



Contents lists available at ScienceDirect

Journal of Financial Economics

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

Hidden liquidity: An analysis of order exposure strategies in electronic stock markets [☆]

Hendrik Bessembinder ^{a,*}, Marios Panayides ^a, Kumar Venkataraman ^b

^a David Eccles School of Business, University of Utah, Salt Lake City, UT 84112, USA

^b Cox School of Business, Southern Methodist University, Dallas, TX 75275, USA

ARTICLE INFO

Article history:

Received 21 April 2008

Received in revised form

16 December 2008

Accepted 19 February 2009

Available online 13 August 2009

JEL classification:

G11

G12

G14

G24

Keywords:

Limit order market

Trading strategies

Hidden liquidity

Dark pools

Iceberg orders

ABSTRACT

Many stock exchanges choose to reduce market transparency by allowing traders to hide some or all of their order size. We study the costs and benefits of order exposure and test hypotheses regarding hidden order usage using a sample of Euronext-Paris stocks, where hidden orders represent 44% of the sample order volume. Our results support the hypothesis that hidden orders are associated with a decreased probability of full execution and increased average time to completion, and fail to support the alternate hypothesis that order exposure causes defensive traders to withdraw from the market. However, exposing rather than hiding order size increases average execution costs. We assess the extent to which non-displayed size is truly hidden and document that the presence and magnitude of hidden orders can be predicted to a significant, but imperfect, degree based on observable order attributes, firm characteristics, and market conditions. Overall, the results indicate that the option to hide order size is valuable, in particular, to patient traders.

© 2009 Elsevier B.V. All rights reserved.

[☆] We thank Alex Butler, Robert Battalio, Charles Jones, Chitru Fernando, Amar Gande, Pengjie Gao, Michael Goldstein, Rick Harris, Joel Hasbrouck, Ohad Kadan, Qin Lei, Maureen O'Hara, Rajesh Narayanan, Michael Pagano, Venkatesh Panchapagesan, Paul Schultz, George Sofianos, Rex Thompson, Sunil Wahal, Andrew Waisburd, Jennifer Warren, Kelsey Wei, Avi Wohl, Harold Zhang, and, in particular, Avanidhar Subrahmanyam (the referee), as well as seminar participants at the American Economic Association Conference, Financial Intermediation Research Society Conference, National Bureau of Economic Research Market Microstructure Conference, Western Finance Association Conference, Brigham Young University, Goldman Sachs Algorithmic Trading Group, University of Central Florida, University of Cyprus, University of Delaware, Notre Dame University, University of Texas at Dallas, University of Oklahoma, Ohio University, Southern Methodist University, and Texas Tech University for valuable comments and discussion, Machiko Hollifield for programming assistance, and Marianne Demarchi of Euronext-Paris for assistance in obtaining the data.

* Corresponding author.

E-mail address: finhb@business.utah.edu (H. Bessembinder).

1. Introduction

Electronic limit order markets, which automatically execute traders' orders on the basis of specified priority rules, account for a large and increasing percentage of global financial and commodity trading.¹ As a consequence, understanding the effect of electronic market design features on traders' order submission decisions and market quality is becoming increasingly important to investors, as well as to those who regulate and design automated markets. While electronic limit order markets are generally designed to be transparent, a number of

¹ In a study on stock exchanges around the world, Jain (2005) reports that electronic trading is the leading stock market structure in 101 of the 120 countries that the study investigates. Furthermore, of these 101 exchanges, 85 are fully electronic, with no floor trading.

markets—including the Toronto Stock Exchange, Euronext, the Swiss Stock Exchange, the Madrid Stock Exchange, the Australian Stock Exchange (ASX), various Electronic Communications Networks (ECNs) in the US, and, most recently, the New York Stock Exchange—have elected to reduce transparency by providing traders the option to hide some or all of the size of their limit order.²

Despite the growth in electronic stock trading around the world, there have been few studies that examine traders' order submission strategies in electronic markets, particularly when they have an option to hide their orders. In this paper, we use data drawn from the Euronext-Paris market to investigate how traders use the option to hide order size and to study the various costs and benefits of order exposure. Specifically, we quantify the effects of order exposure on the likelihood that the order executes in full, on the expected time-to-completion for the order, and on order execution costs as measured by the implementation shortfall method. We also study the empirical determinants of traders' decision to hide their order size and assess whether market participants can use observable information to detect the presence and the magnitude of hidden size.

Traders considering whether to expose the full size of their orders face both costs and benefits of doing so. To encourage order exposure, many stock exchanges specify that marketable orders will execute against hidden quantities only after exhausting all the displayed size at the same price. Thus, hidden orders maintain price priority, but lose time priority to displayed orders at the same price. Exposing an order increases the chance that it attracts a counterparty who is sufficiently interested in trading to monitor the market, but who has not yet revealed herself ["reactive" traders, as per Harris (1996)]. On the other hand, exposing an order could cause other traders to withdraw liquidity or trade ahead of the order if they infer that the limit order submitter may have access to private information regarding security value. The inferences drawn by market participants regarding the information content of an order are likely to be related to the price aggressiveness of the order, since informed traders have an interest in assuming positions before their private information becomes public. Therefore, we anticipate that the decision to expose order size depends in part on the limit price selected, so that limit price and order exposure are optimally selected simultaneously.

We study a sample of 100 stocks that are traded on Euronext-Paris during April 2003. The sample includes a broad cross section of stocks ranging from the most actively traded to illiquid stocks that trade less than once per day on average. Hidden orders are used extensively on Euronext. For the full sample, 18% of the incoming orders include some hidden size and 44% of the order volume are

hidden. The usage of hidden orders is more prevalent for less liquid firms, increasing from 30% of order volume for firms in the most liquid quintile to around 50% for firms in the less liquid quintiles. For larger orders, the usage of hidden orders is increasing from 5% of order volume for order sizes less than €5,000 to over 70% for order sizes greater than €50,000.

Since exposed orders gain time priority versus hidden orders at the same price and could be more effective in drawing trading interest from reactive traders, we test the hypothesis that exposing an order increases the likelihood of full execution and reduces the delay between order submission and execution. We find that fully displayed orders are indeed more likely to execute completely and that, after controlling for other order attributes, exposed orders are associated with shorter time-to-completion. We also assess whether execution costs, measured by the Perold (1988) implementation shortfall method, are affected by traders' decisions to hide or display orders. The results indicate that hidden orders are associated with smaller opportunity costs and lower implementation shortfall costs. Finding lower opportunity costs for hidden orders despite a decreased likelihood of full execution is consistent with the reasoning that hidden orders tend to be used by traders who do not possess information regarding future changes in the prices of securities.

Our study contributes to the small but growing literature on the determinants of order exposure. De Winne and D'Hondt (2007) also study hidden orders on Euronext Paris, documenting that the decision to hide order size is related to the prevailing market conditions such as depth in the limit order book, bid-ask spread, and time of the day, as well as to order characteristics such as price aggressiveness, the total order size, and whether the order is principal- or agency-based. They also show that the revelation of hidden depth at the best quote by a prior trade leads to aggressive subsequent orders. However, our study differs from De Winne and D'Hondt (2007) in important ways.

We study a number of issues related to order exposure that they do not. In particular, we assess the costs and benefits of order exposure. We document that, *ceteris paribus*, fully exposing an order increases the probability of complete execution and decreases expected time to execution. To our knowledge, this offers the first evidence of a tangible benefit to exposing order size in the presence of an option to hide size. However, we also find that orders that hide a portion of their total size experience smaller implementation shortfall (trading) costs, *ceteris paribus*, implying a clear cost-benefit trade off in the order exposure decision.

We also study the extent to which non-displayed size is truly hidden, by assessing the ability of market participants other than the order submitter to detect *ex ante* both the presence and the magnitude of hidden depth. In contrast, the analysis in De Winne and D'Hondt focuses on traders' reactions to the revelation by a trade execution of hidden depth, not *ex ante* detection of hidden orders. The results reveal that market participants can infer the presence and magnitude of hidden depth to a significant but imperfect degree, based on observable

² Some markets, such as US-based INET, allow limit orders to be fully hidden, while other markets, such as Euronext, require limit orders to display a minimum size, in which case the orders are also referred to as "iceberg" orders. For expositional ease, we refer to orders with any unexposed quantity as hidden. Orders with zero displayed size were introduced as part of a pilot program on the New York Stock Exchange in October 2008.

order and stock characteristics and market conditions. One interesting result is that the decision to hide order size is *negatively* related to the revealed size of the order, while the earlier literature has shown that the decision to hide size is *positively* related to unobservable total order size.

Further, De Winne and D'Hondt treat traders' choice of limit price and order exposure as independent decisions. Since theory (e.g. Buti and Rindi, 2008) implies that traders are likely to select order attributes simultaneously, we also model traders' joint decisions regarding these order attributes using a simultaneous equation framework. Some inferences differ meaningfully across specifications. For example, while the estimation without correction for endogeneity indicates that more aggressively priced orders contain more hidden shares, *ceteris paribus*, the simultaneous equation estimation does not indicate a significant effect of price aggressiveness on the number of hidden shares contained in the order.³

To assess the ability of market participants to detect hidden depth on the basis of observable market conditions and order attributes, we implement logistic and Tobit regressions. We find that incoming orders are more likely to be hidden when order arrival rates are low, when spreads are wide, when competition from the limit order book on the same side is small, and when recent transactions are large. With regard to the endogenously determined order attributes, we document that traders are more likely to hide a portion of the order when they also select a more aggressive limit order price.

We examine the out-of-sample predictive ability of the Logistic model, which assesses the presence of hidden order size, and the Tobit model, which estimates the magnitude of hidden size. To do so, we construct a pair of variables that we term the "H-Score" (for Hidden Score) and the "H-Size" (for Hidden Size) for each order. In terms of interpretation, an H-Score of 1.00 indicates that the order has the same likelihood of containing hidden size as the overall sample, and H-Scores less (greater) than one indicate lower (higher) probabilities of containing hidden size. This terminology closely follows Dechow, Ge, Larson, and Sloan (2007), who implement a logistic model and create an "F-score" to detect accounting fraud.

The results indicate substantial out-of-sample forecast power. On average, 78% of all orders containing hidden size have an H-score greater than the full sample average. Actual hidden size for the order (relative to stock-specific mean hidden size) on average increases from 0.70 for orders in the lowest quintile of H-Size to 1.6 in the highest quintile of H-Size. We also find that the model based on displayed order size is less informative than the model

based on total order size in detecting hidden size. This indicates that, while market participants can detect hidden size to a substantial extent, the option to hide size still enables large traders to partially conceal the magnitude of their order. We present a simple application on how the H-Score can inform market participants, by reporting the frequency of *Type I* (the model predicts that the order has hidden size when it does not) and *Type II* (the model fails to detect hidden size when it exists) errors.

Our findings have implications for market regulators and designers of trading systems. While the prior literature reports that increased transparency improves market quality, the substantial usage of hidden orders on Euronext indicates that the option to hide order size is valuable to many market participants. In the absence of such an option, market participants may choose alternative means to complete their transactions while concealing the magnitude of their overall trading programs, e.g., by splitting their orders across markets, by relying on upstairs markets, or by trading in "dark pools," which are non-transparent markets that are limited to certain institutional investors, as described by Abrokwhah and Sofianos (2006). The resulting fragmentation of markets could potentially lower market quality and reduce price efficiency. The New York Stock Exchange made its limit order book transparent to investors in 2002. However, prior to October 2008, investors (except for floor brokers) were not permitted to use hidden orders on the NYSE. In contrast, many US-based ECNs, which compete with the NYSE for institutional order flow, allowed hidden limit orders. In late October 2008 the NYSE announced the introduction of a pilot program allowing investors to place orders with displayed size of zero shares.⁴ Our results suggest that this reduction in transparency is likely to be effective in consolidating order flow at the NYSE.

Our findings are also useful for institutional trading desks responsible for executing block orders received from portfolio managers. By modeling the hidden dimension of liquidity for firms with varying liquidity characteristics and by relating the order exposure decision to market conditions, our results provide insights on the circumstances when liquidity is likely to be hidden and when the search for hidden liquidity is likely to be more effective.

This paper is organized as follows. Section 2 reviews the prior literature on price aggressiveness and order exposure, and presents our testable hypotheses. Section 3 describes the sample and presents descriptive statistics on hidden order usage in the Euronext-Paris market. We examine the impact of order exposure on execution time in Section 4 and on implementation shortfall costs in Section 5. In Section 6, we empirically model the determinants of the decision to hide order size and describe our development of the H-Score to detect hidden orders. Section 7 presents extensions and our main conclusions.

³ De Winne and D'Hondt focus on a sample of large (greater than median) hidden limit orders and an equal number of random non-hidden large limit orders, while our analysis applies to orders of all sizes. We find that some results are sensitive to sample composition. For example, De Winne and D'Hondt find that traders tend to hide (expose) orders when spreads are narrow (wide), which is inconsistent with their hypothesis, and during the last trading hour of the day. In contrast, we find that traders tend to hide (expose) orders when spreads are wide (narrow), reflecting higher adverse selection risk, and expose orders closer to the market close, suggesting an urgency in executing the order.

⁴ The press release is available at <http://www.nyse.com/press/1224844297395.html>.

2. Literature review

2.1. The literature on price aggressiveness

The prior literature on order submission strategies has mainly focused on the determinants of limit order price aggressiveness. Griffiths, Smith, Turnbull, and White (2000) and Rinaldo (2004) find that traders submit aggressively priced orders when there is more depth on the opposite side (at the bid price for sales and at the ask price for purchases), resulting in improved execution probabilities in these more competitive market states, which is consistent with the crowding out hypothesis formally developed by Parlour (1998). Rinaldo (2004) finds that increased recent volatility is associated with more aggressive orders, while Handa and Schwartz (1996) and Ahn, Bae, and Chan (2001) find that increased recent volatility induces more liquidity provision, consistent with the theoretical prediction in Foucault (1999).⁵

Biais, Hillion, and Spatt (1995) study the Paris Bourse (one of the three markets that subsequently merged to form Euronext) and report that traders submit limit orders rather quickly when the bid–ask spread widens or the depth declines, which they attribute to the motivational effect of time priority rules. They also find that a large fraction of the limit orders submitted are at prices at or within the quotes, which they attribute to the competition stemming from price priority rules.

2.2. The literature on hidden orders

Harris (1996, 1997) has articulated some important economic reasoning relevant to understanding hidden order usage. He observes that some traders follow a passive strategy, where they wait for other traders to indicate their interest in trading on favorable terms. The presence of these passive or “reactive” traders increases the attractiveness of publicly displaying one’s own interest in trading, in order to draw out the passive traders. Other traders, in contrast, follow what Harris terms “defensive” and “parasitic” strategies. Parasitic traders seek to exploit the existence of large buy orders by “front running” the order or by using “order matching” strategies, i.e., by posting a limit order at a price one tick more favorable than the existing order.⁶ Harris (1996) finds that traders on the Paris Bourse are more likely to display their orders when the tick size is larger, which increases the cost of order matching strategies. Or, in those cases where the display of trading interest (e.g., the posting of a large buy limit order) conveys that the limit order trader may possess positive private information regarding security value, defensive traders may react by

ceasing to submit market sell orders or canceling existing limit sell orders. Building on this reasoning, Moinas (2006) presents a theoretical model where the display of a large limit order can decrease the execution probability, as defensive traders withdraw their trading interests. Buti and Rindi (2008) present a model where traders use both limit price and order size to compete in supplying liquidity. They note that hidden orders can be used to obscure the intensity of size competition, thereby reducing incentives for other traders to undercut the hidden order.

Aitken, Berkman, and Mak (2001) study the Australian Stock Exchange, reporting that the price impact of hidden orders does not differ from that of other limit orders, and that hidden order usage is negatively related to tick size and positively related to volatility and order size.⁷ Hasbrouck and Saar (2002) study the Island ECN, documenting the extensive use of limit orders that are canceled within a few seconds of order submission. These fleeting orders are likely used by aggressive traders searching for hidden orders, which on Island need not display any size. Tuttle (2006) notes that Nasdaq market makers may hide a portion of their quotation size on Nasdaq’s SuperSOES system and that they make use of hidden quotation size in more risky stocks. Bessembinder and Venkataraman (2004) study the Paris Bourse, and find that the implied transaction costs for block-sized marketable orders walking up the limit order book were on average only half as large when hidden orders were considered, as compared with costs that would have been incurred had the limit order book contained only the displayed liquidity. Pardo and Pascual (2003) examine stocks traded on the Madrid Stock Exchange, documenting that spreads do not widen and depth does not shrink after hidden order executions. Anand and Weaver (2004) examine the abolition in 1996 and reintroduction in 2002 of hidden orders on the Toronto Stock Exchange. They find that the size of the publicly displayed orders at the inside quote did not change after either event, implying that total order size decreased when orders could not be hidden.

2.3. Our contributions and testable predictions

Our paper is related to the literature on the determinants of order submission strategies. It is distinguished from the existing literature in part because the order exposure decision has been relatively unstudied but also because (a) we accommodate in some specifications the fact that order price aggressiveness and order exposure decisions are made simultaneously, (b) we relate order submission strategies to market conditions and observable order attributes, and (c) we examine whether market participants can detect hidden orders, which is informative about the extent to which hidden orders allow traders

⁵ Other studies, such as Harris and Hasbrouck (1996), Chakravarty and Holden (1995), Bae, Jang, and Park (2003), Anand, Chakravarty, and Martell (2005), and Ellul, Holden, Jain, and Jennings (2007), examine the traders’ choice of market versus limit orders.

⁶ The order matching strategy relies on the fact that if the buy limit order is executed the quote matcher will capture any upward movement in prices, while if prices fall she can sell to the party that posted the original buy limit order and lose only one tick.

⁷ Hidden orders on the ASX are displayed to the public as having size “U” (for undisplayed). Hence, market participants can identify with certainty the presence of all hidden orders on this market. In contrast, orders that include a hidden quantity are not labeled as such on most other markets that allow them.

to conceal trading intentions. Further, we examine the effects of order exposure on several dimensions of execution quality, including the likelihood of achieving full execution, the expected time-to-completion and the implementation shortfall cost.

Limit order traders can draw reactive traders either by posting an attractive price or by exposing order size. The two methods of attracting reactive traders differ in their relative costs and benefits. A more aggressive order gains price priority over orders at inferior prices, while a fully exposed order gains time priority versus hidden orders at the same price. Further, the relative costs and benefits are likely to depend on other order attributes. The models presented by Easley and O'Hara (1987) and Moinas (2006) imply that, other things equal, informed traders are likely to submit larger and more aggressive orders, because they typically have an interest in assuming large positions before their information becomes public. Large, aggressively priced orders are therefore likely to be perceived as originating from informed traders, which can cause defensive traders to exit the market or parasitic traders to indulge in front running strategies. A limit order trader may be able to counteract this effect by hiding a portion of their trading interest. For these reasons, Moinas (2006) notes that exposing a large limit order may decrease the probability of order execution and increase the time-to-completion. On the other hand, the arguments in Harris (1996) suggest that order exposure would likely attract trading interests from reactive traders, thus increasing probability of full execution and decreasing the time-to-completion.

Liquidity demanders' response to order exposure is related, at least in part, to perceptions on who (informed or uninformed traders) exposed their orders. Moinas (2006) presents a theoretical framework where an informed trader with long-lived private information on security value chooses to supply liquidity using large limit orders with hidden depth. In her model, these informed traders use a camouflage strategy that attempts to mimic the behavior of uninformed liquidity suppliers. In contrast, Harris (1996) argues that uninformed traders use hidden orders to mitigate the option value of limit orders that are expected to remain standing on the book. Since aggressively priced orders provide more valuable options to other traders, such orders are more likely to be hidden. We assess the extent to which traders who hide size are informed or uninformed based on the price movements observed subsequent to the order submission. If traders motivated by information submit hidden orders, we expect stock prices to move further (rise for buy orders and fall for sell orders) after order submission.

These discussions support the following testable hypothesis:

Hypothesis IA. Order exposure increases execution probability and decreases time-to-completion (Harris, 1996).

Hypothesis IB. Order exposure decreases execution probability and increases time-to-completion (Moinas, 2006).

Hypothesis IIA. Hidden orders are used primarily by (uninformed) traders to mitigate the option value of a standing limit order (Harris, 1996).

Hypothesis IIB. Hidden orders are used primarily by (informed) traders to protect themselves against defensive trading strategies (Moinas, 2006).

Hypothesis III. Hidden order usage is more for larger orders, aggressively priced orders and when adverse selection risk is high.

3. Sample selection and descriptive statistics on order exposure

3.1. Institutional features

Order precedence rules on Euronext are price, exposure, and time. Specifically, an aggressively priced incoming buy (sell) order will first exhaust the depth on the best offer (bid) and walk up (down) the book. At any price, the hidden portion is filled only after an incoming order has exhausted all displayed size at the price, including orders that have arrived after the hidden order was submitted. When the displayed size of a hidden order is filled, the system immediately replenishes in full the initial disclosed quantity specified during order submission and positions the order behind displayed quantities at the same price. Thus, exposed orders gain time priority over hidden orders at the same price. While some markets, such as US-based INET, allow limit orders to be fully hidden, Euronext requires that each order must display at least ten shares. Hidden orders on Euronext are therefore what have been described as "iceberg" orders; a portion of the order size is exposed, while a portion unobservable to other traders remains hidden.

3.2. Sample selection

We examine traders' order submission strategies for a broad cross section of Euronext-Paris firms, using the Base de Données de Marche (BDM) database from April 2003. The BDM database contains information on all orders for all stocks, including the firm symbol; the date and time of order submission; a buy or sell indicator; the total size of the order (in shares); the displayed size (in shares); an order type indicator for identifying market, open or limit orders; a limit price in the case of a limit order; and instructions on when the order expires.⁸ In addition, each order contains fields that allow tracking of any modifications made to the order prior to the expiration, except that the dataset does not record cancellations.⁹

⁸ Until April 23, 2001, brokerage firms on Euronext-Paris could observe the identification codes for broker-dealers submitting limit orders. Since then, the limit order book is anonymous. Foucault, Moinas, and Theissen (2007) find that concealing liquidity suppliers' identities can help improve market liquidity and price efficiency.

⁹ The database contains fields that allow us to track any modifications made to the order (typically order size and limit price) with complete accuracy. Canceled orders can be identified as of the end of the

Our initial sample consists of all stocks that are listed on Euronext-Paris ($N = 1,109$) in the BDM database in April 2003. We retain common stocks that have listed “France” as the home country, as prior research shows that home country stocks exhibit trading patterns that differ significantly from cross-listed stocks. We eliminate the less-liquid stocks on Euronext that trade in a call auction and focus on stocks traded continuously, so that the analysis can capture the decision to make or take liquidity at the time of order submission.¹⁰ Prior research also suggests that initial public offerings (IPOs) exhibit unusual trading patterns in the initial months after listing, partly reflecting the market making activity of the underwriting syndicate. We therefore eliminate stocks that appear for the first time in the BDM database after December 2002. We also eliminate stocks that switched from continuous trading to call auctions (or vice versa) or were delisted from the exchange in 2003. These screens reduce the sample size to 320 firms.

We sort sample firms into liquidity quintiles based on the number of trades in the BDM database during January 2003, which precedes our main April 2003 study period. We then randomly select 20 firms from each liquidity quintile, resulting in a final sample size of 100 firms. Panel A of Table 1 presents summary statistics for the sample, and Panel B presents the statistics by liquidity quintiles and order size. For the full sample, the mean (median) stock price and market capitalization in April 2003 are €54 (€43) and €2,990 million (€386 million), respectively. The market activity (measured as the number of monthly trades, quote updates, incoming orders, or cumulative trading volume) exhibits wide variation across sample firms. The average firm in the sample reported 4,920 trades, 6,475 quote updates, 20,840 order submissions, and a cumulative trading volume of 3.5 million shares during April 2003.

3.3. Univariate analysis of firm liquidity, order size, and order exposure

Panel B.1 of Table 1 presents the average across sample firms of the percentage of orders that were submitted with some hidden size. For the full sample, 18% of orders include non-zero hidden size. The usage of hidden orders is more prevalent for less liquid firms, increasing from 9% of orders for firms in the most liquid quintile to over 20% of orders for firms in the less liquid quintiles. This pattern in hidden order usage could reflect the longer expected waiting time until execution for limit orders in less liquid

firms, due to lower order arrival rates. A strong monotonic relation exists between hidden order usage and total order size. For the full sample, only 1% of orders with size less than €1,000 have a hidden size. In contrast, over 75% of orders with size greater than €50,000 have a hidden component. Controlling for order size, hidden orders are used more frequently in less liquid firms.

Panel B.2 of Table 1 presents statistics on the percentage of order volume that is hidden. Remarkably, 44% of the incoming sample order flow in shares are hidden. The percentage of order volume that is hidden increases from 30% for firms in the most liquid group to approximately 50% for firms in the less liquid groups. The greater use of hidden orders for less liquid firms highlights the desirability of studying a broad cross section of firms, as opposed to focusing exclusively on the more liquid CAC-40 firms, as in De Winne and D'Hondt (2007). Hidden order volume increases with order size, with approximately 70% of share volume in orders greater than €250,000 being hidden.

Panel B.3 of Table 1 presents statistics on hidden volume for those orders that include some hidden size. For the full sample, the percentage of order volume that is hidden, conditional on some hidden size, is 75%. Consistent with earlier results, the percentage of hidden volume is higher for larger orders. However, the percentage of hidden order volume, conditional on a hidden size, does not differ significantly across liquidity groups.

3.4. Univariate analysis of price aggressiveness, order size, and order exposure

We follow Biais, Hillion, and Spatt (1995) in defining categories of price aggressiveness. To do so, we reconstruct from the BDM data estimates of the limit order book, including liquidity that is publicly displayed and liquidity that is hidden, at the time of each order submission. Our reconstruction of the limit order book closely follows the approach described in Appendix B of Bessembinder and Venkataraman (2004).

The first four categories represent orders that demand liquidity from the book and the last three categories represent orders that supply liquidity to the book. The most aggressive orders (category 1) represents buy (sell) orders with order size greater than those displayed in the inside ask (bid) and with instructions to walk up (down) the book until the order is fully executed. Category 2 represents buy (sell) orders with order size greater than those displayed in the inside ask (bid) and with instructions to walk up (down) the book, but the order specifies a limit price such that the order is not expected to execute fully based on displayed book. Category 3 represents buy (sell) orders with the limit price equal to the inside ask (bid) and with order sizes greater than those displayed in the inside ask (bid). Orders in categories 2 and 3 may execute fully due to hidden liquidity but may also clear the book and convert to a standing limit order. Category 4 represents buy (sell) orders with the limit price equal to the inside ask (bid) and with order size less than those displayed in the inside ask (bid). These orders are

(footnote continued)

day with complete accuracy, but cannot always be identified intraday. We are able in many instances to infer the exact time when an order has been canceled, based on quote updates that do not reflect completed trades or order modifications, as in Bessembinder and Venkataraman (2004). Since the database identifies the cancellation date, any minor errors in the reconstructed limit order book attributable to undetected order cancellations do not accumulate across trading days.

¹⁰ For the same reason, we examine only orders that arrive during regular trading hours, thereby excluding orders submitted for the opening and closing batch auction. However, the limit order book is constructed using all orders submitted for the stock.

Table 1

Sample descriptive statistics and hidden order usage, by firm liquidity and order size.

Reported in Panel A are the market capitalization, stock price, daily return volatility, monthly trading activity, trade size, and order size in April 2003 for 100 Euronext-Paris firms. The data are obtained from the Base de Donnees de Marche database made available by Euronext-Paris. Euronext stocks that trade in the continuous market are sorted into liquidity quintiles based on trading activity in January 2003. We randomly select 20 firms from each liquidity quintile, resulting in the final sample of 100 firms. Panel B reports on hidden order usage for sample firms by liquidity quintiles and order size. The relevant statistic is calculated for each firm during April 2003, and the table reports the (cross-sectional) average across sample firms. Panel B.1 presents statistics on the percentage of orders that are submitted with a hidden size; Panel B.2 on the percentage of order volume that is hidden; and Panel B.3 on percentage of order volume that is hidden for orders with a hidden size.

	Number of firms	Mean	Median	Standard deviation	Maximum	Minimum
Panel A: Descriptive statistics based on firm averages, full sample						
Average stock price (euros)	100	54	43	48	235	1
Market capitalization (millions of euros)	100	2,990	386	7,821	65,121	3
Number of monthly trades	100	4,920	325	10,137	44,267	12
Number of monthly quote updates	100	6,475	379	13,253	58,309	15
Number of monthly orders	100	20,840	1,273	42,312	210,444	28
Cumulative monthly trading volume (shares)	100	3,512,852	54,619	11,394,139	98,362,569	723
Daily return volatility (percent)	100	2.67	2.43	2.24	21.45	0.68
Average trade size (shares)	100	397	204	652	4,323	20
Average order size (shares)	100	676	400	883	5,821	26
Panel B: Hidden orders by liquidity group and order size						
		By order size (euros)				
	All orders	Less than 1,000	1,000–5,000	5,000–50,000	50,000–250,000	Greater than 250,000
Panel B.1: Percentage of orders with a hidden size (based on firm average)						
Full sample	18	1	5	34	75	76
Least liquid quintile	21	1	6	46	87	80
Quintile 2	23	2	10	44	87	92
Quintile 3	21	1	6	46	88	75
Quintile 4	15	0	2	27	81	80
Most liquid quintile	9	0	1	7	43	69
Panel B.2: Percentage of order volume that is hidden (based on firm average)						
Full sample	44	1	4	35	69	72
Least liquid quintile	45	0	5	48	82	73
Quintile 2	48	1	7	43	79	90
Quintile 3	53	1	5	46	80	74
Quintile 4	43	0	2	29	74	78
Most liquid quintile	30	0	0	7	39	62
Panel B.3: Conditional on a hidden size, the percentage of order volume that is hidden						
Full sample	75	15	46	71	87	90
Least liquid quintile	79	9	37	79	92	92
Quintile 2	74	23	49	70	92	98
Quintile 3	75	9	49	74	90	88
Quintile 4	75	15	46	71	88	89
Most liquid quintile	72	20	49	62	78	90

expected to immediately execute in full. Category 5 represents orders with limit prices that lie within the inside bid and ask prices. Category 6 represents buy (sell) orders with limit price equal to the inside bid (ask). Finally, category 7 represents buy (sell) orders with limit price less (greater) than the inside bid (ask).¹¹

Panel A of Table 2 reports statistics on the percentages of orders with hidden size and the percentage of order volume that is hidden, by price aggressiveness groups. Orders that are expected to execute fully based on

displayed depth, categories 1 and 4, are least likely to hide order size, as only 1% of orders in category 4 and 7% of orders in category 1 include hidden depth. In contrast, traders are more likely to hide orders and also hide larger order volume for orders that are likely to remain standing in the book. For orders that are not expected to execute immediately (categories 5, 6 and 7), almost 20% of orders and 50% of order volume are hidden. Similarly, orders that are expected to be remain standing in the book after partial execution, categories 2 and 3, are more likely to be hidden. Panel A.3 reports similar patterns for the subset of orders that include a hidden size. These patterns support the reasoning that traders use hidden orders to reduce the option value of standing limit orders in the book.

¹¹ Biais, Hillion, and Spatt (1995) define six categories of orders, as they combine categories 1 and 2 into a single category. Our definitions are consistent with Biais, Hillion, and Spatt for the other categories.

Table 2

Hidden order usage and empirical probabilities of full execution, by price aggressiveness and order size.

Panel A presents descriptive statistics on hidden order usage in April 2003 by price aggressiveness and order size groups. The data are obtained from the Base de Donnees de Marche database made available by Euronext–Paris. The relevant statistic is calculated for each firm during April 2003, and the table reports the (cross-sectional) average across sample firms. Panel A.1 presents statistics on the percentage of orders submitted with a hidden size; Panel A.2 on the percentage of order volume that is hidden; and Panel A.3 on percentage of order volume that is hidden for orders with a hidden size.

The most aggressive category (category 1) represents buy (sell) orders with order size greater than those displayed in the inside ask (bid) and with instructions to walk up (down) the book until the order is fully executed. Category 2 represents buy (sell) orders with order size greater than those displayed in the inside ask (bid) and with instructions to walk up (down) the book, but the order specifies a limit price such that the order is not expected to execute fully based on displayed book. Category 3 represents buy (sell) orders with the limit price equal to the inside ask (bid) and with order sizes greater than those displayed in the inside ask (bid). Category 4 represents buy (sell) orders with the limit price equal to the inside ask (bid) and with order size less than those displayed in the inside ask (bid). Category 5 represents orders with limit prices that lie within the inside bid and ask prices. Category 6 represents buy (sell) orders with limit price equal to the inside bid (ask). Category 7 represents buy (sell) orders with limit price less (greater) than the inside bid (ask).

Panel B presents the empirical probabilities of full execution for orders characterized by whether a portion of order size is hidden. Orders are further classified by price aggressiveness and order size. The empirical probabilities of full execution are defined as the ratio of the number of orders that are completely executed over the total number of orders submitted. The ratio is calculated separately for each type of order (displayed, hidden), by price aggressiveness and order exposure. 'NA' denotes categories with no corresponding observations in the data. ***, **, and * denote statistical significance of difference across order types at the 1%, 5% and 10% level, respectively.

Panel A: Hidden order usage, by price aggressiveness and order size						
Variable	All orders	By order size (euros)				
		Less than 1,000	1,000–5,000	5,000–50,000	50,000–250,000	Greater than 250,000
<i>Panel A.1: Percentage of orders with a hidden size (based on firm average)</i>						
Most aggressive	7	2	2	8	17	43
Category 2	18	0	2	15	30	47
Category 3	13	1	4	15	37	63
Category 4	1	0	1	3	13	10
Category 5	19	1	6	41	80	80
Category 6	26	0	8	47	83	88
Least aggressive	21	1	4	35	84	84
<i>Panel A.2: Percentage of order volume that is hidden (based on firm average)</i>						
Most aggressive	15	0	1	7	15	41
Category 2	25	0	1	14	30	44
Category 3	25	0	3	15	35	61
Category 4	2	0	1	3	10	11
Category 5	48	1	5	41	74	76
Category 6	50	0	5	43	74	84
Least aggressive	45	1	3	35	76	78
<i>Panel A.3: Conditional on a hidden size, the percentage of order volume that is hidden</i>						
Most aggressive	43	0	12	40	58	77
Category 2	61	0	3	51	67	73
Category 3	67	2	26	60	74	83
Category 4	33	6	20	32	54	48
Category 5	75	11	48	72	89	91
Category 6	68	6	32	67	86	94
Least aggressive	72	5	34	68	86	90
Panel B: Execution probability of hidden and non-hidden orders, by price aggressiveness and order size						
Most aggressive						
Non-hidden	98***	100	88	82	95***	100***
Hidden	84***	NA	100	96	74***	78***
Category 2						
Non-hidden	83	75	81	75	87	93
Hidden	72	NA	NA	59	79	76
Category 3						
Non-hidden	85***	92	90***	82***	80***	71*
Hidden	44***	100	67***	48***	48***	48*
Category 4						
Non-hidden	62	64	72	60	72	27
Hidden	60	33	78	73	71	38
Category 5						
Non-hidden	40***	47	43	35*	37*	43**
Hidden	26***	47	44	27*	25*	19*

Table 2. (continued)

Category 6						
Non-hidden	40 ^{***}	46	44	36 ^{**}	42 ^{***}	36 ^{**}
Hidden	22 ^{***}	67	46	25 ^{**}	20 ^{***}	16 ^{**}
Least aggressive						
Non-hidden	18 [*]	24 [*]	20 [*]	17	19	10
Hidden	12 [*]	41 [*]	30 [*]	14	12	11

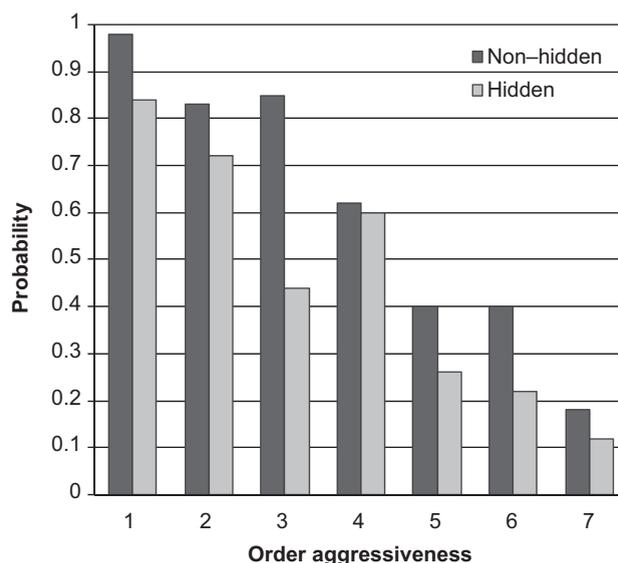


Fig. 1. Probability of full execution. This figure shows the empirical probabilities of full execution for orders characterized by whether a portion of the order size is hidden. Orders are further classified by price aggressiveness. For each order, price aggressiveness is defined as a discrete value between 1 and 7 by comparing the order's limit price with the price of the opposite quote at the time of submission, similar to Biais, Hillion, and Spatt (1995). The first four categories represent orders that demand liquidity (categories 1 to 4) from the book and the last three categories represent orders that supply liquidity to the book (categories 5 to 7). The empirical probabilities of full execution are defined as the ratio of the number of orders that are fully executed over the total number of orders submitted. The ratio is calculated separately for each type of order (displayed, hidden) and each price aggressiveness categories 1 to 7.

4. Order exposure and execution time

While exposing an order could cause other traders to withdraw liquidity or employ front-running strategies (*Hypothesis 1B*), exposed orders gain time priority versus hidden orders at the same price and may be effective in drawing trading interest from passive traders (*Hypothesis 1A*). The latter reasoning suggests that exposing an order can increase the likelihood of order execution. In Fig. 1, we report the empirical probability of full execution for orders, depending on whether order size is hidden or not, grouped by price aggressiveness. Panel B of Table 2 presents a more detailed analysis that is grouped by price aggressiveness and order size. Within each category of price aggressiveness, fully displayed orders have higher execution probabilities than hidden orders. The difference is statistically significant in all except aggressiveness categories 2 and 4. Further, as reported in Panel B of Table 2, the higher execution probabilities for exposed orders is observed for almost all order size and price aggressiveness categories. These results indicate that order exposure is effective in drawing trading interest from reactive traders.

Hypothesis 1A predicts that exposing an order reduces the elapsed time from order submission to execution, after

controlling for the effects of order size and price aggressiveness. To test this reasoning formally, we build a model of limit order time-to-completion using survival analysis, as described in Lo, MacKinlay, and Zhang (2002). Explanatory variables identical to Lo, MacKinlay and Zhang include the proportional difference between the midpoint and the limit price as a measure of price aggressiveness; a buy indicator variable that equals one if the prior trade is buyer-initiated and equals zero otherwise; trade frequency, which is the number of trades in the last hour; and relative trade frequency, which is the number of trades in the last half hour divided by the number of trades in the last hour.

Similar to Lo, MacKinlay and Zhang, we also include a set of variables that capture the state of the limit order book and proxy for competition from traders on the same and opposite side of the market. These variables, defined following the price aggressiveness literature, include same side depth, measured as the displayed depth at the best bid (ask) for a buy (sell) order (normalized); the square of the previous measure to allow for non-linearity in the relation; opposite side depth, measured as the displayed depth at the best ask (bid) for a buy (sell) order; and total (exposed plus hidden) size of the order. We supplement

these variables with a hidden order indicator variable, which equals one if the order has hidden size and equals zero otherwise.¹²

4.1. Cross-sectional aggregation

We estimate the survival model and all subsequent analyses on a firm-by-firm basis. In the interest of parsimony, we present results that are aggregated across firms. Harris and Piwowar (2006) emphasize the desirability of assigning larger weights in cross-sectional aggregation to those securities whose parameters are estimated more precisely. To do so, we assess statistical significance relying on a Bayesian framework attributable to DuMouchel (1994) and also implemented by Panayides (2007). The method assumes that, for each estimated firm i coefficient, β_i :

$$\hat{\beta}_i | \beta_i \sim \text{i.i.d. } N(\beta_i, s_i^2) \quad (1)$$

and

$$\beta_i \sim \text{i.i.d. } N(\beta, \sigma^2) \quad (2)$$

where N is the Gaussian distribution. The standard errors, s_i , are estimated by use of the Newey-West method to correct for autocorrelation and heteroskedasticity, and σ^2 is estimated by maximum likelihood.¹³ The aggregated β estimate is obtained from the N individual firm estimates as

$$\hat{\beta} = \frac{\sum_{i=1}^N \frac{\hat{\beta}_i}{(s_i^2 + \hat{\sigma}_{m.l.e}^2)}}{\sum_{i=1}^N \frac{1}{(s_i^2 + \hat{\sigma}_{m.l.e}^2)}} \quad (3)$$

Assuming independence across firms, the variance of the aggregate estimate is:

$$\text{Var}(\hat{\beta}) = \frac{1}{\sum_{i=1}^N \frac{1}{(s_i^2 + \hat{\sigma}_{m.l.e}^2)}} \quad (4)$$

where $\hat{\sigma}_{m.l.e}^2$ is the maximum likelihood estimator of σ^2 . The aggregate t -statistic is based on the aggregated coefficient estimate relative to the standard error of the aggregate estimate. This method allows for variation across stocks in the true β_i and also for cross-sectional differences in the precision with which β_i is estimated. The key feature of the aggregation method is that it places more weight on those coefficients that are estimated more precisely.

However, the method assumes independence of estimation errors across stocks, which may be incorrect. We therefore augment the approach by also implementing the

correction recommended by Chordia, Roll, and Subrahmanyam (2000, 2005, hereafter CRS), which allows for cross-correlation in the individual stock regression residuals. Assuming that the residual cross-correlation is constant across pairs of stocks, CRS show the standard error of the aggregated estimate is inflated by the factor $[1+(N-1)\rho]^{1/2}$, where N is the number of regressions and ρ is the common cross-correlation in the residuals. Since order arrival times differ across sample firms, the regression residuals are not naturally matched in time. We therefore measure the average residual for each firm over each half-hour period, and estimate ρ as the average cross-correlation (across all 4,950 unique pairs of stocks) of these average residuals. The final t -statistic is based on DuMouchel (1994) but also incorporates the CRS correction for cross-correlation:

$$t = \frac{\sum_{i=1}^N \frac{\hat{\beta}_i}{(s_i^2 + \hat{\sigma}_{m.l.e}^2)}}{\left[(1 + (N - 1)\rho) \sum_{i=1}^N \frac{1}{(s_i^2 + \hat{\sigma}_{m.l.e}^2)} \right]^{1/2}} \quad (5)$$

4.2. Empirical results regarding time-to-completion

Columns 1 and 2 of Table 3 report results of the time-to-completion model for buy and sell limit orders, respectively. The parameter estimates are generally consistent with those reported in Lo, MacKinlay and Zhang. The positive (negative) sign on the difference between the mid-quote and the limit price for buy (sell) orders indicates that the time-to-completion is longer for less aggressively priced orders. The positive estimated coefficients on book depth on the same side of the book, *same side depth* and *same side depth squared*, indicates that the time-to-completion increases when the competition for liquidity provision is high. The negative estimated coefficient on *opposite side depth* indicates that the expected time-to-completion decreases when book depth on the opposite side is deeper. A possible explanation is that an aggressively priced sell (buy) order is more likely when the book depth on the Offer (Bid) side of the book is deep, which decreases the execution time for a buy (sell) order. The positive coefficient on *order size* indicates that the time-to-completion is higher for larger orders, on both the buy and sell sides. This result can be contrasted with the puzzling lack of a relation between order size and time to completion that was reported by Lo, MacKinlay and Zhang. Although the negative coefficient on *trade frequency* indicates that both buy and sell orders execute more quickly during active market conditions, the variable is marginally significant only for buy orders.¹⁴

¹² Following Lo, MacKinlay and Zhang, we estimate the survival function assuming that the distribution of failure times follows a generalized gamma distribution and using the accelerated failure time approach. Lo, MacKinlay and Zhang implement a pooled analysis and include a number of variables to control for differences across stocks. We omit these variables because our analysis is performed on a stock-by-stock basis.

¹³ To estimate the Newey-West corrected standard errors, we use the Generalized Method of Moments (GMM) with a Bartlett kernel and a maximum lag length of 10.

¹⁴ The estimated *shape parameters* are statistically significantly different from one, the value consistent with simple distributions, indicating that the generalized gamma distribution is an appropriate assumption for the survival analysis. For robustness purposes, we have also assessed results using the simple rather than weighted average cross-sectional parameter, as in CRS (2005). Results are quite similar, and available from the authors on request.

Table 3

Order submission strategies and execution time: survival analysis.

The table reports parameter estimates of an econometric model of limit order time-to-execution using survival analysis, following Lo, MacKinlay, and Zhang (2002). The model describes an accelerated failure time specification of limit order execution times under the generalized gamma distribution for a sample of 100 Euronext stocks from April 2003. The explanatory variables are: the distance in basis points of the order's limit price from the quote midpoint (midquote—limit price); an indicator variable that equals one if the prior trade is buyer-initiated and equals zero otherwise (*last trade buy indicator*); the displayed depth at the best bid (ask) for a buy (sell) order (*same side depth*); the square of the previous measure to account for non-linearity (*same side depth squared*); the displayed depth at the best ask (bid) for a buy (sell) order (*opposite side depth*); the total (exposed plus hidden) size of the order (*order size*); the number of trades in the last half hour divided by the number of trades in the last hour (*rel. trade frequency*); the number of trades in the last hour (*trade frequency*); an indicator variable that equals one if the order has hidden size and equals zero otherwise (*hidden order*). Coefficients are estimated on a firm-by-firm basis. Reported results are aggregated across firms using the Bayesian framework of DuMouchel (1994), with *t*-statistics corrected for cross-correlation in the individual regression residuals following Chordia, Roll, and Subrahmanyam (2005).

Variable	Firm-by-firm regressions	
	Buy limit order model (1)	Sell limit order model (2)
Intercept (<i>t</i> -statistic)	10.5523 (8.81)	12.8761 (12.38)
Midquote—limit price (<i>t</i> -statistic)	3.0481 (2.86)	−0.6268 (−1.57)
Last trade buy indicator (<i>t</i> -statistic)	0.1259 (0.98)	−0.2091 (−1.46)
Same side depth (norm) (<i>t</i> -statistic)	0.0722 (2.34)	0.0226 (0.72)
Same side depth squared (<i>t</i> -statistic)	0.0295 (0.22)	0.0063 (2.06)
Opposite side depth (norm) (<i>t</i> -statistic)	−0.3278 (−3.13)	−0.3283 (−5.51)
Order size (<i>t</i> -statistic)	0.1448 (2.62)	0.1540 (2.46)
Rel. trade frequency (<i>t</i> -statistic)	0.3331 (0.28)	1.6105 (1.36)
Trade frequency (<i>t</i> -statistic)	−0.2972 (−1.75)	−0.2088 (−1.36)
Hidden order indicator (<i>t</i> -statistic)	1.8155 (2.33)	0.9678 (2.02)
Scale (fitted distribution) (<i>t</i> -statistic)	3.5802 (5.19)	1.6997 (2.95)
Shape (fitted distribution) (<i>t</i> -statistic)	0.2250 (0.21)	3.7503 (3.17)

Most important for our investigation, the results include a significantly positive coefficient estimate for *Hidden Order* for both buy limit orders (*t*-statistic = 2.33) and sell limit orders (*t*-statistic = 2.02). These results imply that, after controlling for price aggressiveness, order size, and market conditions, the choice to hide a portion of order size is associated with a longer time-to-completion and an increase in investors' price risk due to a delayed trade. Conversely, exposing size shortens the time-to-

completion by providing time priority over hidden orders at the same price and by attracting passive traders (Harris, 1996). To our best knowledge, this is the first documentation of a tangible benefit to traders of exposing order size in markets that provide the option to hide size.

The effect of hiding order size on expected time to completion is substantial. Evaluating expression 6.9 in Lo, MacKinlay, and Zhang (2002), the expected time to completion conditional on a vector of explanatory variables X_1 , relative to the expected time to completion conditional on a vector of explanatory variables X_2 , is given as $\exp(X_1 - X_2)\beta$, where β denotes the vector of coefficients obtained by estimating the survival model. The coefficient estimate of 1.82 reported on the hidden order indicator for buy limit orders in Table 3 therefore implies that the expected time-to-completion for a buy order containing hidden size is increased by a factor of $\exp(1.82) = 6.1$ times, relative to an otherwise similar order with fully exposed size. For sell limit orders, the relative increase in expected time-to-completion for orders containing hidden size is $\exp(0.97) = 2.6$ times.

5. Order exposure and implementation shortfall costs

The evidence reported in Section 4 indicates that order exposure increases the probability of full execution and reduces the anticipated time from order submission to execution. However, almost 18% of the incoming orders include a hidden size, implying that at least some market participants perceive tangible benefits to limiting order exposure. In this section, we investigate whether execution costs are affected by the trader's decision to hide or display orders.

To measure execution costs, we rely on the implementation shortfall approach proposed by Perold (1988), which incorporates not only the *price impact* on the portion of the order that is filled but also imputes a penalty, or *opportunity cost*, for any portion of the order that goes unfilled. Following Harris and Hasbrouck (1996) and Griffiths, Smith, Turnbull, and White (2000), we calculate the two components of implementation shortfall as follows. The *price impact* is the appropriately signed difference between the fill price and the quote midpoint at the time of order submission, in Euros. It is expected to be positive for orders that demand liquidity (aggressiveness groups = 1, 2, 3, and 4) and is expected to be negative for orders that supply liquidity (aggressiveness groups = 5, 6, and 7). For a passive order that goes unfilled (fill rate = 0%), the price impact is zero. For orders that are not completely filled due to an order cancellation or expiration, the *opportunity cost* is the appropriately signed difference between the closing price on the order expiration or cancellation date and the quote midpoint at the time of order submission.¹⁵ If prices move away (rise for buy orders or fall for sell orders) after order submission,

¹⁵ For NYSE SuperDot orders, Harris and Hasbrouck (1996) assume that an expired buy (sell) order is filled at the closing ask (bid) price on expiration date. Because Euronext implements a closing call auction for our sample stocks, we assume that both expired buys and sells are executed at the closing (call auction clearing) price.

Table 4

Regressions of implementation shortfall, price impact, and opportunity cost on order attributes and market conditions.

The table reports on regression coefficients of execution costs on order attributes and market conditions for a sample of Euronext-Paris stocks during April 2003. Execution costs are based on the implementation shortfall approach proposed by Perold (1988), defined as follows. For a buy order, *price impact* is defined as the difference between the filled price of each submitted order and the mid-quote price at the time of order submission. *Opportunity cost* is defined as the difference between the closing price on the day of order cancellation or expiration and the quote midpoint at the time of order submission. Each cost is regressed with respect to four variables that represent order attributes (price aggressiveness, order size, buyer order indicator, and hidden order indicator) and two variables that represent market conditions during the trading hour prior to order submission (trading frequency and return volatility). For *price impact*, we report regression results conditional on either partial or full order execution (*price impact* ≠ 0, Column 3). For *opportunity cost*, we report regression results conditional on either partial or full non-execution (*opportunity cost* ≠ 0, Column 5). The time series coefficients are estimated on a firm-by-firm basis. Reported results are aggregated across firms using the Bayesian framework of DuMouchel (1994), with *t*-statistics corrected for cross-correlation in the individual regression residuals following Chordia, Roll, and Subrahmanyam (2005).

Variable	Implementation shortfall		Price impact		Opportunity cost	
	All orders	All orders	If fill rate > 0%	All orders	If fill rate < 100%	
	Coefficient (1)	Coefficient (2)	Coefficient (3)	Coefficient (4)	Coefficient (5)	
Intercept (<i>t</i> -statistic)	−0.0083 (−0.42)	−0.0607 (−5.45)	0.0872 (6.33)	0.0505 (2.42)	0.0855 (3.27)	
Price aggressiveness (<i>t</i> -statistic)	0.8452 (4.67)	−0.0001 (−1.36)	29.1275 (7.17)	0.4713 (3.20)	0.6245 (4.48)	
Order size (million shares) (<i>t</i> -statistic)	1.98 (2.19)	3.78 (2.53)	0.43 (0.57)	0.42 (0.44)	−0.44 (−0.52)	
Buy order (dummy) (<i>t</i> -statistic)	0.1175 (2.69)	0.0170 (2.50)	0.0175 (0.98)	0.0917 (2.02)	0.1445 (2.39)	
Hidden order (dummy) (<i>t</i> -statistic)	−0.0231 (−2.62)	−0.0245 (−3.63)	0.0007 (0.63)	−0.0151 (−1.66)	−0.0360 (−2.67)	
Trading frequency (<i>t</i> -statistic)	−0.0083 (−2.40)	0.0058 (3.15)	−0.0052 (−3.56)	−0.0127 (−3.75)	−0.0182 (−4.81)	
Volatility (<i>t</i> -statistic)	0.0012 (0.81)	−0.0107 (−1.61)	0.1690 (5.61)	0.0058 (1.71)	0.0114 (1.84)	

the *opportunity cost* will be positive, reflecting the cost of delayed execution. The *opportunity cost* for a fully executed order (fill rate = 100%) is zero. The *implementation shortfall cost* for an order is the weighted sum of the *price impact* and the *opportunity cost*, where the weights are the proportion of the order size that is filled and unfilled, respectively.

Table 4 presents the coefficients obtained in the regressions of *implementation shortfall*, *price impact*, and *opportunity cost*, respectively, on order characteristics and market conditions. The coefficients are estimated for each firm and aggregated across firms using the approach described in Section 4.1. For the *price impact* measure, Column 2 presents coefficients based on all orders and Column 3 presents coefficients obtained when the sample includes only orders with either partial or full execution (that is, fill rates > 0%, *price impact* ≠ 0). Similarly, for the *opportunity cost* measure, Column 4 presents coefficients based on all orders and Column 5 presents coefficients estimated for orders with either partial or full non-execution (that is, fill rates < 100%, *opportunity cost* ≠ 0). Note that the measures of *price impact* conditional on execution are effectively measures of the aggressiveness of the order's limit price.

As expected, *price impact* is larger for aggressively priced orders, conditioned on order execution. Focusing on Column 2, *price impact* is greater for large orders, for buy orders, and for orders submitted when markets are

more active. However, each of these results can be attributed to variation in execution rates. Coefficient estimates in Column 3 indicate that, conditional on execution, order size, and order direction do not affect the *price impact*. *Price impact* increases with recent volatility and declines with market activity, conditional on execution.

Focusing on Columns 4 and 5, *opportunity costs* are higher for more aggressive orders and for buy orders, and they are lower for orders submitted when markets are more active. To the extent that the information that motivates informed traders becomes public before the close of trading, these results suggest that aggressively priced orders and buy orders tend to be placed by informed traders.

We are most interested in coefficients estimated on the hidden size indicator. The coefficient estimates in Columns 2 and 4 indicate lower price impact and lower opportunity costs for orders containing a hidden component. However, the estimate reported in Column 3 suggests no significant effect of hiding size on price impact, conditional on execution. Equivalently, the negative coefficient on the hidden indicator in Column 2 simply reflects the lower execution rate for hidden orders (see Fig. 1), not more favorable execution prices.

In contrast, comparing results across Columns 4 and 5 reveals a stronger negative effect of the hidden indicator on *opportunity costs* when we condition on the

non-execution of the order. Other things equal, non-execution should imply larger *opportunity costs*. Finding smaller *opportunity costs* associated with hidden orders even conditional on non-execution therefore implies less adverse movement in market prices from order submission to the close of trading for those orders with a hidden component. This evidence is strongly consistent with the reasoning that orders which are fully displayed execute quickly, either by taking liquidity from the book or by drawing out passive traders, and tend to be placed by informed traders. In contrast, hidden orders are used primarily by uninformed traders, consistent with the reasoning that hidden orders mitigate the option value of standing limit orders (Harris, 1996), as implied by *Hypothesis IIA*.

Finally, Column 1 presents coefficients when the *implementation shortfall*, which is the sum of *price impact* and *opportunity cost*, is the dependent variable. As might be expected, the implementation shortfall is smaller when markets are more active. Consistent with prior literature, implementation shortfall costs are higher for aggressively priced orders, for larger order sizes, and for buyer initiated orders. Most important for this analysis, implementation shortfall costs are lower for orders that hide a portion of the order size. Panel B of Table 1 shows that hidden order usage is concentrated in large orders. The point estimate of -0.023 reported in Column 1 of Table 4 implies an implementation shortfall cost savings of €23 for an order of 1,000 shares that contains hidden size, as compared to a similar fully exposed order.

These empirical results indicate both cost-benefit tradeoffs and self-selection in the order exposure decision. On average, exposing an order increases the likelihood of full execution and lowers the time between order submission and execution. However, despite more rapid executions and higher execution rates, exposed orders have higher opportunity costs and larger implementation shortfall. These results likely reflect self-selection by which informed traders tend to expose orders and uninformed traders tend to hide orders. Explicitly incorporating trader self-selection in the econometric analysis is beyond the scope of this paper due to the lack of empirical proxies for trader motivation, but presents an important and interesting avenue for future research.

6. Determinants of the decision to hide order size

In this section, we examine the empirical determinants of order exposure from two perspectives. First, we examine the determinants of order exposure from the perspective of the initiating trader, who chooses the order's limit price, order size, and the portion of the order size to hide. Second, we examine the problem from the perspective of market participants who seek to detect the presence and the quantity of hidden size. The crucial distinction between examining order exposure from the perspective of the initiating trader versus the market participants lies in the information set: Market participants observe only the displayed order size, while the order initiator knows also the total size of the order.

Table 5 presents results of specifications where we treat the order attributes as independent choices. However, because theory (e.g. Buti and Rindi, 2008) suggests that traders may well choose the order attributes simultaneously, we also examine an alternative specification in Table 6 where the price aggressiveness and order exposure decisions are treated as endogenous variables.¹⁶ The choice of the explanatory variables in each analysis is guided by the prior theoretical and empirical literature, as described in Section 2.

In Table 5, we examine the extent to which a logistic model that conditions on the information set of market participants is successful in detecting the presence of orders containing some hidden size (Model 1) and compare the predictive ability to that of a second model (Model 2) that conditions on the richer information set of the initiating trader. In a similar spirit, Models 3 and 4 report results obtained from implementing a Tobit specification, focusing on the quantity of shares that are hidden, as opposed simply to the presence or absence of hidden shares.

The empirical specification includes variables that capture (1) the state of the limit order book, including the bid-ask spread and the displayed depth at the inside quotes, cumulative order book imbalance, standing limit orders at the same price as the incoming order, and revelation of hidden orders at the inside quotes by the most recent transaction; (2) trading conditions for the stock, such as recent volatility, the trading frequency, and the waiting time between recent order arrivals; (3) order attributes, such as price aggressiveness and order size; and (4) control variables such as recent industry volatility, overall market volatility, and time-of-the-day effects. The inclusion of industry and market volatility variables helps to control for any commonality in economic fundamentals, as in CRS (2000). A detailed description of the variables is provided in the Appendix. To render results more comparable across stocks, we normalize some variables. The depth and spread variables are each normalized by dividing the actual observation by the median for that stock during the month, while order size and trade size are normalized by dividing the actual observations by the stock's average daily trading volume.

6.1. Determinants of order exposure when treating order attributes as independent choices

Table 5 reports regression coefficients along with corresponding *t*-statistics, estimated on a firm-by-firm basis and aggregated across firms using the approach described in Section 4.1. The analysis uses data from the first 15 trading days of the sample, while the remaining five days are reserved for examining out-of-sample predictive ability. The dependent variable for results reported in Columns 1 and 2 equals one for orders that

¹⁶ We recognize that traders may also select simultaneously other order attributes such as the order size. In an earlier draft of the paper, we model the limit order traders' choice of price aggressiveness, order exposure, and order size in a simultaneous equation framework. Results are available upon request.

Table 5

Modeling the decision to hide order size and the magnitude of the hidden size.

The table reports coefficients (and *t*-statistics) of logistic models predicting the presence of hidden orders and Tobit models predicting the magnitude of hidden size, on order attributes and market conditions. For the logistical models, the dependent variable equals one if the order has a hidden quantity and equals zero otherwise. For the Tobit models, the dependent variable equals the quantity of hidden size (in shares) normalized by the stocks' average daily trading volume. Detailed definitions of the explanatory variables are provided in the Appendix. Results are reported from the perspective of market participants (displayed size as explanatory variable) and from the perspective of the order initiator (total size as explanatory variable). Time series coefficients are estimated on a firm-by-firm basis. Reported results are aggregated across firms using the Bayesian framework of DuMouchel (1994), with *t*-statistics corrected for cross-correlation in individual regression residuals following Chordia, Roll, and Subrahmanyam (2005).

Variable	Decision to hide size		Magnitude of hidden size	
	Estimate (1)	Estimate (2)	Estimate (3)	Estimate (4)
Intercept (<i>t</i> -statistic)	−1.5281 (−25.50)	−2.6772 (−15.19)	−0.1679 (−6.41)	−0.0192 (−5.03)
<i>Order attributes:</i>				
Price aggressiveness (<i>t</i> -statistic)	0.5141 (3.68)	3.3207 (7.01)	0.2377 (3.17)	−0.0001 (−2.80)
Displayed order size (norm) (<i>t</i> -statistic)	−0.3521 (−2.00)		0.1600 (2.21)	
Total order size (norm) (<i>t</i> -statistic)		17.8051 (25.07)		0.7597 (27.64)
<i>Market conditions:</i>				
Bid–ask spread (norm) (<i>t</i> -statistic)	11.0993 (8.09)	8.6708 (4.34)	0.3681 (3.63)	−0.0023 (−0.28)
Depth—same side (norm) (<i>t</i> -statistic)	−0.0272 (−2.65)	−0.0724 (−6.53)	−2.4807E−04 (−3.25)	−0.1981E−04 (−3.74)
Depth—opposite side (norm) (<i>t</i> -statistic)	0.0017 (0.43)	−0.0125 (−1.92)	0.2641E−04 (1.11)	−0.1146E−04 (−3.40)
Conditional volatility (previous hour) (<i>t</i> -statistic)	−0.0192 (−0.28)	−0.1996 (−1.21)	0.0034 (1.47)	−0.0002 (−0.29)
Waiting time (<i>t</i> -statistic)	0.1711 (2.00)	0.8100 (2.87)	0.0559 (2.87)	0.0021 (2.03)
Trade frequency (last hour) (<i>t</i> -statistic)	0.0471 (1.15)	−0.0517 (1.55)	0.2740E−04 (0.10)	−0.1121E−04 (−2.13)
HiddenSameSide (norm) (<i>t</i> -statistic)	0.1513 (0.54)	3.7664 (4.99)	−0.0311 (−0.41)	0.0722 (3.84)
Same price book displayed depth (norm) (<i>t</i> -statistic)	0.5462 (2.46)	−0.2876 (−0.56)	0.1488 (2.08)	−0.0133 (−1.93)
Book order imbalance (norm) (<i>t</i> -statistic)	−0.0115 (−0.22)	−0.2151 (−2.59)	−0.0003 (−1.95)	−0.0005 (−3.72)
Last trade size (norm) (<i>t</i> -statistic)	0.6811 (4.69)	−1.5098 (−3.16)	0.3618 (5.96)	−0.0471 (−5.16)
Market volatility (previous hour) (<i>t</i> -statistic)	−0.0186 (−2.51)	−0.0153 (−2.58)	−0.5203E−04 (−1.63)	−0.4492E−06 (−0.93)
Industry volatility (previous hour) (<i>t</i> -statistic)	−0.0275 (−5.00)	−0.0517 (−2.84)	−0.9463E−04 (−2.04)	0.1303E−06 (0.88)
Last trading hour indicator (<i>t</i> -statistic)	0.0449 (0.90)	−0.1915 (−3.49)	−0.0001 (−0.26)	−0.0002 (−5.37)

contain hidden size and zero for orders that do not. We focus here on limit orders that supply liquidity (order aggressiveness categories 5, 6 and 7), since the exposure decision is more relevant for traders submitting limit orders that stand for some period of time in the limit order book.

In Column 1 we report results from the specification that includes displayed, not total, order size. These

confirm that order attributes and market conditions are useful in explaining traders' exposure decisions, which is broadly consistent with the findings of De Winne and D'Hondt (2007). The positive coefficient on waiting time between order arrivals indicates that hidden orders are more likely when markets are less active. A slower order arrival rate implies a decreased likelihood that a subsequent limit order arrives at the same price, meaning that

the loss of time priority due to hiding a portion of the order is less costly. Hidden order usage is less likely when depth at the best quote on the same side is greater but is more likely when displayed depth at the same price as the limit order is greater. The former result likely reflects that the loss in time priority is costly when competition from traders on the same side is high, while the latter result is unexpected. Somewhat surprising, greater industry and market return volatility is associated with a lower likelihood of hidden order usage.

The results also indicate that traders choose to hide more of their orders when the bid–ask spread is wide, consistent with the reasoning that traders use the option to hide order size during periods when adverse selection risk is high. In contrast, De Winne and D'Hondt (2007) report a negative relation between bid–ask spread and hidden order usage, a finding that is inconsistent with their stated hypothesis. Similarly, De Winne and D'Hondt find that hidden orders are more likely during the last trading hour of the day, while we find the opposite result. Our results indicate that traders choose to expose orders prior to the market close, likely due to improved execution probabilities.

An important determinant of order exposure identified by prior literature is total order size (see Harris, 1996; Aitken, Berkman, and Mak, 2001; De Winne and D'Hondt, 2007). Consistent with the prior work, we find (Column 2 of Table 5) a positive and highly significant (t -statistic = 25.1) coefficient on total order size, implying that traders who select a large order size are much more likely to hide a portion of the size. However, other market participants have information only on displayed order size. The estimated coefficient on displayed order size reported in Column 1 of Table 5 is significantly *negative*, implying that orders with larger displayed sizes are *less* likely to contain a hidden component.

With regard to order attributes, we obtain a positive and significant coefficient on limit price aggressiveness, consistent with De Winne and D'Hondt (2007), indicating that hidden order usage increases as prices become more aggressive. Since the value of the free option increases with the limit order's price aggressiveness, these findings suggest that traders use hidden orders to control the cost of order exposure.

To assess the economic significance of the coefficient estimates reported on Table 5 we rely mainly on evidence regarding the ability of the model to predict the presence of hidden size out-of-sample. As a supplement, we compute the change in the predicted probability that an order contains hidden size for a one standard deviation change in each independent variable, following Wooldridge (2002, Eq. (15.29)).¹⁷ The results confirm the dominant effect of total order size, for which a one standard deviation increase implies a 0.891 increase in the probability that the order contains hidden size. A one standard deviation increase in price aggressiveness is

associated with an increase of 0.099 in the probability that the order contains hidden size, while a one standard deviation increase in the bid–ask spread is associated with an increase of 0.021 in the probability that the order contains hidden size.

In addition to detecting the presence of hidden orders, market participants are interested in predicting the quantity of hidden shares. Columns 3 and 4 of Table 5 report regression coefficients and t -statistics obtained from the firm-by-firm Tobit analysis, aggregated across firms as discussed in Section 4.1. The dependent variable is the number of hidden shares in the order, normalized by the stock's daily average trading volume. We use a Tobit specification because the number of hidden shares is subject to a lower limit of zero. In fact, as reported on Table 1, over 80% of the incoming orders do not contain a hidden quantity.

The signs of the regression coefficients for market condition variables obtained in the Tobit analysis are generally the same as those reported for the logistic analysis; the exceptions being the insignificant coefficients in Column 4 on bid–ask spread and return volatility. A number of explanatory variables, including depth on the same side, depth at the same price on the same side, and the size of the most recent trade, are statistically significant for each Tobit specification. Interestingly, different insights emerge regarding order attributes. With regard to price aggressiveness, the estimated coefficient for the Tobit analysis (Column 4) is negative, which contrasts with the positive sign reported for the logistic analysis (Column 2). With regard to displayed order size, the estimated coefficient for the Tobit analysis (Column 3) is positive, which contrasts with the negative sign reported for the logistic analysis (Column 1). The evidence therefore implies that an order with larger displayed size is less likely to contain a hidden component, but that, conditional on the presence of some hidden size, the quantity (the number of shares) of hidden size increases with the displayed order size. These findings could explain the traders' use of "pinging" strategies, first submitting a marketable order with size slightly larger than that displayed and then following with a larger order if hidden size is detected.

6.2. The endogenous choice of price aggressiveness and order exposure

Since theory suggests that the choice of limit price and order exposure may well be simultaneous decisions, we report in Table 6 the results of a simultaneous equation model where we treat price aggressiveness and order exposure as endogenous variables. We employ two-stage least squares (2SLS) to estimate the simultaneous equation model. In the absence of fully developed economic models of optimal order submission strategies, we rely in part on the prior empirical evidence to select instrumental variables, as well as on the standard technique of designating lagged variables as instruments.

De Winne and D'Hondt (2007) report that the revelation of hidden depth on the opposite side quote (i.e., at the

¹⁷ Probabilities are assessed conditional on cross-sectional mean outcomes on other explanatory variables, and the standard deviation employed is the cross-sectional average of firm-specific standard deviations.

Table 6

Simultaneous equation estimates of price aggressiveness and order exposure.

Reported are the regression coefficients for simultaneous equations of price aggressiveness and order exposure. The presence of hidden orders is modeled using logistical models and the magnitude of hidden size is modeled using Tobit models. For the logistical models, the dependent variable equals one if the order has a hidden quantity and equals zero otherwise. For the Tobit models, the dependent variable equals the quantity of hidden size (in shares) normalized by the stocks' average daily trading volume. Detailed definitions of the explanatory variables are provided in the Appendix. The time series coefficients are estimated on a firm-by-firm basis. Reported results are aggregated across firms using the Bayesian framework of DuMouchel (1994), with *t*-statistics corrected for cross-correlation in the individual regression residuals following Chordia, Roll, and Subrahmanyam (2005).

Variable	Decision to hide size	Price aggressiveness	Magnitude of hidden size	Price aggressiveness
	Estimate (1)	Estimate (2)	Estimate (3)	Estimate (4)
Intercept (<i>t</i> -statistic)	-1.4337 (-9.38)	-0.0208 (-11.40)	-0.0315 (-5.33)	-0.0212 (-11.79)
<i>Order attributes:</i>				
Price aggressiveness (endogenous) (<i>t</i> -statistic)	8.4067 (2.03)		0.0236 (0.67)	
Order exposure (endogenous) (<i>t</i> -statistic)		0.0025 (2.05)		-0.0019 (-0.43)
Total order size (norm) (<i>t</i> -statistic)	10.4493 (15.72)	0.0004 (1.35)	1.0760 (92.01)	0.0043 (1.14)
<i>Market conditions:</i>				
Bid-ask spread (norm) (<i>t</i> -statistic)	7.4466 (2.32)	-0.5408 (-16.35)	0.1236 (1.77)	-0.5438 (-18.43)
Depth—same side (norm)(*10 ²) (<i>t</i> -statistic)	-2.4137 (-7.05)	0.0051 (1.66)	-0.0342 (-3.05)	0.0068 (1.88)
Depth—opposite side (norm) (<i>t</i> -statistic)	-0.0037 (-0.97)	-1.50E-05 (-0.67)	-5.63E-06 (-0.41)	-0.0001 (-0.65)
Conditional volatility (previous hour) (<i>t</i> -statistic)	-0.0400 (-0.51)	-0.0033 (-5.67)	0.0002 (0.19)	-0.0034 (-5.99)
Waiting time (<i>t</i> -statistic)	0.0248 (0.21)	-0.0027 (-1.97)	0.0108 (1.90)	-0.0030 (-2.10)
Trade frequency (last hour)(*10 ²) (<i>t</i> -statistic)	-0.2422 (-0.14)	0.1086 (5.12)	-0.0073 (-0.89)	0.1074 (5.68)
HiddenSameSide (norm) (<i>t</i> -statistic)	3.2222 (3.42)		0.0813 (5.15)	
HiddenOppSide (norm) (<i>t</i> -statistic)		0.0126 (3.24)		0.0084 (1.75)
Same price book displayed depth (norm) (<i>t</i> -statistic)	-0.3314 (-0.52)		-0.0677 (-2.20)	
Book order imbalance (norm) (<i>t</i> -statistic)	-0.0934 (-2.18)	-0.0005 (-1.37)	-0.0002 (-2.79)	-0.0004 (-1.24)
Lag (price aggressiveness) (<i>t</i> -statistic)		0.0419 (5.66)		0.0397 (5.59)
Lag (displayed order size) (<i>t</i> -statistic)		0.0004 (0.29)		-0.0006 (-0.48)
Last trade size (norm) (<i>t</i> -statistic)	-0.8248 (-2.00)	0.0018 (1.21)	-0.0213 (-2.23)	0.0020 (1.04)
Market volatility (previous hour) (*10 ⁴) (<i>t</i> -statistic)	-8.5955 (-3.35)	0.2070 (1.34)	-0.1768 (-1.83)	-1.4229 (-0.07)
Industry volatility (previous hour)(*10 ⁴) (<i>t</i> -statistic)	-289.7625 (-2.77)	-0.1588 (-0.38)	-0.7449 (-2.27)	4.6040 (2.32)
Last trading hour indicator (<i>t</i> -statistic)	-0.1085 (-3.39)	0.0007 (2.71)	-0.0005 (-3.68)	-0.0212 (-11.79)

ask quote for an incoming buy order) is an important determinant of price aggressiveness. However, in the absence of any reason to anticipate that hidden depth on the opposite side should affect the exposure decision

(after accounting for information contained in other exogenous variables), we exclude *HiddenOppSide*, the quantity of hidden shares revealed by a prior execution on the opposite side of the book, from the order exposure

equation. In addition, we exclude the lagged values of the two order attributes, price aggressiveness and displayed order size.¹⁸

To identify the price aggressiveness equation we rely on similar reasoning. In particular, we exclude *HiddenSameSide*, the size of the hidden depth on the same side quote revealed in the prior transaction (i.e., bid quote for a buy order and vice versa), and *SamePriceBookDispDepth*, the displayed depth at the same price as the specified limit price of the incoming order, from the price aggressiveness equation. *HiddenSameSide* is expected to have a first-order effect on the exposure decision but not for price aggressiveness, after accounting for information contained in other exogenous variables. Since *SamePriceBookDispDepth* is estimated conditional on the choice of the limit price, we expect it to directly effect the exposure decision, but not the price aggressiveness decision.

For the sake of brevity, our discussion focuses on the sensitivity of findings to specification. Most coefficients are similar to those reported in Table 5 without correction for endogeneity. However, order arrival rates (average waiting time) has an insignificant estimated coefficient in the simultaneous specification reported in Column 1, while it was significant in Column 2 of Table 5. Also, the puzzling positive relation between displayed depth at the same price and hidden order usage reported in Column 1 of Table 5 is no longer observed after controlling for endogeneity.

With regard to price aggressiveness, the results in Column 1 of Table 6 indicate that traders are more likely to hide a portion of order size when they also select a more aggressive limit price, and the coefficient is substantially larger (8.41 versus 3.32) than in the absence of controls for endogeneity. That traders are more likely to hide order size when they also choose more aggressive order prices is consistent with the reasoning that traders are concerned about strategies that exploit the option value of standing limit orders. These trading options are more valuable for aggressively price orders. The value of the free trading option can be reduced by hiding size, as traders are unsure of how many shares can be traded at the indicated price.

Turning to the price aggressiveness equations (Columns 2 and 4 of Table 6), we find that, consistent with Griffiths, Smith, Turnbull, and White (2000), orders are less aggressively priced when the spread is wide, as limit order traders prefer to provide liquidity rather than take liquidity from the book. We estimate a significant negative coefficient on firm volatility, implying that orders are less price-aggressive when market conditions are turbulent. This likely reflects limit order traders concerns that their orders may be “picked off” by better-informed traders during times of greater uncertainty. The positive coefficient on relative trading frequency and the negative coefficient on waiting time suggest that orders are more price-aggressive when recent trading activity has been

high, and when the order arrival rate and trading activity are slow. Consistent with De Winne and D’Hondt (2007), we find that traders submit aggressively priced orders when the prior trade reveals the existence of hidden depth on the opposite side of the market. Focusing on the endogenously determined order exposure variables, we find a significant positive coefficient on the decision to hide variable in Column 2, indicating that traders who choose to hide their orders tend to also use aggressively priced limit orders, while traders who choose to expose their orders tend to submit limit orders that are placed away from the best quotes. However, after allowing for endogeneity, results reported in Columns 3 and 4 of Table 6 do not indicate a significant effect of price-aggressiveness on the number of shares hidden. We next quantify the extent to which the models have an ability to detect the presence of and quantity of shares hidden, both in- and out-of-sample.

6.3. Evaluating predictability: the *H-Score* and *H-Size* for an order

To evaluate the predictive ability of the logistic and Tobit models, we sort and rank each order into quintiles based on the predicted probability that the order contains a hidden component or the predicted number of shares that are hidden. Predicted probabilities (shares hidden) are based on coefficients estimated during the first 15 days of sample as reported in Table 5 and on order attributes and market conditions that correspond to each order.¹⁹ We then calculate a *Hidden-Score* (*H-Score*) for each order, by dividing the predicted probability for the order by the unconditional probability of hidden order usage, which for each firm is the number of orders with non-zero hidden size as a proportion of total orders. Analogous, we define a *Hidden-Size* (*H-Size*) for each order as the predicted (normalized) hidden size for the order divided by the average normalized hidden order size for the firm. Our approach follows that of Dechow, Ge, Larson, and Sloan (2007), who implement a similar design and designate an “F-score” for detecting accounting fraud. An *H-Score* of 1.00 indicates that the order has the same likelihood of containing hidden size as a random order in the firm, and *H-Scores* less (greater) than one indicate lower (higher) probabilities of containing hidden size. For example, an *H-Score* of 2.00 indicates that the order is twice as likely to contain some hidden size as a random order for the firm. Analogously, an *H-Size* of 2.00 indicates that the hidden quantity for the order is expected to be twice as large as for a random order for the firm.

¹⁸ Results are broadly similar when we use only the two lagged order attributes as instruments. The models are tested for over-identification restrictions using Basmann’s (1960) test. The tests produce an overall rejection rate of around 15% at the 5% significance level.

¹⁹ The out-of-sample predictions based on the 2SLS analysis in Table 6 are similar to but slightly weaker than those based on Models 2 and 4 in Table 5. The 2SLS analysis uses predicted (from the first-stage regression) rather than actual outcomes on the endogenous price aggressiveness variable. While this is necessary to obtain unbiased estimates of the structural parameters, market participants observe and can take into account for forecasting purposes actual outcomes as opposed to outcomes predicted on the basis of a first stage regression on instrumental variables. For this reason, we evaluate predictability based on Table 5 coefficients.

Table 7

Goodness of fit and predictive ability in-sample and out-of-sample of the logistical and tobit models.

The table presents results regarding the predictive ability of the logistic and Tobit models, as reported in Tables 5 and 6, for in-sample (first 15 days of April 2003) and out-of-sample (last five days of April 2003) periods for standing limit orders.

In Panel A, the detection rates of logistical Models 1 and 2 in Table 5 are reported by ordered quintile groups of *H-Scores*. For each order, we calculate the predicted probability based on the logistical regression coefficients, the order attributes and the market conditions that correspond to each order. We then calculate a Hidden-Score (*H-Score*) for each order by dividing the predicted probability for the order by the unconditional probability of hidden order usage, which is the number of orders for the firm with non-zero hidden size divided by the total number of orders during the in-sample period. An *H-Score* of 1.00 indicates that the order has same likelihood of containing hidden size as the overall sample, and *H-Scores* less (greater) than one indicate lower (higher) probabilities of containing hidden size. In Panel A.1, we rank the orders based on *H-Scores* and report the true percentage of hidden and non-hidden orders contained in each ordered quintile groups of *H-Scores*. Panel A.2 reports prediction statistics when the *H-score* cut-off is set to 1. *Correct Classification* represent the total number of correct predictions of hidden and non-hidden orders over the total number of orders, and *Sensitivity* represents the total number of correct predictions of hidden orders over the total number of hidden orders. *Type I* error represents the percentage of non-hidden orders that are misclassified as hidden, and *Type II* error represents the percentage of hidden orders that are misclassified as non-hidden. We report the cross-sectional medians from firm-by-firm estimates.

Panel B reports results for quintile groups ordered by *Hidden-Size*. For each order, we calculate the predicted shares hidden based on the Tobit regression coefficients from Model 3 and 4 in Table 5, the order attributes, and the market conditions that correspond to each order. *Hidden-Size (H-size)* for each order is the predicted (normalized) hidden size divided by the average normalized hidden size for the firm. Reported are the predicted *H-Size*, observed (true) *H-Size*, the mean square error for the model based on predicted *H-Size* and observed *H-Size*, and the mean square error based on the unconditional expectation (sample average) of hidden size.

	Total size		Displayed size	
	In-sample (1)	Out-of-sample (2)	In-Sample (3)	Out-of-Sample (4)
<i>Panel A.1: H-Score quintiles and goodness of fit percentages</i>				
H-Score lowest quintile				
Hidden orders	1%	2%	7%	14%
Non-hidden orders	25%	25%	24%	21%
H-Score quintile 2				
Hidden orders	1%	3%	15%	18%
Non-hidden orders	24%	25%	21%	20%
H-Score quintile 3				
Hidden orders	4%	7%	18%	19%
Non-hidden orders	23%	23%	20%	20%
H-Score quintile 4				
Hidden orders	15%	20%	25%	22%
Non-hidden orders	21%	20%	18%	20%
H-Score highest quintile				
Hidden orders	77%	66%	34%	26%
Non-hidden orders	6%	7%	16%	19%
<i>Panel A.2: In-sample and out-of-sample predictions for H-Score cutoff equal to one</i>				
Correct classification	88%	83%	66%	55%
Sensitivity	85%	82%	69%	61%
Type I error	11%	14%	35%	46%
Type II error	15%	19%	31%	39%

Panel B: Predicted quintiles and goodness of fit measures

	In-sample				Out-of-sample			
	Observed hidden size/ unconditional expectation	Predicted hidden size/ unconditional expectation	MSE model	MSE unconditional	Observed hidden size/ unconditional expectation	Predicted hidden size/ unconditional expectation	MSE model	MSE unconditional
<i>Panel B.1: Quintile portfolios based on displayed order size</i>								
H-Size lowest quintile								
All orders	0.2	0.3	1.1	1.7	0.7	0.5	6.6	5.9
Only hidden orders	4.3	0.5	19.8	21.5	5.2	0.5	48.6	39.4
H-Size quintile 2								
All orders	0.4	0.6	3.1	3.3	0.7	0.8	4.3	5.4
Only hidden orders	4.0	0.7	25.6	21.5	5.0	0.9	33.3	25.7
H-Size quintile 3								
All orders	0.8	0.9	5.7	5.6	0.8	1.1	7.8	8.2
Only hidden orders	4.3	1.0	30.4	30.5	5.4	1.1	56.6	49.1
H-Size quintile 4								
All orders	1.1	1.3	11.2	11.7	0.9	1.3	9.9	7.8
Only hidden orders	4.5	1.3	35.0	37.2	4.6	1.3	38.7	31.6
H-Size highest quintile								
All orders	2.2	2.2	28.5	33.3	1.6	2.2	27.6	17.3
Only hidden orders	6.8	2.3	64.2	93.3	5.6	2.0	54.7	51.9
<i>Panel B.2: Quintile portfolios based on total order size</i>								
H-Size lowest quintile								
All orders	0.0	0.0	0.0	1.0	0.0	0.0	0.0	1.0
Only hidden orders	0.8	0.0	0.8	0.7	1.0	0.0	1.6	1.8
H-Size quintile 2								
All orders	0.0	0.1	0.0	1.0	0.0	0.1	0.0	1.2
Only hidden orders	0.6	0.1	0.4	0.7	0.7	0.1	0.6	0.8
H-Size quintile 3								
All orders	0.0	0.1	0.0	1.0	0.0	0.1	0.1	1.0
Only hidden orders	0.7	0.2	0.4	0.5	1.0	0.2	0.7	0.7
H-Size quintile 4								
All orders	0.1	0.3	0.2	1.0	0.6	0.4	0.4	0.9
Only hidden orders	1.2	0.3	0.7	0.5	1.7	0.4	1.4	1.0
H-Size highest quintile								
All orders	4.7	4.2	4.5	52.6	4.3	3.5	6.1	52.3
Only hidden orders	6.7	5.0	2.5	80.4	7.6	5.5	5.3	94.1

Panel A.1 of Table 7 reports the percentage of actual hidden orders and fully displayed (non-hidden) orders by *H-Score* quintile. These percentages are calculated for each sample firm, and we report the median outcome across firms.²⁰ Results are reported separately for the in-sample period during which the Logit model was estimated (the first 15 trading days of April 2003) and during a later out-of-sample period (the remaining five trading days of April 2003). If the model has no predictive ability then hidden and non-hidden orders will be randomly dispersed across quintiles, implying that 20% of both hidden and non-hidden orders would be allocated to each *H-Score* quintile. In contrast, if the model is effective, the highest percentage of hidden orders should be concentrated in Quintile 5 and the lowest percentage should be found within Quintile 1.

Focusing first on Column 1, which pertains to in-sample results obtained while using total order size, we find that 77% of actual hidden orders are in Quintile 5, compared with only 1% in Quintile 1. Further, the hidden order percentages increase almost monotonically from Quintile 1 to Quintile 5. Out of sample (Column 2) the model assigns 66% of hidden orders to the highest *H-Score* quintile, compared with 2% allocated to the lowest *H-Score* quintile.

Of course, these results were based on a specification that included total order size, which is not observable to market participants other than the order submitter. For comparison, Columns 3 and 4 of Table 7 Panel A, report results obtained when *H-Scores* are computed from the more realistic information set that includes displayed instead of total order size. Not surprisingly, the predictive power of the model is decreased but is still substantial. In-sample, the model allocates 34% of orders containing hidden size to the largest *H-Score* quintile, compared with only 7% allocated to the lowest *H-Score* Quintile. Out-of-sample performance is somewhat weaker but still indicates significant forecast power, with 26% of orders allocated to the largest *H-Score* Quintile, versus 14% allocated to the smallest Quintile.

The superior predictive performance obtained when using total order size implies that hiding size meaningfully reduces the information available to other market participants. However, significant predictive power remains when traders' information sets are restricted to include only displayed order size. We conclude from this analysis that market participants can detect to a significant but imperfect extent the presence of a hidden order based on observable variables.

In Panel B of Table 7, we report the predicted hidden size and the observed (true) hidden size for all orders and for the subset of orders with some hidden size, by *H-Size* quintile. The panel also compares the mean square error (MSE) of the Tobit model with the MSE obtained when comparing the actual hidden shares to the unconditional

(sample) average hidden shares. The out-of-sample forecasting ability of the Tobit model using displayed order size (Panel B.1) is modest for orders in the four lowest *H-Size* quintiles. While the ratio of observed hidden size to unconditional average hidden size is indeed less than one and increases across these four quintiles, magnitudes are moderate, ranging from 0.7 for the lowest *H-Size* Quintile to 0.9 for the fourth *H-Size* Quintile. Further, for these quintiles unconditional averages provide MSEs that are comparable with the Tobit forecasts.

However, the Tobit model exhibits forecast power for the highest *H-Size* quintile, implying that the model succeeds in forecasting hidden size in those cases where the forecasted hidden size is largest. The actual normalized hidden size for the highest *H-Size* quintile is 2.2 in-sample and 1.6 out-of-sample, implying that orders falling in this quintile contain significantly more hidden size than randomly selected orders. In comparison, for the lowest *H-Size* quintile, the actual normalized hidden size averages only 0.2 in-sample and 0.7 and out-of-sample.

Not surprisingly, the performance of the Tobit model would be significantly better if total orders size could be relied on as a predictor. In this case the average normalized hidden size is 4.7 in-sample and 4.3 out-of-sample for orders in the largest *H-Size* Quintile, as compared with 0.1 or less in the other Quintiles. On balance, whether the Tobit model is useful depends on the relative benefits of successfully identifying, on average, the large quantities of hidden size in those cases where the model predicts large hidden size, versus the costs of less precise predictions of hidden size in those cases where the model predicts low or moderate hidden size.

6.4. Model fit and an application

We next assess model performance by reporting on the percentage of orders that are correctly classified as containing hidden size or not containing hidden size. That is, the percentage of orders containing hidden size that are incorrectly classified by the model as not containing hidden size (Type II error) and conversely, the percentage of non-hidden orders that are incorrectly classified by the model as containing hidden size (Type I error). Following Dechow, Ge, Larson, and Sloan (2007), we also report a "Sensitivity" statistic, which is one minus the probability of a Type II error, i.e. the percentage of orders containing hidden size that are correctly detected.

Evaluation of these statistics is intrinsically linked to the costs and benefits of correctly and falsely identifying hidden liquidity. Descriptive statistics reported on Table 2 indicate that 82% of sample orders contain no hidden size. Consequently, in this sample, a naive forecast that no orders contain hidden liquidity would achieve 82% correct classifications. However, this method would yield a sensitivity statistic of zero, as no hidden orders would be detected. The methods used here achieve higher sensitivity (i.e., they succeed in identifying some hidden orders), but at a cost in terms of a lower overall out-of-sample success rate.

²⁰ We also compiled cross-sectional means, which support similar conclusions but are noisier. The additional noise reflects the presence of a few observations with almost no hidden order usage in sample, but some hidden order usage out of sample, leading to exceptionally large ratios of actual to predicted use.

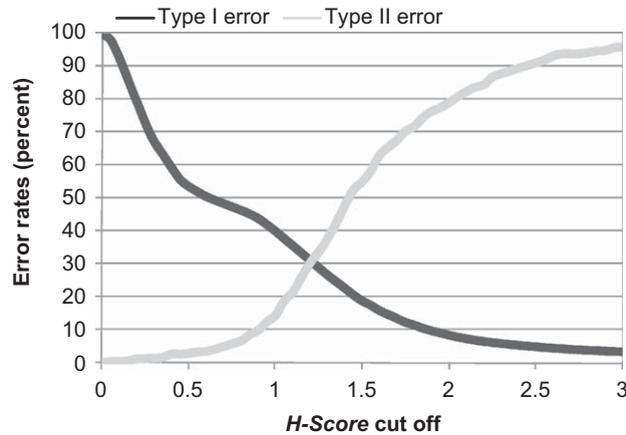


Fig. 2. Error rates for different *H-Score* cutoffs. This figure shows the Type I and Type II error rates with respect to different cutoffs of the *H-Score*, relying on the model predicting non-marketable hidden orders from the perspective of market participants, for a representative company (Rexel Electrical Supplies). Type I error represent the percentage of non-hidden orders that are misclassified as hidden, and Type II error represent the percentage of hidden orders that are misclassified as non-hidden. *H-Score* represents the Hidden-score for each order, obtained by dividing the predicted probability from the model by the unconditional probability of hidden order usage in the sample period, i.e., first 15 days of April 2003. Error rates displayed are calculated from the out-of sample period, i.e., the last five days of April 2003.

The relative costs of Type I and Type II errors may well differ across traders depending on their trading motives. For example, a trader interested in splitting a marketable order across fragmented market venues might wish to minimize *Type II* errors (i.e., hidden depth exists on a market but is not detected) so as to allocate orders to markets with the greatest total depth. On the other hand, a patient trader executing a non-marketable order may prioritize gaining time priority over standing limit orders. She might prefer to reduce *Type I* errors, since the cost of incorrect identification of a hidden order will not be punitive.

Traders can alter the probability of Type I versus Type II errors by selection of an *H-Score* cutoff, whereby orders with *H-Scores* above the cutoff are forecast to contain hidden order size and orders with *H-Scores* below the cutoff are forecast to not include hidden size. As an illustration, Fig. 2 displays rates of Type I and Type II errors for the out-of-sample period for a representative company (Rexel Electrical Supplies), for various possible *H-Score* cutoffs. *H-Score* cutoffs below about 0.5 lead to a very high rate of Type I errors (false identification of hidden size), but a very low rate of Type II errors (failure to detect hidden liquidity). Conversely, *H-Score* cutoffs above about 2.0 lead to very high rates of Type II errors, but infrequent Type I errors. *H-Score* cutoffs between these extremes lead to meaningful tradeoffs between Type I errors (false positives) and sensitivity (ability to detect hidden orders). The optimal *H-Score* can be assessed only in light of information on traders' motives and the relative cost to the individual trader of each type of error.

For illustrative purposes, we report on Panel A.2 of Table 7 results obtained with an *H-Score* cutoff of 1.00, so that all orders with conditional probabilities of containing hidden size greater than the unconditional probability are predicted to contain hidden size. Focusing first on in-sample performance when using the total order size as a predictor (Column 1), we observe very high success rates.

The frequency of correct classifications is 88% and the sensitivity ratio (percentage of hidden orders detected) is 85%. The corresponding out-of-sample performance (Column 2) is only slightly weaker, as the frequency of correct classifications falls to 83% and sensitivity falls to 82%.

As anticipated, success rates when relying on displayed size are somewhat lower. In-sample, the overall success rate is 66%, and the sensitivity is 69%. Out of sample, the overall success rate falls to 55%, while the sensitivity rate is 61%. While the assessment is necessarily subjective, it seems striking that a logistic model based on observable market and order characteristics can succeed in successfully identifying out-of-sample over 60% of orders containing a hidden size.

That the ability to detect hidden orders would hypothetically be improved if total order size were observable indicates that hiding size succeeds in withholding some information from other market participants. That is, market participants have considerable ability to detect hidden size based on observable order attributes and market conditions. However, the findings also indicate that hiding order size allows large traders to partially conceal their trading intentions and lower the option value of the limit order. These findings imply that hidden orders remain an important tool for market participants in controlling the risks of order exposure.

7. Conclusions and extensions

Many electronic stock exchanges choose to limit the transparency of their market by allowing traders to hide a portion of their order size. Yet, few studies have examined the costs and benefits of order exposure and described the trader's order exposure decision in markets where they have an option to hide their orders. This study empirically examines the order exposure decision for a sample of 100 stocks traded on Euronext-Paris during April 2003.

We assess the relative costs and benefits of order exposure by examining its impact on several dimensions of execution quality, including the likelihood of full execution, the time-to-execution, and implementation shortfall (transaction) costs. We find that hidden orders are associated with lower implementation shortfall cost and smaller opportunity cost; i.e., smaller price drifts subsequent to order submission time, suggesting that hidden orders are used primarily by traders without superior information on future security price movements. On the downside, even after controlling for price aggressiveness, order size, and market conditions, hidden orders take longer to execute and have larger non-execution rates. Thus, traders select the optimal exposure strategies on the basis of both their private trading motives and the tradeoffs involved in selecting more aggressive prices and exposing their orders. To our knowledge, this is the first study that attempts to quantify the costs and benefits of exposing order size in markets that provide the option to hide size.

We also investigate the factors that govern the trader's decision to hide order size and the selection of the hidden quantity, and assess whether market participants can detect the presence (existence and size) of hidden orders based on observable order attributes, market conditions, and firm characteristics. We find that prevailing market conditions, including order arrival rates, bid–ask spread, displayed book depth on same and opposite side, book imbalance, recent transaction size, and time of the day help to detect the presence of hidden orders. Hidden order usage is also related to the revelation of hidden liquidity by recent transactions, implying a degree of momentum in order exposure, and to the existence of standing limit orders at the same price in the book, thus reflecting the effect of the competition for liquidity provision. While we find that, consistent with prior work, total order size is positively related to hidden order usage, we also find that displayed order size is negatively related to the presence of hidden size. Further, displayed order size is less informative than total order size in detecting hidden liquidity, indicating that the option to hide order size enables large limit order traders to conceal at least the magnitude of their trading programs.

These findings have important implications for market centers that are moving toward implementing fully automated trading systems, such as the New York Stock Exchange, and for market centers that currently operate automated trading systems but require traders to fully display orders, such as the Hong Kong Stock Exchange. The set of order types that traders can submit represents an important dimension of trading system design. That hidden orders are used extensively and strategically on Euronext Paris implies that these orders are valuable to traders. Thus, market centers may be more successful in attracting large orders if they allow traders to hide order size. Our findings are suggestive that the recent fragmentation of orders in NYSE-listed stocks across ECNs and institution-oriented “dark pools” can be attributable in part to the inability to hide order size on the NYSE limit order book. Interestingly, the NYSE announced in October 2008 a pilot program that will allow traders to enter orders with a displayed size of zero.

These findings are of interest to market regulators, academics, and institutional trading desks. A better understanding of trader behavior in electronic limit order markets would enable regulators to more accurately assess the impact of new regulation on market liquidity. The empirical evidence on order submission strategies, and, in particular, order exposure, could help theorists in developing comprehensive models on trader behavior. Finally, institutional trading desks, responsible for executing block orders received from portfolio managers, are facing new challenges in the search for liquidity pools in an increasingly fragmented and automated US marketplace (see [Abrokwah and Sofianos, 2006](#)). By modeling the hidden dimension of liquidity for firms with differing liquidity characteristics and by relating order exposure to market conditions, these results provide insights on the circumstances in which the search for hidden liquidity is likely to be most successful.

Appendix A. Variable definitions

We examine a set of models related to the decision to hide order size (Logistic) and the quantity of hidden size (Tobit) for standing limit orders that supply liquidity. Specifically, we estimate

$$\begin{aligned}
 \text{OrderExposure}_{it} = & \\
 = & \alpha_0 + \alpha_1 \text{PriceAggressive}_{it} + \alpha_2 \text{OrderSize}_{it} \\
 & + \alpha_3 \text{Spread}_{it} + \alpha_4 \text{DepthSame}_{it} + \gamma_5 \text{DepthOpp}_{it} \\
 & + \alpha_6 \text{Volatility}_{it} + \alpha_7 \text{WaitTime}_{it} + \alpha_8 \text{TradeFreqHour}_{it} \\
 & + \alpha_9 \text{HiddenSameSide}_{it} + \alpha_{10} \text{SamePriceDisplayedDepth}_{it} \\
 & + \alpha_{11} \text{BookOrderImbalance}_{it} + \alpha_{12} \text{TradeSize}_{it-1} \\
 & + \alpha_{13} \text{MktVolatility}_{it-1} + \alpha_{14} \text{Ind.Volatility}_{it-1} \\
 & + \alpha_{14} \text{LastHour}_{it} \tag{6}
 \end{aligned}$$

The subscript i, t refers to the time t order in stock i . For the Logistic analysis, the dependent variable is an indicator variable that equals one if the order has a hidden size and equals zero otherwise. For the Tobit analysis, the dependent variable is the quantity of hidden shares divided by the average daily trading volume in the stock over the prior 30 trading days.

The remaining variables are defined as follows. *PriceAggressive* is the distance of the order's limit price from the opposite quote price, suitably signed (a higher value indicates a more aggressively priced order) divided by the quote midpoint. *DisplayedOrderSize* is exposed size of the order divided by average daily trading volume. *TotalOrderSize* is total (displayed plus hidden) size of the order divided by average daily trading volume. *Spread* is the percentage bid–ask spread. *DepthSame* is the displayed depth at the best bid (ask) for a buy (sell) order divided by the monthly median. *DepthOpp* is the displayed depth at the best ask (bid) for a buy (sell) order divided by the monthly median. *Volatility* is the standard deviation of quote midpoint returns over the preceding hour. *WaitTime* is the average elapsed time between the prior three order arrivals on the same side, refreshing the time clock each day. *TradeFreqHour* is the number of transactions in the last hour. *HiddenSameSide* is the number of hidden shares

at the best quote on the same side (i.e., at the bid side for a buy order) revealed by the most recent transaction. *HiddenOppSide* is the number of hidden shares at the best quote on the opposite side revealed by the most recent transaction. *TradesSize* is the size of the most recent transaction divided by the average daily trading volume. *SamePriceDisplayedDepth_{it}* is the current depth at the price level of the limit order that is submitted divided by the average daily trading volume. *BookOrderImbalance* is the percentage difference between the displayed liquidity in the best five prices on the buy and sell side of the book, suitable signed (i.e., the variable is positive when same size liquidity exceeds opposite side liquidity). *Ind.Volatility* is the return volatility of a portfolio of stocks in the same industry during the prior hour. *Mkt.Volatility* is the return volatility of the CAC40 Index during the prior hour. *Last Hour* is an indicator variable that equals one for orders submitted in the last hour of the trading day and is zero otherwise.

We also employ two-stage least squares to estimate the simultaneous equation model (Table 6), treating *Price Aggressiveness* and *Order Exposure* as endogenous variables. The *Price Aggressiveness* equation is defined as

$$\begin{aligned} \text{PriceAggressive}_{it} &= \gamma_0 + \gamma_1 \text{OrderExposure}_{it} + \gamma_2 \text{PriceAggressive}_{it-1} \\ &+ \gamma_3 \text{HiddenOppSide}_{it} + \gamma_4 \text{DisplayedSize}_{it-1} \\ &+ \gamma_5 \text{OrderSize}_{it} + \gamma_6 \text{Spread}_{it} + \gamma_7 \text{DepthSame}_{it} \\ &+ \gamma_8 \text{DepthOpp}_{it} + \gamma_9 \text{Volatility}_{it} + \gamma_{10} \text{WaitTime}_{it} \\ &+ \gamma_{11} \text{TradeFreqHour}_{it} + \gamma_{12} \text{BookOrderImbalance}_{it} \\ &+ \gamma_{13} \text{TradeSize}_{it-1} + \gamma_{14} \text{MktVolatility}_{it-1} \\ &+ \gamma_{15} \text{Ind.Volatility}_{it-1} + \gamma_{16} \text{LastHour}_{it} \end{aligned} \quad (7)$$

For the system of Eqs. (6) and (7), we rely on the exclusion restrictions on *PriceAggressive_{it-1}*, *HiddenOppSide_{it}*, and *DisplayedSize_{it}* to identify *PriceAggressive_{it}* and *HiddenSameSide_{it}* and *SamePriceBookDispDepth_{it}* to identify *OrderExposure_{it}*.

References

- Abrokwha, K., Sofianos, G., 2006. Accessing displayed and non-displayed liquidity. *Journal of Trading* Fall 47–57.
- Ahn, H., Bae, K., Chan, K., 2001. Limit orders, depth, and volatility. *Journal of Finance* 56, 767–788.
- Aitken, M., Berkman, H., Mak, D., 2001. The use of undisclosed limit orders on the Australian Stock Exchange. *Journal of Banking and Finance* 25, 1589–1603.
- Anand, A., Chakravarty, S., Martell, T., 2005. Empirical evidence on the evolution of liquidity: choice of market versus limit orders. *Journal of Financial Markets* 8, 265–287.
- Anand, A., Weaver, D., 2004. Can order exposure be mandated? *Journal of Financial Markets* 7, 405–426.
- Bae, K., Jang, H., Park, K., 2003. Traders' choice between limit and market orders: evidence from NYSE stocks. *Journal of Financial Markets* 6, 517–538.
- Basmann, R.L., 1960. On finite sample distributions of generalized classical linear identifiability test statistics. *Econometrica* 28, 939–952.
- Bessembinder, H., Venkataraman, K., 2004. Does an electronic stock exchange need an upstairs market? *Journal of Financial Economics* 73, 3–36.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris bourse. *Journal of Finance* 50 (5), 1655–1689.
- Buti, S., Rindi, B., 2008. Hidden orders and optimal submission strategies in a dynamic limit order market. Unpublished working paper, University of Toronto, Toronto, Canada.
- Chakravarty, S., Holden, C., 1995. An integrated model of market and limit orders. *Journal of Finance* 4, 213–241.
- Chordia, T., Roll, R., Subrahmanyam, A., 2000. Commonality in liquidity. *Journal of Financial Economics* 56, 3–28.
- Chordia, T., Roll, R., Subrahmanyam, A., 2005. Evidence on the speed of convergence to market efficiency. *Journal of Financial Economics* 76, 271–292.
- Dechow, P., Ge, W., Larson, C.R., Sloan, R.G., 2007. Predicting material accounting manipulations. Unpublished working paper, University of California, Berkeley, CA.
- De Winne, R., D'Hondt, C., 2007. Hide-and-seek in the market: placing and detecting hidden orders. *Review of Finance* 11, 663–692.
- DuMouchel, W., 1994. Hierarchical Bayes linear models for meta-analysis. Technical Report 27, National Institute of Statistical Sciences.
- Easley, D., O'Hara, M., 1987. Price, trade size and information in securities markets. *Journal of Financial Economics* 18, 69–90.
- Ellul, A., Holden, C., Jain, P., Jennings, R., 2007. Order dynamics: recent evidence from the NYSE. *Journal of Empirical Finance* 14, 636–661.
- Foucault, T., 1999. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets* 2, 193–226.
- Foucault, T., Moinas, S., Theissen, E., 2007. Does anonymity matter in electronic limit order markets? *Review of Financial Studies* 20 (5), 1707–1747.
- Griffiths, M., Smith, B., Turnbull, D., White, R., 2000. The costs and determinants of order aggressiveness. *Journal of Financial Economics* 56, 65–88.
- Handa, P., Schwartz, R., 1996. Limit order trading. *Journal of Finance* 51, 1835–1861.
- Harris, L., 1996. Does a minimum price variation encourage order exposure? Unpublished working paper, University of Southern California, Los Angeles, CA.
- Harris, L., 1997. Order exposure and parasitic traders. Unpublished working paper, University of Southern California, Los Angeles, CA.
- Harris, L., Hasbrouck, J., 1996. Market vs. limit orders: the SuperDOT evidence on order submission strategy. *Journal of Financial and Quantitative Analysis* 31, 213–231.
- Harris, L., Piwowar, M., 2006. Secondary trading costs in the municipal bond market. *Journal of Finance* 61, 1361–1397.
- Hasbrouck, J., Saar, G., 2002. Limit orders and volatility in a hybrid market: the Island ECN. Unpublished working paper, New York University, New York, NY.
- Jain, P., 2005. Financial market design and the equity premium: electronic versus floor trading. *Journal of Finance* 60 (6), 2955–2985.
- Lo, A.W., MacKinlay, C., Zhang, J., 2002. Econometric models of limit-order executions. *Journal of Financial Economics* 65, 31–71.
- Moinas, S., 2006. Hidden limit orders and liquidity in limit order markets. Unpublished working paper, Toulouse Business School, Toulouse, France.
- Pardo, A., Pascual, R., 2003. On the hidden side of liquidity. Unpublished working paper, University of Valencia, Valencia, Spain.
- Panayides, M., 2007. Affirmative obligations and market making with inventory. *Journal of Financial Economics* 86, 513–542.
- Parlour, C., 1998. Price dynamics in limit order markets. *Review of Financial Studies* 11, 789–816.
- Perold, A., 1988. The implementation shortfall: paper versus reality. *Journal of Portfolio Management* 14, 4–9.
- Rinaldo, A., 2004. Order aggressiveness in limit order book markets. *Journal of Financial Markets* 7, 53–74.
- Tuttle, L., 2006. Hidden orders, trading costs, and information. Unpublished working paper, University of Kansas, Lawrence, KS.
- Wooldridge, J.M., 2002. *Econometric Analysis of Cross Section and Panel data*. MIT Press, Cambridge, MA.