

On the Hidden Side of Liquidity*

Ángel Pardo

Departamento de Economía Financiera
Universidad de Valencia
Facultad de Economía
Edificio Departamental Oriental
46022 Valencia, Spain
e-mail: angel.pardo@uv.es
Phone: +34-963-82-83-77
Fax: +34-963-82-83-70

Roberto Pascual

Departamento de Economía de la Empresa
Universidad de las Islas Baleares
Ctra. Valldemossa, km. 7.5
07122 Palma de Mallorca, Baleares, Spain
e-mail: rpascual@uib.es
Phone: +34-971-17-13-29
Fax: +34-971-17-23-89

This version: November 2003

First version: September 2003

(Comments are welcome)

*

We wish to thank Yakov Amihud, Julio J. Lucia and Ingo Saar for their suggestions and comments. We gratefully acknowledge helpful discussions with ten practitioners from Spanish fund managers and financial services organizations (Bancaja, CAM, Fibanc, Gesamed, Inverseguros and Morgan Stanley Dean Winter). We are grateful for the financial support of the Spanish DGICYT projects BEC2001-2552-C03-03 and BEC2000-1388-C04-04. Roberto Pascual also acknowledges the financial sponsorship of the Fulbright Grant and the Spanish Ministry of Education, Culture and Sports. The first version of this paper was finished while Roberto Pascual was a Visiting Scholar at the New York University Salomon Center. Corresponding author: Roberto Pascual, Departamento de Economía de la Empresa, Universidad de las Islas Baleares, Ctra. Valldemossa Km. 7.5, 07122 Palma de Mallorca, Baleares, Spain. Phone: +34-971-17-13-29, e-mail: rpascual@uib.es.

On the Hidden Side of Liquidity

ABSTRACT

An important number of stock exchanges allow market participants to enter limit orders without revealing the full size. However, there is a lot of controversy over the use and consequences of hidden orders, since they embrace a complex interaction between order exposure risk, market liquidity and transparency. Our study focuses on the motives of submitting undisclosed limit orders to trade as well as on the market response when the presence of these orders is publicly revealed. Using data from the Spanish Stock Exchange, we find that hidden orders emerge in periods of intense trading activity and extremely high liquidity. Our results find no evidence that the undisclosed volume is used as a defensive strategy against parasitic traders. On the contrary, we provide support to the notion that liquidity suppliers use hidden orders to mitigate adverse selection costs. We also report that hidden orders temporarily increase the aggressiveness of traders when they are revealed to the marketplace but, as opposed to the widespread opinion among practitioners, they have no relevant price impact.

JEL: G1

1. Introduction

“Pure” order driven markets are characterized by the absence of market makers. These markets rely upon limit orders to replace market makers for the provision of liquidity. Until recently, the empirical evidence about limit-order trading focused on the limit order book of the NYSE, which complements the liquidity provided by the specialist (e.g., Kavajecz, 1999). However, the patent trend in market design toward electronic limit order book systems (see Domowitz and Wang, 1994) together with the increasing availability of order book data have yielded a growing interest on the liquidity provision in “pure” order driven markets.

Empirical research has highlighted the capability of pure order driven venues to provide liquidity when it is valuable for the marketplace.¹ Several studies have compared electronic limit order books with other market structures, reporting a good relative performance in terms of liquidity.² Although these papers suggest that limit order traders substitute market makers fairly well in providing immediacy and depth to the market, there may be significant differences in terms of permanence and visibility. First, market makers are forced to continuously posting bid and ask quotes whilst limit order traders are not required to be always providing liquidity and have the freedom to choose whether to submit limit orders to buy or to sell. In this regard, Hasbrouck and Saar (2002) evidence that an important proportion of limit orders submitted to the Island ECN are cancelled between two seconds of their submission. Second, order-driven markets usually provide facilities that allow traders to submit (partially) undisclosed limit orders, also known as iceberg orders or hidden orders. By submitting an iceberg order, the trader only displays a fraction of the total quantity she wishes to buy or sell. Hence, investors do not necessarily know the exact quantity of shares offered or demanded at the posted quotes.

Some recent studies provide evidence about the usage of iceberg orders. Degryse (1999) estimates that hidden orders account for about 16% of the entire limit order book of the Brussels CATS system. Hasbrouck and Saar (2002) report that almost 12% of all order executions and executed shares of the Island ECN involve hidden orders. D'Hondt, De Winne and François-Heude (2001) obtain that iceberg orders represent 14% of limit orders and 45% of the shares offered or demanded of the Paris Euro-NM. Tuttle (2002) obtains that hidden liquidity accounts for 22% of the inside depth in a sample of Nasdaq stocks after the Super SOES implementation. Finally, D'Hondt et al. (2003) report that half of the depth available in the five best levels of the Euronext is undisclosed. Up to date, however, very little is known about the motives and consequences of submitting undisclosed orders. In fact, there is a lot of controversy among market regulators over the possibility to partially hide the size of the orders, since it creates a real trade-off between two desirable features for any market: transparency and liquidity. We have conducted a small informal survey of ten Spanish fund managers and financial services traders. All of them agree that iceberg orders involve large-sized orders. Practitioners argue that iceberg orders are used to avoid unfavorable price movements, obscuring the trading strategy of large investors and/or preventing from being front-run. However, there is no consensus concerning when and for what stocks iceberg order placement is more likely. Moreover, practitioners disagree regarding the market response when hidden volume is detected.

In this paper, we cope with hidden orders in a continuous electronic pure order-driven market: the Stock Exchange Interconnection System (henceforth SIBE) of the Spanish Stock Exchange (SSE). We use six months of limit order book and transaction data on the 36 SSE-listed stocks with the highest activity rates and liquidity levels during the year 2000. We identify all market or marketable limit orders that picked off hidden limit orders, and we use this information to accomplish two goals: firstly, we study the effective contribution of hidden

limit orders to the trading process; namely, we analyze the probability of an order being executed against an iceberg order conditional on different information sets: the trading hour, the state of the limit order book and the information intensity. We evidence that hidden volume concentrates towards the end of the trading session, when the NYSE and the SSE sessions overlap, a period of intense trading activity and extremely high liquidity. Our findings strongly support the hypothesis that liquidity suppliers use hidden orders to mitigate adverse selection costs. In contrast, our findings do not support the hypothesis that hidden orders are used as a mechanism to reduce the option value of limit orders for parasitic traders.

Secondly, we investigate the market response when hidden volume is exposed to all traders; namely, we evaluate the information content of iceberg orders by studying its impact on the stock returns and volatility. We also examine the incidence of iceberg orders detection on trader's strategies by analyzing its impact on the composition of the order flow. Our results suggest that hidden orders have no relevant price impact but they temporally increase the aggressiveness of traders when revealed to the marketplace.

The paper proceeds as follows. Section 2 discusses the reasons why limit order traders may choose not to display the whole size of their orders. Section 3 details the structure of the SIBE with special attention to hidden orders. Section 4 describes the database and provides some statistics about the 36 stocks in the sample. Section 5 studies the probability of hitting iceberg orders depending on the trading hour. Section 6 analyzes the probability of hitting iceberg orders conditional on the limit order book and several proxies for information intensity. Section 7 evaluates the information content of iceberg orders. Finally, section 8 concludes.

2. The rationale for hidden order placement

The desirability of hidden orders embraces a complex interaction between order exposure risk, market liquidity and transparency. In a pure order driven market, traders choose to submit either limit orders or market orders. This decision is relevant since market orders consume posted liquidity and limit orders supply liquidity. If a trade occurs, limit orders result in better execution prices than market orders. The inconvenience is that limit order traders run various risks.

Firstly, the execution of a limit order is not guaranteed. Favorable (unfavorable) news can result in unexecuted limit orders to buy (sell), missing the investment opportunity. The time to execution may depend on many factors, like the limit price, the state of the limit order book, the market conditions, etc. Lo, MacKinlay and Zhang (2002) evidence that execution times of limit orders submitted are very sensitive to the limit price but not to the limit size. Parlour (1998), Foucault (1999), and Handa, Schwartz and Tiwari (2003) model the trader's decision as to whether to place a limit order or a market order. These models suggest that variables like the imbalance between potential buyers and sellers and the volatility of the asset determine the non-execution risk of a limit order and, hence, the mix between market and limit orders. Consequently, the lower the non-execution risk, the higher the expected proportion of limit orders in total order flow. These models do not offer any theoretical prediction regarding the mix between disclosed and (partially) undisclosed limit orders. What is more, assuming that Lo et al.'s (2002) findings can be extended to other trading venues, reducing the time to execution should not be a reason for keeping in secret part of the order size. Nonetheless, intuition suggests that given a non-execution risk level, the probability of hidden order placement should increase with other exposure risks. Thus, whether to submit undisclosed or

disclosed limit orders emerges as a second stage decision after choosing between ordinary limit orders and market orders.

Secondly, limit order traders face adverse selection costs. Because the limit price is fixed over time, limit orders can become mispriced when new public information arrives. This situation creates a winner's curse problem for limit order traders since their orders are more likely to be executed (at a loss) when they become mispriced (Foucault, 1999). Copeland and Galai (1983) characterized limit orders to buy (sell) as free put (call) options for the entire market. When a limit order executes against an informed agent, it does it in the money. In Glosten's (1994) model, limit order traders implicitly gain from liquidity driven price changes but lose from information driven price changes. Handa and Schwartz (1996) examine the profitability of limit order trading. In their model, the trader's choice between limit and market orders depends on the probability of the limit order being executed against an informed trader since executing orders against permanent price changes is undesirable. Hence, this model predicts exposure costs to decrease with liquidity-driven volatility and to increase with information-driven volatility. In Foucault (1999), however, higher volatility arising from information causes limit order traders to set prices less aggressively, increasing the relative risk of market orders and the proportion of limit orders in the incoming order flow.

In a recent empirical paper, Bae, Jang and Park (2003) decompose the price volatility in volatility arising from noise or liquidity trading and volatility arising from information to test the two competing predictions in Foucault's and Handa and Schwartz' models. They evidence that NYSE traders place more limit orders through SuperDot when the bid-ask spread is large and they expect high transitory price volatility. The impact of informational volatility is indeterminate, a finding that supports neither Foucault's nor Handa and Schwartz' arguments. These authors do not deal with undisclosed limit orders. However, our intuitive hypothesis is

that limit order traders prefer to submit partially undisclosed orders whenever the perceived risk of trading against informed traders is high.

Since losses with informed traders would be increasing with the limit order size, iceberg orders are likely to be associated with large buyers and sellers that become less willing to display their trading. If informed traders only know the displayed depth available at prices where they can profitably trade, they may submit orders for fewer shares than they would if they could also see the hidden side of depth. In this manner, undisclosed order traders may reduce their adverse selection costs. Because news may always happen, there would be a strictly positive probability of hidden depth in the book, and the size of the hidden depth (relative to the displayed depth) would increase with the probability of such an event.³

Moreover, informed traders are usually characterized as impatient traders. Patient traders place limit orders and supply liquidity whereas the impatient ones place market orders and consume liquidity (e.g., Glosten, 1994). Informed traders should prefer to submit market orders since their information is short lived and immediate execution of these orders is guaranteed. Nonetheless, Kaniel and Liu (2001) argue that informed traders would prefer to submit limit orders when their information is long-lived or their valuation is close to the current market quotes. A patient informed trader would be more reluctant to submit large market orders since by doing so she would signal that the stock is mispriced. If she submits an aggressive market order when her valuation is close to the current quotes, she would bear price risk. Since large limit orders might also reveal her information advantage, the patient informed trader would be more willing to submit partially undisclosed limit orders.

Finally, limit order traders face parasitic traders, also known as front-runners or quote-matchers. Even though limit order traders are not informed, the limit order book may be informative for other traders. Parasitic traders try to infer information about the security's

value from the exposed liquidity and use front-running strategies to exploit that information. These traders are professional ones, with some kind of technological, informational or geographical advantage that allows them to issue and cancel orders more quickly than other traders (Harris, 1999). Since an asymmetric limit order book may reflect the market sentiment or the presence of informed traders, parasitic traders might want to step in front of the heavy side of the book. Even if an asymmetric book is uninformative about future price changes, this trading strategy might be profitable. If prices move against the front-runner, she may limit losses by trading with the heavy side of the book (Harris and Panchapagesan, 1999). Consequently, if parasitic traders act frequently, limit order traders would be less willing to display their trading interests. Therefore, we should expect limit order traders to place more iceberg orders when the risk of being front-run is high.

The design of the trading system may aggravate the exposure costs of limit order traders. First, the parasitic traders' activity is facilitated when the complete limit order book is accessible in real time for all public traders, as it is usually the case in pure electronic order driven markets. Therefore, parasitic traders benefit from pre-trade transparency (see Madhavan, 2000). Second, the size of the minimum price increment (tick) determines the cost of obtaining precedence through price priority in a financial market and, hence, the profitability of front-running strategies. Several papers evidence that decreasing the NYSE minimum price variation reduces the quoted size throughout the limit order book (e.g., Goldstein and Kavajecz, 2000) and induces the specialist to step ahead of the heavy side of the book more frequently (e.g., Harris and Panchapagesian, 1999).

Moreover, order driven markets use secondary order precedence rules, like the time precedence rule and the display precedence rule, to encourage limit order traders to display orders. These secondary order precedence rules, however, are only meaningful when

protected by an economically significant tick (Harris, 1999). Harris (1996) evidences that in the Paris Bourse and the Toronto Stock Exchange large ticks are associated with greater order display. Similarly, D'Hondt et al. (2001) for the Paris Euro-NM and Aitken, Berkman and Mak (2001) for the Australian Stock Exchange provide cross-sectional evidence that the use of hidden orders is negatively related to the relative tick size. These findings support the hypothesis that undisclosed limit orders are used to reduce the option value of limit orders as a defensive strategy against parasitic traders.

In summary, iceberg order placement is expected to increase whenever limit order traders perceive an increase in the exposure risk, either due to higher adverse selection costs or higher activity by parasitic traders. By submitting undisclosed limit orders, they reduce the option value of their orders, mitigate information asymmetry risk and prevent unfavorable price reactions.

From a market design point of view, iceberg orders represent a real trade-off between liquidity and transparency. On the one hand, trading systems need to attract liquidity and they want traders to display that liquidity. Iceberg orders encourage traders to supply liquidity when they might be reluctant to fully disclose their trading interests. On the other hand, iceberg orders diminish the benefits of a transparent limit order book. The dissemination of order book information is assumed to sharply reduce the costs of monitoring the market, permits real-time assessment of liquidity, reduces information asymmetries and improves price efficiency (e.g., Madhavan, 2000; Bloomfield and O'Hara, 2000). Therefore, order-driven systems that allow traders to submit iceberg orders impose certain degree of opacity in an apparently transparent trading mechanism. Iceberg orders increase the total available liquidity, but liquidity demanders only learn about the undisclosed size when these iceberg orders execute. Certainly, limit order traders might protect themselves from an elevated

exposure risk using other strategies rather than submitting hidden orders. For example, by breaking up their large orders into small ones and spreading them over time, the limit order trader may reduce the probability of being front-run and the price impact of their orders. They can also cancel and modify orders more frequently or simply switch to market orders. These alternative mechanisms of protection, however, might increase the trader's transaction costs and either reduce or consume liquidity.

3. Institutional background

Since 1995, the Spanish Stock Exchange Interconnection System (SIBE) is the computer-assisted trading platform that communicates the four Spanish Stock Exchanges (Madrid, Barcelona, Bilbao and Valencia). This electronic system holds the trading activity of the most active and liquid stocks of the SSE. Drawing on its leading-edge technology, the SIBE enables large trading volume to be handled efficiently and transparently, providing real time information and immediate dissemination of trading data.

The SIBE is organized as an order-driven market with a continuous trading session from 9:00 a.m. to 5:30 p.m. and two call auctions, the first one determines the opening price (8:30-9:00 a.m.) and the second one determines the official closing price (5:30-5:35 p.m.). Investor's orders are placed in the electronic system through brokers. During the pre-opening auction, orders can be entered, altered or cancelled, but no trade occurs. The limit order book is partially visible since tentative equilibrium auction prices and volumes are publicized and revised continuously. The auction finishes with a 30-seconds random-end period that fixes the definitive opening price. During the continuous trading session, orders can be submitted, modified or cancelled. A trade takes place whenever a counterpart order hits the quotes. The open market is governed by a strict price-time priority rule: limit orders at the highest bid or the lowest ask take priority over the other orders. At a given price, orders are executed in the

sequence they are submitted to the system. An order might lose priority when modified. Stocks in the SSE are quoted in euros. The minimum price variation (tick) depends on the trade price. The tick equals €0.01 for prices below €50 and €0.05 for prices above €50.

Three basic types of orders are allowed. Market orders are executed against the best prices on the opposite side of the book. These orders walk up the book until they are fulfilled. Market to limit orders do not specify a limit price but are limited to the best opposite-side price on the book at the time of entry; the non-executed part is stored in the book as a limit order at that price. Finally, limit orders are to be executed at the limit price or better. For all of these orders and during the continuous market, brokers may specify special execution restrictions: “execute and eliminate”, “minimum execution” and “fill or kill”. By default, orders expire at the end of the session. Nonetheless, the broker can enter a specific expiration date for each order submitted. The maximum validity period is 90 calendar days. The minimum order size is one share.

The SIBE allows submitting partially undisclosed limit orders. Only the supervisor of the SIBE and the broker that submits the order know of its presence. Managers of the SIBE consider iceberg orders as “interesting for large orders, so as to avoid adverse price movements” (see SIBE, 2001). The investor chooses the “displayed volume unit” of the order, with a minimum of 250 shares. A new displayed volume unit emerges as soon as the current one is executed. The hidden part of the order loses, however, its time precedence. During the call auctions, hidden orders take part at their total volume.

In the SIBE, iceberg orders might be used to camouflage large pre-arranged trades by submitting, in sequence, two matching orders. During the Special Operations Market (SOM) which opens after the closing call auction (5:40 to 8:00 p.m.), a member of the SSE can cross his/her clients’ buyer and seller-initiated orders. This pre-arranged trades or “applications”,

however, are subject to very restrictive price and minimum size conditions. Moreover, the broker cannot accumulate orders from different clients. During the open trading session this kind of pre-matched trades are forbidden. Nonetheless, the iceberg order placement might facilitate members to match large-sized trades at the current quotes or at prices inside the current bid-ask spread and to pull together orders from different clients. In addition to hidden orders and the SOM, brokers may also manage large volume orders through the Block Market (BM). Trading hours are from 9:00 a.m. to 5.30 p.m. The BM handles pre-arranged trades and competitive large orders but, again, under very rigid price and minimum size conditions.

The SIBE includes a computerized Dissemination Information System (DIS) designed to distribute trade, order book and index information in real time. This system guarantees a high degree of pre-trade and post-trade transparency (Madhavan, 2000). The most important indicator of how the SSE is performing is the IBEX-35 index. It is composed of the 35 most liquid and active SIBE securities during the most recent six-months control period. The composition is ordinarily revised twice a year, but extraordinary revisions are possible due to major events like mergers or new stock issues. The IBEX-35 is computed as a cross-stock average trade price weighted by market capitalization.⁴

4. Data

The database includes all the information provided in real time by the SIBE DIS to any entity connected to the system. Therefore, we have the same information that any broker member of the SSE receives from her vendor feed. The limit order book (LOB) data contain the five best buy and sell positions, including quotes, depths and number of orders. The book is updated each time there is a change in these best positions. Additionally, the trade related information details the price, the volume and the counter-parties of the trade. A new register enters each time a trade is carried out on the system. We consider the sample period July-

December 2000, 124 trading days. The database includes information on all SIBE-listed stocks. However, we keep only those stocks that traded the whole year and were included in the IBEX-35 index at least 6 months (36 stocks). We only exclude one stock that disappeared due to a merger. Table I provides some descriptive statistics on the 36 stocks, including the average IBEX-35 weight and measures of liquidity, trading activity and volatility averaged per 30 minutes intervals. Albeit these are the most liquid and active stocks of the system and they cover the major portion of the market activity, Table I evidences remarkable differences among them in terms of market capitalization, liquidity, activity and volatility. Usually, the most liquid stocks are the most active but not necessarily the less volatile ones. Using Spearman rank correlations, we obtain a significant correlation between market capitalization and trading activity (0.65 with share volume and 0.61 with the number of trades) at the 1% level and between market capitalization and liquidity (0.27 with depth and -0.25 with spread) at the 10% level. Market capitalization and volatility are not significantly correlated.

[Table I]

We have perfectly matched both the book data and the trade data. A direct comparison of both files allows us to classify all the movements of the LOB into cancellations, modifications, market or marketable limit orders, market-to-limit orders and limit orders. Buyer and seller-initiated trades are labeled as aggressive whenever they exhaust, at least, the best quote on the opposite side of the book. This implies that an aggressive order consumes the total (both disclosed and undisclosed) depth available at that quote.

Since hidden orders are not marked as such on the SIBE DIS screens, we cannot identify all iceberg order submissions and cancellations. However, we can detect executed iceberg orders by comparing the reported trade size with the associated change in the LOB. That is, we do not have information on all the hidden orders submitted and the specific times of

submission; we are making inferences only with the hidden orders that become partially or totally executed. Consequently, our findings are conditional on the implicit assumption that the subsample of executed iceberg orders is representative of the whole sample of iceberg orders submitted. All the theoretical arguments revised in section 2 suggest that iceberg orders should be mainly competitive limit orders, that is, limit orders inside the best quotes or limit orders hitting the best quotes. If hidden order traders submit non-competitive limit prices, there would be scarce arguments to motivate a partial concealing of the order size. Firstly, front-runners would be a minor problem since we would be dealing with patient traders that do not worry about the time to execution. Secondly, the larger premiums/discounts associated with non-competitive limit order prices could be enough to compensate the higher adverse selection costs associated with large limit orders. Therefore, we believe that our sample of executed hidden orders is a good approximation of the whole set of hidden orders submitted. For the same motive, we do not expect the time of execution to be far apart from the time of submission. Roughly 60% of all trades involving hidden volume are fulfilled when there is only one limit order supporting the best quote on the LOB. Following backwards the evolution of these limit orders, it is possible to find the exact time of submission for 104.452 iceberg orders. Around 93% of these orders were placed inside the best quotes. The median time-to-execution for these orders is 3 minutes. Finally, since we focus on the effective contribution of hidden orders to trading and the impact of hidden volume detection, the information on hidden orders executions is enough to accomplish our goals. Table I reports that, on average, 26.26% of all trades in our database involve hidden volume, 19.70% of all non-aggressive trades and 41.68% of all aggressive trades. We do not find remarkable differences between buyer and seller initiated trades.

5. The intra-daily distribution of hidden orders

In this section we study the intra-daily distribution of hidden orders. We divide the continuous trading session of the SSE (from 9:00 a.m. to 17:30 p.m.) into seventeen half-hour intervals. For each interval we compute the number of trades comprising hidden orders. We also compute the time-weighted average bid-ask spread, the time-weighted average cumulative depth at the five best ask and bid quotes, the volume (number of shares) executed and the number of trades completed. Volatility is the ratio between the highest and the lowest quote midpoint in each interval minus one. To facilitate comparisons, each of the previous variables has been normalized and standardized for every stock as follows. We divide each observation by the average of the corresponding day. Then, we compute the mean of these normalized observations for each half-hour interval. Finally, we subtract one from each half-hour mean and divide it by the standard deviation of the 17 half-hour intervals. Figure 1 represents the equally weighted average of these normalized and standardized measures for the 36 stocks.⁵

[Figure 1]

This figure shows the familiar U-shaped patterns in volatility and trading activity but not in quoted spreads and depths. Spreads achieve their maximum at the opening interval but do not increase towards the end of the session. Depths achieve their minimum at the beginning of the day and their maximum at the end of the day. Therefore, the SSE attains its highest liquidity levels towards the end of the day. For the hidden volume series, we obtain average correlations of 75.24% with the number of trades, 89.95% with the share volume, -17.76% with the time weighted average bid-ask spread, 42.58% with the time weighted average quoted depth, and 36.23% with volatility.

Figure 1 also reports a new finding: the effective contribution of hidden orders to the trading activity increases towards the end of the trading session, reaching its maximum during the last intervals. The patterns indicate that hidden orders emerge in periods of intense trading activity and extremely high liquidity. The Kruskal-Wallis test for equality of medians confirms all these patterns are statistically significant.

Figure 2 represents the median percentage of trades involving hidden volume per half-hour interval. There is a remarkable and statistically significant jump just 1 hour before the opening of the NYSE (15:30 Spanish time) and the highest percentages are achieved during the next trading hour (15:30-16:30). Only during the 14:30-close interval the effective contribution of iceberg orders is significantly above the median (13%).

[Figure 2]

Finally, we study whether trade size matters in explaining the previous findings. We classify all trades into 7 size categories created in the following way. For each firm in the sample, we separate buyer from seller-initiated trades. Trades are ordered by their size in shares. Trades are then partitioned into the following seven percentile categories: less than 25% (S1), 25-50% (S2), 50-75% (S3), 75-90% (S4), 90-95% (S5), 95-99% (S6) and greater than 99% (S7). The size cut-offs depend on each firm's trade size distribution. Hence, large and small trades are defined for each firm relative to its own order flow experience over the entire sample period. Figure 3 presents the equally weighted average of the normalized and standardized number of trades in each size category comprising hidden volume.

[Figure 3]

This figure reports similar patterns to those already seen in previous figures. Hidden volume happens to be especially relevant towards the end of the trading session, and the inflexion point is the opening of the NYSE. The relevant point here is that this pattern does not depend

on the trade size category. In addition, Figure 3 evidences an increase in the average trade size after the opening of the NYSE, suggestive of a concentration of large traders, like institutional investors, towards the end of the trading day.

The patterns shown in hidden volume suggests non-uniform exposure costs in the course of a session. During the opening period, liquidity traders may perceive a higher information asymmetry risk, reflected in the lowest liquidity levels, which accentuates the winner's curse problem of submitting limit orders. In addition, during the first half-hour of trading the volatility achieves its maximum, 263% larger than its mean (see Figure 1). In this context, liquidity providers should prefer to submit (partially) undisclosed limit orders to mitigate adverse selection costs. Both informed and uninformed investors may prefer to hide large limit orders as to avoid unfavorable price movements and reduce the option value of their orders.

Similarly, the increase in hidden volume just about the opening of the NYSE might be explained by an upward revision in the information-asymmetry risk around the initial trading intervals at the US market.⁶ Public announcements in the SSE are specially frequent and regular between 14:30 and 16:30. In addition to all the information about the opening of the NYSE, unemployment rates, consumption price indexes, GDP data, the University of Michigan's consumer confidence survey, the European Commission's Economic Sentiment Indicator (ESIN) and profit announcements of many firms are usually disseminated during this trading interval.

Previous figures show that the hidden volume is far less frequent during the opening than between 14:30-16:30. Therefore, in addition to adverse selection costs, other arguments must explain the dramatic ultimate increase in iceberg orders detection. D'Hondt et al. (2001) report a large number of hidden orders submitted in the Paris Euro-NM just before the end of

each auction. These authors hypothesize that the informational value of an order for parasitic traders increases as the end of an auction approaches and, hence, exposure risk becomes higher. Admati and Pfleiderer's (1988, 1989) work suggests an alternative explanation, that hidden order traders may be trying to take advantage of the benefits of concentrated trading.⁷ In fact, during the last trading hours at the SSE, the average trade size reaches maximums and trading activity is more intense than in any other moment (see Pascual, Pascual-Fuster and Climent, 2003). Hidden order executions should be more difficult to perceive when a large number of limit orders are submitted simultaneously and trading is heavy. These periods may be also particularly appropriate to hide large applications.

Finally, our previous discussion suggests that market participants whose quotes contribute the greatest information to the market are more likely to use hidden size. Chakravarty (2001) evidences that the so-called "stealth trading" hypothesis, where privately informed traders concentrate their trades in medium-sizes, appears to be primarily driven by medium-sized trades initiated by institutions. Since their quotes may have a greater signaling effect, these traders will be more prone to use undisclosed orders in order to reduce free-rider costs of displaying size. It is a shared opinion among the SSE traders that the increase in the average trade-size during the last trading hours reflects the presence of large institutional investors trying to balance their portfolios. These traders may retain their orders until all the pending news is finally disseminated around 14:30-16:30.

6. Hidden orders, the limit order book and the information intensity

As it was discussed in section 2, the decision of (partially) not disclosing the size of a limit order depends on the non-execution risk, which determines the decision of submitting either an ordinary limit order or a market order (e.g., Foucault, 1999), and on the order exposure risk, which determines whether to submit an ordinary limit order or a hidden limit order.

Whenever liquidity providers believe that the order exposure risk has augmented, either because they perceive an increase in the information asymmetry risk or a higher risk of being front-run, the probability of them choosing undisclosed limit orders also increases.

In this section we study the probability of undisclosed limit order execution by analyzing the determinants of the effective contribution of hidden orders to trading. To achieve this aim, we define H_t as a binary variable that equals one whenever a market (or marketable limit) order is executed against hidden volume and zero otherwise. Assume there is an underlying response variable that measures the liquidity provider's judgment about the order exposure risk, OER_t , defined by the linear regression $OER_t = X_t' \beta + u_t$, where X_t represents a set of control variables, β is a vector of unknown parameters and u_t is the noise term. Whenever the perceived level of order exposure risk is high, say $OER_t > \delta$, the liquidity provider does not reveal the total size of his/her position. In this way, hidden volume concentration patterns are the result of a shared perception among liquidity providers about the market signals; that is, hidden volume clusters occur when many traders coincide in their judgment about the order exposure risk. Whilst OER_t is not observable, we do can observe the outcome H_t such that $H_t = I(OER_t > \delta)$, where $I(\cdot)$ is an indicator function taking the value 1 if the condition within parenthesis is satisfied and zero otherwise. If we assume that u_t is normally distributed, we get a Probit model,

$$E(H_t | X_t) = \Pr(H_t = 1 | X_t) = \Phi(X_t' \beta), \quad [1]$$

with $\Phi(z)$ being the standard normal cumulative distribution function.

We assume that liquidity traders evaluate the order exposure risk based on two information sets, the status of the limit order book and the information intensity. All variables are computed over a 5-minutes window before each trade. Book information consists of the

quoted bid-ask spread (SPR_t), the number of ticks between the best bid and ask quotes; the net depth, the difference between the number of shares on the offer and bid sides of the book, distinguishing between net depth at the best quotes (ND_t^1) and net depth between the 2nd and the 5th best quotes (ND_t^{2-5}); the net number of limit orders supporting the posted quotes, distinguishing between net orders at the best quotes (NLO_t^1) and net orders between the 2nd and the 5th best quotes (NLO_t^{2-5}); and the net “dispersion” of the book ($NDISP_t$), measured as the depth-weighted average distance (in ticks) from the five best ask quotes to the quote midpoint minus the equivalent measure for the bid side of the book.⁸ For all these book indicators, we compute time-weighted averages over the 5-minutes window.

The second set of variables signifies the intensity of information arrival to the market. The information intensity set consists on the net number of trades (NT_t), computed as the difference between the number of buyer-initiated and seller-initiated trades; the net share volume (NV_t), defined as the difference between the buyer-initiated share volume and the seller-initiated share volume; the net number of limit orders, distinguishing between limit orders inside the best quotes ($NLOIQ_t$), limit orders hitting the best quotes ($NLOBQ_t$), and limit orders with worse limit-prices than the best quotes (NLO_t), all them computed as the difference between the number of limit orders to sell and the number of limit orders to buy; and the net number of cancellations, defined as the difference between the number of cancellations at the ask side and the bid side of the book, setting apart cancellations at the best quotes ($NCBQ_t$) from cancellations at worst prices than the best quotes (NC_t). All these measures characterize the recent order flow history before each trade.

Finally, we consider two additional variables: the cumulative mid-quote returns (CMR_t), computed as the signed change in the quote midpoint, and the ratio between the highest and

the lowest quote midpoint (QV_t). The first measure captures the recent price trend and the second is a proxy of the mid-quote volatility.⁹

All previous variables have been standardized for each of the seventeen half-hour intervals of the trading session. As a result of the large number of observations, most of the correlations (not reported) among the variables are found to be statistically different from zero, even though most of them seem to be negligible. The highest average correlation is found between ND_t^{2-5} and NLO_t^{2-5} (0.7), far more larger than that between ND_t^1 and NLO_t^1 (0.26). The recent price trend (CMR_t) is fairly correlated with measures of the information intensity set, like NT_t (0.37) or $NLOIQ_t$ (-0.36). These correlations are suggestive of multicollinearity problems. Some preliminary tests, however, show no relevant changes in the coefficients estimates depending on which set of variables is included in the Probit model.

Table II summarizes the estimation results of the model [1] for the 36 stocks in the sample. For each explanatory variable, we report the median, maximum and minimum of the slope of the probability function, evaluated at the average value of the explanatory variables and computed using the estimated β coefficients of the estimated model [1],

$$\frac{\partial}{\partial x_t^i} \Phi(\beta x_t^i) = \phi(\hat{\beta} \bar{x}^i) \hat{\beta},$$

where $\phi(\cdot)$ is the standard normal density function. We also provide the number of statistically significant coefficients at the 1% level and their sign.

[Table II]

The discussion in section 2 suggests several testable hypotheses regarding the probability of hidden order execution. On the one hand, if undisclosed limit orders are used to reduce the option value of limit orders as a defensive strategy against parasitic traders, then:

Hypothesis 1: Liquidity providers will hide more volume on the heavy side of the book, since parasitic traders are expected to infer the market sentiment from asymmetric limit order books.

Hypothesis 2: Liquidity providers will hide more volume during periods of intense price improvement in the same side of the book, because this may be a signal of front-running activity.

Hypothesis 3: Liquidity providers will hide less volume when the bid-ask spread is tight, because the likelihood of being front-run decreases and the time precedence rule is more effective.

On the other hand, liquidity providers would be averse to submit large limit orders whenever the perceived risk of trading against informed traders is high. Sellers (buyers) will raise their perceived information-asymmetry risk whenever the book accumulates more orders on the bid (ask) side, the length of the book on the ask (bid) side is larger than on the bid (ask) side, price improvement is more common on the bid (ask) side, or seller (buyer) initiated trades are less frequent than buyer (seller) initiated trades. Therefore,

Hypothesis 4: Limit order traders hide more volume when providing liquidity against the market trend because they face a higher information asymmetry risk.

The empirical evidence in Table II provides little support to the argument that undisclosed volume is used as a mechanism to reduce the option value of limit orders for parasitic traders. Indeed, hypotheses 1 to 3 are rejected for the greater part of the sample. First, we found a strong and negative (positive) relationship between net depth and hidden volume on the ask (bid) side of the book, rejecting the hypothesis 1. This relationship is stronger for ND_t^1 than for ND_t^{2-5} . Table II also reports a positive (negative) association between NLO_t^1 and hidden

volume on the ask (bid) side of the book. Altogether, hidden volume on the ask (bid) side of the book is more frequent when the market is heavier on the bid (ask) side and the average size of the limit orders to buy (sell) is larger. Second, we found a strong negative (positive) relationship between $NLOIQ_t$ and hidden volume on the ask (bid) side of the book; that is, when price improvement is more frequent on the bid (ask) side of the book, hidden volume is more common on the ask (bid) side of the book. This result contradicts hypothesis 2. Finally, the relationship between the bid-ask spread and hidden volume detection is negative, which disagrees with hypothesis 3.

On the contrary, the results in Table II are more supportive of the hypothesis that undisclosed limit orders are used to mitigate adverse selection costs. Table II suggests that hidden volume is placed on the side of the book with higher adverse selection costs. Undisclosed limit orders to buy (sell) are more frequent when the net depth is positive (negative) and the average size of the posted limit orders to sell (buy) at the best quotes is larger. An asymmetric book may reflect the traders' sentiment, the likely direction of future price changes, or even informed traders. A larger average order size on the opposite side of the book may also reveal the presence of institutional investors or informed traders. Moreover, a higher degree of dispersion on the ask (bid) side of the book, $NDISP_t > 0$ (< 0), increases the probability of hitting iceberg orders on that side of the book. Therefore, hidden orders are found in the side of the book where the expected price impact of trades is larger, traders are less willing to provide liquidity and there is an apparent weaker consensus about the true value of the stock.¹⁰ The variables in the information arrival set provide additional support. The probability of hidden volume on the ask (bid) side of the book is positively (negatively) related with NT_t and negatively (positively) associated with $NLOIQ_t$. This finding means that undisclosed limit orders are more likely used when providing liquidity

against the market trend, which is consistent with previous findings. Notice that NT_t is much more informative about iceberg orders than NV_t .

Mid-quote volatility (QV_t) has a strong negative effect on the use of hidden volume. Handa and Schwartz (1996) suggest that a rise in volatility due to informed trading discourages the placement of limit orders. Foucault (1999), however, predicts just the opposite: when volatility due to informed trading increases, limit order traders' risk of being picked off by informed traders increases. Consequently, traders post less attractive quotes and this causes market order trading to be more costly. As a consequence, the proportion of limit orders in the order flow increases.

How can these theoretical predictions be reconciled with our empirical finding? On the one hand, if information-induced volatility causes limit order traders to post less aggressive prices so that higher margins compensate higher adverse selection costs, the use of hidden orders will be needless. This argument is also consistent with the weaker but negative relationship between SPR_t and the probability of picking off hidden orders. On the other hand, information-driven volatility may discourage all kind of large limit order traders either disclosed or undisclosed. But, what if QV_t is not information-driven. Both Glosten (1994) and Handa and Schwartz (1996) predict that transitory volatility triggered by noise traders encourages limit order placement. This is because, other things being equal, the greater the volatility, the greater the probability of a limit-order being picked off and, consequently, the shorter the expected time to execution. Therefore, if QV_t only reflected transitory volatility there would be no reason to expect a positive relationship with the probability of hidden order detection.¹¹

7. The information content of hidden orders

In this section, we investigate the market response when hidden volume is exposed to all traders. We evaluate the information content of iceberg orders by studying its impact on the stock returns. We also study how it influences trading strategies by analyzing its effect on liquidity and the composition of the order flow.

For each minute in a five-minute window before each trade, we compute the following variables, already defined in the previous section: the bid-ask spread (SPR_t) and the net depth ($ND_t^1 + ND_t^{2-5}$) both weighted by time, the net number of trades (NT_t), the net share volume (NV_t), and the net number of limit order submissions, pooling $NLOIQ_t$, $NLOBQ_t$ and NLO_t in previous section. We also compute the cumulative mid-quote returns (CMR_t) and the mid-quote volatility (QV_t) 5, 10 and 30 minutes after the trade and until the closing of the trading session. These variables summarize the aspects of the market this section focus on: liquidity, order flow, volatility and returns. We separate buyer from seller-initiated trades and standardize the variables for each of the seventeen half-hour intervals of the trading session. Finally, we bring together all trades belonging to the same trade-size category (S1 to S7). We use this information to construct matched samples of trades evidencing hidden volume and ordinary trades. Namely, we construct two matched samples for each of the S1-S7 trade-size categories, one for buyer-initiated trades and another for seller-initiated trades.^{12, 13} Table III provides some details on the resulting matched samples. This table shows that the matching becomes more difficult as the trade size augments and the pre-trade matching interval increases from 1 to 5 minutes. We discard S7 trades from the posterior analysis because of the small number of matched observations. Table III also evidences that there are no remarkable trade size differences between the matched samples.

[Table III]

Next, for each variable and for each trade size category, we construct 1000 subsamples of 1000 matched trades randomly selected from the respective samples in Table III. For each subsample, we compute the non-parametric Wilcoxon test of equality of medians. Tables IV and V report the proportion of cases for which the null hypothesis of equality of medians is rejected and the proportion of cases for which the null is rejected against the alternative hypothesis that the median of the corresponding variable for the 1000 trades revealing hidden liquidity is larger than the median for the matched ordinary trades.

Table IV focuses on liquidity and order flow in the 10-minute interval centered on the execution of each trade. This table reports the results for the S4 trade-size category, trades with size between the 75% and 90% percentiles of the empirical trade-size distribution of the corresponding stock, and for a 1-minute pre-matching interval. The results for the other trade size categories and for the 5-minute pre-matching interval are identical and available upon request from the authors.

[Table IV]

This table shows that the hidden volume detection has a remarkable impact neither on the net limit order book depth nor on the net limit order submission. Nonetheless, there is evidence of a narrower bid-ask spread and a more intense buyer (seller) initiated trading activity after a trade that discovers the presence of hidden volume on the ask (bid) side of the limit order book than after an equally-sized and pre-trade matched ordinary trade, at least up to one minute after. The presence of hidden volume reduces the impact of posterior trades and this may explain the narrower bid-ask spread with respect to an ordinary trade of equal size. The increased trading activity together with the irrelevant impact on the submission of limit orders indicates that the aggressiveness of traders temporally increases when the presence of

hidden volume is exposed to the market. The possibility of buying or selling larger quantities at the same cost induces traders to submit more aggressive orders to consume the hidden part of the available liquidity. These effects, however, dissipate quickly.

Table V provides the findings for the cumulative mid-quote returns and the mid-quote volatility. If hidden orders are informative, we should observe some positive (negative) impact on stock returns when a hidden buy (sell) volume is exposed to the market. Indeed, Tuttle (2002) observes that the presence of hidden depth on the relevant side of the market indicates that the price will move against the trade. Additionally, she detects that price reactions are stronger when market participants quoting hidden size are those working large institutional orders. Our findings in Table V, however, do not support the hypothesis that hidden orders are information-motivated. There are no statistically significant difference between the return of a trade revealing hidden volume and the return of an equally sized and matched ordinary trade, independently of the time horizon considered. The results with a 5-minute pre-trade matching interval (not reported) are exactly equal. Regarding volatility, it is weakly lower in the short-run after a trade revealing hidden volume but, again, this effect dissipates in a few minutes. Therefore, these findings indicate that the market does not attribute any unknown information content to the hidden side of liquidity.

8. Summary, conclusions and final comments

In this paper, we have studied hidden orders in a continuous electronic pure order-driven market, the Stock Exchange Interconnection System of the Spanish Stock Exchange. Using six months of limit order book and transaction data on the stocks with higher activity rates and liquidity levels, we have identified all market (or marketable) orders that picked off hidden limit orders, and we have used this information to study the effective contribution of hidden limit orders to the trading process.

We have shown that hidden volume concentrates towards the end of the trading session, when the NYSE and the SSE sessions overlap. Our findings strongly support the hypothesis that liquidity suppliers use iceberg orders to mitigate adverse selection costs. There is no evidence that corroborates the alternative hypothesis that hidden orders are submitted to reduce the option value of limit orders for parasitic traders. Finally, our results suggest that hidden orders have no relevant impact on returns or volatility but, when they are revealed to the marketplace, they temporally increase the aggressiveness of traders.

Our findings contrast with the practitioners' beliefs about the use of hidden orders. We provide evidence suggesting that iceberg orders are not mainly used to avoid unfavorable price movements since no relevant permanent price impact is associated with the discovery of hidden volume. Our results also contradict the idea that hidden orders are used to prevent from being front-run since all the theoretical arguments based on this hypothesis have been statistically rejected. On the contrary, our findings indicate that hidden orders are a vehicle to obscure the trading strategy of large and uninformed investors aimed to reduce adverse selection costs. This behavior would explain why the market does not attribute any relevant information content to the discovery of undisclosed liquidity. Therefore, market regulators could draw two basic implications from the SSE experience: the main motive of hiding volume is protection against informed traders and iceberg orders have distorting effects neither on prices nor on volatility.

This study suggests additional questions for further empirical research. Our findings report that traders submit iceberg orders to manage the information asymmetry risk. Therefore, we should expect the use of iceberg orders to be more frequent in illiquid and infrequently traded stocks, foreign stocks and highly volatile stocks (like companies of new technologies or R&D

firms), usually associated with higher adverse selection costs. A cross-sectional analysis of these and alternative sets of stocks could provide further insights in the use of hidden orders.

Footnotes

1. For instance, Biais, Hillion and Spatt (1995) report that investors in the Paris Bourse place limit orders when the spread is large or the order book is thin. Limit order placements that improve upon the best quotes occur in succession and on the same side of the book, suggesting competition in the supply of liquidity. Ahn, Bae and Chan (2001) observe that investors in the Hong Kong Stock Exchange submit more buy (sell) limit orders relative to market orders when liquidity-driven transitory volatility rises from the ask (bid) side (see Hamao and Hasbrouck, 1995, Danielson and Payne, 2001, and Rinaldo, 2003, for further evidence). Coppejans, Domowitz and Madhavan (2002) evidence that volatility shocks in the electronic market for Swedish stocks index futures (OMX) reduce liquidity. Shocks to liquidity, however, dissipate quickly, indicating a high degree of resiliency. Finally, Coppejans and Domowitz (1999) also evidence that an automated market can operate well in a relatively illiquid setting.
2. Seppi (1997) proposes a theoretical model of liquidity provision where a specialist competes against a limit order book. In this model, a hybrid market, like the NYSE, provides better liquidity to small retail and institutional trades, but a pure order driven market may offer better liquidity on mid-sized orders. Degryse (1999) compares the costs of trading Belgian shares on both the order-driven Brussels CATS system and the London SEAQ dealership market. He concludes that total trading costs on the CATS market are smaller for small trade sizes. de Jong et al. (1995) obtain similar findings in a comparison of the Paris Bourse and the SEAQ. Huang (2002) compares the liquidity provision on both the ECNs and the market makers of the Nasdaq. In general, ECNs quoted spreads are smaller than dealer quoted spreads. Pirrong (1996), Frino et al. (1998), and Kofman and Moser (1997) evidence nearly equal bid-ask spreads, adverse selection components, and pricing efficiency in automated limit order books relative to floor trading venues.
3. Aitken et al. (2001) report a cross-sectional positive and significant relationship between the relative price volatility and the use of undisclosed limit orders. Tuttle (2000) also finds the residuals of the market model, as a measure of the idiosyncratic risk, to be positively and significantly related to the use of hidden size. However, she also observes that the mid-quote volatility is negatively associated with the proportion of hidden-size and the market-model beta has no explanatory power.
4. For more complete and detailed information on the SSE regulation, organization and trading procedures please visit <http://www.sbolsas.es>.
5. The intra-daily patterns for capitalization-weighted averages are similar. This information is available upon request from the authors.
6. Pascual, Pascual-Fuster and Climent (2003) study the contribution of the NYSE to the price discovery process of the Spanish stocks cross-listed at the NYSE. All these dually listed stocks are among the stocks in this paper's sample.
7. In their 1988 model, trading clusters result from the strategic behavior of traders. Admati and Pfleiderer argue that liquidity traders have strong incentives to concentrate their trading in order to minimize the price impact of their trades. Even though liquidity-trading concentration attracts informed traders, the competition between informed traders intensifies the forces leading to the cluster of trading by liquidity traders since prices are more informative in those periods. In their 1989 model, trading clusters result from a "divide and conquer" strategy by market makers. Volume concentration serves in this model to discourage short-lived information acquisition by informed traders while giving liquidity traders more favorable terms of trade.
8. A high absolute net dispersion value may be interpreted as a signal of little consensus regarding the true value of the stock; it may be informative about future price changes or it may indicate the presence of informed traders; it may also be interpreted as a measure of the willingness of traders to provide liquidity at a given side of the limit order book; finally, it may also be a measure of the expected price impact of large market orders. Coppejans and Domowitz (1999) construct a similar measure they call the "length" of the book.
9. We have also considered other volatility proxies, like the standard deviation of the quote midpoint, the cumulative squared relative changes of the quote midpoint and even asymmetric volatility, computed as the difference between the cumulative squared relative changes of the best ask and bid quotes. A priori, there is no reason to prefer one proxy to the others.
10. Even if large absolute values of $NDIS_t$ are not associated with future price changes, it may still provide relevant information to traders. Ahn et al. (2001) and Bae et al. (2003) show that traders place more limit orders relative to market orders when they expect high transitory price volatility. Ahn et al. argue that transitory volatility is induced by temporary order imbalances. This liquidity-

driven price volatility “will attract public traders to submit limit orders rather than market orders, as the gains from supplying liquidity can more than offset the potential loss from trading with informed traders” (Handa and Schwartz, 1996). Therefore, $NDIS_t > 0$ (< 0) might precede a raise in transitory volatility induced by the ask (bid) side of the book. If this is the case, $NDIS_t > 0$ (< 0) should encourage the placement of limit orders on the ask (bid) side of the book. Our results do not contradict this argument; they rather suggest that traders prefer to submit partially undisclosed that totally disclosed limit orders.

11. As a matter of robustness, we have also estimated the Probit model [1] for each of the trade-size categories S1-S7 defined in section 5 and for each of the 17 half-hour trading intervals of the SSE continuous session. Explanatory variables are standardized accordingly and the results are consistent.

12. A matched sample design is justified in this case by the empirical findings in the previous section. To ensure that any difference in stock price movements, order flow or trading activity is only attributable to the detection of hidden volume, the situation before a trade detecting hidden volume takes place must be akin to the situation before an equally-sized ordinary trade occurs.

13. The matching procedure for each trade-size category is as follows: We eliminate trades preceded by some other trade detecting hidden volume within the last 5 minutes. For the ordinary trades, we also discard trades followed by some other trade detecting hidden volume within 5 minutes after their execution. These filters reduce the risk of market conditions being affected by previous or posterior orders evidencing undisclosed volume. For the remaining trades, we compute the percentiles 10%, 30%, 50%, 70% and 90% of the empirical distribution of each the variables previously defined. These percentiles are the thresholds of the 5 categories used to characterize the level of all the variables during the five-minute window before each particular trade. For each trade that detects hidden volume, we look for an ordinary trade with perfectly matching prior market conditions. If no matching occurs, the trade of reference is eliminated. If multiple matching occurs, we choose the ordinary trade with the minimum sum of the absolute deviations in prior market conditions with respect to the trade of reference. Two alternative matching periods have been considered: 1 minute before each trade and 5 minutes before each trade. Our main results are independent of the matching interval considered.

References

- Admati, Anat R., and Paul Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies*, 1, 1, 3-40.
- Admati, Anat R., and Paul Pfleiderer, 1989, Divide and conquer: A theory of intraday and day-of-the-week mean effects, *Review of Financial Studies*, 2, 2, 189-223.
- Ahn, Hee-Joon, Kee-Hong Bae, and Kalok Chan, 2001, Limit orders, depth and volatility: evidence from the stock exchange of Hong Kong, *Journal of Finance*, 56, 2, 769-790.
- Aitken, Michael J., Henk Berkman and Derek Mak, 2001, The use of undisclosed limit orders on the Australian Stock Exchange, *Journal of Banking and Finance*, 25, 1589-1603.
- Bae, Kee-Hong, Hasung Jang and Kyung Suh Park, 2003, Traders' choice between limit and market orders: evidence from NYSE stocks, *Journal of Financial Markets*, 6, 4, 517-538.
- Biais, Bruno, Pierre Hillion, and Chester Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *The Journal of Finance*, 50, 1655-1689.
- Bloomfield, Robert, and Maureen O'Hara, 2000, Can transparent markets survive?, *Journal of Financial Economics*, 55, 425-459.
- Chakravarty, Sugato, 2001, Stealth trading: Which traders' trades move stock prices? *Journal of Financial Economics*, 61, 2, 289-307
- Copeland, Thomas E., and Dan Galai, 1983, Information effects on the bid-ask spread, *The Journal of Finance*, 38, 1457-1469.
- Coppejans, Mark, Ian Domowitz, 1999, Screen information, trader activity, and bid-ask spreads in a limit order market, *Working Paper*, Duke University.
- Coppejans, Mark, Ian Domowitz, and Ananth Madhavan, 2002, Dynamics of liquidity in an electronic limit order book market, *Working Paper*, Duke University.
- Danielsson, Jón, and Richard Payne, 2001, Measuring and explaining liquidity on an electronic limit order book: evidence from Reuters D200-2, *Working Paper*, London School of Economics.
- de Jong, Frank, Theo Nijman and Aisla Roell, 1995, A comparison of the costs of trading French shares on the Paris Bourse and SEAQ International, *European Economic Review*, 39, 1277-1301.
- Degryse, Hans, 1999, The total costs of trading Belgian shares: Brussels versus London, *Journal of Banking and Finance*, 23, 1331-1355.
- Domowitz, Ian, and Jianxin Wang, 1994, Auctions as algorithms: Computerized trade execution and price discovery, *Journal of Economic Dynamics and Control*, 18, 29-60.
- D'Hondt, C., R. De Winne and A. François-Heude, 2001, Hidden orders: an empirical study on the French segment of Euro.NM, *Working Paper*, University of Perpignan and FUCaM.
- D'Hondt, C., R. De Winne and A. François-Heude, 2003, Hidden orders on Euronext: Nothing is quite as it seems, *Working Paper*, University of Perpignan and FUCaM.
- Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market, *Journal of Financial Markets*, 2, 99-134.
- Frino, Alex, Tom McNish, and Martin Toner, 1998, The liquidity of automated exchanges: New evidence from German bond futures, *Journal of International Financial Markets, Institutions and Money*, 8, 225-241.
- Glosten, Lawrence, 1994, Equilibrium in an electronic open limit order book, *Journal of Finance*, 49, 1127-1161.
- Goldstein, M.A. and Kenneth A. Kavajecz, 2000, Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the NYSE, *Journal of Financial Economics*, 56, 125-149.
- Hamao, Yasushi, and Joel Hasbrouck, 1995, Securities trading in the absence of dealers: trades, and quotes on the Tokyo Stock Exchange, *Review of Financial Studies* 8, 3, 849-878.
- Handa, Puneet, and Robert A. Schwartz, 1996, Limit order trading, *The Journal of Finance*, 51, 1835-1861.
- Handa, Puneet, Robert A. Schwartz, and Ashish Tiwari, 2003, Quote setting and price formation in an order driven market, *Journal of Financial Markets*, 6, 4, 461-489.
- Harris, Lawrence, 1996, Does a minimum price variation encourages order exposure?, *Working Paper*, Marshall School of Business.
- Harris, Lawrence, 1997, Order exposure and parasitic traders, *Working Paper*, Marshall School of Business.
- Harris, Lawrence, 1999, Trading in pennies, *NYSE conference "U.S. equity markets in transition"*, Scottsdale, Arizona.
- Harris, Lawrence E., and Venkatesh Panchapagesan, 1999, The information-content of the limit order book: evidence from NYSE specialist actions, *Working Paper*, Marshall School of Business, University of Southern California.
- Hasbrouck, Joel, and Gideon Saar, 2002, Limit orders and volatility in a hybrid market: the Island ECN, *Working Paper*, Stern School of Business, New York University.

- Hasbrouck, Joel, George Sofianos and Deborah Sosebee, 1993, New York Stock Exchange systems and trading procedures, *NYSE Working Papers*, #93-01.
- Huang, Roger D., 2002, The quality of ECN and Nasdaq market maker quotes, *Journal of Finance*, 57, 3, 1285-1319.
- Kaniel, Ron, and Hong Liu, 2001, So what orders do informed traders use?, *Working Paper*, University of Texas at Austin.
- Kavajecz, Kenneth A., 1999, A specialist's quoted depth and the limit order book, *The Journal of Finance*, 54, 747-771.
- Kofman, Paul, and James T. Moser, 1997, Spreads, information flows, and transparency across trading systems, *Applied Financial Economics*, 7, 281-294.
- Lo, Andrew W., A. Craig MacKinlay, and June Zhang, 2002, Econometric models of limit-order executions, *Journal of Financial Economics*, 65, 31-71.
- Madhavan, Ananth, 2000, Market microstructure: a survey, *Journal of Financial Markets*, 3, 205-258.
- Parlour, Christine A., 1998, Price dynamics IN limit order markets, *Review of Financial Studies*, 11, 789-816.
- Pascual, Roberto, Bartolome Pascual-Fuster, and Francisco Climent, 2003, Cross-listing, price discovery and the informativeness of the trading process, *Working Paper*, Universidad de las Islas Baleares.
- Pirrong, S. Craig, 1996, Market liquidity and depth on computerized and open outcry trading systems: A comparison of DTB and LIFFE bund contracts, *Journal of Futures Markets*, 16, 519-543.
- Ranaldo, Angelo, 2003, Order aggressiveness in limit order book markets, *Journal of Financial Markets*, 7, 53-74.
- Seppi, Duane, 1997, Liquidity provision with limit orders and a strategic specialist, *Review of Financial Studies*, 10, 103-150.
- Tuttle, Laura, 2002, Hidden orders, trading costs and information, *Working Paper*, Fisher College of Business.

TABLE I
Sample: descriptive statistics

This table reports some descriptive statistics for the 36 most liquid stocks in the SSE in July-December 2000. The table reports the average weight of each stock during the year 2000. The other statistics are computed for a 30-minutes time resolution. Both spread and depth are averages weighted by time. The spread is the difference between the best ask and bid quotes. The tick is 0.01€ for all the stocks. The depth is the total shares offered at the five best ask and bid quotes. The minimum depth per quote is one share. Volume is the number of shares transferred. The minimum trade size is one share. Volatility is the ratio between the max and the min of the quote midpoint minus one. Finally, this table provides the percentage of (all, non aggressive and aggressive) trades involving hidden volume. In each column, we remark in bold format the ten stocks with higher capitalization, narrower spread, larger depth, higher trading activity, lower volatility and higher proportions of hidden volume detection, respectively.

Stock	IBEX-35						Hidden Volume Detection (%)		
	Average Weight (%)	Spread	Depth	Volume	Trades	Volatility	Total Trades	Non-Aggres. Trades	Aggressive Trades
ACR	0.3745	0.0295	25271.85	15883.48	21.83	0.00463	23.75	14.88	41.32
ACS	0.5897	0.111	8355.51	8163.52	14.7	0.00431	21.81	17.07	29.01
ACX	0.5529	0.114	7490	11063.12	19.38	0.00509	25.80	18.85	35.96
AGS	0.6047	0.0607	10514.07	10079.68	13.75	0.00493	28.07	26.85	29.63
ALB	0.6859	0.1694	6019.58	7891.9	11.17	0.00475	25.26	23.59	27.18
ALT	1.6125	0.0395	22629.36	78207.32	45.52	0.00591	26.64	20.08	42.05
AMS	2.2001	0.0268	28813.95	80405.57	54.44	0.00853	30.49	22.13	48.69
ANA	0.8247	0.134	5693.53	7876.4	16.99	0.00423	27.11	21.73	34.71
AUM	0.2915	0.0763	12930.9	6975.17	4.6	0.0014	26.05	24.53	28.81
BBVA	15.3956	0.0163	132761.51	456637.09	119.09	0.00486	24.19	14.92	56.51
BKT	1.1942	0.1117	7075.9	14850.56	32.93	0.00519	24.81	14.96	42.77
CAN	0.7757	0.1142	14283.19	9607.05	8.86	0.0034	25.46	26.24	24.37
CTG	2.6269	0.0629	13412.56	25156.06	21.76	0.00536	33.20	31.32	35.87
DRC	0.4038	0.0342	28075.33	44219.33	25.94	0.0062	30.69	26.56	38.52
ELE	6.7036	0.0233	45427.93	214150.42	88.05	0.00466	26.34	17.61	50.3
FCC	0.7421	0.0762	10302.08	11728.69	19.04	0.00532	26.34	20.77	34.85
FER	0.5673	0.0453	12364.11	11272.55	23.4	0.00506	24.67	15.95	42.08
GPP	0.1478	0.0157	84238.83	68225.26	36.51	0.00982	24.38	9.65	60.34
IBE	3.798	0.0246	55397.08	143763.62	44.63	0.00464	29.22	21.42	45.96
IDR	0.5529	0.0547	10999.82	18657.93	31.27	0.00674	26.93	16.96	44.73
MAP	0.2506	0.1083	11305.52	9195.94	6.93	0.00446	28.76	26.64	31.66
NHH	0.3083	0.0588	19638.23	23728.58	13.71	0.00492	26.59	27.3	25.48
POP	2.1956	0.0819	12404.15	26498.34	27.31	0.00412	30.54	27.2	35.91
PRS	1.6592	0.0741	10117.32	27249.8	32.67	0.00578	22.44	14.78	41.54
REE	0.5065	0.0427	16547.39	9656.78	16.51	0.00473	22.20	13.88	36.88
REP	7.7296	0.023	47288.48	247408.24	97.42	0.00509	26.73	18.26	50.31
SCH	14.162	0.013	283176.86	654880.1	174.31	0.00531	18.31	11.12	61.04
SGC	1.1269	0.1069	6018.71	12844.13	32.14	0.00698	27.33	18.31	40.45
SOL	0.6264	0.0443	15454.02	18492.25	17.25	0.00543	27.57	22.69	35.71
TEF	23.0108	0.0149	117941.66	1335221.13	388.8	0.00705	22.10	11.6	63.16
TPI	1.1352	0.0223	31878.48	78430.54	59.38	0.00892	29.45	19.44	52.21
TPZ	0.4117	0.0157	108745.55	93100.63	49.89	0.00753	21.43	13.02	55
TRR	3.5387	0.0517	15281.65	212069.22	270.45	0.01257	28.10	19.64	51.41
UNF	1.8017	0.0418	33398.14	47491.26	24.38	0.00363	28.23	21.38	39.37
VAL	0.2517	0.0287	25980.4	20988.67	15	0.00506	27.43	23.24	35.4
ZEL	0.6406	0.0829	12138.53	58128.22	104.54	0.00924	26.77	14.73	51.28

TABLE II
Probit Analysis

This table provides, for each standardized explanatory variable, the median, maximum and minimum across stocks of the slope of the probability function of the Probit model [1]. The number of significant coefficients at the 1% level and its sign is also reported. The dependent variable (H_i) is the probability of detecting hidden volume on a given side of the limit order book. For a detailed description of the explanatory variables see section 6. All variables are defined over a 5-minutes interval before each trade. Variables in the “status of the book” set: SPR = bid-ask spread, ND^1 = net depth at the best quotes, ND^{25} = net depth between the 2nd and the 5th best quotes, NLO^1 = net number of limit order supporting the best quotes, NLO^{25} = net number of limit orders between the 2nd and the 5th best quotes, $NDISP$ = net dispersion of the book. Variables in the “information intensity” set: NV = net share volume, NT = net number of trades, $NLOBQ$ = net number or limit orders hitting the best quotes, NLO = net number of limit orders with worst limit prices than the best quote, $NLOIQ$ = net number of limit orders inside the best quotes, $NCBQ$ = net number of cancellations at the best quotes, NC = net number of cancellations between the 2nd and the 5th quote, CMR = cumulative mid-quote returns, QV = mid-quote volatility.

Panel A: Ask Side of the Book

	SPR	ND^1	ND^{25}	NLO^1	NLO^{25}	$NDISP$	NV	NT	$NLOBQ$	NLO	$NLOIQ$	$NCBQ$	NC	CMR	QV	Psd. R2	Obs.
Median	-0.0116	-0.0237	-0.0127	0.0162	0.0063	0.0111	0.0082	0.0248	0.0083	0.0049	-0.0270	0.0068	0.0069	-0.0074	-0.0198	0.0249	23556
Max	0.0145	-0.0063	0.0186	0.0241	0.0189	0.0210	0.0164	0.0408	0.0172	0.0102	-0.0076	0.0168	0.0133	0.0089	0.0054	0.0441	496390
Min	-0.0368	-0.0392	-0.0329	0.0036	-0.0188	0.0033	0.0048	-0.0248	-0.0081	-0.0099	-0.0472	-0.0081	-0.0187	-0.0278	-0.0438	0.0090	3529
Signif. (1%)	30	35	17	25	16	26	11	36	13	9	36	11	11	12	35		
Positive	3	0	3	25	10	26	11	34	11	5	0	9	9	3	1		

Panel B: Bid Side of the Book

	SPR	ND^1	ND^{25}	NLO^1	NLO^{25}	$NDISP$	NV	NT	$NLOBQ$	NLO	$NLOIQ$	$NCBQ$	NC	CMR	QV	Psd. R2	Obs.
Median	-0.0105	0.0194	0.0132	-0.0118	-0.0083	-0.0121	-0.0078	-0.0214	-0.0088	-0.0053	0.0224	-0.0074	-0.0101	0.0095	-0.0139	0.0280	30676
Max	0.0075	0.0437	0.0355	0.0073	0.0235	-0.0027	-0.0033	0.0059	-0.0043	0.0089	0.0618	-0.0016	0.0039	0.0160	0.0054	0.0549	309672
Min	-0.0293	0.0057	-0.0094	-0.0291	-0.0413	-0.0208	-0.0126	-0.0460	-0.0232	-0.0088	0.0063	-0.0167	-0.0182	-0.0076	-0.0372	0.0044	5580
Signif. (1%)	23	36	19	32	18	30	13	34	15	8	36	13	12	17	29		
Positive	3	36	18	2	6	0	0	1	0	1	36	0	1	11	3		

TABLE III
Matched Samples

This table provides descriptive statistics on the matched samples of ordinary trades and trades revealing the presence of hidden volume resulting from the matching procedure described in footnote #12. The table reports the average trade size (standard deviation in parenthesis) the number of matched observations for the seven trade-size percentile categories S1-S7: less than 25% (S1), 25-50% (S2), 50-75% (S3), 75-90% (S4), 90-95% (S5), 95-99% (S6) and greater than 99% (S7). We also report the proportion all trades executed against hidden volume that passed the filters described in footnote #12 that become matched. The first panel provides the results when a 1-minute pre-trade interval is used to do the matching and the second panel provides the results when a five-minute pre-trade interval is used. BIT (SIT) stands for buyer (seller) initiated trades.

			Size category						
1 minute-matching			S1	S2	S3	S4	S5	S6	S7
Hidden	BIT	Avg. Size	96.74 (63.6)	325.13 (210.2)	770.78 (640.0)	1911.26 (1727.8)	3163.22 (2744.0)	6725.89 (6147.6)	33491.87 (49817.0)
No Hidden	BIT	Avg. Size	90.90 (63.2)	289.80 (200.2)	726.16 (629.4)	1881.28 (1800.2)	3246.10 (2872.5)	6906.07 (6296.4)	51333.46 (176309.7)
		Obs.	7122	11418	18066	12478	3302	2699	394
		% matched	(85.5)	(88.6)	(88.5)	(80.8)	(60.3)	(54.1)	(35.8)
Hidden	SIT	Avg. Size	86.59 (58.5)	241.67 (127.2)	559.91 (336.8)	1525.48 (1146.3)	3065.51 (2162.3)	6364.15 (4817.3)	35324.54 (69741.9)
No Hidden	SIT	Avg. Size	78.17 (56.4)	219.51 (128.8)	539.90 (339.9)	1502.63 (1088.8)	3163.88 (2171.6)	6401.08 (4633.6)	39591.32 (73194.3)
		Obs.	7836	10053	18153	14562	3902	3485	473
		% matched	(92.1)	(91.0)	(90.5)	(84.5)	(65.8)	(65.0)	(39.0)
5 minutes-matching									
Hidden	BIT	Avg. Size	96.88 (63.5)	323.98 (208.9)	763.40 (631.0)	1869.42 (1688.7)	3111.67 (2636.1)	6504.93 (5991.2)	33696.39 (64833.9)
No Hidden	BIT	Avg. Size	88.54 (64.0)	292.03 (202.0)	716.39 (617.2)	1846.42 (1776.5)	3233.14 (2834.4)	6590.25 (6061.8)	37714.01 (75450.0)
		Obs.	6803	10334	16486	10723	2418	1989	203
		% matched	(81.6)	(80.2)	(80.7)	(69.4)	(44.1)	(39.8)	(18.5)
Hidden	SIT	Avg. Size	86.25 (58.4)	243.03 (127.7)	558.74 (336.4)	1518.58 (1141.1)	3001.48 (2157.7)	6161.47 (4753.3)	29847.45 (39867.4)
No Hidden	SIT	Avg. Size	79.10 (58.1)	220.81 (129.1)	536.68 (340.1)	1497.48 (1080.5)	3143.37 (2168.2)	6159.03 (4584.9)	35542.13 (54563.7)
		Obs.	7181	9204	16661	12718	2902	2458	218
		% matched	(84.4)	(83.3)	(83.0)	(73.8)	(49.0)	(45.8)	(18.0)

TABLE IV
Liquidity and order flow

This table summarizes the impact of hidden volume on liquidity and the order flow when its presence is revealed to the market. Liquidity is measured by the bid-ask spread (*SPR*) and the net limit order book depth ($ND^l + ND^{25}$). The order flow is characterized by the net share volume (*NV*), the net number of trades (*NT*) and the net number of limit order submitted ($NLOIQ + NLOBQ + NLO$). For a detailed description of these variables see section 6. The table reports the results for the S4 trade size category, that is, trades with size between the 75% and 90% percentiles of the trade-size empirical distribution of the corresponding stock. We have constructed 1000 subsamples of matched trades randomly extracted from the samples described in Table III. This table reports the percentage of subsamples for which the median of the corresponding variable for the ordinary trades is found to be statistically different at the 1% level to the median of the matched trades revealing hidden volume. We also provide the percentage of subsamples for which the null is rejected against the alternative that the median of the variable for the trades revealing hidden volume is larger than the median for the matched ordinary trades. We report the results for the ten-minute interval centered on the execution of the trades. To test the equality of medians we use the non-parametric Wilcoxon test. Trades are matched using the one-minute pre-trade interval (for further details on the matching procedure, see footnote #12).

	Ask side of the book										Bid side of the book									
<u>Spread</u>	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5
Rejections at 1%	1.20	1.10	1.40	1.30	1.30	75.40	58.90	57.10	42.60	31.80	0.60	0.60	0.70	0.70	1.31	55.20	38.90	35.00	21.30	21.42
Hidden > No Hidden at 1%	0.50	0.60	0.40	0.90	0.90	0.00	0.00	0.00	0.00	0.00	0.10	0.10	0.00	0.20	1.20	0.00	0.00	0.00	0.00	0.00
<u>Net depth</u>	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5
Rejections at 1%	1.20	0.60	1.10	1.30	2.30	0.60	1.40	0.30	1.00	0.90	0.80	0.80	0.80	0.80	0.80	1.10	0.80	0.50	0.40	0.60
Hidden > No Hidden at 1%	0.90	0.30	0.00	0.30	0.00	0.50	1.30	0.10	0.30	0.10	0.00	0.20	0.50	0.50	0.80	1.10	0.70	0.30	0.20	0.40
<u>Net share volume</u>	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5
Rejections at 1%	1.30	1.50	3.81	1.10	3.70	79.20	5.12	14.50	6.91	8.01	0.70	1.50	1.30	1.40	6.41	53.20	17.92	8.70	4.80	3.41
Hidden > No Hidden at 1%	0.50	1.10	0.10	0.30	3.60	79.20	5.12	14.50	6.91	7.91	0.00	1.10	0.20	0.30	0.00	0.00	0.00	0.00	0.00	0.10
<u>Net number of trades</u>	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5
Rejections at 1%	1.80	0.60	1.61	0.90	10.80	83.50	16.53	21.66	8.41	39.90	1.41	2.61	4.11	1.51	3.02	65.50	4.31	1.10	0.90	1.00
Hidden > No Hidden at 1%	1.40	0.10	0.10	0.70	10.80	83.50	16.53	21.66	8.41	39.90	1.00	2.61	4.01	1.51	0.10	0.00	0.10	0.00	0.10	0.30
<u>Net number of limit orders</u>	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5	T-5	T-4	T-3	T-2	T-1	T+1	T+2	T+3	T+4	T+5
Rejections at 1%	4.71	3.81	1.50	1.70	6.31	0.80	1.11	0.90	1.71	0.81	2.61	7.12	2.21	2.01	21.00	0.70	1.21	0.91	2.20	1.81
Hidden > No Hidden at 1%	4.71	3.81	1.30	1.60	0.00	0.60	0.50	0.50	0.91	0.20	0.00	0.00	0.00	1.91	21.00	0.70	0.50	0.60	0.40	0.20

TABLE V
Returns and volatility

This table summarizes the impact of hidden volume on returns and volatility when its presence is revealed to the market. We compute the cumulative mid-quote returns (*CMR*) and the mid-quote volatility (*QV*) 5, 10 and 30 minutes after the trade and until the closing of the trading session. For a detailed description of these variables see section 6. We have constructed 1000 subsamples of matched trades randomly extracted from the samples described in Table III. This table reports the percentage of subsamples for which the median of the corresponding variable for the ordinary trades is found to be statistically different at the 1% level to the median of the matched trades revealing hidden volume. We also provide the percentage of subsamples for which the null is rejected against the alternative that the median of the variable for the trades revealing hidden volume is larger than the median for the matched ordinary trades. To test the equality of medians we use the non-parametric Wilcoxon test. Trades are matched using the one-minute pre-trade interval (for further details on the matching procedure, see footnote #12).

Panel A: Returns								
Size category	Ask side of the book				Bid side of the book			
	t+5	t+10	t+30	Closing	t+5	t+10	t+30	Closing
S1								
Rejections (1% lev.)	1.10	1.40	1.10	3.10	3.21	2.20	0.90	2.60
Hidden > No Hidden	0.00	0.00	0.30	3.00	0.00	0.40	0.50	2.60
S2								
Rejections (1% lev.)	0.70	1.40	1.70	2.60	1.10	1.20	1.60	2.30
Hidden > No Hidden	0.60	1.30	1.50	2.50	0.80	1.00	1.50	2.30
S3								
Rejections (1% lev.)	2.60	3.61	5.90	14.90	1.50	3.40	5.40	13.90
Hidden > No Hidden	2.60	3.61	5.80	14.90	1.20	3.40	5.40	13.90
S4								
Rejections (1% lev.)	2.21	1.30	1.00	1.00	2.71	2.20	0.50	1.00
Hidden > No Hidden	0.10	0.20	0.20	0.70	0.10	0.00	0.10	0.40
S5								
Rejections (1% lev.)	1.90	1.80	1.30	9.70	1.30	1.40	1.20	9.90
Hidden > No Hidden	0.00	0.10	1.10	9.70	0.00	0.10	0.90	9.90
S6								
Rejections (1% lev.)	1.70	2.11	2.40	6.70	3.21	1.50	1.31	3.90
Hidden > No Hidden	1.70	1.31	2.10	6.70	2.81	0.90	0.90	3.90

Panel B: Volatility								
Size category	Ask side of the book				Bid side of the book			
	t+5	t+10	t+30	Closing	t+5	t+10	t+30	Closing
S1								
Rejections (1% lev.)	1.00	1.30	1.10	15.30	1.70	1.10	1.10	13.70
Hidden > No Hidden	0.00	0.10	0.40	15.30	0.20	0.40	0.50	13.70
S2								
Rejections (1% lev.)	5.10	1.50	0.50	19.60	4.60	1.10	1.00	19.80
Hidden > No Hidden	0.00	0.00	0.20	19.60	0.00	0.00	0.20	19.80
S3								
Rejections (1% lev.)	21.80	18.30	6.90	2.00	22.90	17.90	6.50	2.90
Hidden > No Hidden	0.00	0.00	0.00	1.60	0.00	0.00	0.00	2.90
S4								
Rejections (1% lev.)	35.30	15.80	7.70	2.70	31.30	13.60	6.70	3.20
Hidden > No Hidden	0.00	0.00	0.00	2.60	0.00	0.00	0.00	3.20
S5								
Rejections (1% lev.)	18.50	5.50	4.30	2.00	17.20	3.60	2.80	2.00
Hidden > No Hidden	0.00	0.00	0.10	1.90	0.00	0.00	0.10	1.90
S6								
Rejections (1% lev.)	25.80	14.00	6.00	5.90	22.50	12.90	5.10	5.00
Hidden > No Hidden	0.00	0.00	0.00	5.90	0.00	0.00	0.00	5.00

FIGURE 1
Intra-daily Patterns

This figure provides the average intra-daily distribution of hidden order detection compared with liquidity, trading activity and volatility. The continuous trading session of the SSE (from 9:00 a.m. to 17:30 p.m.) is divided into seventeen 30-minutes intervals. For each interval we compute the number of trades comprising hidden orders, the time-weighted average bid-ask spread, the time-weighted average cumulative depth at the five best ask and bid quotes, the number of shares executed (volume) and the number of trades completed. Volatility is the ratio between the highest and the lowest quote midpoint in each interval minus one. We represent the equally weighted average of these measures normalized and standardized for the 36 stocks in the sample. For normalization and standardization details see section 5.

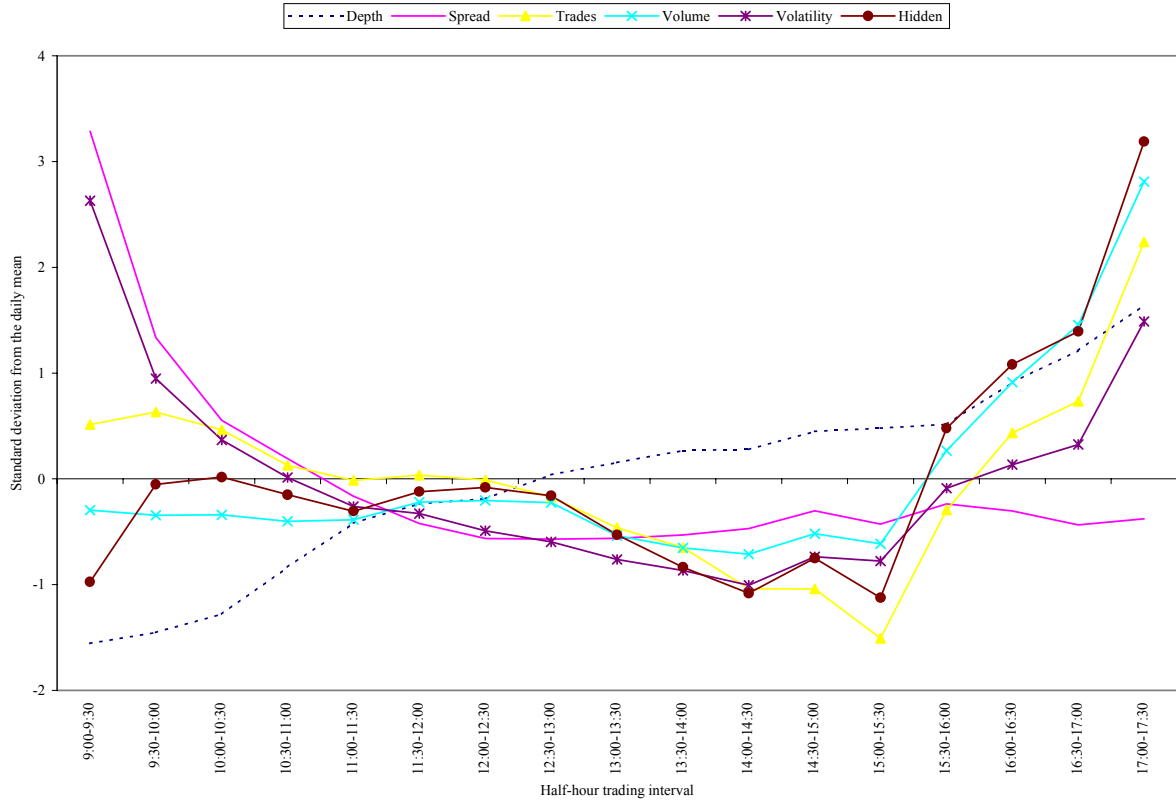


FIGURE 2
Proportion of trades involving hidden volume

This figure provides median for the 36 stocks in the sample of the percentage of trades completed against the undisclosed side of the quoted depth. The continuous trading session of the SSE (from 9:00 a.m. to 17:30 p.m.) is divided into seventeen 30-minutes intervals.

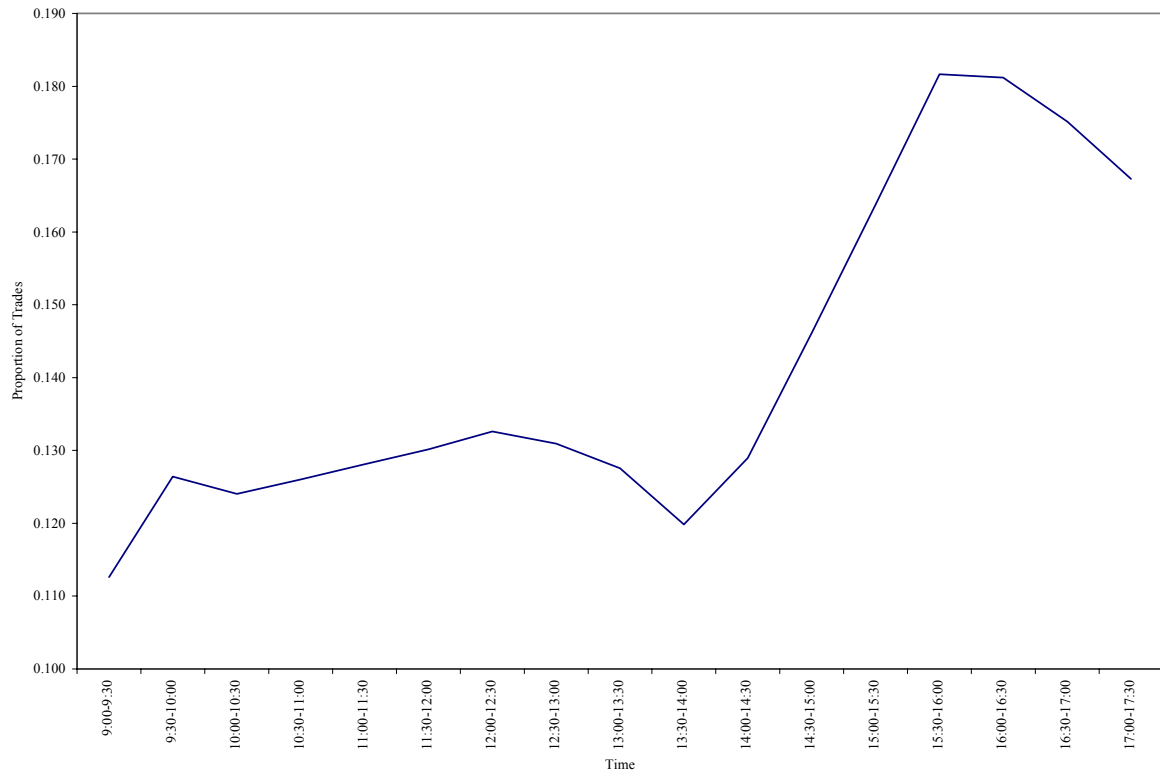


FIGURE 3
The intra-daily contribution of hidden orders conditional on trade size

This figure provides the average intra-daily distribution of hidden order detection conditional on each firm's trade size distribution. The continuous trading session of the SSE (from 9:00 a.m. to 17:30 p.m.) is divided into seventeen 30-minutes intervals. For each firm, trades are partitioned into seven percentile categories: less than 25% (S1), 25-50% (S2), 50-75% (S3), 75-90% (S4), 90-95% (S5), 95-99% (S6) and greater than 99% (S7). The cut-offs are also different for buyer and seller initiated trades. For each interval we compute the number of trades in each trade size category comprising hidden orders. We represent the equally weighted average of these measures normalized and standardized for the 36 stocks in the sample. For normalization and standardization details see section 5.

