

This manuscript has been accepted for publication in *Journal of Financial Economics*. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all disclaimers that apply to the journal apply to this manuscript.

Shorting at close range: a tale of two types *

Carole Comerton-Forde
Department of Finance, University of Melbourne

Charles M. Jones[#]
Columbia Business School

Tālis J. Putniņš
University of Technology Sydney and Stockholm School of Economics in Riga

Abstract

We examine returns, order flow, and market conditions in the minutes before, during, and after NYSE and Nasdaq short sales. We find two distinct types of short sales: those that provide liquidity, and those that demand it. Liquidity-supplying shorts are strongly contrarian at intraday horizons. They trade when spreads are unusually wide, facing greater adverse selection. Liquidity-demanding shorts trade when spreads are narrow and tend to follow short-term price declines. These results support a competitive rational expectations model where both market-makers and informed traders short, indicating that these two shorting types are integral to both price discovery and liquidity provision.

JEL classification: G14, G19

Keywords: short selling, information content, market quality, high-frequency trading

* We thank the Australian Research Council (ARC Linkage Project LP0455536) and the Paul Woolley Centre for Capital Market Dysfunctionality for funding. We thank the Securities Industry Research Centre of Asia-Pacific and Reuters for providing access to data used in this study. We are grateful for the comments of an anonymous referee, Henk Berkman, Veljko Fotak, Larry Glosten, Bruce Grundy, Larry Harris, Terry Hendershott, Albert Menkveld, Erik Theissen, Andriy Shkilko, Josef Zechner, and seminar participants at the Paul Woolley Centre for Capital Market Dysfunctionality 2009 annual conference, the European Finance Association 2010 annual conference, the Financial Management Association 2010 European conference, the Financial Management Association 2010 annual meeting, the IX Madrid Finance Workshop, the Institut Louis Bachelier market microstructure conference, the University of Technology Sydney market microstructure conference, the American Finance Association 2012 annual meeting, Baruch, the Hanken School of Economics, University of Warwick, Vienna Graduate School of Finance, the Stockholm School of Economics / Stockholm Institute for Financial Research, Hong Kong University and the University of Western Australia. This paper succeeds one by Comerton-Forde and Putniņš entitled “Short selling: information or manipulation?”

[#] Corresponding author at: Columbia Business School, 801 Uris Hall, 3022 Broadway, New York, NY, 10027; Tel.: +1 212 854 4109; Fax: +1 212 662 8474.

E-mail addresses: cj88@columbia.edu (C.M. Jones), carole.comerton-forde@unimelb.edu.au (C. Comerton-Forde), and talis.putnins@uts.edu.au (T.J. Putniņš).

1. Introduction

Short sellers are at the center of an intense debate among regulators, politicians, the media and academics. The key questions in this debate concern the role of short sellers in stock markets and elsewhere. In particular, do short sellers improve market efficiency? Do short sellers destabilize stock prices in any way? Do short sellers improve or worsen market quality?

On one side of the debate, industry observers, issuers, and some of the popular media argue that short sellers employ abusive trading strategies, damage investor confidence and market quality, and amplify price declines.¹ Company directors, shareholders and the media have even gone so far as to blame short sellers for the sharp price declines or collapses of companies such as Bear Stearns, Halifax Bank of Scotland, Lehman Brothers and Merrill Lynch.² Regulators have responded with several new rules to limit or discourage some short sales. For example, within one week of Lehman's collapse in September 2008, the U.S. Securities and Exchange Commission instituted a temporary emergency ban on short sales in all financial stocks, stating that "unbridled short selling is contributing to the recent, sudden price declines in the securities of financial institutions unrelated to true price valuation".³

On the other hand, most academic research argues that short sellers are relatively informed, improve market efficiency, and generally stabilize share prices by identifying and then leaning against overvalued stocks.⁴ Much of this academic literature analyzes the association between various measures of shorting activity and stock returns, typically at horizons of days, weeks, or months.

We contribute to this debate by taking a fine-toothed comb and analyzing the behavior of stock returns, order flow, and market conditions in the minutes before, during, and after short sales. We use trade-level information on all short sales executed on the

¹ For examples see "There's a Better Way to Prevent 'Bear Raids'" by R. Pozen and Y. Bar-Yam, *The Wall Street Journal*, 18 November 2008, "Anatomy of the Morgan Stanley Panic" by S. Pulliam et al., *The Wall Street Journal*, 24 November 2008.

² For example, Richard Fuld Jr., the former CEO of Lehman Brothers, during hearings on the bankruptcy filing by Lehman Brothers and bailout of AIG alleged that a host of factors including naked short selling attacks followed by false rumors contributed to both the collapse of Bear Stearns and Lehman Brothers (<http://oversight.house.gov/documents/20081006125839.pdf>).

³ SEC press release 2008-211 (<http://www.sec.gov/news/press/2008/2008-211.htm>).

⁴ See, for example, Dechow, Hutton, Meulbroek and Sloan (2001), Abreu and Brunnermeier (2002), Alexander and Peterson (2008), Boehmer, Jones, and Zhang. (2008), Boehmer and Wu (2013) and Diether, Lee and Werner (2009).

NYSE and Nasdaq during the first eight months of 2008 for a sample of 350 stocks. Our analysis highlights the fact that there are two very distinct types of short sales: those that provide liquidity, and those that demand it. The heterogeneity we find suggests that researchers, regulators, and market participants should not view short sellers as monolithic.

The high level of granularity in our analysis is important for several reasons. First, it allows us to separate liquidity-demanding short sellers from liquidity-supplying short sellers. These two different types of short sales are likely to originate from different types of traders and therefore may have different effects on liquidity and price discovery. Liquidity-demanding short sales are likely to arise from informed traders, whereas liquidity-supplying short sales can arise when a market-maker with zero or negative inventory provides liquidity to an incoming buy order. Aggregating these two categories could mask important features of their contribution to liquidity. For example, if shorts contribute liquidity on average, but tend to take liquidity during times of market stress, it is not clear that they benefit markets. Our analysis allows us to examine this possibility and thereby provide a more complete characterization of the effects of short sellers.

Second, previous studies clearly document that shorts tend to be contrarian, but it is less clear *why* shorts are contrarian. Partitioning these two types of shorts can shed light on this question. In some sense, liquidity-supplying short sellers are inherently contrarian, because they sell when there is buying pressure. However, if short sellers are able to detect and actively trade against mispricings, we should also see contrarian behavior by liquidity-demanding shorts.

Third, if short sellers are conducting ‘bear raids,’ they are likely to use aggressive liquidity-demanding short sales. By comparing price behavior around liquidity-demanding short sales vs. long sales, we provide evidence on whether the data contain patterns consistent with short seller bear raids. Again, disaggregating short sale types is important, as liquidity-supplying short sales could mask the patterns around liquidity-demanding short sales.

Finally, given that many short sellers employ high-frequency trading strategies, understanding the behavior of short sellers in today's markets, and how their trades impact the market, requires observations at intraday horizons.⁵

We find that seller-initiated short sales that demand liquidity are quite distinct from passive, liquidity-supplying, buyer-initiated shorting. Shorts that supply liquidity do so when spreads are unusually wide, which is when liquidity is most highly valued by market participants. These short sellers are also strongly contrarian, stepping in to initiate or increase a short position after fairly sharp share price rises over the past hour or so. In contrast, shorts that demand liquidity are not contrarian on average. Especially in smaller stocks, these aggressive short sellers tend to be momentum traders, as their shorting activity tends to follow a price decline over the previous 24 hours. In addition, we find that aggressive short sales have significantly bigger price impacts at short horizons. Our results are consistent with the predictions of a competitive rational expectations model in which both market-makers and informed traders use short sales. In addition, the finding that only liquidity-supplying shorts are contrarian on average is consistent with the hypothesis that the dynamics of liquidity provider inventories are responsible for a large part of the contrarian behavior of short sellers.

We find that at an intraday level, liquidity-supplying short sales are clearly a stabilizing force in stock markets. The evidence strongly indicates that they help to narrow spreads, limit price spikes, and provide liquidity at important times such as when it is scarce. These results provide an insight into why restrictions imposed on short selling harm market quality (Boehmer, Jones, and Zhang (2013) and Beber and Pagano (2013)). We also find that aggressive order flow from short sellers is not very different from aggressive order flow that originates from long sellers. Based on our close-in examination of the data, the evidence provides no reason for regulators to target short sellers over other sellers.

⁵ For example, Jones (2012) finds that 'in-and-out shorting' (short selling and covering the position before the end of the day as in the first scenario) represented about 5% of total daily volume (and a much bigger, but unknown, fraction of short selling activity) in the early 1930s. It is reasonable to expect this fraction to be higher in today's markets given the increases in automation, algorithmic trading, statistical arbitrage and turnover. The argument that short sellers employ rapid trading strategies is also consistent with the finding of Diether, Lee and Werner (2009) that short sales represent on average 23.9% of NYSE and 31.3% of Nasdaq volume.

2. Related literature

Much of the literature on short selling is devoted to examining the informativeness of short sales. The large increase in the amount of short selling in recent years is likely due to the growth in algorithmic and in particular high-frequency trading (HFT), which uses short sales in automated market making and high-frequency arbitrage strategies. Therefore, the emerging literature on HFT is relevant to understanding the use of short sales in recent times.

2.1. The informativeness of short sales and granularity of data

There are strong theoretical reasons to expect some types of short sellers to be relatively informed. For example, market participants do not short sell for liquidity reasons because they do not have use of the sale proceeds. Theory predicts that prices diverge from fundamental values when short selling is constrained (e.g., Miller, 1977; Duffie, Garleanu, and Pedersen, 2002; Hong, Scheinkman, and Xiong, 2006), and this prediction is supported by empirical evidence (e.g., Danielsen and Sorescu, 2001; Jones and Lamont, 2002; Cohen, Diether, and Malloy, 2007; and Saffi and Sigurdsson, 2011). Beber and Pagano (2013) show that bans on short selling, imposed in markets around the world during the financial crisis, slowed down price discovery. Evidence on the relation between short selling and future returns is not uniform, but is increasingly moving towards the consensus that, on average, short sellers predict future returns. This trend is particularly true in the more recent work that uses data on short sale flows.

Empirical studies that use monthly short interest data (total outstanding short positions for each stock, measured in shares, at a particular point in time each month) find mixed results on the informativeness of short sales. For example, Brent, Morse, and Stice (1990) and Lamont and Stein (2004) find that short interest is positively related to past returns but does not predict future returns in cross-section or time-series. Asquith, Pathak, and Ritter (2005) find return predictability only in the smallest stocks and report that the effect is stronger in stocks with low institutional ownership. In contrast, Desai, Ramesh, Thiagarajan, and Balachandran (2002) find that high short interest predicts negative returns in Nasdaq stocks, and Dechow, Hutton, Meulbroek, and Sloan (2001) find that short interest is highest in firms that are overpriced according to fundamental ratios.

Boehmer, Huszár, and Jordan (2010) find that low short interest predicts high future returns, but the relationship between high short interest and future returns is much weaker.

Some recent studies use more granular data. Boehmer, Jones, and Zhang (2008) and Diether, Lee, and Werner (2009) construct portfolios of stocks with high and low daily short-sale flows. Heavily shorted stocks underperform lightly shorted stocks over the following month. Christophe, Ferri, and Angel (2004) and Boehmer, Jones, and Zhang (2012) find that daily flows of short sales are concentrated prior to disappointing earnings announcements, which suggests short sellers have access to private information. Engelberg, Reed, and Ringgenberg (2012) and Fox, Glosten, and Tetlock (2010) find that short sellers trade around negative news releases. All of these papers aggregate shorting flow to the daily level; none of them focus on intraday returns, spreads, or trading behavior. Similarly, none of these papers distinguish between liquidity-supplying and liquidity-demanding short sales.

Only a few papers employ intraday data to study short sales. Boehmer and Wu (2013) find that high-frequency informational efficiency of prices improves with greater daily shorting flow. Aromi and Caglio (2008) show that extreme negative returns during September 2008 are not the result of short selling activity. Boehmer, Jones, and Zhang (2013) use intraday data to measure market quality and the aggressiveness of short sellers around the 2008 temporary ban on shorting U.S. financial stocks. Jain, Jain, and McNish (2012) use intraday data to study the effects of the ‘alternative uptick rule’ on price stability and liquidity and find that this rule did not improve the return dynamics of target stocks. Chakrabarty, Moulton, and Shkilko (2012) use data from INET (an ECN that was integrated with Nasdaq in 2006) for a sample of 200 Nasdaq-listed stocks during two months in 2005. While the focus of their study is the performance of trade classification algorithms for classifying the initiator of short sales, they document that trade aggressiveness of short sellers is similar to that of long sellers and that short sellers tend to supply liquidity when spreads are wider, consistent with our findings. The present paper examines more broadly the market conditions at different intraday horizons before and after liquidity-demanding and liquidity-supplying shorts, including returns, and uses a more encompassing sample of NYSE and Nasdaq trades. We also provide an interpretation for the intraday patterns that is consistent with a rational expectations model that includes long and short sellers.

2.2. *Short selling and high-frequency trading*

Many HFT strategies, such as market making and high-frequency arbitrage across markets or securities, use short sales. Given that HFT may account for the majority of volume in today's US equity markets,⁶ a considerable share of short sales are executed by high-frequency algorithms. Thus, some of the more granular studies of HFT are relevant to the present paper.

Menkveld (2013) examines the trading of a high-frequency market-maker in 12 Dutch stocks in 2007. A substantial majority of the market maker's trades (80%) result from passive orders. The market maker's inventory position in any given stock typically flips between long and short tens of times during a trading day. Therefore, passive shorting is an integral part of the high-frequency market-maker's business model.

Brogaard, Hendershott, and Riordan (2014) study a dataset of Nasdaq trades during 2008 and 2009, in which the trades of 26 proprietary trading firms that are classified as HFTs are flagged. They decompose price changes into permanent and temporary components using a state-space model and find that liquidity-demanding trades by HFT firms tend to be in the direction of permanent price changes and in the opposite direction to transitory pricing errors. HFT liquidity-supplying trades are adversely selected: they tend to occur in the opposite direction to permanent price changes. In related work on algorithmic trading, Hendershott and Riordan (2013) show that algorithmic traders supply liquidity when it is expensive and consume liquidity when it is cheap.

3. Short sales in a sequential trade model

To guide our empirical examination, we first introduce a sequential trade model that distinguishes between long sales and short sales. In the model, shorting can occur for two reasons. First, an informed trader may want to short sell if she receives a negative signal about the value of the asset. Second, an uninformed market-maker may sell short as a consequence of supplying liquidity.

⁶ The Tabb Group, a market structure consultancy firm, estimates that HFT accounted for 52% of trading activity in US equity markets in 2008. Tabb Group also estimates that the level of HFT activity in US equity markets peaked at 61% in 2009.

The basic structure of the model is based on Glosten and Milgrom (1985), and is illustrated in Fig. 1. Nature chooses the fundamental value of the stock v to be the high value v^H with probability p_0 , and the low value v^L with probability $(1 - p_0)$, where $v^H > v^L$. The probabilities and possible fundamental values are common knowledge at time $t = 0$. Trading involves the sequential arrival of T traders at times $t = 1, 2, 3, \dots, T$. In each trading round t , a competitive, risk-neutral representative liquidity provider posts bid and ask quotes for a unit trade. Due to competition, these quotes must be equal to the expected fundamental value conditional on the arrival of a given market order. That is, denoting the direction of the trade as $D_t = -1$ for a seller-initiated trade and $D_t = +1$ for a buyer-initiated trade, we have $Bid_t = E[v | \text{history}_t, D_t = -1]$ and $Ask_t = E[v | \text{history}_t, D_t = +1]$. In other words, the liquidity provider sets ‘no regret’ quotes with zero expected profits. The liquidity provider cannot observe whether a seller-initiated trade is from a short-seller. After every trade the liquidity provider Bayesian updates her beliefs about the probability that the fundamental value is high, p_t . For simplicity, the liquidity provider starts with an initial inventory of zero, $I_0^{LP} = 0$, and her inventory at any point in time is equal to the net number of sells up to that point, $I_t^{LP} = -\sum_{\tau=1}^t D_\tau$.⁷

The probability of informed trading, π , is drawn from a uniform distribution ranging from 0 to π^{MAX} and remains constant throughout the T trading rounds. The value of π is common knowledge. The informed trader’s starting inventory is chosen from a uniform distribution ranging from $\underline{I^{INF}}$ to $\overline{I^{INF}}$. In each trading round, with probability π an informed trader arrives at the market and submits a unit size market order, and with probability $(1 - \pi)$ an uninformed trader arrives and submits a unit size market order.⁸ Uninformed traders submit buy market orders (direction, $D_t = +1$) and sell market orders

⁷ The assumption that the liquidity provider starts trading with zero inventory is consistent with the evidence that both formal designated market makers and informal liquidity providers avoid the inventory risk associated with carrying a positive or negative inventory overnight, and therefore tend to start and end each trading day with an inventory close to zero (see, for example, Menkveld, 2013). However, this assumption is not crucial, and our main results can be obtained with a stochastic initial inventory.

⁸ In our model, a trade occurs in every trading round. This follows from the simplifying assumption that informed and uninformed trading demands are inelastic and thus are not sensitive to the terms of trade (e.g., the spread). Although Glosten and Milgrom (1985) do not make this assumption, it is common in subsequent implementations of Glosten-Milgrom type models. An alternative approach is to assume that traders have private valuations that are drawn from some distribution (e.g., Bessembinder, Hao and Lemmon, 2011). Under this approach, ‘no-trade’ outcomes (trading rounds in which the arriving trader chooses not to trade) occur when a trader’s valuation lies within the spread.

(direction, $D_t = -1$) with equal probability, and never short sell because they trade for exogenous liquidity needs, not information. The informed trader knows the fundamental value and therefore buys when $v = v^H$ and sells when $v = v^L$. If the informed trader has positive inventory at the time she sells ($I_t^{INF} > 0$) the trade is a seller-initiated long sale (SIL), but if she does not hold the stock ($I_t^{INF} \leq 0$), she short sells, leading to a seller-initiated short sale (SIS). Similarly, if the liquidity provider has positive inventory when she receives a market buy order ($I_t^{LP} > 0$) she sells from her inventory resulting in a buyer-initiated long sale (BIL), whereas if the liquidity provider has zero inventory or a short position when she receives a market buy order ($I_t^{LP} \leq 0$) she short sells resulting in a buyer-initiated short sale (BIS).

< FIGURE 1 HERE >

Diamond and Verrecchia (1987) also examine short sales in a Glosten and Milgrom model with sequential trade. Their focus is on shorting constraints and the revelation of private information, and they assume that short sales only arise from informed traders with negative information. In our model, by contrast, market-makers also short sell. Thus, in Diamond and Verrecchia, short sales always demand liquidity, while some shorts provide liquidity in our model.

Perhaps the most important parameter in the model is the probability of informed trading, π . With more informed trading, spreads are relatively wider in early trading rounds, order flow is more one-sided, and there is faster convergence towards the fundamental value. Adverse selection clearly changes over the course of the T trading rounds, but adverse selection also depends on π , and variation in π turns out to provide interesting comparative statics on the relationship between bid-ask spreads and the prevalence of shorting.

Four key results from the sequential trade model provide empirical predictions of the market conditions around liquidity-supplying and liquidity-demanding short sales. In what follows, we describe these results and provide the underlying intuition.

Result 1: Liquidity-supplying short sales are contrarian, in an absolute sense and relative to liquidity-supplying long sales.

Proof. See the appendix.

< FIGURE 2 HERE >

To provide an example of this result, Fig. 2 plots the average cumulative midquote returns before and after each of the four trade types for the following parameter values: (i) $T = 35$ trading periods; (ii) the fundamental value is either $v^L = \$50$ or $v^H = \$51$ with equal probability ($p_0 = 0.5$); and (iii) the probability of informed trade (π) is drawn from a uniform distribution ranging from 0 to $\pi^{MAX} = 0.5$, and the informed trader's initial inventory is uniformly distributed over the integers from -5 to $+5$. The figure is based on a simulation of 500,000 trading sequences. The plot shows a distinct pattern: liquidity-supplying short sales (the light-weight solid red line) tend to occur after a period of price increases. Midquotes tend to rise before liquidity-supplying short sales, while midquotes tend to fall before liquidity-supplying long sales (the light-weight dashed blue line).

The intuition is relatively straightforward. Buyer-initiated trades tend to move prices up, and seller-initiated trades tend to move prices down, because the trade direction provides a noisy signal about the fundamental value. At the same time, buyer-initiated trades deplete any positive inventory that the liquidity provider might have, making it more likely that she will short sell to an incoming buy order. Therefore, liquidity-supplying short sales are more likely to occur following increases in the price. The pattern of increasing prices preceding liquidity-supplying short sales extends back in time beyond the previous trade because the more buys (relative to sells) received by the liquidity provider in the past, the larger the past price increase and the more likely the liquidity provider does not have positive inventory. When the fundamental value is high ($v = v^H$) we tend to see a sequence of buys, which drives prices up and causes the liquidity provider to go short, resulting in liquidity-supplying short sales following price rises. In contrast, liquidity-supplying long sales occur when the liquidity provider has built up a positive inventory. This occurs when the liquidity supplier has received more sell orders than buy orders since the start of trading, which means that the price has declined since the start of trading. This is why we observe the opposite price pattern (decreasing prices) preceding liquidity-supplying long sales.

Result 2: Liquidity-demanding short sales tend to follow price declines, in an absolute sense and relative to liquidity-demanding long sales.

Proof. See the appendix.

Fig. 2 also illustrates Result 2 for the set of parameter values outlined above. Comparing liquidity-demanding short sales (the heavy-weight solid red line) and liquidity-demanding long sales (the heavy-weight dashed blue line) we see a clear difference. Liquidity-demanding short sales tend to occur following a period of falling prices, whereas prices preceding liquidity-demanding long sales are approximately flat on average. Liquidity-demanding short sellers may appear to be short-term momentum traders, but this is just due to the fact that these informed investors trade multiple times on their information. They are not selling just because the price has declined.

Furthermore, the price patterns *following* liquidity-demanding short and long sales reveal a second regularity: prices tend to continue to decline following liquidity-demanding short sales, whereas after the immediate impact of a liquidity-demanding long sale there is no clear subsequent trend. Thus, liquidity-demanding short sales also tend to predict future price patterns.

Again, the intuition for these patterns is relatively straightforward. The patterns are driven by the fact that liquidity-demanding short sales only arise from informed traders when they have bad news (i.e., when $v = v^L$). In environments where informed traders have bad news, prices decline on average as they impound the informed traders' information and converge to the low fundamental value. It is precisely in such environments that liquidity-demanding short sales will occur, resulting in patterns of declining prices before and after liquidity-demanding short sales (SIS trades).

Result 3: Short sellers provide liquidity at wider spreads than long sellers.

Although it is relatively straightforward to obtain an analytic expression for the spread at a point in time as a function of the underlying parameters, the expression does not readily allow computation of expected spreads conditional ex post on the different

trade types of interest and integrated over time. We therefore obtain Result 3 using numerical techniques, under a range of reasonable parameter values. For example, consider the parameter values introduced earlier, with $T = 35$, v^L and v^H differing by \$1 and equally likely, $\pi \sim U[0, 0.5]$, and the informed trader's initial inventory $\sim U[-5, +5]$. We simulate 500,000 trading sequences using these parameters and record the quotes and inventory positions at every point in time as well as every trade and its type. The average effective spread when short sellers provide liquidity (BIS trades) is 19.2 basis points, whereas when long sellers provide liquidity it is 17.3 basis points (BIL trades).

The intuition is as follows. Consider the likely state of the liquidity provider's inventory when she is providing liquidity to incoming buys in a relatively high adverse selection environment. In this case, there is a high proportion of informed traders and spreads are wide. Order flow is relatively one-sided, and prices tend to converge rapidly towards the fundamental value. If the fundamental value is high ($v = v^H$) a large number of marketable buy orders arrive, which causes the liquidity provider to be short most of the time, and therefore many BIS trades occur. If the fundamental value is low ($v = v^L$) a large number of marketable sell orders arrive and very few buyer-initiated trades. Therefore, when the probability of informed trade is high, most of the liquidity provision to incoming buys occurs with short sales. In contrast, if the probability of informed trade is low, spreads are relatively narrow and, because order flow is more balanced, the liquidity provider's inventory tends to fluctuate around zero. Consequently, in such low adverse selection environments, the proportions of long and short selling by the liquidity provider are approximately equal. Because a disproportionately large number of liquidity-providing short sales (BIS trades) occur in high adverse selection situations when spreads are wide, the average spread at the time of a BIS trade tends to be larger than at the time of a liquidity-providing long sale (BIL trade). Ultimately, short sales (compared to long sales) tend to provide liquidity when adverse selection risk is higher.

Result 4: Short sellers take liquidity at narrower spreads than long sellers.

As before, due to the difficulty in obtaining closed form expressions for the expected spreads conditional on the trade types, we obtain Result 4 using numerical techniques. For example, using the previous parameter values and a simulation of

500,000 trading sequences, we find that the average effective spread when short sellers take liquidity (SIS trades) is 17.3 bps, whereas when long sellers take liquidity it is 18.9 bps (SIL trades).

The intuition behind this result is as follows. Liquidity-demanding sales originate from informed traders when they have bad news (i.e., when $v = v^L$). At the start of trading this informed trader can either be long the stock, short the stock or have no position in the stock. If the informed trader is initially short or has no position, she will short sell the stock throughout the course of trading. If the informed trader is initially long the stock, she will initially execute long sales, deplete her position down to zero and then turn to short selling. Consequently, liquidity-demanding shorts are more likely to occur after a period of informed selling. This result does not require that every time an informed trader starts with a long position she eventually turns to short selling; there may be situations in which the informed trader's inventory remains positive. As long as there are some occasions on which an informed trader starts long and ends short, it will be true that liquidity-demanding shorts are relatively more likely to occur after a longer period of informed trading (i.e., in later trading rounds). Importantly, spreads tend to narrow over the course of trading because the informed trader's information is gradually revealed through trade and information asymmetry decreases. As a result, liquidity-demanding shorts tend to occur when spreads are narrower.

Next we turn to the data to see if these predictions of the model are borne out empirically. Results 1 and 2 can be confirmed or refuted by looking at returns before, during, and after different kinds of sales. Results 3 and 4 require an examination of the spreads in effect for short vs. long sales that are buyer- vs. seller-initiated.

4. Data and summary statistics

In the U.S., the SEC requires a broker-dealer to mark a short sale when the order is sent to a trading venue. Short sale orders are not publicly disclosed as such in real time, but for audit and compliance purposes exchanges are required to track all short sale orders. Thus, we are able to obtain data from the respective exchanges on all short sales executed at the NYSE or Nasdaq from January 1, 2008 through August 31, 2008 for a sample of 350 stocks. Sample stocks are selected by sorting all NYSE-listed common stocks (Center for Research in Securities Prices (CRSP) share codes 10 and 11) into deciles based on

their market capitalization as of 31 December 2007 and then randomly sampling 35 stocks from each size decile.

We use data from 2008 for three reasons. First, this is a period when short selling received considerable regulatory attention. Regulators in the US and other jurisdictions imposed bans on short selling shortly after our sample period. Therefore, this time period seems to be the most likely to exhibit negative or harmful short selling activity. Second, the NYSE and Nasdaq still had substantial market share during this period, therefore short selling activity on these venues will be most representative of the market as a whole. Subsequent periods with greater fragmentation would require additional data from other trading venues such as BATS and Direct Edge. Finally, the market structure in place in 2008 remains largely unchanged today.

We couple the short sales data with all trades and quotes during the same period, obtained from the Thomson Reuters Tick History database maintained by the Securities Industry Research Centre of Asia-Pacific.^{9,10} We add company level variables, including book-to-market ratio, market capitalization, number of shares on issue, and short interest from the CRSP database, Compustat and Thomson Reuters Datastream.

< TABLE 1 HERE >

Table 1 reports summary statistics about the number, volume and size of short sales relative to long sales on the two exchanges, as well as for the pooled sample. There

⁹ Similar to the TAQ database, the Thomson Reuters database is a record of the consolidated tape. We retain all trades executed at the NYSE or Nasdaq, as well as the national best bid and offer (NBBO) prevailing at the time of the trade. We apply a number of filters to screen for errors and outliers in the raw data, similar to those commonly applied to TAQ data. For trades: (i) the trade must be time-stamped between 9:30:00 and 16:00:00, (ii) price and volume must be strictly positive, (iii) the date and time must be non-missing, and (iii) error/correction records are removed. Trades at prices that are at least 20% away from the previous trade price, next trade price or midquote are manually inspected. When it is clear that a trade record contains an error, the error is corrected if possible or the record is removed. When it is not clear that a trade record contains an error, the trade is left in the sample and is subject to winsorization. For NBBO quotes: (i) the quote must be recorded between 9:30:00 and 16:00:00 inclusive, (ii) quotes are deleted if the ask or bid is zero or negative or if the ask is less than the bid, and (iii) the date and time must be non-missing. Quotes that are at least 20% away from either the previous quote or the midquote are manually inspected and any clear errors are corrected or removed.

¹⁰ We match short sales to trades by exchange, symbol, date, price, size and time. We first look for exact matches, down to the second. To correct for small differences in the timestamps recorded in each of the databases, we then look for matches where time stamps differ by one second, and finally we look for matches with two-second timestamp differences. For our sample, 98.2% of the short sales match to trades and we discard the remaining 1.8%.

are more than 103 million short sales in these 350 stocks over the eight-month sample period, and almost 170 million long sales. About 60% of the short sales in these NYSE-listed stocks take place on the NYSE, with the remaining 40% executing on Nasdaq.

Short sales constitute 40.2% and 39.2% of the total dollar volume sold on the NYSE and Nasdaq, respectively. Shorting is a similar fraction of the overall number of reported trades on those two venues. For about a quarter of the stocks in our sample, shorting actually accounts for a majority of all selling. The proportion of short selling has been increasing over time. Boehmer, Jones, and Zhang (2008) find that short selling represents about 13% of all selling in their 2001-2005 sample, and Diether, Lee, and Werner (2009) report that short selling represents an average of 24% of shares sold on the NYSE in 2005. Note that our sample ends before the SEC ban on short selling in financial stocks and is thus unaffected by that regulatory change.

On the NYSE, short sales and long sales are distributed quite similarly. The mean size of a short sale is \$6,581, while the mean size of a long sale is \$6,754. The quartile points are also virtually identical. The 25th percentile is \$1,710 for short sales vs. \$1,721 for long sales, the medians are \$3,046 vs. \$3,065, and the 75th percentiles are \$4,928 for short sales vs. \$4,953 for long sales. The similarity in distribution makes sense if short sellers and long sellers use similar order-splitting algorithms, as both types of sellers could be expected to have similar incentives to blend in and not telegraph their order flow. Short sales on Nasdaq have a very similar distribution to the short sales and long sales on the NYSE, with a mean size of \$6,551. Long sales on Nasdaq tend to be somewhat smaller, with a mean size of \$5,860.

< TABLE 2 HERE >

Table 2 reports how average short selling activity varies with firm characteristics. In Panel A, we sort stock-days into quintiles based on both size and book-to-market ratios. There are only very modest patterns across these two characteristics. Shorting is a bit more prevalent in small stocks (41.2% of dollar volume) than it is in the largest cap quintile (38.0% of dollar volume). On average, short sales account for 36.5% of dollar volume for the lowest book-to-market firm quintile vs. 41.3% of dollar volume for the highest book-to-market firm quintile.

In Panel B, we sort into quintiles based on share price and short interest at the end of the previous month, where short interest is calculated as a fraction of shares outstanding. The pace of shorting activity does not appear to be related to share price. Not surprisingly, the shorting flow measure used here (shorting as a fraction of overall trading volume) is positively correlated with short interest (short positions outstanding at a given point in time), which is a stock measure. But it is worth noting that there is still substantial shorting activity even in the stocks with the smallest short interest. Short sales account for 34.2% of dollar volume for the low short interest quintile, compared to 43.4% of dollar volume for the high short interest quintile.

< TABLE 3 HERE >

Table 3 reports the aggressiveness of executed short and long sell trades relative to the prevailing quotes. Short sales and long sales display similar aggressiveness, indicating nothing unusual about shorting activity relative to other selling. About 40% of sales occur at the best ask quote (and thus are liquidity-supplying limit sell orders), about 35% are at the best bid quote (indicating that they are liquidity-demanding market sell orders), and about one-sixth execute between the best bid and best ask quotes. A small fraction (3.8% to 4.2%) execute at prices above the ask quote, and a similar number (3.7% to 3.9%) execute at prices below the bid quote. Short sales and long sales are virtually identical in terms of aggressiveness, indicating that long sellers and short sellers are similarly patient.

When we partition the sample by market cap, the most notable result is that there are relatively few executions within the quotes for the largest cap quintile. Only about 13% of sales (long or short) take place between the quotes, while the corresponding figure is more than 20% for most of the other size quintiles. This is probably due to the fact that the inside spread in most large-cap stocks is usually equal to the minimum tick of one cent. Finally, when we compare sell orders executed on the NYSE and the Nasdaq, there are only very small differences in the distribution of executions relative to the prevailing quotes, and no discernible cross-exchange differences in the aggressiveness of short sales.

5. Two types of short sales

The heart of the paper is a comparison at close range of the dynamics and characteristics of aggressive (seller-initiated) vs. passive (buyer-initiated) short sales. We assign short sales to one of these two categories based on the transaction price compared to the midpoint between the prevailing bid and ask price at the time of the trade.¹¹ Short sales that take place below the prevailing midpoint are considered seller-initiated, while short sales that take place above the prevailing midpoint are considered buyer-initiated. We use the Lee and Ready (1991) method to assign trades that are at the prevailing midpoint.¹² As a practical matter, during our sample period the quoted spread on most U.S. stocks is the minimum tick of one cent most of the time. In most cases, then, trades are at either at the prevailing bid or the prevailing ask, and short sales at the bid are considered seller-initiated, while short sales at the ask are buyer-initiated. Note that we will often refer to seller-initiated shorts as aggressive or liquidity-demanding, while buyer-initiated shorts are often referred to passive or liquidity-supplying, because in an electronic limit order book market, these trades result from a marketable buy order trading with a standing (non-marketable) limit order to short sell.

For comparison, we use the same methods to partition long sales into buyer-initiated and seller-initiated. Thus we have four different types of trades: buyer-initiated trades involving a long seller (BIL), seller-initiated trades involving a long seller (SIL), buyer-initiated trades involving a short seller (BIS), and seller-initiated trades involving a short seller (SIS).

It is important to emphasize that our partition is based on trade-level data; a particular seller might find her individual trades in more than one of our categories. For example, some high-frequency traders move back and forth rapidly between long sales and short sales. Menkveld (2013) provides evidence on one such high-frequency market-maker, who switches from a long position to a short position many times in a single day. Along the same lines, a large institutional order – sometimes called the parent order – is

¹¹ The prevailing bid and ask are taken from the most recent consolidated quote with a time-stamp prior to that of the transaction.

¹² Chakrabarty, Moulton, and Shkilko (2012) use Nasdaq order book data to show that the Lee-Ready algorithm misclassifies 21% to 31% of short sales. This adds noise to our estimates and reduces statistical power, but Chakrabarty, Moulton, and Shkilko (2012) find that short sale misclassification does not introduce significant bias. Similarly, the presence of hidden orders on the exchange may also add noise, but not bias. Rosenblatt Securities estimates the volume of executions arising from hidden orders on exchanges during 2008 to be approximately 5% of consolidated volume.

likely to be broken up and executed in many small pieces via child orders. Algorithms typically submit (and often cancel and resubmit at different prices) these child orders over time, and typically they use a mix of limit orders and marketable orders. Thus, a single seller would in part supply liquidity and in part demand liquidity. However, impatient or better-informed sellers would still be expected to use more marketable orders.

5.1. Effective spreads, realized spreads, and price impacts

We begin by examining whether market conditions are associated with the likelihood of these four different transaction types. In particular, we focus on the variation in liquidity over time, to see whether long and short sales of various types are more or less likely to take place when spreads are relatively wide. To do this, we calculate the effective half-spread associated with a transaction in stock i at time t (ES_{it}) as the distance, in basis points, between the transaction price and the prevailing quote midpoint.

< TABLE 4 HERE >

Table 4 has the mean effective spreads for each transaction type, partitioned by market-cap quintile as well as pooled across all stocks. Comparing the two buy trade types, we find that effective spreads for buyer-initiated short sales (BIS) average 6.03 basis points for the whole sample vs. 5.74 basis points for buyer-initiated long sales (BIL), and this difference is statistically distinguishable from zero. There are similar results for each of the market-cap quintiles: compared to long sellers that supply liquidity, short sellers that supply liquidity do so at wider average effective spreads. This is consistent with Result 3 from our sequential trade model in Section 3.

Turning to seller-initiated trades, we find the reverse result. Short-sellers that demand liquidity (SIS) do so at average effective spreads which are narrower than average effective spreads for liquidity-demanding long sellers (SIL). For example, using the pooled sample, the average spreads are 5.89 basis points for SIL and 5.55 basis points for SIS, and the average spreads in the largest market-cap quintile are 1.85 basis points for seller-initiated longs vs. 1.79 basis points for seller-initiated shorts. This is consistent with Result 4 from our sequential trade model.

There are several possible explanations for these findings. It could be that our four trade types tend to concentrate in different stocks. It could be that short sellers and long sellers trade at different times in the same stocks. To sort out these explanations, we estimate the following fixed-effect panel regression on the 270 million trades in our sample:

$$ES_{it} = \mu_i + \beta_1 D_{it}^{BI} + \beta_2 D_{it}^{BIL} + \beta_3 D_{it}^{SIL} + \beta_4 D_{it}^{NYSE} + \varepsilon_{it} \quad (1)$$

where there are fixed effects for each stock, indicator variables for each transaction in stock i at time t if the transaction is associated with a buyer-initiated trade (D_{it}^{BI}), buyer-initiated long sale (D_{it}^{BIL}), or a seller-initiated long sale (D_{it}^{SIL}), and an indicator variable for trades executed on the NYSE (D_{it}^{NYSE}). The indicator variables D_{it}^{BI} , D_{it}^{BIL} , D_{it}^{SIL} , and D_{it}^{NYSE} take a value of one for buyer-initiated trades, buyer-initiated long sales, seller-initiated long sales, and NYSE trades, respectively, and zero otherwise. The indicator variables are configured so that differences between long and short sales, both buyer- and seller-initiated, can be read straight from the coefficient estimates. For example, the coefficient on the indicator variable D_{it}^{SIL} reflects the differential effective spread of a seller-initiated long sale compared to a seller-initiated short sale. Similarly, the coefficient on D_{it}^{BIL} reflects the differential effective spread of a buyer-initiated long sale compared to a buyer-initiated short sale. Additionally, the coefficient on the indicator variable D_{it}^{BI} reflects the differential effective spread of a buyer-initiated short sale compared to a seller-initiated short sale. We estimate the model on the pooled sample as well as on five market-cap quintiles, and in each case we report standard errors that are clustered by date.

< TABLE 5 HERE >

The results are in Panel A of Table 5. The fixed-effect regression results are consistent with the overall means from the earlier table, confirming that the effects are not the result of short sellers and long sellers concentrating their activity in different stocks. For a given stock, average effective spreads are narrowest when seller-initiated short sales take place. In the pooled sample, buyer-initiated long sales take place at effective spreads that are 0.07 basis points wider on average (0.193-0.126), buyer-initiated short sales

average 0.19 basis points wider, and seller-initiated long sales average 0.16 basis points wider. The differences between SIS and the other trade types (BIS, BIL and SIL) are statistically significant. The results are broadly similar across the market-cap quintiles. Overall these results indicate that liquidity-demanding shorts (SIS) are more prevalent when spreads are unusually narrow.¹³

One other possibility is that these four transaction types – BIL, BIS, SIL, and SIS – differ in the distribution of trade size. For example, if seller-initiated long sales (SIL) are larger than seller-initiated short sales (the omitted category in the regression), this could account for the observed positive coefficient on the SIL dummy. Given that Table 1 showed that short sales and long sales have a very similar distribution of sizes, we do not expect this to account for the effective spread differences. However, to investigate this, we estimate a variant of the effective spread regression above that includes a linear and a quadratic term for the dollar volume of the trade. We also include the inverse of share price in the regression to control for differences caused by price level and tick size. The results are in Panel B of Table 5. The trade size coefficients capture important variation in effective spreads, but the coefficients on the trade type dummies are virtually unchanged. Thus, trade size does not explain the observed differences between aggressive short sellers and other seller types. Similarly, the inverse of share price is also highly significant but does not change the significance of the trade type coefficients. Including these controls, the evidence continues to support the predictions of Results 3 and 4.

As noted above, short sellers that supply liquidity do so when effective spreads are an average of 0.19 basis points wider, compared to spreads when short sellers demand liquidity. Furthermore, across all market-cap quintiles, when a short seller provides liquidity spreads are statistically wider than when a long seller provides liquidity. For example, in the smallest quintile, spreads are 0.69 basis points wider for a buyer-initiated short sale compared to buyer-initiated long sale, and in the pooled sample the difference is 0.13 basis points.

Do passive short sellers (BIS) actually capture more spread than passive long sellers (BIL)? It could be that passive short sellers face greater adverse selection and are not able to pocket the wider effective spreads. To investigate this possibility, we measure one-minute realized spreads for the four types of trades. Realized spreads are defined for

¹³ We find very similar results for seller-initiated short sales across share price quintiles.

buys as the transaction price minus the quote midpoint one minute later, and are defined analogously for sells. Realized spreads are a proxy for liquidity supplier profits under the assumption that these liquidity suppliers hold onto their position for one minute and are then able to unwind at the then-prevailing midpoint. Realized spreads can be equivalently defined as the effective spread less the one-minute price impact, so if realized spreads do not reconcile with effective spreads, it must be due to differences in the magnitude of adverse price moves following a transaction.

Average realized spreads are also in Table 4 for both the pooled sample and the five market-cap quintiles. In the pooled sample, passive short sellers (BIS) trade at effective spreads that are 0.38 basis points wider than passive long sellers (BIL), but the realized spread they earn is only 0.12 basis points larger. The realized spread differential is smaller because passive short sellers face greater adverse selection – trades for which they provide liquidity have a statistically significant 0.32 basis points greater one-minute price impact than trades for which long sellers provide liquidity. Thus, on average, passive short sellers are relatively more willing to step in when conditions are more adverse, indicating that short sellers are adding to overall market quality.

The adverse selection results are also consistent with our sequential trade model. In the model, passive short sellers trade at times when there is more adverse selection, so they require wider spreads, and those wider spreads are completely offset by larger post-trade price moves, so that realized spreads should be identical for the two types of passive sales. Empirically, there is very little difference in observed realized spreads between passive short sellers and passive long sellers, consistent with the model.

Price impacts associated with liquidity demanders are also consistent with the sequential trade model. In the pooled sample, for example, one-minute price impacts associated with short sellers (SIS) are 4.12 basis points, compared to 3.95 basis points for liquidity-demanding long sellers (SIL). The difference is statistically significant. This is also consistent with the intuition behind Result 2 of the sequential trade model (and with Diamond and Verrecchia (1987)). Liquidity-demanding short sales are predominantly informed (in our model, they are always informed), whereas liquidity-demanding long sales can occur for non-informational reasons, such as liquidity needs. Thus, liquidity-demanding short sales predict future negative returns and their price impacts tend to be larger than those of liquidity-demanding long sales.

It is worth noting how small the realized spreads are in general. Across all stocks in our sample and all four trade types, realized spreads average only about 2.03 basis points, or about one-half cent on a typical \$30 stock. In fact, realized spreads are less than one basis point in the largest-cap quintile.¹⁴ Note that these are gross realized spreads that do not take into account rebates paid to liquidity suppliers. During our sample period, the NYSE did not pay these rebates, while Nasdaq rebates were 0.28 cents per share for high-volume liquidity suppliers, which amounts to about one additional basis point on a typical \$30 stock. Even though realized spreads do not directly measure trading revenue earned by liquidity suppliers, these results indicate that supplying liquidity is very competitive during our sample period.

In sum, our evidence so far indicates that compared to passive long sellers, passive short sellers are more willing to step in and supply liquidity in adverse market conditions, and short sellers demanding liquidity do so when liquidity is relatively cheap. These results are consistent with our sequential trade model. In the next section, we investigate price behavior just prior to a trade in order to further characterize the behavior of our four types of sellers.

5.2. Stock returns just before and after various types of sales

There is some prior evidence that short sellers are contrarian, taking short positions in stocks that have experienced recent stock price increases. For example, Diether, Lee, and Werner (2009) use daily data on shorting activity from 2005 and find that short sales increase the day after a share price rise. In this section, we investigate whether this is also true at finer time intervals. We also examine whether aggressive short sellers and passive short sellers are similarly contrarian, and we compare short sellers to long sellers to see if they are qualitatively similar.

To do this, we take our four types of executed order flow – BIS, SIS, BIL, and SIL – and for each stock we aggregate each type of order flow over five-minute intervals. Since U.S. stock markets are open for 6½ hours a day, from 9:30am to 4:00pm, there are 78 five-minute intervals each day. For calculating five-minute returns, we use the midquote that is in effect at the end of each five-minute interval. In our plots, the

¹⁴ This differs from Brogaard, Hendershott and Riordan (2014), who find that liquidity supply on Nasdaq by high-frequency traders during 2008-2009 actually loses money excluding rebates.

overnight return period from the 4:00pm close to the 9:30am open on the next trading day is considered a separate five-minute interval.

We then examine cumulative stock returns up to one day before and after transactions of various types using an event study methodology. The results are in Fig. 3. We show average cumulative stock returns pooled over all stocks and five-minute intervals, weighted either by the dollar volume of the particular sale type (Panel A) or by the number of trades of the relevant sale type (Panel B). Dollar volume weighting is dominated by trading in large-cap stocks, so the equal-weighting in Panel B gives more influence to trades in smaller-cap names. Fig. 3 Panel C shows the results for each weighting scheme for each of the five market-cap quintiles.

< FIGURE 3 HERE >

The results for price moves leading up to a sale are quite revealing. Not much happens to prices before a long sale. For SIL and BIL, average returns are generally close to zero up to one trading day (78 five-minute periods) before the relevant sale. Aggressive short sales (SIS) tend to follow short-term price declines, particularly in the smallest two market-cap quintiles. For example, when short sales are equal weighted, the average return over the 24 hours prior to an aggressive short sale is -0.43% for the smallest quintile and -0.32% for the next-smallest quintile. Standard errors clustered by stock and by date confirm that these average returns are statistically different from the returns prior to the other three transaction categories. This is consistent with Result 2 of the sequential trade model, and indicates that this pattern is simply due to informed short sellers trading gradually over time on their information.

The strongest regularity is for short sales that provide liquidity. The pooled dollar-volume weighted results in Panel A of Fig. 3 illustrate that the average share prices rise a full ten basis points in the 100 minutes (from five-minute period -20 to period 0) prior to a buyer-initiated short sale (BIS). The effect is even steeper as we get closer to event time zero. This result is not limited to large-cap stocks; the share price rise is statistically significant and of very similar magnitude in four out of the five market-cap quintiles. Passive short-sellers are sharply contrarian, responding by adding sell-side liquidity following a sharp price rise in the previous few minutes. It is important to emphasize the

magnitude of this result. On a typical \$30 stock, ten basis points means an increase of \$0.03 in the quote midpoint, averaged across every short sale that takes place at the ask.

This result is completely consistent with Result 1 of the sequential trade model. To see this, suppose that market-makers begin a given period with a long inventory position. A sequence of buy orders arrive. Prices rise, and market-makers reduce their long position. If there are enough buys in the sequence, eventually the market-makers must go short to provide liquidity. The data record a price rise before the market-making short sales as the competitive response to persistent buyer-initiated order flow.

It turns out that this price rise is permanent, not temporary. Of course, there is an immediate price impact associated with the buyer-initiated order flow. Looking out further after the passive short sale, prices remain more or less flat. There is no further upward drift and no downward price reversal, consistent with markets that efficiently incorporate trade-related information. This is also consistent with our rational expectations sequential trade model. The pooled results indicate that three of the four sale types – BIS, BIL, and SIL – do not show appreciable drift in the midpoint following the trade. Prices continue to erode slightly following aggressive short sales (SIS), as average quote midpoints drift down an additional two basis points or so over the 24 hours following these liquidity-demanding short sales, but this additional decline is not statistically significant.

When we disaggregate the results by market-cap quintile, we find that the results are not uniform in the cross-section. Passive shorts demonstrate distinctive contrarian behavior most strongly in large-cap stocks, particularly Quintiles 4 and 5. They are not at all contrarian for stocks in Quintile 2, and they are not as contrarian as passive long sellers in the smallest-cap quintile. Thus, the evidence is most consistent with the sequential trade model in the large-cap quintiles

The graphs also suggest that price impacts are bigger in magnitude when a short seller is involved. Compared to a seller-initiated long sale (SIL), for example, prices fall further following a seller-initiated short sale (SIS). To determine whether the patterns in the graphs are statistically significant, we turn to a panel regression approach that measures share price moves immediately following the four sale types.

For a time horizon k following a trade in stock i at time t , we regress the cumulative post-trade midpoint return $r_{i,t+k}$ (the subscript i is suppressed) on indicator

variables for buyer-initiated trades (D_{it}^{BI}), buyer-initiated long sells (D_{it}^{BLL}), and seller-initiated long sells (D_{it}^{SLL}), and possible fixed effects:

$$r_{t,t+k} = \mu_i + \gamma_t + \beta_1 D_{it}^{BI} + \beta_2 D_{it}^{BLL} + \beta_3 D_{it}^{SLL} + \varepsilon_{it} \quad (2)$$

Again, the indicator variables are configured to allow direct tests of differences between long and short sellers, both liquidity-demanding and liquidity-supplying. Model 1 has an intercept and no fixed effect, Model 2 has a stock fixed effect, and Model 3 has a fixed effect for each five-minute interval. The time dummies in Model 3 take out the average return across all sample stocks for that interval. Thus, this model is equivalent to using excess returns vs. the market as the dependent variable, thereby accounting for relative rather than raw stock price behavior. As in earlier regressions, the omitted category is aggressive shorting (SIS), so that the intercept in Model 1 is the average value for liquidity-demanding short sales (SIS), and the coefficients on indicator variables measure the incremental effect relative to the SIS case. Standard errors are clustered by stock and by date. We examine post-trade time horizons of five minutes, 30 minutes, two hours, and 24 hours. The five-minute and 30-minute horizons are common in empirical microstructure work, but we also examine returns up to one day out to connect to some of the lower-frequency information content literature. We estimate a version using the entire pooled sample as well as five market-cap quintile subsamples.¹⁵ We also estimate the models where each trade is weighted by dollar volume and where each trade is equal weighted. The results for these alternative weights are very similar. We focus on the dollar volume-weighted results because these better reflect liquidity suppliers' aggregate dollar losses to better-informed traders.

< TABLE 6 HERE >

The pooled results in Panel A of Table 6 show that at the five-minute horizon, we cannot statistically distinguish between the price impact of a seller-initiated short sale

¹⁵ Note that because all trades are grouped into five-minute intervals, the return period depends on a trade's location within such an interval. For example, if a trade takes place at 10:32:00am, the five-minute post-trade midpoint return would be measured from 10:30am to 10:35am. The same five-minute return would be used for any trade in that stock that occurs between 10:30:00am and 10:34:59am inclusive. Also, except for the 24 hour time horizon, the time horizons count only time when the market is open. For example, for a trade that takes place at 3:31pm, the two-hour return would be calculated from 3:30pm to 11:00am the next trading day.

(2.26 basis points on average, based on Model 1) and the price impact of a seller-initiated long sale ($2.26 - 0.14 = 2.12$ basis points). This result holds in Panel B across all five market-cap quintiles. At longer horizons there is also not a great deal of difference between long and short liquidity-demanding sells and any differences are only marginally statistically significant. Aggressive short sellers and aggressive long sellers have virtually the same average price impact.¹⁶ Initially, this might seem incompatible with results by others showing that short sellers have superior information. However, those results are at longer horizons, with holding periods measured in weeks or months. The results here simply indicate that if this group of short sellers is indeed better informed, they possess long-lived information that takes more than one day to be impounded into prices.

In contrast, for buyer-initiated trades the price impact is statistically different when a short seller is on the other side. At a five-minute horizon, a passive short sale (BIS) faces an adverse price impact of $-2.26 + 4.89 = 2.63$ basis points, while a long seller providing liquidity (BIL) experiences an adverse price impact of only $-2.26 + 4.89 - 0.72 = 1.91$ basis points. The difference of 0.72 basis points (the estimated coefficient on D_{it}^{BIL}) is statistically significant at the 1% level. This result is also present in Model 2 with stock fixed effects, which rules out the possibility that passive short sellers are simply concentrated in stocks with less adverse selection. As noted in the previous subsection, this result is consistent with the sequential trade model, because passive short sales tend to occur at times when there is more adverse selection.¹⁷ This result is consistent across all five size quintiles, although in the smallest quintile BIL and BIS five-minute price impacts are statistically indistinguishable.

¹⁶ In fact, Model 3 shows that price impacts are actually *smaller* when it is a short seller that demands liquidity. Market-adjusted prices actually move down 0.70 basis points further for SIL trades compared to similar trades involving a short seller. This difference is significant at the 1% level.

¹⁷ Interestingly, this difference between BIS and BIL is fairly short-lived. At the 30-minute horizon, price impacts associated with the two types of trades differ by only 0.26 basis points on average using Model 1. This price impact difference is not statistically significant.

5.3 The joint evolution of long and short sales over time

In the preceding analysis, we have conditioned on the single occurrence of a certain type of trade. However, we know from Biais, Hillion and Spatt (1995) and others that order flow is persistent. It is therefore likely that our different order flow types may themselves persist and may induce persistence in other types of order flow.

To investigate the dynamic relationships among our four sale types and between order flow and returns, we adapt the Hasbrouck (1991) vector autoregressive framework to our order flow partition. We classify trades into four categories (SIS, SIL, BIS, and BIL) to create order flow variables x_{it}^{SIS} , x_{it}^{SIL} , and so on, measured in number of shares. For each stock over the whole sample period, we estimate the following system:

$$\begin{aligned}
 x_t^{BIL} &= \mu^{BIL} + \sum_{i=1}^{10} \phi_i^r r_{t-i} + \sum_{i=1}^{10} \phi_i^{BIL} x_{t-i}^{BIL} + \sum_{i=1}^{10} \phi_i^{BIS} x_{t-i}^{BIS} + \sum_{i=1}^{10} \phi_i^{SIL} x_{t-i}^{SIL} + \sum_{i=1}^{10} \phi_i^{SIS} x_{t-i}^{SIS} + \varepsilon_t^{BIL} \\
 x_t^{BIS} &= \mu^{BIS} + \sum_{i=1}^{10} \theta_i^r r_{t-i} + \sum_{i=1}^{10} \theta_i^{BIL} x_{t-i}^{BIL} + \sum_{i=1}^{10} \theta_i^{BIS} x_{t-i}^{BIS} + \sum_{i=1}^{10} \theta_i^{SIL} x_{t-i}^{SIL} + \sum_{i=1}^{10} \theta_i^{SIS} x_{t-i}^{SIS} + \varepsilon_t^{BIS} \\
 x_t^{SIL} &= \mu^{SIL} + \sum_{i=1}^{10} \lambda_i^r r_{t-i} + \sum_{i=1}^{10} \lambda_i^{BIL} x_{t-i}^{BIL} + \sum_{i=1}^{10} \lambda_i^{BIS} x_{t-i}^{BIS} + \sum_{i=1}^{10} \lambda_i^{SIL} x_{t-i}^{SIL} + \sum_{i=1}^{10} \lambda_i^{SIS} x_{t-i}^{SIS} + \varepsilon_t^{SIL} \\
 x_t^{SIS} &= \mu^{SIS} + \sum_{i=1}^{10} \beta_i^r r_{t-i} + \sum_{i=1}^{10} \beta_i^{BIL} x_{t-i}^{BIL} + \sum_{i=1}^{10} \beta_i^{BIS} x_{t-i}^{BIS} + \sum_{i=1}^{10} \beta_i^{SIL} x_{t-i}^{SIL} + \sum_{i=1}^{10} \beta_i^{SIS} x_{t-i}^{SIS} + \varepsilon_t^{SIS} \\
 r_t &= \mu + \sum_{i=1}^{10} \gamma_i^r r_{t-i} + \sum_{i=0}^{10} \gamma_i^{BIL} x_{t-i}^{BIL} + \sum_{i=0}^{10} \gamma_i^{BIS} x_{t-i}^{BIS} + \sum_{i=0}^{10} \gamma_i^{SIL} x_{t-i}^{SIL} + \sum_{i=0}^{10} \gamma_i^{SIS} x_{t-i}^{SIS} + \varepsilon_t^r,
 \end{aligned} \tag{3}$$

where t indexes trades and individual stock subscripts are suppressed.¹⁸ Therefore, for each trade t , one of $\{x_{it}^{SIS}, x_{it}^{SIL}, x_{it}^{BIS}, x_{it}^{BIL}\}$ will be equal to the volume of the trade and the others will be zero.¹⁹ r_t is the log-midquote change subsequent to the t^{th} trade.

For each equation, we estimate coefficients on ten lags of each variable. In addition to the ten lags, midquote returns are also determined by contemporaneous order flow. A separate estimation is performed for each stock using the entire eight-month sample. Impulse responses are then calculated for a shock to each of the order flow types,

¹⁸ There is no separate overnight return in this analysis. The overnight return (adjusted for distributions) is included as part of the return from the last trade for the day to the first trade of the following day.

¹⁹ An exception occurs for the small number of trades that are made up partly from short volume and partly from long volume. Such trades arise when one order executes simultaneously against two or more orders and a single trade is recorded on the tape. If, for example, a market buy order for 500 shares executes against a 200 share long sale and a 300 share short sale and is recorded as a single trade, then the volume variables would take the values $x_{it}^{BIL} = 200$ and $x_{it}^{BIS} = 300$.

holding all other types of volume equal to their unconditional means. To make the impulse responses comparable across order flow types, the magnitude of each volume shock is set equal to the standard deviation of unanticipated buyer-initiated long (BIL) volume. We report the equal-weighted average impulse response up to 20 trades ahead, averaging across all stocks in a given size quintile.

< FIGURE 4 HERE >

Fig. 4 contains the results for the cumulative return response to various kinds of volume shocks. Midpoints adjust fully to a trade innovation within approximately ten trades. This partial adjustment is consistent with the effects of discreteness, as well as the standard practice of reporting multiple trades to the consolidated tape when a single marketable order interacts with multiple standing limit orders of smaller size. To confirm this, we also conduct the VAR in calendar time (results available on request), and we find that quote midpoints fully adjust to order flow innovations within one minute, consistent with market efficiency.

The price impacts are generally quite symmetric. To be precise, the price impact for a buyer-initiated short sale (BIS) is similar in magnitude to the price impact for a seller-initiated short sale (SIS), and the price impact for a buyer-initiated long sale (BIL) is similar in magnitude to the price impact for a seller-initiated long sale (SIL). Most interesting is that price impacts associated with long sales tend to be smaller than price impacts associated with short sales. The relative gap between long sale and short sale price impacts is widest for large stocks. In fact, for the large-cap quintile, the price impacts associated with short sales are about twice as large as the price impacts associated with similarly-sized long sales. The difference cannot be attributed to trade size. As noted earlier, long sales and short sales have very similar trade size distributions. Instead, it appears that short sellers tend to trade at times when information asymmetries are more severe. This is true for short sellers that demand liquidity as well as short sellers that supply liquidity. This result is consistent with recent empirical evidence that shorting is concentrated around news events which give rise to higher information asymmetries (e.g., Engelberg, Reed, and Ringgenberg (2012), Fox, Glosten, and Tetlock (2010), Christophe, Ferri, and Angel (2004), Boehmer, Jones, and Zhang (2012)).

< FIGURE 5 HERE >

Fig. 5 shows the cumulative response of each type of volume to its own shock, as inferred from the estimated VAR. All four types of volume are persistent, and for the three small-cap quintiles, there are no discernible differences in persistence across order-flow types. The results for the two large-cap quintiles are different, as long sales are considerably less persistent than short sales for this group. We conjecture that the order flow persistence arises from working a large parent order via a sequence of child orders. It could also be the case that multiple traders receive similar signals about valuation and tend to trade in the same direction at approximately the same time. The difference in persistence could arise from differences in average parent order size. Long sale parent orders are likely to be smaller on average, because the long sale categories mix large institutional parent orders and small individual orders, while shorting tends to be dominated by institutions. The institutional long sales would generate persistence, but long sales by individuals would not, and on average at the trade level, long sales would appear less persistent.

We also use the VAR to examine cross-volume effects. For each volume type, the cross effects are virtually zero and therefore are not reported.

6. Concluding discussion

In this paper, we take a magnifying glass to short sales and characterize what happens just prior to, at the same time as, and just after a short sale on the Nasdaq or NYSE during the first two-thirds of 2008. We partition by trade initiator, distinguishing more aggressive seller-initiated short sales from buyer-initiated shorting. We find that these two types of shorting are quite distinct. Shorts that supply liquidity, step in when spreads are unusually wide. They are also strongly contrarian, stepping in to short sell after fairly sharp share price rises over the past hour or so. They face greater adverse selection when they do step in, as prices tend to move up more following a passive short sale compared to times when a long seller is providing liquidity to a buyer-initiated trade. These results are consistent with a simple model of competitive liquidity supply in the

presence of long-lived private information. Thus, it appears that liquidity-supplying short sales are important contributors to market quality.

Who are these short sellers who provide liquidity? We suspect many are high-frequency traders who make markets algorithmically, a number of large hedge funds²⁰, as well as the market-making operations at many large broker-dealers. These liquidity suppliers may be formal, registered market-makers, or they may simply provide liquidity informally.

The popular press tends to focus on short sellers who gather and trade on fundamental information, but the evidence here indicates that shorting is also an important component of liquidity supply and overall market quality. Our results suggest that regulators and policymakers should encourage these particular market participants. More specifically, rules targeting short sellers should take care not to undermine this particular shorting function. For instance, a complete ban on short selling would make it impossible for many of these participants to provide liquidity, and the evidence in Boehmer, Jones, and Zhang (2013) and Beber and Pagano (2013) indicates that the bans on short selling in the US and other world markets during the financial crisis severely damaged market quality in banned stocks. In contrast, the SEC's more recent regulatory efforts are more carefully crafted. In early 2010, the SEC adopted Rule 201 which prohibits short sellers from demanding liquidity in a stock once it has fallen at least 10% in one trading day.²¹ However, short sellers that supply liquidity are unaffected. Our theoretical and empirical results suggest that these restrictions are unlikely to damage market quality, consistent with the results in Jain, Jain, and McInish (2012).

In contrast, shorts that demand liquidity are not contrarian on average, and actually these short sellers tilt slightly in the opposite direction. However, the evidence associated with these short sales is consistent with a standard model that involves gradual trading over time on negative information.

The slight tilt toward momentum trading by liquidity-demanding short sellers is the only piece of evidence reported here that warrants a closer look. It could be that

²⁰ One of the common strategies employed by high frequency trading firms is electronic market-making. Like traditional market-makers this involves the firms posting two-sided quotes, however, they typically enter and exit positions over very short time horizons.

²¹ SEC filing 34-61595, adopted February 26, 2010, effective date May 10, 2010 and compliance date February 28, 2011.

aggressive short sellers possess important negative information about fundamentals or future order flow, and this information is simply revealed over the short-term. A potentially less benign alternative is that aggressive short sellers are actually driving prices down by their continued selling activity. However, we do not see any evidence of this in the VAR analysis. Aggressive shorting is persistent, but no more so than any of the other types of trading that we identify. There is no evidence to indicate that seller-initiated shorting follows any sort of unusual pattern, and there is no price reversal after a sequence of short sales. Therefore, the results point away from “bear raid” and other manipulative explanations.

Of course, reality may be more complex than our model. In the model, informed traders with a negative signal submit liquidity-demanding orders, shorting if they do not own the stock. Liquidity-supplying shorts only originate from market-makers. In reality, informed traders that have a bearish view and do not hold the stock are likely to use a mix of liquidity-demanding and liquidity-supplying shorts managed through an execution algorithm. This is likely to bias our analysis against finding differences between liquidity-demanding and liquidity-supplying short sales predicted by our model; the distinction in the data between the two types of shorts should be less pronounced than in the model. Our results suggest that there is still a strong tendency for different traders to use different types of short sales. Future work might use parent and child order-level data to examine the extent to which different types of short sellers use combinations of market and limit orders and how market conditions influence the mix.

It is important to note that these are average results, cumulated over eight months of trading and across a large number of stocks. It is possible that this aggregation masks some interesting heterogeneity across stocks or over time. Seller behavior might be different around certain corporate events. It might differ near the open or the close of the market, and it might be somewhat different for stocks that are in financial distress. There could also be further heterogeneity within the two types that we study. For example, liquidity-demanding short sellers may consist mainly of informed traders who improve price efficiency, but could also include market manipulators and front-runners. In future work, various data subsets can be investigated to see whether the aggregate findings are broadly applicable.

In sum, the paper identifies an important heterogeneity among short sales. On close examination, it is clear that short sales that provide liquidity are qualitatively very different from short sales that demand liquidity. However, the model and the evidence indicate that both types are important to well-functioning markets: passive short sales are a key contributor to competitive, liquid markets, while liquidity-demanding short sales are important to price discovery. There is no evidence that shorting causes any deleterious effects. As a corollary, policymakers should remember that constraints or restrictions on shorting can cause deleterious effects via two separate channels: impeding liquidity supply as well as hindering price discovery.

APPENDIX

Preliminaries

Without loss of generality assume that after T rounds of trading, v will be either zero or one (corresponding to v^L and v^H), and that the two outcomes are equally likely ($p_0 = \frac{1}{2}$).

Recall that p_t is the liquidity provider's post-trade time t expectation of v . That is, the expectation is based on all information available at time t , including the direction of the trade at time t . This is also the conditional probability that the asset takes on the high value. At the start of trading, $p_0 = \frac{1}{2}$.

Recall that the probability of an informed trade in any given trading round is π . If selected, it turns out that the informed trader always buys if she knows that $v = 1$ (because the ask price will always be strictly less than one), and the informed trader always sells if she knows that $v = 0$. Uninformed liquidity traders are equally likely to buy and sell. Thus, we have:

$$\text{Prob}[buy_t | v = 1] = \pi + \frac{1}{2}(1 - \pi) \quad (4)$$

$$\text{Prob}[sell_t | v = 1] = \frac{1}{2}(1 - \pi). \quad (5)$$

Because p_t is the liquidity provider's conditional probability (given the history of observed trades) that the informed trader has received the high signal, we have:

$$\text{Prob}[buy_t | \text{history}_t] = \pi p_{t-1} + \frac{1}{2}(1 - \pi) \quad (6)$$

$$\text{Prob}[sell_t | \text{history}_t] = \pi (1 - p_{t-1}) + \frac{1}{2}(1 - \pi). \quad (7)$$

The ask price, Ask_t , should give zero profits and be the conditional expectation of v given all the trades up to $t - 1$ and conditional on the time t trade being a buy. By Bayes' rule:

$$Ask_t = E[v | buy_t] = \text{Prob}[v = 1 | buy_t] = \frac{\text{Prob}[buy_t | v = 1] \text{Prob}[v = 1]}{\text{Prob}[buy_t]}, \quad (8)$$

where all of these are also conditional on the history up to $t - 1$. This history can be summarized in the current expectation, p_{t-1} . This gives:

$$Ask_t = p_{t-1} \left(\pi + \frac{1}{2}(1 - \pi) \right) / \left(p_{t-1} \pi + \frac{1}{2}(1 - \pi) \right). \quad (9)$$

An even more compact expression is available for the odds ratio associated with the ask price:

$$Ask_t/(1 - Ask_t) = \frac{(1+\pi)}{(1-\pi)} \frac{p_{t-1}}{(1-p_{t-1})} \quad (10)$$

There is an analogous expression for $(1 - Bid_t)/Bid_t$.

Since bids and offers are just updated expectations in response to sells and buys, respectively, the above expressions show that the dynamics of expectations take a convenient form. Recall that the indicator variable, as $D_t = -1$ for a seller-initiated trade at time t and $D_t = +1$ for a buyer-initiated trade. Eq. (10) can be rewritten in terms of updated expectations:

$$p_t/(1 - p_t) = C^{D_t} p_{t-1}/(1 - p_{t-1}), \quad (11)$$

where $C = (1 + \pi)/(1 - \pi)$. Iterating, we get:

$$\ln(p_t/(1 - p_t)) = \ln(p_0/(1 - p_0)) + (\ln(C))N_t, \quad (12)$$

where $N_t = \sum_{\tau=1}^{t-1} D_\tau$ is the number of buys less the number of sells seen up to time t . The term N_t also captures the change in the liquidity provider's inventory from its starting level of zero. More details on these preliminary results can be found in Foucault, Pagano, and Roell (2013). Given the initial conditions, this reduces to:

$$\ln(p_t/(1 - p_t)) = (\ln(C))N_t. \quad (13)$$

Thus the liquidity provider will be short-selling if $p_t \geq \frac{1}{2}$ and long-selling otherwise.

Define the return $R_{0,t}$ to be the change in the expected value of the asset from the start of trading ($t = 0$) to time t . Then we have $R_{0,t} = p_t - \frac{1}{2}$. Now we are ready to derive the results stated in the text.

Result 1. Liquidity-supplying short sales (BIS) are contrarian over short horizons, in an absolute sense and relative to liquidity-supplying long sales (BIL).

Proof. The liquidity provider is shorting if and only if $p_t \geq \frac{1}{2}$, and since the lagged return from the start of trading is given by $R_{0,t} = p_t - \frac{1}{2}$, this is clearly non-negative if the liquidity provider is shorting (BIS). Similarly, since BIL requires $p_t < \frac{1}{2}$, the lagged return since the start of trading is strictly negative, assuming that $T \geq 2$ so that a BIL is possible. Thus $[R_{0,t}|BIS_t] \geq 0 > [R_{0,t}|BIL_t]$.

QED.

Result 2. Liquidity-demanding short sales (SIS) tend to follow price declines, in an absolute sense and relative to liquidity-demanding long sales (SIL).

Proof. Consider first the situation in which an informed trader starts with zero or negative inventory. In that case, an informed trader who is selling is always shorting, and SIL trades only come from uninformed traders. Since SIL arrivals do not depend on the history we have $E[R_{0,t} | SIL_t] = 0$ for this case.

A similar expression conditional on an SIS trade is given by:

$$E[R_{0,t} | SIS_t] = E[p_t - \frac{1}{2} | SIS_t] = \frac{\sum_j (p_{jt} - \frac{1}{2}) \text{Prob}(SIS_{jt})}{\sum_j \text{Prob}(SIS_{jt})}, \quad (14)$$

where the summation is over all nodes j (all possible states, which can be summarized by all possible values of p_t) at time t . At any such node, $\text{Prob}(SIS_{jt}) = \pi (1 - p_{jt})$. Choose a node such that $p_t = \frac{1}{2} + e$ and a corresponding node such that $p_t = \frac{1}{2} - e$. Then this pair of nodes contributes the following to the numerator of the conditional expected return:

$$p_t = \frac{1}{2} + e: \quad e\pi \left(\frac{1}{2} - e\right) \quad (15)$$

$$p_t = \frac{1}{2} - e: \quad -e\pi \left(\frac{1}{2} + e\right) \quad (16)$$

and the sum of these two terms is $-2\pi e^2$ which is negative. This is true of all such node pairs, so the overall conditional expectation must be negative: $E[R_{0,t} | SIS_t] < 0$.

Suppose now that the informed trader's initial inventory is positive. As long as there are enough trading periods, it will be possible for the informed trader to reach a zero position and short sell thereafter. For any such starting inventory $I_0^{INF} > 0$, SIS trades can only be observed once the informed trader has sold an amount I_0^{INF} . At and beyond this crossover point, we expect negative net order flow overall, because we have observed sales of at least I_0^{INF} and the expected uninformed order flow is zero. Thus, the expected price is below $\frac{1}{2}$ at that crossover point and any later time that an SIS occurs:

$$E[p_t | I_0^{INF} > 0, SIS_t] < \frac{1}{2}. \quad (17)$$

From there, our earlier results show that:

$$E[R_{0,t} | SIS_t] < 0. \quad (18)$$

QED.

References

- Abreu, D., Brunnermeier, M.K., 2002. Synchronization risk and delayed arbitrage. *Journal of Financial Economics* 66, 341-360.
- Alexander, G.J., Peterson, M.A., 2008. The effects of price tests on trader behavior and market quality: an analysis of Reg SHO. *Journal of Financial Markets* 11, 84-111.
- Aromi, D., Caglio, C., 2008. Memorandum re: analysis of short selling activity during the first weeks of September 2008, U.S. Securities and Exchange Commission Office of Economic Analysis, December 16, available at <http://www.sec.gov/comments/s7-08-09/s70809-369.pdf>.
- Asquith, P., Pathak, P.A., Ritter, J.R., 2005. Short interest, institutional ownership, and stock returns. *Journal of Financial Economics* 78, 243-276.
- Beber, A., Pagano, M., 2013. Short-selling bans around the world: evidence from the 2007-09 crisis. *Journal of Finance* 68, 343-381.
- Bessembinder, H., Hao, J., Lemmon, M., 2011. Why designate market makers? Affirmative obligations and market quality. Unpublished manuscript.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of limit order book and the order flow in the Paris Bourse. *Journal of Finance* 50, 1655-1681.
- Boehmer, E., Huszár, Z., Bradford, J.D., 2010. The good news in short interest. *Journal of Financial Economics* 96, 80-97.
- Boehmer, E., Jones, C.M., Zhang, X., 2008. Which shorts are informed?. *Journal of Finance* 63, 491-527.
- Boehmer, E., Jones, C.M., Zhang, X., 2012. What do short sellers know?. Unpublished manuscript.
- Boehmer, E., Jones, C.M., Zhang, X., 2013. Shackling short sellers: the 2008 shorting ban. *Review of Financial Studies* 26, 1363-1400.
- Boehmer, E., Wu, J., 2013. Short selling and the price discovery process. *Review of Financial Studies* 26, 287-322.
- Brent, A., Morse, D., Stice, E.K., 1990. Short interest - explanations and tests. *Journal of Financial and Quantitative Analysis* 25, 273-289.
- Brogaard, J., Hendershott, T., Riordan, R., 2014. High-frequency trading and price discovery. *Review of Financial Studies* 27, 2267-2306.
- Chakrabarty, B., Moulton, P.C., Shkilko, A., 2012. Short sales, long sales, and the Lee-

- Ready trade classification algorithm revisited. *Journal of Financial Markets* 15, 467-491.
- Christophe, S., Ferri, M., Angel, J., 2004. Short-selling prior to earnings announcements. *Journal of Finance* 59, 1845-1875.
- Cohen, L., Diether, K.B., Malloy, C., 2007. Supply and demand shifts in the shorting market. *Journal of Finance* 62, 2061-2096.
- Danielsen, B., Sorescu, S., 2001. Why do option introductions depress stock prices? An empirical study of diminishing short sale constraints. *Journal of Financial and Quantitative Analysis* 36, 451-484.
- Dechow, P., Hutton, A., Meulbroek, L., Sloan, R.G., 2001. Short-sellers, fundamental analysis, and stock returns. *Journal of Financial Economics* 61, 77-106.
- Desai, H., Ramesh, K., Thiagarajan, S., Balachandran, B., 2002. An investigation of the informational role of short interest in the Nasdaq market. *Journal of Finance* 57, 2263-2287.
- Diether, K.B., Lee, K-H., Werner, I.M., 2009. Short-sale strategies and return predictability. *Review of Financial Studies* 22, 575-607.
- Duffie, D., Garleanu, N.B., Pedersen, L.H., 2002. Securities lending, shorting, and pricing. *Journal of Financial Economics* 66, 307-339.
- Engelberg, J., Reed, A.V., Ringgenberg, M., 2012. How are shorts informed? Short sellers, news, and information processing. *Journal of Financial Economics* 105, 260-278.
- Fox, M.B., Glosten, L.R., Tetlock, P.C., 2010. Short selling and the news: a preliminary report on an empirical study. *New York Law School Law Review* 54, 645-686.
- Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 71-100.
- Hasbrouck, J., 1991. Measuring the information content of stock trades. *Journal of Finance* 46, 179-207.
- Hendershott, T., Riordan, R., 2013. Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis* 48, 1001-1024.
- Hong, H., Scheinkman, J., Xiong, W., 2006. Asset float and speculative bubbles. *Journal of Finance* 61, 1073-1117.

- Jain, C., Jain, P.K., McInish, T.H., 2012. Short selling: the impact of SEC Rule 201 of 2010. *Financial Review* 47, 37-64.
- Jones, C.M., 2012. Shorting restrictions: revisiting the 1930's. *Financial Review* 47, 1-35.
- Jones, C.M., Lamont, O.A., 2002. Short sale constraints and stock returns. *Journal of Financial Economics* 66, 207-239.
- Lamont, O.A., Stein J.C., 2004. Aggregate short interest and market valuations. *American Economic Review* 94, 29-32.
- Lee, C.M., Ready, M.J., 1991. Inferring trade direction from intraday data. *Journal of Finance* 46, 733-746.
- Menkveld, A.J., 2013. High frequency trading and the new-market makers. *Journal of Financial Markets* 16, 712-740.
- Miller, E.M., 1977. Risk, uncertainty, and divergence of opinion. *Journal of Finance* 32, 1151-1168.
- Saffi, P.A.C., Sigurdsson, K., 2011. Price efficiency and short selling. *Review of Financial Studies* 24, 821-852.

Table 1

Summary statistics.

This table reports summary statistics on the number, volume and size of short sales relative to long sales executed on the NYSE (Panel A), Nasdaq (Panel B) and the pooled sample consisting of NYSE and Nasdaq trades. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008. 25%, 50% and 75% refer to distribution quartiles, n refers to the total number of short sales and long sales in the sample. To calculate daily number of trades, dollar volume per stock and short proportion, we first sum the number of trades and dollar volume and calculate the short proportion for each stock-day and then calculate the mean and quartiles for the stock-days in the sample. To calculate short and long trade sizes, we first find the mean (median) short and long trade sizes for each stock-day and then calculate the mean (quartiles) for the stock-days in the sample.

	Daily number of short sales per stock	Shorting share of trades	Daily dollar volume of short sales per stock (\$'000)	Shorting share of dollar volume	Size of a short sale (\$)	Size of long sale (\$)
Panel A: NYSE trades						
Mean	1,123	38.0%	11,047	40.0%	6,580	6,754
25%	268	29.2%	939	30.2%	1,710	1,721
50%	682	38.2%	3,280	40.7%	3,046	3,065
75%	1,449	47.7%	10,659	50.9%	4,927	4,953
n					60,848,720	94,090,063
Panel B: Nasdaq trades						
Mean	788	38.6%	8,940	39.1%	6,587	5,892
25%	105	31.1%	354	30.8%	1,767	1,704
50%	318	39.6%	1,524	40.1%	3,171	3,039
75%	866	47.8%	6,443	49.0%	5,102	4,884
n					42,675,697	75,503,593
Panel C: Pooled sample						
Mean	1,911	39.2%	19,988	38.7%	6,602	6,414
25%	390	31.4%	1,380	30.0%	1,713	1,700
50%	1,016	40.3%	4,925	39.6%	3,057	3,026
75%	2,325	48.9%	17,371	48.8%	4,960	4,889
n					103,524,417	169,593,656

Table 2

Dollar volume of short sales.

This table reports mean daily dollar volumes (\$ '000s) of short sales per stock and the percentage of total daily dollar volume per stock made up by short sales. The dollar volumes and percentages are reported for stocks grouped in quintiles of four variables: book-to-market ratio (*B/M*); market capitalization (*Size*); closing price (*Price*); and the open short interest as a percentage of total shares on issue (*Short Interest*). *Pooled* indicates results not grouped by quintiles of the corresponding variable. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.

Panel A: Size and book-to-market quintiles						
B/M	Size					Pooled
	Small	2	3	4	Big	
Low	841	1,854	8,135	20,585	72,910	31,314
	40.5%	40.9%	42.0%	37.3%	36.0%	36.5%
2	811	2,615	7,233	13,825	52,410	17,430
	43.3%	38.8%	40.6%	37.8%	41.0%	39.9%
3	596	2,993	8,567	17,616	54,965	17,154
	41.7%	41.4%	39.8%	39.3%	36.5%	37.6%
4	542	2,293	5,844	16,505	69,951	17,251
	40.5%	41.2%	42.2%	42.6%	38.9%	39.5%
High	1,252	2,980	7,056	18,055	221,123	15,558
	41.1%	40.3%	41.2%	42.7%	41.0%	41.3%
Pooled	902	2,591	7,464	16,625	71,470	19,730
	41.2%	40.7%	41.0%	39.3%	38.0%	38.5%

Panel B: Price and short interest quintiles						
Price	Short Interest					Pooled
	Low	2	3	4	High	
Low	2,968	1,713	997	2,037	6,208	3,213
	36.6%	38.8%	41.6%	42.0%	39.8%	39.6%
2	39,685	11,500	3,386	3,314	7,032	8,478
	29.8%	41.3%	38.5%	39.5%	42.3%	35.2%
3	7,535	46,467	4,159	25,145	6,793	21,168
	35.4%	41.6%	38.5%	45.0%	45.5%	42.0%
4	44,926	27,865	22,891	16,107	10,289	25,703
	35.8%	40.3%	38.7%	43.5%	45.0%	38.8%
High	45,657	40,680	39,876	45,787	16,166	39,838
	34.2%	38.5%	41.9%	37.6%	46.0%	37.3%
Pooled	31,382	31,790	13,126	14,031	8,366	19,730
	34.2%	40.3%	40.3%	41.4%	43.4%	38.5%

Table 3

Trade aggressiveness.

This table reports the aggressiveness of executed short and long sale trades relative to the prevailing quotes. Size quintiles are calculated using market capitalization. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.

		NYSE		Nasdaq		Pooled	
		Long sales	Short sales	Long sales	Short sales	Long sales	Short sales
Panel A: Pooled							
	Behind-the-ask limit order	4.0%	4.1%	3.8%	4.2%	3.9%	4.2%
	At-the-ask limit order	40.8%	39.9%	39.9%	39.5%	40.4%	39.8%
	Within the quotes	15.5%	15.6%	17.2%	17.8%	16.3%	16.5%
	At-the-bid market order	35.8%	36.6%	35.2%	34.8%	35.5%	35.9%
	Below-the-bid market order	3.9%	3.7%	3.9%	3.7%	3.9%	3.7%
Panel B: By size quintile							
Small	Behind-the-ask limit order	2.0%	2.3%	1.9%	2.3%	2.0%	2.3%
	At-the-ask limit order	41.1%	39.3%	39.0%	37.0%	40.3%	38.6%
	Within the quotes	20.1%	20.3%	22.3%	23.3%	20.9%	21.2%
	At-the-bid market order	34.8%	36.3%	34.7%	35.4%	34.8%	36.0%
	Below-the-bid market order	2.0%	1.8%	2.1%	1.9%	2.1%	1.8%
2	Behind-the-ask limit order	2.0%	2.2%	1.9%	2.2%	2.0%	2.2%
	At-the-ask limit order	41.0%	39.1%	38.5%	37.2%	40.1%	38.5%
	Within the quotes	19.4%	19.2%	23.3%	23.9%	20.8%	20.7%
	At-the-bid market order	35.5%	37.7%	34.2%	34.8%	35.1%	36.8%
	Below-the-bid market order	2.1%	1.8%	2.1%	1.8%	2.1%	1.8%
3	Behind-the-ask limit order	2.6%	2.6%	2.5%	2.7%	2.5%	2.7%
	At-the-ask limit order	41.2%	39.7%	39.5%	38.5%	40.5%	39.2%
	Within the quotes	18.0%	17.5%	20.3%	20.9%	18.9%	18.7%
	At-the-bid market order	35.7%	37.9%	35.1%	35.6%	35.4%	37.0%
	Below-the-bid market order	2.6%	2.3%	2.6%	2.3%	2.6%	2.3%
4	Behind-the-ask limit order	3.3%	3.6%	3.2%	3.6%	3.3%	3.6%
	At-the-ask limit order	38.4%	37.4%	36.5%	36.0%	37.6%	36.9%
	Within the quotes	20.2%	20.2%	24.2%	24.3%	21.8%	21.7%
	At-the-bid market order	34.7%	35.6%	32.8%	32.9%	33.9%	34.5%
	Below-the-bid market order	3.4%	3.3%	3.4%	3.2%	3.4%	3.2%
Big	Behind-the-ask limit order	5.1%	5.4%	4.5%	5.0%	4.8%	5.2%
	At-the-ask limit order	41.6%	41.4%	41.2%	41.2%	41.4%	41.3%
	Within the quotes	11.9%	11.9%	13.8%	14.2%	12.8%	12.9%
	At-the-bid market order	36.4%	36.5%	36.0%	35.2%	36.2%	35.9%
	Below-the-bid market order	5.0%	4.8%	4.5%	4.4%	4.8%	4.6%

Table 4

Mean spreads and price impacts.

This table reports means of effective half-spreads, one-minute midquote price impacts, and one-minute realized spreads, measured in basis points. Trades are classified into four groups: buyer initiated trades with long sell side (BIL), buyer initiated short sales (BIS), seller initiated long sales (SIL) and seller initiated short sales (SIS). Means of the three variables for each trade type and differences in means (BIL-BIS and SIL-SIS) are calculated in three steps: (i) a simple average across trades in each stock-day, (ii) calculate the differences in stock-day means (BIS-BIL and SIS-SIL), and (iii) a simple average of the stock-day means and differences in means. Prior to calculating means the variables are winsorised within each stock at the 1st and 99th percentile for price impact and realized spread and the 99th percentile for effective spreads. *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. Standard errors are reported in parentheses. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.

		BIL	BIS	BIS-BIL	SIL	SIS	SIS-SIL
Panel A: Pooled							
	Effective spread	5.74 (0.04)	6.03 (0.04)	0.38*** (0.02)	5.89 (0.03)	5.55 (0.03)	-0.23*** (0.03)
	Price impact	3.89 (0.03)	4.16 (0.04)	0.32*** (0.03)	3.95 (0.04)	4.12 (0.04)	0.25*** (0.04)
	Realized spread	2.00 (0.03)	2.07 (0.04)	0.12*** (0.03)	2.10 (0.03)	1.59 (0.04)	-0.47*** (0.04)
Panel B: By size quintile							
Small	Effective spread	13.17 (0.14)	14.01 (0.15)	1.17*** (0.09)	13.49 (0.13)	12.53 (0.12)	-0.56*** (0.12)
	Price impact	8.28 (0.13)	9.61 (0.19)	1.48*** (0.17)	9.03 (0.15)	9.19 (0.19)	0.46** (0.19)
	Realized spread	5.22 (0.15)	4.97 (0.17)	-0.05 (0.16)	4.80 (0.14)	3.75 (0.17)	-0.93*** (0.17)
Q2	Effective spread	6.50 (0.04)	6.87 (0.04)	0.38*** (0.01)	6.69 (0.04)	6.45 (0.04)	-0.24*** (0.04)
	Price impact	4.69 (0.04)	4.82 (0.05)	0.13*** (0.04)	4.58 (0.05)	4.92 (0.05)	0.33*** (0.05)
	Realized spread	1.95 (0.04)	2.23 (0.05)	0.28*** (0.04)	2.27 (0.04)	1.67 (0.05)	-0.60*** (0.05)
Q3	Effective spread	4.27 (0.03)	4.45 (0.03)	0.18*** (0.01)	4.38 (0.03)	4.22 (0.03)	-0.16*** (0.03)
	Price impact	3.02 (0.03)	3.06 (0.03)	0.04 (0.02)	2.91 (0.03)	3.07 (0.04)	0.16*** (0.04)
	Realized spread	1.34 (0.03)	1.50 (0.04)	0.16*** (0.02)	1.58 (0.03)	1.25 (0.04)	-0.33*** (0.04)
Q4	Effective spread	2.86 (0.01)	2.99 (0.01)	0.13*** (0.01)	2.93 (0.01)	2.81 (0.01)	-0.13*** (0.01)
	Price impact	2.09 (0.02)	2.09 (0.02)	0.00 (0.02)	1.96 (0.02)	2.13 (0.02)	0.17*** (0.02)
	Realized spread	0.84 (0.03)	0.98 (0.02)	0.14*** (0.02)	1.05 (0.03)	0.74 (0.02)	-0.31*** (0.02)
Big	Effective spread	1.82 (0.01)	1.87 (0.01)	0.05*** (0.00)	1.85 (0.01)	1.79 (0.01)	-0.06*** (0.01)
	Price impact	1.27 (0.02)	1.25 (0.02)	-0.01 (0.01)	1.18 (0.02)	1.31 (0.02)	0.13*** (0.02)
	Realized spread	0.62 (0.02)	0.68 (0.02)	0.07*** (0.01)	0.74 (0.02)	0.53 (0.02)	-0.21*** (0.02)

Table 5

Effective spread regressions.

This table reports trade-level regression estimates for which the dependent variable is effective half-spread measured in basis points. D_{BI} , D_{BIL} , and D_{SIL} are dummy variables that take the value of one for buyer initiated trades, buyer initiated trades with long sell side, and seller initiated long sales, respectively (short seller initiated trades are the base case), and zero otherwise. D_{NYSE} is a dummy variable for trades executed on the NYSE (Nasdaq is the base case). $dVol$ and $Price$ are the dollar volume and price of the trade, respectively. Regressions include fixed effects for each stock. *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. Standard errors are clustered by date. T-statistics are reported in parentheses. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.

	Pooled	By size quintile				
		Small	Q2	Q3	Q4	Big
Panel A: Model 1						
Intercept	9.789*** (39.40)	24.395*** (9.63)	9.698*** (38.23)	2.448*** (35.13)	2.194*** (30.37)	1.750*** (40.64)
D_BI	0.193*** (12.45)	0.645*** (10.41)	0.484*** (9.84)	0.243*** (12.08)	0.173*** (15.24)	0.117*** (6.24)
D_BIL	-0.126*** (-5.54)	-0.691*** (-12.91)	-0.347*** (-8.90)	-0.147*** (-6.08)	-0.148*** (-7.11)	-0.048** (-2.12)
D_SIL	0.156*** (4.88)	0.404*** (5.34)	0.368*** (6.40)	0.214*** (4.82)	0.121*** (4.22)	0.109*** (3.55)
D_NYSE	0.151*** (5.79)	0.399*** (6.21)	0.079 (1.46)	0.181*** (3.67)	0.163*** (6.86)	0.134*** (5.32)
Panel B: Model 2						
Intercept	0.907 (0.67)	-5.002 (-1.22)	3.001*** (3.59)	-0.573 (-0.80)	0.498*** (3.99)	0.630*** (3.35)
D_BI	0.186*** (11.89)	0.783*** (12.09)	0.454*** (9.61)	0.207*** (8.56)	0.180*** (15.89)	0.117*** (6.22)
D_BIL	-0.134*** (-5.90)	-0.701*** (-13.85)	-0.382*** (-9.81)	-0.150*** (-6.60)	-0.150*** (-7.19)	-0.051** (-2.22)
D_SIL	0.145*** (4.59)	0.488*** (7.98)	0.317*** (5.76)	0.194*** (4.36)	0.126*** (4.42)	0.107*** (3.46)
D_NYSE	0.194*** (6.95)	0.473*** (7.78)	0.271*** (5.51)	0.228*** (4.36)	0.178*** (7.38)	0.145*** (5.51)
dVol	0.019*** (11.19)	0.969*** (5.51)	0.364*** (9.90)	0.192*** (8.87)	0.069*** (12.13)	0.012*** (8.14)
dVol ²	0.000*** (-6.91)	-0.001*** (-4.25)	0.000*** (-8.23)	0.000*** (-7.77)	0.000*** (-5.07)	0.000*** (-5.03)
1/Price	66.969*** (6.54)	96.165*** (7.25)	49.017*** (7.73)	65.722*** (4.06)	47.698*** (22.43)	39.904*** (6.85)

Table 6

Price impact measured through varying post-trade horizons.

This table reports interval-level regression estimates, for which the dependent variable is the midquote price impact measured in basis points over various length periods forward in time. Each trade is weighted by dollar volume. D_{BI} , D_{BIL} , and D_{SIL} are dummy variables that take the value of one for buyer initiated trades, buyer initiated trades with long sell side, and seller initiated long sales, respectively (short seller initiated trades are the base case), and zero otherwise. *Small* to *Big* are size quintiles measured by average market capitalization. Model 1 does not have fixed effects, Model 2 includes fixed effects on stocks and Model 3 includes fixed effects on five-minute intervals. *, ** and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. Standard errors are clustered by stock and by date. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.

	5 mins			30 mins			2 hours			24 hours			
	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3	
Panel A: Pooled													
Intercept	-2.26***			-2.66***			-3.75**			-3.46			
D_BI	4.89***	4.88***	3.54***	4.92***	4.93***	3.68***	5.11***	5.14***	3.88***	5.55***	5.57***	4.05***	
D_BIL	-0.72***	-0.68***	-0.99***	-0.26	-0.22	-0.51***	0.20	0.21	0.06	0.29	0.43	0.41	
D_SIL	0.14	0.18	-0.32***	0.40**	0.44**	-0.04	0.99**	1.00**	0.62**	1.21	1.35	1.00*	
Panel B: By size quintile													
Small	Intercept	-7.15***		-7.89***			-7.80**			-8.30			
	D_BI	16.1***	16.0***	9.19***	19.3***	19.3***	11.9***	18.2***	18.2***	11.2***	20.3***	20.6***	12.4***
	D_BIL	-1.27	-1.20	-0.37	-1.75	-1.71	-1.39**	-1.72	-1.59	-1.78**	-3.52	-3.35	-2.71
	D_SIL	0.74	0.72	-0.15	-0.55	-0.59	-1.28*	-2.09	-2.00	-2.58***	-4.42	-3.96	-3.42
2	Intercept	-4.56***		-6.14***			-9.68***			-13.1***			
	D_BI	11.1***	11.1***	6.77***	12.9***	12.9***	8.51***	14.0***	14.0***	9.40***	14.6***	14.5***	9.60***
	D_BIL	-2.52**	-2.44**	-1.89***	-1.91*	-1.75*	-1.91***	-0.40	-0.20	-1.53	0.21	0.73	-1.70
	D_SIL	0.14	0.28	-0.60**	0.61	0.80*	-0.56	2.32**	2.50**	0.51	3.17	3.60	0.26
3	Intercept	-3.46***		-3.66***			-4.95***			-6.03*			
	D_BI	7.66***	7.65***	4.94***	7.70***	7.68***	5.36***	7.30***	7.29***	5.39***	7.70***	7.69***	5.49***
	D_BIL	-0.91**	-0.92**	-0.92***	-0.39	-0.43	-0.89***	0.60	0.51	-0.23	0.73	0.56	0.25
	D_SIL	0.32	0.33	-0.26*	0.54	0.49	-0.51**	1.18	1.07	0.20	1.90*	1.65	1.08
4	Intercept	-2.52***		-3.24***			-5.00***			-8.29*			
	D_BI	5.49***	5.49***	3.84***	5.77***	5.77***	4.12***	6.05***	6.11***	4.35***	6.06***	6.20***	4.37***
	D_BIL	-0.68**	-0.66**	-0.96***	-0.45	-0.35	-0.70***	0.10	0.21	-0.49	1.49	1.81	0.33
	D_SIL	0.62	0.63	-0.03	0.93	1.00*	0.27	1.44**	1.46***	0.49	2.67*	2.76*	1.52
Big	Intercept	-1.94***		-2.24***			-3.09*			-1.74			
	D_BI	4.11***	4.10***	2.79***	3.99***	3.99***	2.75***	4.20***	4.23***	2.96***	4.70***	4.71***	3.19***
	D_BIL	-0.60***	-0.56***	-0.81***	-0.08	-0.06	-0.25*	0.25	0.26	0.35	0.00	0.21	0.32
	D_SIL	-0.01	0.03	-0.40***	0.25	0.28	-0.11	0.81	0.84	0.56*	0.70	0.97	0.57




Fig. 1. The sequential trade model. This tree diagram illustrates the parameters of the model and the process by which the type of trade in each round, t , is determined. The variables are defined as follows: v is the fundamental value, which takes the high value v^H with probability p_0 , and the low value v^L with probability $(1 - p_0)$, where $v^H > v^L$. Through the course of trading, the liquidity provider Bayesian updates the probability of the fundamental value taking the high value, p_t , from her initial belief, p_0 . The liquidity provider starts with an initial inventory of zero, $I_0^{LP} = 0$, and her inventory at any point in time is equal to the net number of sells up to that point, $I_t^{LP} = -\sum_{\tau=1}^t D_\tau$. Nature picks the probability of informed trading (the probability that in any given trading round the informed trader submits an order), π , from a uniform distribution ranging from 0 to π^{MAX} , as well as the informed trader's inventory from a uniform distribution ranging from \underline{I}^{INF} to \overline{I}^{INF} . In each trading round the liquidity provider receives a unit size market order from the informed trader with probability π or the uninformed traders with probability $(1 - \pi)$. Uninformed traders submit buy market orders (direction, $D_t = +1$) and sell market orders (direction, $D_t = -1$) with equal probability, and never short sell because they trade for exogenous liquidity needs, not information. The informed traders know the fundamental value and therefore they buy when $v = v^H$ and sell when $v = v^L$. If the informed trader has positive inventory at the time she sells ($I_t^{INF} > 0$) the trade is a seller-initiated long sale (SIL), but if she does not hold the stock ($I_t^{INF} \leq 0$), she short sells, leading to a seller-initiated short sale (SIS). Similarly, if the liquidity provider has positive inventory when she receives a market buy order ($I_t^{LP} > 0$) she sells from her inventory resulting in a buyer-initiated long sale (BIL), whereas if the liquidity provider has zero inventory or a short position when she receives a market buy order ($I_t^{LP} \leq 0$) she short sells resulting in a buyer-initiated short sale (BIS).






Fig. 2. Average cumulative log midquote returns (CR) before and after short sales (solid red line) that are buyer initiated (light-weight line) and seller initiated (heavy-weight line) and long sales (dashed blue line) that are buyer initiated (light-weight line) and seller initiated (heavy-weight line). The horizontal axis measures time (trading periods) before and after the trade (which takes place in interval zero). The returns in this figure are averages from 500,000 simulations of the sequential trade model summarized in Fig. 1 and detailed in the text.

Fig. 3. Average cumulative raw log midquote returns (CR) before and after short sales (solid red line) that are buyer initiated (light weight line) and seller initiated (heavy weight line) and long sales (all trades other than those in which the sell side is short – dashed blue line) that are buyer initiated and seller initiated. The horizontal axis measures the number of five-minute intervals from the trade (which takes place in interval zero). 78 intervals equals a whole trading day, i.e., for a trade at 13:05 on Wednesday, the cumulative return to interval 78 is the return from 13:05 on Wednesday to 13:05 on Thursday. In Panel A trades are weighted by their dollar volume, in Panel B trades are equal weighted, and Panel C provides separate results for the five market-cap quintiles. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.

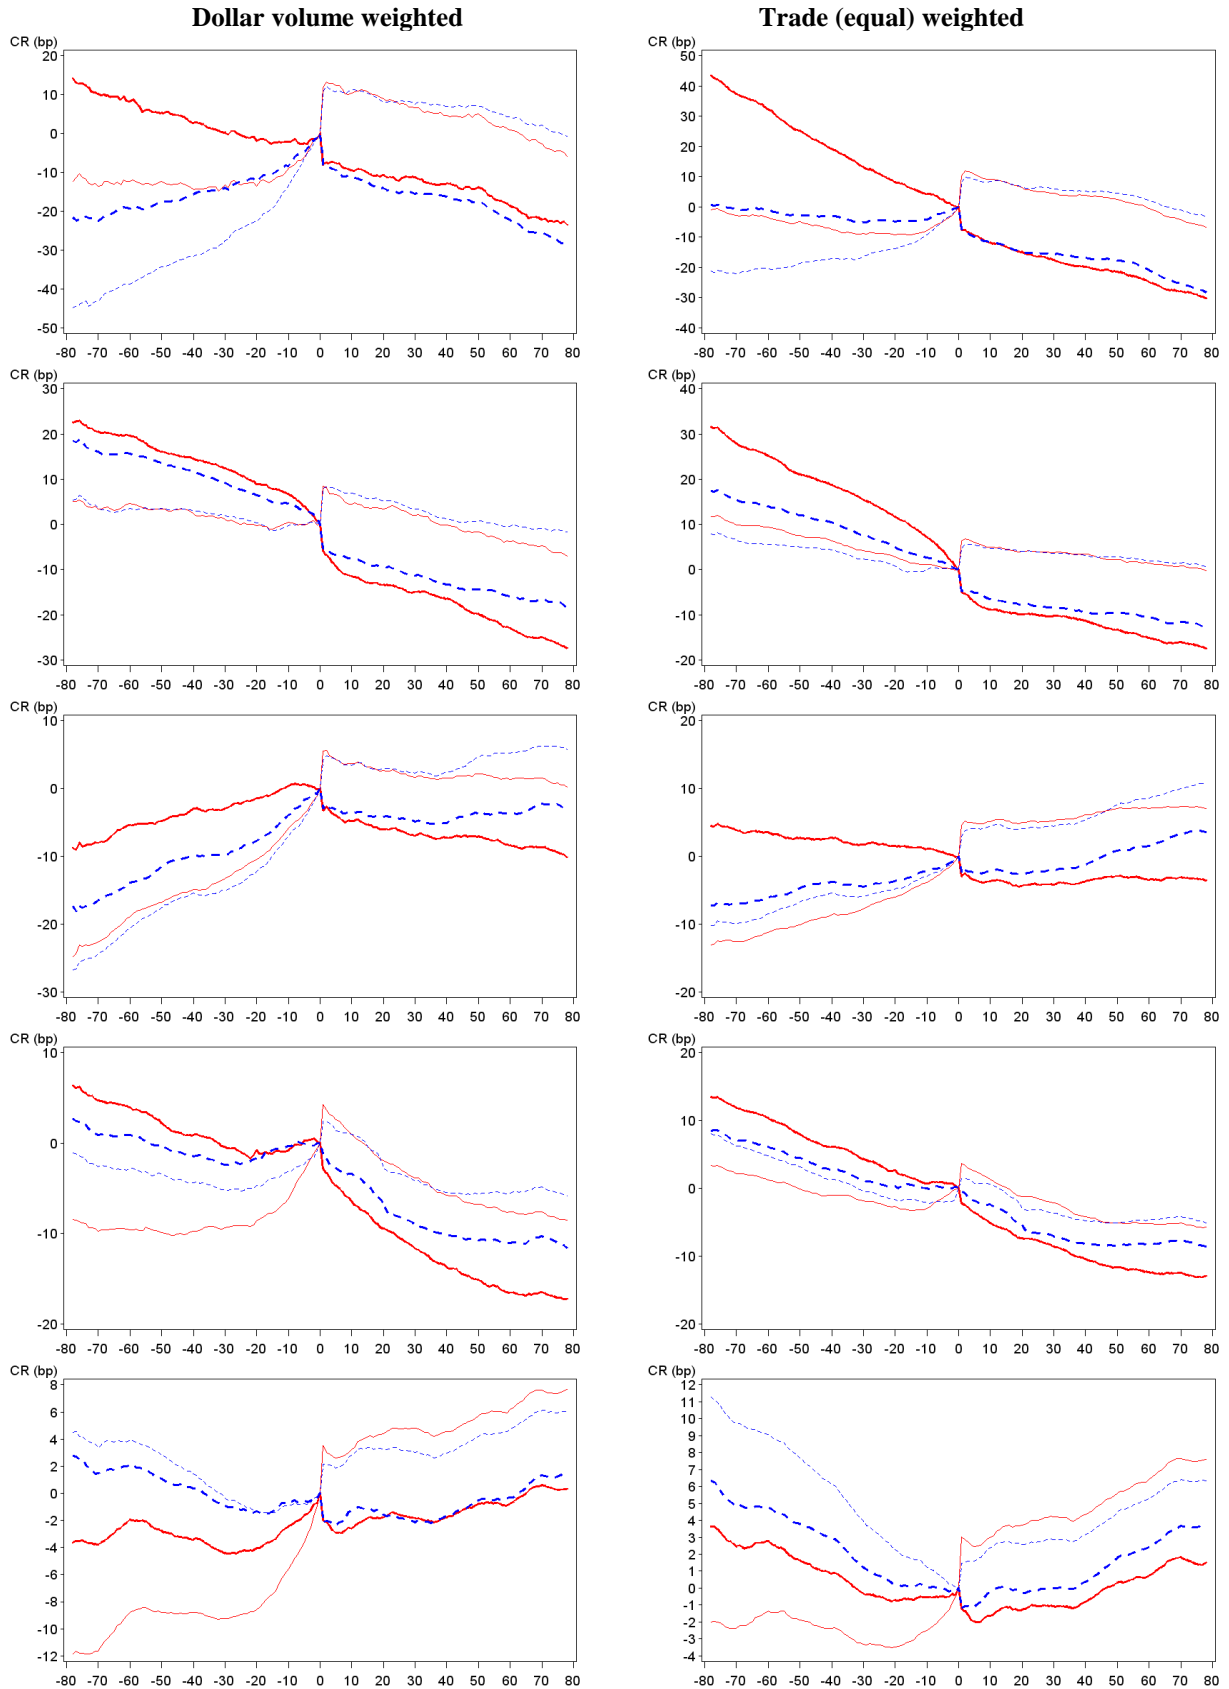
Panel A: Dollar volume weighted



Panel B: Trade (equal) weighted



Panel C. Results for each market-cap quintile, with the smallest quintile at the top.



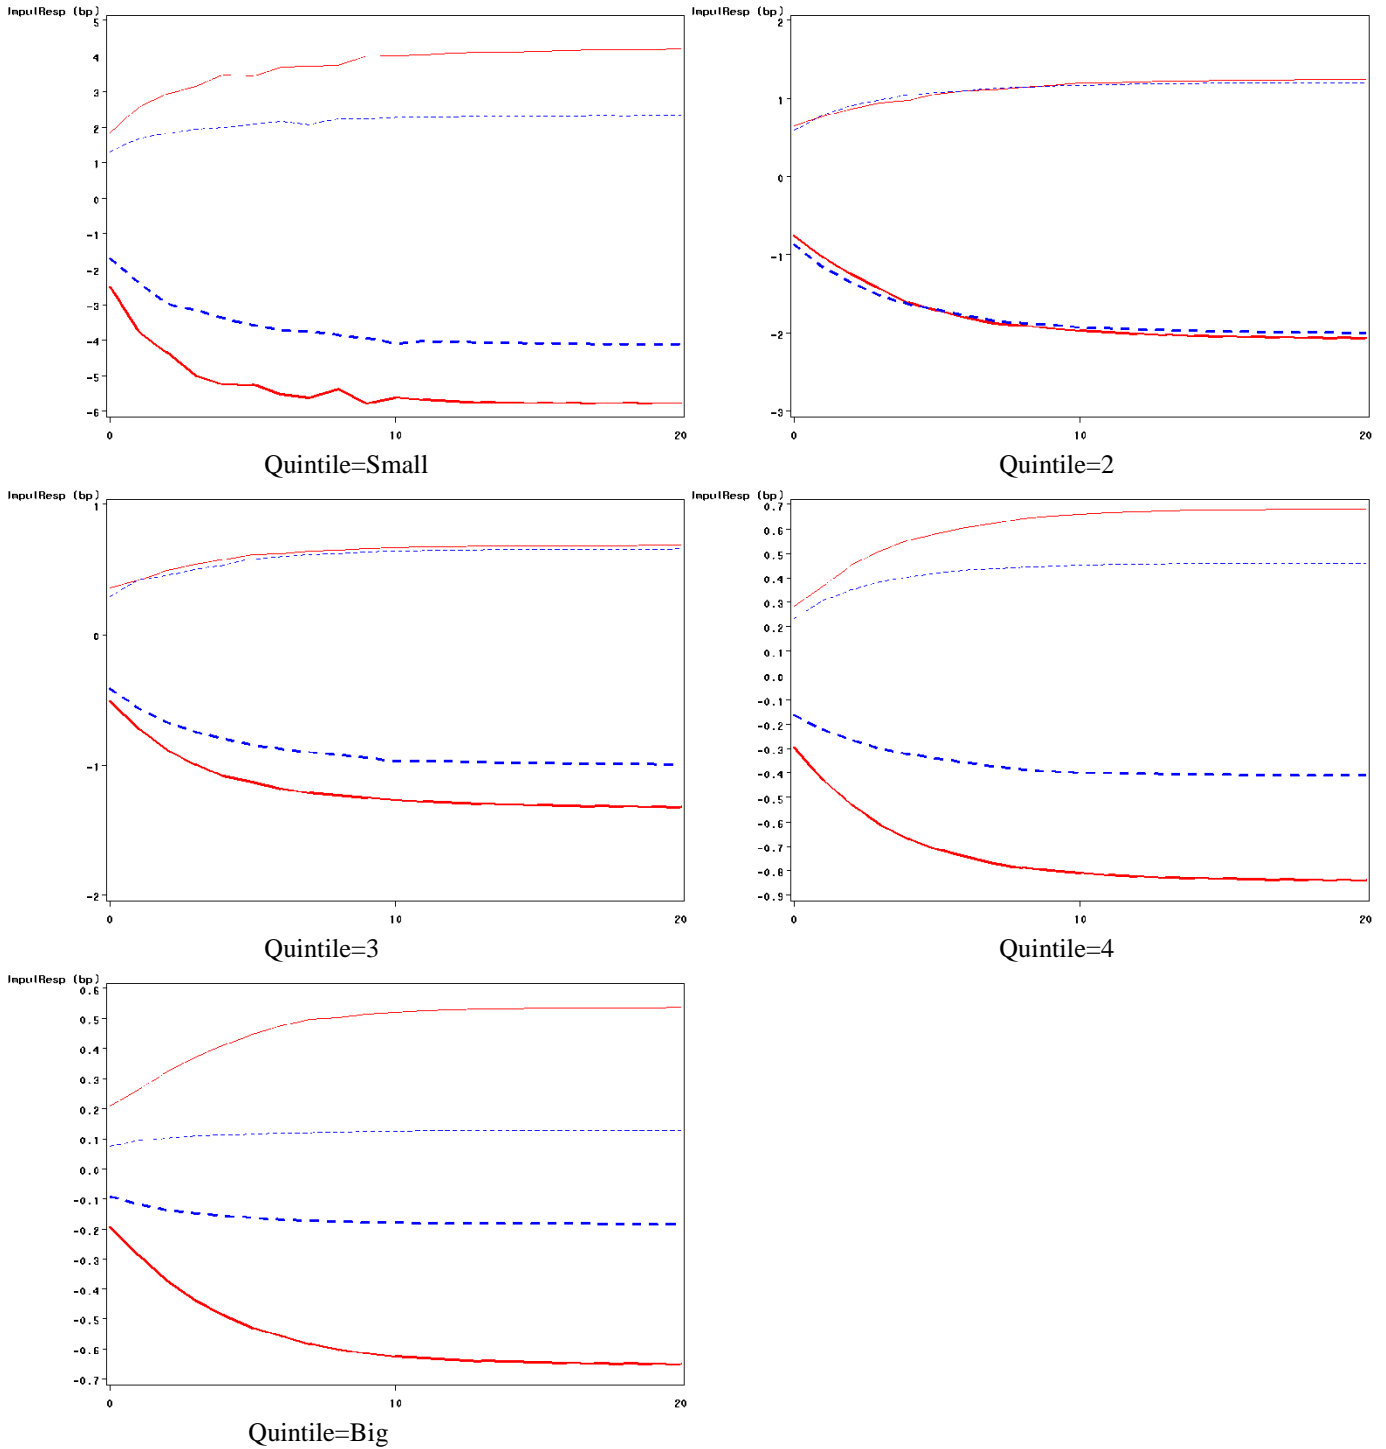


Fig. 4. Cumulative return response to volume shocks, estimated from a VAR model. Solid red (dashed blue) lines represent shocks to short (long) volume, with heavyweight (lightweight) lines for seller initiated (buyer initiated) volume. The horizontal axis measures the number of trades from the shock (which takes place at $t = 0$). The VAR coefficients and impulse response are estimated separately for each stock and then the estimates averaged across stocks within size quintiles (measured by market capitalization). The magnitude of each volume shock is equal to the standard deviation of unanticipated buyer-initiated long volume. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.




Fig. 5. Cumulative dollar volume response to same volume type shock, estimated from a VAR model. Solid red (dashed blue) lines represent shocks to short (long) volume, with heavyweight (lightweight) lines for seller initiated (buyer initiated) volume. The horizontal axis measures the number of trades from the shock (which takes place at $t = 0$). The VAR coefficients and impulse response are estimated separately for each stock and then the estimates averaged across stocks within size quintiles (measured by market capitalization). The magnitude of each volume shock is equal to the standard deviation of unanticipated buyer-initiated long volume. The sample comprises 350 NYSE-listed stocks during Jan-Aug 2008.