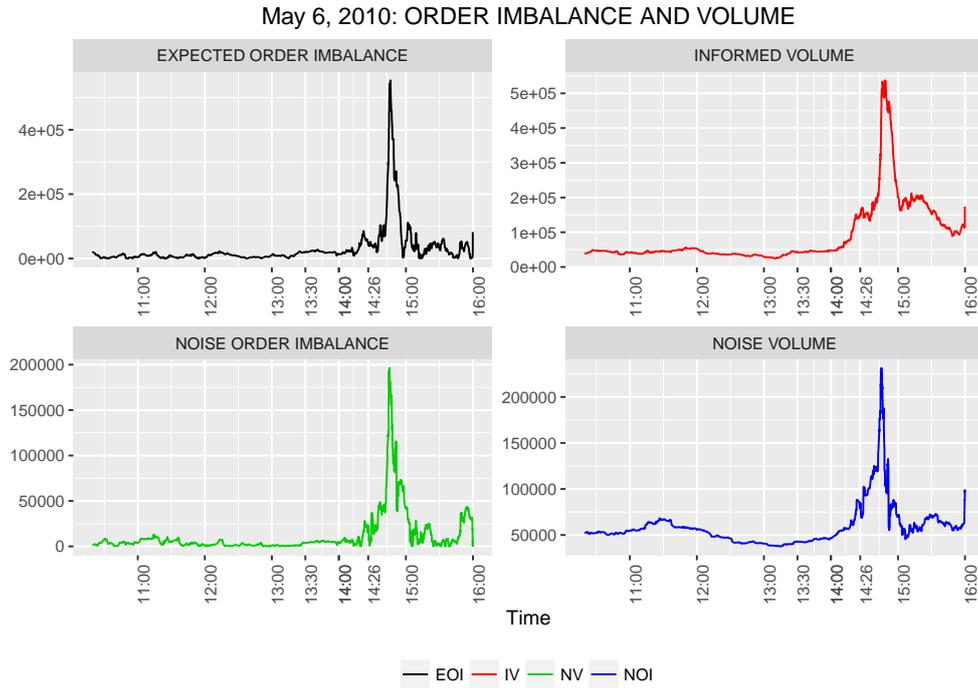


Figure 4: Left panel: The top plot shows the expected order imbalance obtained from estimates of the model parameter, whereas the order imbalance solely generated by uninformed traders ($|\epsilon_B - \epsilon_S|$) is illustrated in the bottom plot. Right panel: Informed volume (μ) is represented by the red line in the upper plot, while the blue line indicates volume from uninformed traders ($\epsilon_B + \epsilon_S$).

8. Conclusion

This paper examines whether maximum likelihood estimation of the VPIN metric improves its predictive power for short-term return volatility. The MLE method, which captures information in volume time, was recently developed by Lin and Ke (2017), and enables the estimation of a modified PIN model in high-frequency market microstructure. Not only VPIN but also PIN may be computed, which is typically difficult to estimate in highly active markets.

The main contribution of this paper is to conduct a large empirical investigation on the MLE of VPIN. The empirical study is concentrated on the SPRD S&P500 ETF (SPY) traded from 2007-2015, with the VPIN computations done out-of-sample. So far, studies in the literature, whether supporting or disputing the usefulness of VPIN as a real-time indicator of toxic order flow, have been based solely on the moment estimation of VPIN. In contrast, this paper directly compares the two VPIN estimator

performances as predictor of short-term return volatility. Moreover, this work contributes with a similar investigation of the usefulness of the PIN metric.

The findings show that maximum likelihood estimation of VPIN is superior to the moment estimation of VPIN for predicting large intraday price movements. Under moment estimation, toxic order flow reflected by high VPIN readings would in 77% – 85% of the cases — depending on the event horizon — lead to large price movements. With maximum likelihood estimation of VPIN, the corresponding precision (positive predictive rate) was between 85% and 90%. In addition, the moment estimation of VPIN became unstable for small volume buckets that are typically optimal in empirical applications.

The results also reveal that VPIN does not approximately measure PIN, emphasizing the theoretical disagreement between the two measures. Only the MLE of the model parameter should be used to compute PIN.

As regards the PIN's ability to anticipate peaks in return volatility, the results are striking. When the toxic order flow was measured by extreme PIN values, the positive predictive rate was about 50%, and thus inferior to VPIN. This observation reinforces the argument that flow toxicity is more accurately measured on the volume clock.

This paper can motivate future research through a minor case study illustrating a performance trade-off between the positive predictive rate and type II errors captured by the F_1 -score. The practitioner's key preference may be to avoid type II errors. However, the performance evaluation of VPIN in this work, which is similar to the methodology used by Wu et al. (2013a), neglects false negatives. An optimal positive predictive rate may lead to a suboptimal F_1 -score. Thus, it is pertinent to conduct a larger study investigating this decision problem from the practitioner's perspective.

Appendix A Cleaning High-Frequency Trade Data

Data cleaning follows the procedure by Holden and Jacobsen (2013). Only trades with time stamps between 9:30 am and 4:00 pm are considered. The entries with transaction price less than or equal to zero are deleted. The corrected trades and trades with abnormal sale condition are also removed from the dataset. These restrictions correspond to the following conditions in the datafile:

- Only trades during the core trading session: $utcsec \in [34000, 56000]$ (seconds after midnight).
- Only trades with positive prices: $price > 0$.
- Only trades with normal sales conditions: $cond \in [",", "@", "E", "F", "I", "FI"]$.
- Only trades that are not corrected: $corr \in ["0", "00"]$.

For more details see https://www.nyse.com/publicdocs/nyse/data/Daily_TAQ_Client_Spec_v3.0.pdf.

Appendix B Empirical Findings: Figures and Tables

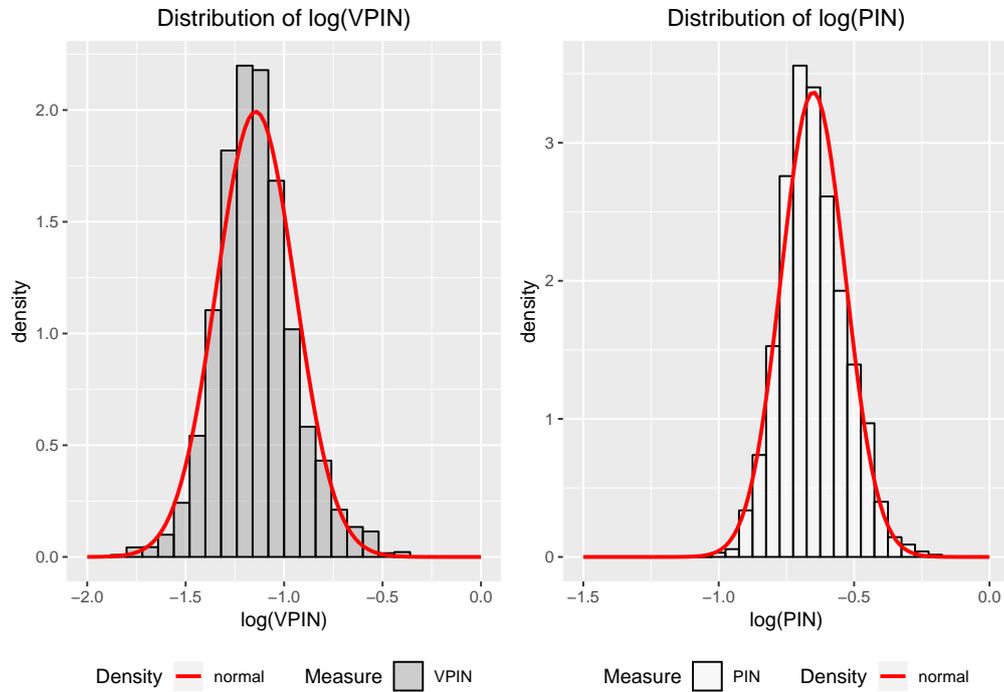


Figure 5: Left plot: Histogram of the logarithm of VPIN under maximum likelihood estimation for a given trading year. Right plot: Histogram of the logarithm of PIN under maximum likelihood estimation. In both panels, the red line represents the density from the normal distribution (with mean and variance equal to the sample mean and sample variance, respectively, of the relevant series).

Table 6: MLE — SUMMARY STATISTICS ($\beta = 100$)(a) $(\beta, \sigma) = (100, 1/4)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.465 (0.11)	0.516 (0.18)	2.292×10^5	6.161×10^4	6.154×10^4
2009	0.416 (0.11)	0.506 (0.2)	5.762×10^5	5.099×10^4	5.083×10^4
2010	0.425 (0.11)	0.496 (0.2)	2.817×10^5	4.289×10^4	4.372×10^4
2011	0.417 (0.11)	0.497 (0.19)	2.773×10^5	4.518×10^4	4.579×10^4
2012	0.371 (0.11)	0.502 (0.21)	4.554×10^5	2.668×10^4	2.645×10^4
2013	0.358 (0.11)	0.512 (0.22)	5.222×10^5	2.467×10^4	2.521×10^4
2014	0.444 (0.11)	0.509 (0.2)	5.822×10^4	2.101×10^4	2.103×10^4
2015	0.43 (0.11)	0.526 (0.19)	7.712×10^4	2.307×10^4	2.357×10^4

(b) $(\beta, \sigma) = (100, 1/2)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.458 (0.074)	0.512 (0.12)	2.539×10^5	5.923×10^4	5.913×10^4
2009	0.411 (0.076)	0.515 (0.13)	6.740×10^5	4.684×10^4	4.689×10^4
2010	0.428 (0.077)	0.495 (0.13)	2.389×10^5	3.861×10^4	3.939×10^4
2011	0.422 (0.075)	0.494 (0.13)	2.140×10^5	4.143×10^4	4.196×10^4
2012	0.38 (0.075)	0.506 (0.15)	4.431×10^5	2.343×10^4	2.335×10^4
2013	0.329 (0.073)	0.507 (0.15)	6.308×10^5	2.177×10^4	2.209×10^4
2014	0.442 (0.078)	0.505 (0.14)	5.079×10^4	1.893×10^4	1.903×10^4
2015	0.427 (0.075)	0.52 (0.13)	7.833×10^4	2.086×10^4	2.139×10^4

(c) $(\beta, \sigma) = (100, 1)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.467 (0.052)	0.512 (0.086)	1.722×10^5	5.528×10^4	5.530×10^4
2009	0.407 (0.053)	0.522 (0.1)	8.722×10^5	4.411×10^4	4.422×10^4
2010	0.443 (0.054)	0.497 (0.093)	1.583×10^5	3.541×10^4	3.597×10^4
2011	0.426 (0.053)	0.493 (0.093)	2.104×10^5	3.805×10^4	3.855×10^4
2012	0.374 (0.051)	0.504 (0.1)	4.569×10^5	2.253×10^4	2.248×10^4
2013	0.308 (0.049)	0.515 (0.12)	6.515×10^5	2.020×10^4	2.040×10^4
2014	0.445 (0.054)	0.5 (0.096)	4.655×10^4	1.746×10^4	1.757×10^4
2015	0.428 (0.053)	0.517 (0.092)	7.662×10^4	1.901×10^4	1.944×10^4

Description: Maximum likelihood estimates of the model parameter across trading years for the three different support window choices. Standard errors are computed using the BHHH-estimator. The bucket size (β) is equal to 100.

Table 6: CONTINUED ($\beta = 200$)(d) $(\beta, \sigma) = (200, 1/4)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.477 (0.072)	0.512 (0.12)	2.674×10^5	5.697×10^4	5.677×10^4
2009	0.422 (0.073)	0.504 (0.12)	6.319×10^5	4.853×10^4	4.822×10^4
2010	0.426 (0.073)	0.496 (0.13)	3.123×10^5	4.066×10^4	4.124×10^4
2011	0.413 (0.071)	0.497 (0.13)	3.131×10^5	4.263×10^4	4.303×10^4
2012	0.355 (0.071)	0.502 (0.15)	5.022×10^5	2.591×10^4	2.572×10^4
2013	0.346 (0.072)	0.51 (0.15)	5.692×10^5	2.380×10^4	2.428×10^4
2014	0.457 (0.074)	0.506 (0.12)	6.227×10^4	1.975×10^4	1.988×10^4
2015	0.433 (0.072)	0.527 (0.12)	8.586×10^4	2.154×10^4	2.204×10^4

(e) $(\beta, \sigma) = (200, 1/2)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.475 (0.051)	0.509 (0.082)	2.608×10^5	5.457×10^4	5.467×10^4
2009	0.421 (0.052)	0.511 (0.088)	6.642×10^5	4.472×10^4	4.461×10^4
2010	0.441 (0.052)	0.501 (0.087)	2.539×10^5	3.630×10^4	3.690×10^4
2011	0.427 (0.051)	0.495 (0.089)	2.383×10^5	3.900×10^4	3.939×10^4
2012	0.368 (0.05)	0.501 (0.1)	4.795×10^5	2.284×10^4	2.267×10^4
2013	0.331 (0.048)	0.513 (0.11)	5.807×10^5	2.095×10^4	2.135×10^4
2014	0.462 (0.052)	0.505 (0.087)	5.328×10^4	1.782×10^4	1.797×10^4
2015	0.438 (0.051)	0.521 (0.085)	8.328×10^4	1.952×10^4	2.006×10^4

(f) $(\beta, \sigma) = (200, 1)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.485 (0.036)	0.509 (0.056)	1.657×10^5	5.122×10^4	5.123×10^4
2009	0.421 (0.036)	0.509 (0.063)	7.837×10^5	4.213×10^4	4.203×10^4
2010	0.456 (0.037)	0.5 (0.06)	1.307×10^5	3.341×10^4	3.396×10^4
2011	0.438 (0.036)	0.495 (0.061)	2.010×10^5	3.579×10^4	3.604×10^4
2012	0.372 (0.035)	0.498 (0.067)	3.926×10^5	2.184×10^4	2.169×10^4
2013	0.331 (0.033)	0.516 (0.074)	5.033×10^5	1.926×10^4	1.948×10^4
2014	0.462 (0.037)	0.504 (0.061)	5.232×10^4	1.654×10^4	1.666×10^4
2015	0.441 (0.036)	0.52 (0.059)	8.647×10^4	1.788×10^4	1.833×10^4

Description: Maximum likelihood estimates of the model parameter across trading years for the three different support window choices. Standard errors are computed using the BHHH-estimator. The bucket size (β) is equal to 200.

Table 6: CONTINUED ($\beta = 500$)(g) $(\beta, \sigma) = (500, 1/4)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.514 (0.045)	0.511 (0.068)	2.767×10^5	5.004×10^4	5.021×10^4
2009	0.452 (0.045)	0.507 (0.073)	6.090×10^5	4.399×10^4	4.403×10^4
2010	0.446 (0.044)	0.495 (0.074)	3.054×10^5	3.676×10^4	3.704×10^4
2011	0.444 (0.044)	0.5 (0.076)	2.812×10^5	3.799×10^4	3.858×10^4
2012	0.375 (0.043)	0.506 (0.084)	4.670×10^5	2.420×10^4	2.413×10^4
2013	0.376 (0.044)	0.509 (0.085)	5.317×10^5	2.208×10^4	2.250×10^4
2014	0.485 (0.045)	0.502 (0.07)	6.950×10^4	1.772×10^4	1.785×10^4
2015	0.467 (0.045)	0.523 (0.071)	8.834×10^4	1.907×10^4	1.963×10^4

(h) $(\beta, \sigma) = (500, 1/2)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.517 (0.032)	0.509 (0.048)	2.392×10^5	4.790×10^4	4.825×10^4
2009	0.46 (0.032)	0.506 (0.05)	5.723×10^5	4.051×10^4	4.044×10^4
2010	0.463 (0.032)	0.495 (0.05)	2.130×10^5	3.293×10^4	3.331×10^4
2011	0.457 (0.031)	0.498 (0.051)	2.096×10^5	3.477×10^4	3.529×10^4
2012	0.395 (0.031)	0.504 (0.057)	4.053×10^5	2.125×10^4	2.121×10^4
2013	0.377 (0.031)	0.509 (0.056)	4.293×10^5	1.923×10^4	1.948×10^4
2014	0.49 (0.032)	0.503 (0.049)	6.227×10^4	1.609×10^4	1.625×10^4
2015	0.473 (0.032)	0.52 (0.049)	8.830×10^4	1.743×10^4	1.792×10^4

(i) $(\beta, \sigma) = (500, 1)$

Year	$\hat{\alpha}$	$\hat{\delta}$	$\hat{\mu}$	$\hat{\epsilon}_B$	$\hat{\epsilon}_S$
2008	0.521 (0.022)	0.509 (0.033)	1.908×10^5	4.533×10^4	4.569×10^4
2009	0.462 (0.023)	0.503 (0.035)	5.265×10^5	3.831×10^4	3.816×10^4
2010	0.47 (0.023)	0.496 (0.035)	1.570×10^5	3.064×10^4	3.090×10^4
2011	0.463 (0.022)	0.498 (0.035)	1.645×10^5	3.223×10^4	3.254×10^4
2012	0.398 (0.022)	0.501 (0.038)	2.104×10^5	2.035×10^4	2.031×10^4
2013	0.388 (0.022)	0.512 (0.038)	2.338×10^5	1.756×10^4	1.780×10^4
2014	0.49 (0.023)	0.502 (0.035)	6.180×10^4	1.501×10^4	1.515×10^4
2015	0.472 (0.022)	0.518 (0.035)	9.252×10^4	1.612×10^4	1.655×10^4

Description: Maximum likelihood estimates of the model parameter across trading years for the three different support window choices. Standard errors are computed using the BHHH-estimator. The bucket size (β) is equal to 500.

Table 7: COMPARISON OF VPIN ESTIMATORS ($\beta = 100$)(a) $(\beta, \sigma) = (100, 1/4)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.258 (0.1)	0.216 (0.051)	0.307 (0.05)	0.007	0.013	0.883	1.268
2009	0.215 (0.064)	0.191 (0.052)	0.303 (0.053)	0.006	0.027	0.936	1.452
2010	0.228 (0.076)	0.207 (0.056)	0.324 (0.057)	0.006	0.023	0.965	1.499
2011	0.243 (0.092)	0.205 (0.055)	0.325 (0.056)	0.008	0.022	0.903	1.435
2012	0.213 (0.059)	0.191 (0.059)	0.355 (0.062)	0.007	0.045	0.946	1.697
2013	0.236 (0.084)	0.196 (0.062)	0.398 (0.063)	0.009	0.060	0.878	1.714
2014	0.216 (0.079)	0.206 (0.054)	0.293 (0.054)	0.006	0.016	1.034	1.477
2015	0.241 (0.089)	0.203 (0.053)	0.297 (0.054)	0.007	0.012	0.891	1.304

(b) $(\beta, \sigma) = (100, 1/2)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.258 (0.092)	0.225 (0.037)	0.325 (0.036)	0.007	0.014	0.928	1.345
2009	0.215 (0.051)	0.206 (0.039)	0.336 (0.04)	0.005	0.034	1.005	1.617
2010	0.228 (0.065)	0.229 (0.042)	0.358 (0.042)	0.006	0.029	1.060	1.661
2011	0.243 (0.082)	0.225 (0.041)	0.35 (0.041)	0.007	0.022	0.989	1.542
2012	0.213 (0.042)	0.213 (0.043)	0.392 (0.045)	0.005	0.054	1.028	1.868
2013	0.236 (0.065)	0.2 (0.046)	0.461 (0.05)	0.009	0.086	0.889	2.004
2014	0.216 (0.07)	0.228 (0.041)	0.331 (0.04)	0.005	0.022	1.132	1.657
2015	0.241 (0.08)	0.224 (0.04)	0.336 (0.041)	0.006	0.019	0.981	1.484

(c) $(\beta, \sigma) = (100, 1)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.258 (0.083)	0.243 (0.028)	0.344 (0.026)	0.005	0.014	0.997	1.417
2009	0.215 (0.041)	0.214 (0.029)	0.359 (0.029)	0.004	0.040	1.027	1.721
2010	0.228 (0.058)	0.248 (0.03)	0.372 (0.028)	0.003	0.026	1.126	1.707
2011	0.243 (0.074)	0.24 (0.03)	0.372 (0.029)	0.005	0.024	1.041	1.624
2012	0.213 (0.032)	0.216 (0.03)	0.398 (0.032)	0.004	0.052	1.032	1.888
2013	0.237 (0.05)	0.201 (0.033)	0.478 (0.039)	0.008	0.087	0.867	2.077
2014	0.216 (0.061)	0.242 (0.03)	0.354 (0.028)	0.004	0.024	1.185	1.743
2015	0.241 (0.071)	0.241 (0.03)	0.368 (0.029)	0.004	0.023	1.049	1.609

Description: The left-hand side of the tables reports summary statistics for the three (V)PIN series across trading years for the three different support window choices. For VPIN^{mme}, the standard deviation is reported, whereas standard errors for the two parametric series are computed using the BHHH-estimator. The right-hand side presents the values computed with the two loss functions $L_1(\cdot)$ and $L_2(\cdot)$ for comparison between the (V)PIN series. The bucket size (β) is equal to 100.

Table 7: CONTINUED ($\beta = 200$)(d) $(\beta, \sigma) = (200, 1/4)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.328 (0.12)	0.261 (0.04)	0.383 (0.037)	0.011	0.014	0.830	1.248
2009	0.278 (0.075)	0.225 (0.04)	0.369 (0.04)	0.009	0.027	0.845	1.376
2010	0.291 (0.09)	0.241 (0.042)	0.39 (0.041)	0.009	0.024	0.864	1.410
2011	0.306 (0.11)	0.237 (0.042)	0.393 (0.043)	0.012	0.023	0.809	1.378
2012	0.268 (0.066)	0.211 (0.044)	0.418 (0.047)	0.011	0.046	0.814	1.586
2013	0.295 (0.095)	0.216 (0.047)	0.46 (0.048)	0.015	0.057	0.763	1.607
2014	0.276 (0.095)	0.241 (0.04)	0.349 (0.038)	0.008	0.015	0.930	1.360
2015	0.303 (0.11)	0.237 (0.04)	0.359 (0.04)	0.010	0.012	0.812	1.247

(e) $(\beta, \sigma) = (200, 1/2)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.328 (0.11)	0.269 (0.029)	0.398 (0.027)	0.009	0.015	0.862	1.298
2009	0.278 (0.062)	0.24 (0.03)	0.396 (0.029)	0.007	0.031	0.897	1.478
2010	0.292 (0.079)	0.266 (0.032)	0.422 (0.029)	0.006	0.028	0.948	1.524
2011	0.307 (0.099)	0.257 (0.031)	0.416 (0.03)	0.009	0.024	0.885	1.457
2012	0.268 (0.049)	0.231 (0.032)	0.449 (0.034)	0.007	0.051	0.879	1.701
2013	0.295 (0.075)	0.223 (0.034)	0.512 (0.036)	0.013	0.077	0.782	1.797
2014	0.276 (0.085)	0.263 (0.03)	0.386 (0.027)	0.005	0.021	1.010	1.501
2015	0.303 (0.096)	0.257 (0.03)	0.396 (0.029)	0.007	0.018	0.885	1.380

(f) $(\beta, \sigma) = (200, 1)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.328 (0.1)	0.286 (0.021)	0.415 (0.018)	0.006	0.014	0.916	1.344
2009	0.279 (0.051)	0.249 (0.022)	0.415 (0.021)	0.006	0.034	0.918	1.539
2010	0.292 (0.07)	0.284 (0.023)	0.437 (0.02)	0.003	0.026	1.002	1.558
2011	0.307 (0.091)	0.274 (0.023)	0.437 (0.021)	0.006	0.024	0.933	1.509
2012	0.268 (0.038)	0.239 (0.023)	0.449 (0.023)	0.005	0.046	0.905	1.693
2013	0.295 (0.06)	0.233 (0.024)	0.524 (0.025)	0.010	0.075	0.801	1.831
2014	0.276 (0.076)	0.275 (0.022)	0.409 (0.02)	0.004	0.023	1.045	1.566
2015	0.303 (0.086)	0.274 (0.023)	0.428 (0.02)	0.005	0.023	0.936	1.482

Description: The left-hand side of the tables reports summary statistics for the three (V)PIN series across trading years for the three different support window choices. For VPIN^{mme}, the standard deviation is reported, whereas standard errors for the two parametric series are computed using the BHHH-estimator. The right-hand side presents the values computed with the two loss functions $L_1(\cdot)$ and $L_2(\cdot)$ for comparison between the (V)PIN series. The bucket size (β) is equal to 100.

Table 7: CONTINUED ($\beta = 500$)**(g)** $(\beta, \sigma) = (500, 1/4)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.418 (0.13)	0.329 (0.029)	0.486 (0.023)	0.014	0.017	0.808	1.241
2009	0.376 (0.089)	0.285 (0.029)	0.465 (0.025)	0.014	0.023	0.779	1.288
2010	0.385 (0.11)	0.296 (0.03)	0.49 (0.026)	0.014	0.024	0.787	1.338
2011	0.398 (0.13)	0.293 (0.03)	0.486 (0.026)	0.018	0.023	0.757	1.315
2012	0.355 (0.076)	0.254 (0.03)	0.498 (0.029)	0.017	0.039	0.731	1.434
2013	0.383 (0.11)	0.264 (0.032)	0.537 (0.029)	0.023	0.048	0.708	1.464
2014	0.367 (0.11)	0.297 (0.028)	0.438 (0.024)	0.011	0.015	0.840	1.271
2015	0.39 (0.13)	0.296 (0.029)	0.451 (0.025)	0.014	0.012	0.776	1.219

(h) $(\beta, \sigma) = (500, 1/2)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.418 (0.13)	0.337 (0.021)	0.498 (0.016)	0.012	0.017	0.832	1.268
2009	0.376 (0.076)	0.301 (0.021)	0.488 (0.017)	0.011	0.025	0.822	1.346
2010	0.385 (0.094)	0.319 (0.022)	0.517 (0.018)	0.009	0.027	0.847	1.403
2011	0.398 (0.12)	0.312 (0.022)	0.507 (0.018)	0.013	0.024	0.809	1.361
2012	0.355 (0.057)	0.277 (0.022)	0.523 (0.02)	0.011	0.041	0.793	1.501
2013	0.383 (0.087)	0.278 (0.023)	0.573 (0.02)	0.018	0.057	0.741	1.554
2014	0.367 (0.1)	0.315 (0.021)	0.473 (0.017)	0.008	0.020	0.892	1.364
2015	0.39 (0.12)	0.313 (0.021)	0.484 (0.017)	0.011	0.018	0.825	1.310

(i) $(\beta, \sigma) = (500, 1)$

Year	VPIN ^{mme}	VPIN ^{mle}	PIN ^{mle}	$L_1(\text{VPIN}^{mle})$	$L_1(\text{PIN}^{mle})$	$L_2(\text{VPIN}^{mle})$	$L_2(\text{PIN}^{mle})$
2008	0.419 (0.12)	0.349 (0.015)	0.514 (0.011)	0.009	0.018	0.861	1.302
2009	0.377 (0.065)	0.31 (0.015)	0.501 (0.012)	0.009	0.025	0.840	1.373
2010	0.386 (0.085)	0.332 (0.016)	0.533 (0.012)	0.005	0.027	0.876	1.431
2011	0.398 (0.11)	0.325 (0.016)	0.525 (0.012)	0.010	0.024	0.840	1.394
2012	0.354 (0.046)	0.284 (0.016)	0.524 (0.014)	0.008	0.036	0.807	1.499
2013	0.383 (0.073)	0.293 (0.017)	0.577 (0.014)	0.012	0.050	0.773	1.554
2014	0.367 (0.094)	0.325 (0.015)	0.493 (0.012)	0.005	0.022	0.915	1.408
2015	0.39 (0.1)	0.325 (0.016)	0.512 (0.012)	0.009	0.023	0.854	1.376

Description: The left-hand side of the tables reports summary statistics for the three (V)PIN series across trading years for the three different support window choices. For VPIN^{mme}, the standard deviation is reported, whereas standard errors for the two parametric series are computed using the BHHH-estimator. The right-hand side presents the values computed with the two loss functions $L_1(\cdot)$ and $L_2(\cdot)$ for comparison between the (V)PIN series. The bucket size (β) is equal to 100.

Table 8: MME OF VPIN: FALSE DISCOVERY RATES I ($c = 0.9$)

β	σ	η	False discovery rate
100	1/4	1/10	0.4167
200	1/4	1/10	0.3842
500	1/4	1/10	0.4275
1000	1/4	1/10	0.3666
100	1/2	1/10	0.3918
200	1/2	1/10	0.4074
500	1/2	1/10	0.3518
1000	1/2	1/10	0.4181
100	1	1/10	0.3959
200	1	1/10	0.3384
500	1	1/10	0.3219
1000	1	1/10	0.3522
100	1/4	1/4	0.3720
200	1/4	1/4	0.3529
500	1/4	1/4	0.3776
1000	1/4	1/4	0.3482
100	1/2	1/4	0.3197
200	1/2	1/4	0.3723
500	1/2	1/4	0.2849
1000	1/2	1/4	0.3406
100	1	1/4	0.3485
200	1	1/4	0.2759
500	1	1/4	0.3023
1000	1	1/4	0.3559
100	1/4	1/2	0.3714
200	1/4	1/2	0.3348
500	1/4	1/2	0.3499
1000	1/4	1/2	0.3265
100	1/2	1/2	0.3208
200	1/2	1/2	0.3029
500	1/2	1/2	0.2648
1000	1/2	1/2	0.2979
100	1	1/2	0.3093
200	1	1/2	0.2931
500	1	1/2	0.3333
1000	1	1/2	0.3183

Description: False discovery rates under moment estimation of VPIN given the threshold $c = 0.9$.

Table 9: MME OF VPIN: FALSE DISCOVERY RATES II ($c = 0.99$)

β	σ	η	False discovery rate
100	1/4	1/10	0.2541
200	1/4	1/10	0.2683
500	1/4	1/10	0.2486
1000	1/4	1/10	0.2778
100	1/2	1/10	0.3486
200	1/2	1/10	0.2795
500	1/2	1/10	0.3415
1000	1/2	1/10	0.3388
100	1	1/10	0.3596
200	1	1/10	0.2772
500	1	1/10	0.2299
1000	1	1/10	0.2778
100	1/4	1/4	0.2418
200	1/4	1/4	0.2967
500	1/4	1/4	0.2486
1000	1/4	1/4	0.1597
100	1/2	1/4	0.2800
200	1/2	1/4	0.2547
500	1/2	1/4	0.3089
1000	1/2	1/4	0.3802
100	1	1/4	0.2809
200	1	1/4	0.2376
500	1	1/4	0.1494
1000	1	1/4	0.1667
100	1/4	1/2	0.2336
200	1/4	1/2	0.2683
500	1/4	1/2	0.1618
1000	1/4	1/2	0.1736
100	1/2	1/2	0.2571
200	1/2	1/2	0.2547
500	1/2	1/2	0.2358
1000	1/2	1/2	0.2149
100	1	1/2	0.2247
200	1	1/2	0.1881
500	1	1/2	0.1602
1000	1	1/2	0.2222

Description: False discovery rates under moment estimation of VPIN given the threshold $c = 0.99$.

Table 10: MLE OF VPIN: FALSE DISCOVERY RATES ($c = 0.99$)

β	σ	η	False discovery rate
100	1/4	1/10	0.2655
200	1/4	1/10	0.2182
500	1/4	1/10	0.1993
100	1/2	1/10	0.3065
200	1/2	1/10	0.2340
500	1/2	1/10	0.1538
100	1	1/10	0.3077
200	1	1/10	0.3049
500	1	1/10	0.2230
100	1/4	1/4	0.2566
200	1/4	1/4	0.1879
500	1/4	1/4	0.1624
100	1/2	1/4	0.2903
200	1/2	1/4	0.1809
500	1/2	1/4	0.1538
100	1	1/4	0.2692
200	1	1/4	0.1951
500	1	1/4	0.1007
100	1/4	1/2	0.2035
200	1/4	1/2	0.1636
500	1/4	1/2	0.1476
100	1/2	1/2	0.1935
200	1/2	1/2	0.2021
500	1/2	1/2	0.1674
100	1	1/2	0.2885
200	1	1/2	0.2317
500	1	1/2	0.1583

Description: False discovery rates under maximum likelihood estimation of VPIN given the threshold $c = 0.99$.

Table 11: MLE OF PIN: FALSE DISCOVERY RATES ($c = 0.99$)

β	σ	η	False discovery rate
100	1/4	1/10	0.5193
200	1/4	1/10	0.4987
500	1/4	1/10	0.4430
100	1/2	1/10	0.5364
200	1/2	1/10	0.5417
500	1/2	1/10	0.4979
100	1	1/10	0.5246
200	1	1/10	0.4400
500	1	1/10	0.4459
100	1/4	1/4	0.5401
200	1/4	1/4	0.5113
500	1/4	1/4	0.4693
100	1/2	1/4	0.5182
200	1/2	1/4	0.5250
500	1/2	1/4	0.4772
100	1	1/4	0.4590
200	1	1/4	0.4867
500	1	1/4	0.4268
100	1/4	1/2	0.5579
200	1/4	1/2	0.5263
500	1/4	1/2	0.4934
100	1/2	1/2	0.5409
200	1/2	1/2	0.5167
500	1/2	1/2	0.4606
100	1	1/2	0.4508
200	1	1/2	0.4933
500	1	1/2	0.3822

Description: False discovery rates under maximum likelihood estimation of PIN given the threshold $c = 0.99$.

Table 12: SUMMARY OF SPY TICK DATA: TRADES

Year	#Trades	Average(#Trades)	Min(#Trades)	Max(#Trades)	Shares per Trade
2007	4.98×10^7	1.98×10^5	4.91×10^4	7.45×10^5	679
2008	1.43×10^8	5.64×10^5	1.59×10^5	2.25×10^6	508
2009	1.40×10^8	5.55×10^5	8.25×10^4	1.24×10^6	399
2010	1.25×10^8	4.98×10^5	1.26×10^5	1.50×10^6	365
2011	1.28×10^8	5.09×10^5	1.69×10^5	1.60×10^6	364
2012	7.67×10^7	3.07×10^5	1.00×10^5	5.82×10^5	397
2013	6.79×10^7	2.69×10^5	1.04×10^5	7.37×10^5	381
2014	7.89×10^7	3.13×10^5	1.21×10^5	1.24×10^6	274
2015	9.57×10^7	3.80×10^5	1.21×10^5	2.28×10^6	253

Description: Summary statistics of SPY trade data across trading years. The first column reports total trades/year. Columns 2–4 report, respectively, the average, minimum, and maximum of the daily traded volume across years. The last column shows the shares per trade (average).

Table 13: SUMMARY OF SPY TICK DATA: VOLUME

Year	Total Volume	Average(Volume)	Min(Volume)	Max(Volume)
2007	3.34×10^{10}	1.33×10^8	2.95×10^7	4.86×10^8
2008	6.60×10^{10}	2.61×10^8	5.44×10^7	7.68×10^8
2009	5.44×10^{10}	2.16×10^8	3.46×10^7	4.73×10^8
2010	4.57×10^{10}	1.81×10^8	4.64×10^7	5.65×10^8
2011	4.66×10^{10}	1.85×10^8	5.71×10^7	6.00×10^8
2012	3.04×10^{10}	1.22×10^8	4.78×10^7	2.36×10^8
2013	2.58×10^{10}	1.02×10^8	4.17×10^7	2.78×10^8
2014	2.17×10^{10}	8.60×10^7	3.30×10^7	3.09×10^8
2015	2.39×10^{10}	9.47×10^7	2.81×10^7	3.81×10^8

Description: Summary statistics of SPY trade data across trading years. The first column reports the total traded volume across years. The last three columns report, respectively, the average, minimum, and maximum of the daily traded volume across years.

References

- Abad, D., M. Massot, and R. Pascual (2018). Evaluating VPIN as a trigger for single-stock circuit breakers. *Journal of Banking & Finance* 86(C), 21–36.
- Abad-Diaz, D. and J. Yagüe (2012, 07). From pin to vpin: An introduction to order flow toxicity. *The Spanish Review of Financial Economics* 10, 74–83.
- Andersen, T. G. and O. Bondarenko (2014a). Reflecting on the vpin dispute. *Journal of Financial Markets* 17, 53–64.
- Andersen, T. G. and O. Bondarenko (2014b). Vpin and the flash crash. *Journal of Financial Markets* 17, 1–46.
- Andersen, T. G. and O. Bondarenko (2015). Assessing measures of order flow toxicity and early warning signals for market turbulence*. *Review of Finance* 19(1), 1–54.
- Easley, D., R. Engle, M. O’Hara, and L. Wu (2008). Time-varying arrival rates of informed and uninformed trades. *Journal of Financial Econometrics* 6(2), 171–207.
- Easley, D., S. Hvidkjær, and M. O’Hara (2002). Is information risk a determinant of asset returns? *Journal of Finance* 57(5), 2185–2221.
- Easley, D., S. Hvidkjær, and M. O’Hara (2010). Factoring information into returns. *Journal of Financial and Quantitative Analysis* 45(2), 293–309.
- Easley, D., N. M. Kiefer, M. O’Hara, and J. B. Paperman (1996). Liquidity, information, and infrequently traded stocks. *The Journal of Finance* 51(4), 1405–1436.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2012, 02). Flow toxicity and liquidity in a high frequency world. *Review of Financial Studies* 25(5), 1457–1493.
- Easley, D., M. Lopez de Prado, and M. O’Hara (2013, 03). Bulk classification of trading activity. Working Paper.
- Easley, D., M. M. López de Prado, and M. O’Hara (2010, 11). The microstructure of the ”flash crash”: Flow toxicity, liquidity crashes and the probability of informed trading. *The Journal of Portfolio Management* 37, 118–128.
- Easley, D., M. M. López de Prado, and M. O’Hara (2011). The exchange of flow toxicity. *The Journal of Trading* 6(2), 8–13.
- Easley, D., M. M. López de Prado, and M. O’Hara (2014). Vpin and the flash crash: A rejoinder. *Journal of Financial Markets* 17(C), 47–52.
- Easley, D., M. M. López de Prado, and M. O’Hara (2016, 01). Discerning information from trade data. *Journal of Financial Economics* 120, 269–286.

-
- Eldén, L. and L. Wittmeyer-Koch (1990). *Numerical Analysis: An Introduction*. San Diego, CA, USA: Academic Press Professional, Inc.
- Holden, C. W. and S. Jacobsen (2013). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *The Journal of Finance* 69(4), 1747–1785.
- Lee, C. and M. Ready (1991). Inferring trade direction from intraday data. *Journal of Finance* 46(2), 733–46.
- Leth, L. R. (2019). Delta hedging and the vpin. Working paper.
- Lin, H.-W. W. and W.-C. Ke (2009, 10). A computing bias in estimating the probability of informed trading. *Journal of Financial Markets* 14, 625–640.
- Lin, H.-W. W. and W.-C. Ke (2017). An improved version of the volume-synchronized probability of informed trading. *Critical Finance Review* 6(2), 357–376.
- Mullen, K. M. (2014). Continuous global optimization in r. *Journal of Statistical Software* 60(6).
- Pöppe, T., S. Moos, and D. Schiereck (2016). The sensitivity of vpin to the choice of trade classification algorithm. *Journal of Banking & Finance* 73(C), 165–181.
- Song, J. H., K. Wu, and H. Simon (2014). Parameter analysis of the vpin (volume synchronized probability of informed trading) metric. In C. Zopounidis (Ed.), *Quantitative Financial Risk Management: Theory and Practice*. Wiley.
- Wu, K., E. W. Bethel, M. Gu, D. Leinweber, and O. Ruebel (2013a, 06). A big data approach to analyzing market volatility. *Algorithmic Finance* 2, 241–267.
- Wu, K., E. W. Bethel, M. Gu, D. Leinweber, and O. Ruebel (2013b, 01). Testing vpin on big data response to 'reflecting on the vpin dispute'. *SSRN Electronic Journal*.