

The Impact of Hidden Liquidity in Limit Order Books *

Stefan Frey[†]

Patrik Sandås[‡]

Current Draft: September 7, 2008

First Draft: September 2007

Comments Welcome

*We thank Uday Rajan, Mark Van Achter, Burton Hollifield, Gideon Saar, Rick Harris, Mike Pagano, and seminar participants at the NBER Market Microstructure Meetings, the Center for Financial Studies' Conference on the Industrial Organization of Securities Markets, and the Society for Financial Econometrics's Inaugural Conference for useful suggestions, the German Stock Exchange for providing access to the Xetra order book data and Uwe Schweickert for his help with the order book reconstruction. Frey gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft, and Sandås gratefully acknowledges financial support from the McIntire School of Commerce.

[†]University of Tübingen and Center for Financial Research, E-mail: stefan.frey@uni-tuebingen.de

[‡]University of Virginia and Center for Economic Policy Research, E-mail: patriks@virginia.edu

Abstract

We report evidence that the presence of hidden liquidity is associated with greater visible liquidity in the order books, greater trading volume, and smaller price impact. We construct an algorithm that extracts information about hidden depth from publicly available data. We show that the predicted presence of iceberg orders is associated with larger market orders and market orders that are skewed towards the side of the order book with the iceberg order. We estimate a state-dependent price impact function and the moments of the order flow distributions and use them to calculate the expected surplus to limit and iceberg orders. We find that an iceberg order earns its highest expected surplus when it remains undetected consistent with the ability to hide giving large trader a comparative, but short-lived, advantage. The positive net surplus and the greater trading activity when iceberg orders are present are consistent with uninformed iceberg orders. The expected mid-quote changes conditional on type I and II errors generated by the algorithm provides a measure of the information conveyed by iceberg orders. A ‘surprise’ iceberg order is associated with a mid-quote change over the next 30 traders of the same order of magnitude as the quoted bid-ask spread consistent with informed iceberg orders.

Keywords: Hidden Liquidity; Iceberg Orders; Hidden Orders; Reserve Orders; Limit Order Markets; Limit Order Books; Transparency;

JEL Codes: G10, G14

1 Introduction

Many limit order markets use a market design that allows traders to submit hidden liquidity. The option to submit hidden liquidity alongside the visible liquidity makes the strategic interaction between different market participants more complicated and raises a number of questions. To what extent can market participants detect iceberg orders? If detection is possible, wherein lays the strategic advantage of the hidden orders? Do market participants perceive iceberg orders to be submitted by informed or uninformed traders? How does the presence of hidden liquidity affect the strategies of liquidity suppliers and liquidity demanders? We address these questions using a sample from the German Stock Exchange's Xetra platform that includes iceberg and limit orders.

Exchanges often add the option to submit hidden liquidity by creating a different type of limit order that is known as an *iceberg order*.¹ An iceberg order is a limit order that specifies a price, a total order size, and a visible peak size. The peak size is the maximum number of shares that is displayed to the market at any time. The remainder of the iceberg order is not displayed in the order book. When the first peak size has been fully executed, the visible part is immediately replenished by a size equal to the peak size. At a given price level in the order book all displayed order depth has time priority relative to any hidden depth, irrespective of the order entry times. Because of the replenishment rule, which adds a new peak size immediately after the current visible peak size is executed, an iceberg order is likely to be detected, after its first peak executes, by acute observers of the order book. A sequence of events that includes a trade followed by a new order at the same price with a minimal delay is a signal of an iceberg order.

The above arguments and conventional wisdom suggest that iceberg orders can be and are eventually detected. While the footprint left by a single replenishment is fleeting, repeated patterns of trading volume exceeding the visible volume at a price level provides strong signal of an iceberg order. We develop an algorithm for detecting iceberg orders that uses publicly available information on the order book and price dynamics. We use the predictions generated by the algorithm to

¹This type of order is known as a *reserve order* in some markets. Completely hidden orders can be submitted in some markets and some markets (e.g., BATS Trading) allow both reserve and hidden orders. In fixed-income markets there is a type of reserve order known as a *expandable limit order* that gives the submitter the option but not the obligation to trade more when the initial size has been executed (see Boni and Leach (2004)) but iceberg orders are also used (see Fleming and Mizrach (2008)).

approximate the information set available to liquidity demanders and suppliers. By combining the information generated by the algorithm with the true iceberg state, which is observable to the econometrician, we can approximate the information set available to a large trader with an iceberg order in the order book. We use these two information sets to extract information about how the surplus from providing liquidity is split between limit and iceberg orders and to determine whether iceberg orders convey any information to the market.

Using the signals generated by the algorithm we show that when iceberg orders are predicted order books have greater *visible* depth and narrower inside spreads. The greater liquidity reflects both the choice of the large traders submitting the iceberg order and the reaction of the liquidity suppliers to the predicted presence of an iceberg order.

We estimate a state-dependent affine price impact function which permits the price impact to shift with the iceberg state. The price impact changes significantly for order books with predicted and actual iceberg orders with the salient effect being that the market orders hitting the side with the iceberg order have a significantly smaller fixed and variable price impact over a 30 trader horizon. The distribution of the market order size and the probability of a buy versus a sell market order are also strongly state dependent with larger market orders when iceberg orders are predicted and the market order flow skewed towards the order book side with the iceberg order.

We use the estimates of the price impact function and the moments of the order flow to calculate the expected surplus to limit and iceberg orders in the order book. We show that states in which the iceberg order remains undetected according to the algorithm are associated with the largest expected surplus for both limit and iceberg orders. This evidence suggests that on the one hand iceberg orders work as intended and on the other hand that the efforts exerted by market participants are worthwhile. Liquidity suppliers who submit limit orders alongside the iceberg orders earn a greater expected surplus than limit orders in other states. But even when iceberg orders are detected the expected surplus to liquidity suppliers is positive. The positive surplus combined with the greater liquidity demand suggests that iceberg orders are associated with favorable trading outcome for both liquidity suppliers and demanders.

We quantify the potential information content of iceberg orders by measuring the expected mid-

quote changes conditional on whether or not the iceberg detection algorithm predicts an iceberg and whether or not there actually is an iceberg order. By construction the market will almost certainly learn the true state of the order book after a few rounds of trade. We find that when the order book contains an undetected iceberg order we expect the mid-quote to move in a direction favorable to the iceberg order over the next 30 trades. The expected mid-quote change is of the same order of magnitude as the quoted bid-ask spread and thus appears to be economically significant. On the other hand, when the algorithm generates a false signal of an iceberg order—for example, the order book dynamics suggest an iceberg order is there but it has in fact been canceled—then a reversal is expected over the next 30 trades. Both findings suggest that iceberg orders are perceived to contain some information.

Overall our findings provide some support for two views that have been put forward in the literature. The evidence on mid-quote changes in response to ‘surprise’ and ‘false-call’ iceberg orders suggest that the market perceives the iceberg orders to contain some information. These effects are absent when the algorithm detects an iceberg order or no iceberg order. On the other hand, the greater trading activity and the iceberg order attracting market order flow are both consistent with iceberg orders being perceived as uninformed orders.

There is a large literature on market transparency which is often classified along the pre-trade versus post-trade dimensions.² Hidden liquidity is an example of a pre-trade transparency issue but there are several others. The complexity of the issue of pre-trade transparency arises because the nature of the trade-offs involved change with the trading mechanism, the type of information that is disclosed, and the participants to whom information is disclosed.³ A number of studies focus specifically on hidden orders or iceberg orders (Pardo and Pascual (2007), Tuttle (2006), and Labys (2001)). A number of studies have focused on different aspects of the decision problem faced by a submitter of an iceberg order. Esser and Mönch (2007) focus on the trade-off between

²See the sections on market transparency in O’Hara (1995), Madhavan (2000), and Biais, Glosten, and Spatt (2005) for in-depth discussions.

³Biais (1993), Madhavan (1995), Madhavan (1996), Pagano and Röell (1996), Bloomfield and O’Hara (2000), Baruch (2005), Moinas (2007), Foucault, Moinas, and Theissen (2007) among others develop theoretical models of transparency and Flood, Huisman, Koedijk, and Mahieu (1999) and Bloomfield and O’Hara (1999) carry out experimental studies of trading in different transparency regimes. A number of empirical studies including Anand and Weaver (2004), Boehmer, Saar, and Yu (2005), Foucault, Moinas, and Theissen (2007), Hendershott and Jones (2005), Madhavan, Porter, and Weaver (2005) focus on the impact of changes in market transparency.

price and peak size of the iceberg order. Bessembinder, Panayides, and Venkataraman (2008) and De Winne and D'Hondt (2007) study how the decision to not display the full order interacts with other dimensions of the trader's order choice problem. Harris (1996) studies how variations in the minimum tick size influence the willingness to display larger order quantities. Aitken, Berkman, and Mak (2001) examine variation in the use of hidden orders around a change in the threshold size for such orders. Moinas (2007) and Buti and Rindi (2008) develop theoretical models of hidden liquidity. We add to this literature by quantifying to what extent market participants may detect iceberg orders and how their knowledge of possible iceberg orders change their order submission strategies.

2 Our Sample

Our sample is from the Frankfurt Stock Exchange's electronic trading platform Xetra. On Xetra, traders can, in addition to market and limit orders, submit iceberg orders. An iceberg order specifies a price, a total size, and a peak size. The peak size is the maximum visible volume of the order. When a trader submits an iceberg order, the first peak size is visible in the order book. At that time, the hidden volume of the order is equal to the order's total size minus its peak size. When the first peak size has been fully executed, the visible part is automatically replenished by a number of shares equal to the peak size, and the hidden part is reduced by the corresponding number of shares. The replenishment of the visible part continues automatically until the hidden volume is depleted or the trader cancels the iceberg order.

In the order book, an order is given priority according to price, display condition, and time. A sell order at a lower price has priority relative to any sell orders at higher prices, irrespective of the order's time of submission or display condition. At the same price level, a displayed order has priority relative to any hidden orders regardless of the order's time of submission. Among displayed orders, an order submitted earlier has priority relative to any orders submitted later. When an iceberg order's visible part is replenished and the next peak size converts from hidden to displayed status the newly visible peak size also receives a new time stamp which determines its time priority.

We reconstruct the sequence of order books from the event histories in the sample. The order

records include a flag for an iceberg order which we use to construct complete histories for all limit and iceberg orders. From these histories we reconstruct snapshots of the visible and hidden order books 1/100th of a second before each transaction. We restrict our sample to orders submitted during the continuous trading period. Continuous trading on Xetra starts after an opening auction, ends with a closing auction, and stops for a few minutes, in the middle of the day, for an auction. The reconstruction takes into account the effects that any order submissions, transactions, or cancellations in the auctions have on the state of the order book during continuous trading.

Our sample includes all order entries, trades, and cancellations in the thirty stocks in the DAX-30 German blue chip index for the period January 2nd to March 31st, 2004. Table 1 reports the market capitalization, trading volume, number of trades, and the average mid-quote, trade size, and bid-ask spread for all thirty sample stocks. We sort the stocks into three categories—large, medium, and small—based on trading activity. The large and medium categories have eight and seven stocks, respectively, with the remaining fifteen stocks being assigned to the small category. This sorting scheme ensures that stocks within each category have comparable trading volume and number of trades.

We are interested in how the hidden liquidity of iceberg orders interact with the limit order book, the order flow, and price dynamics. The potential impact that iceberg orders have depends on several factors including the frequency of iceberg orders, their size, and their duration. Table 2 reports some basic statistics for iceberg orders. Panel A of the table shows that the iceberg orders' share of all non-marketable orders submitted ranges from 7 to 11% with an overall average of 9%.⁴ Iceberg orders represent a greater fraction of shares executed with the corresponding percentages ranging from 15 to 20% with an overall average of 16%.

Given the rationale for iceberg orders we expect them to be larger than regular limit orders. Panel B of Table 2 reports the average of the iceberg order's peak size, the ratio of total order size to peak size, and the ratio of executed shares to peak size for all iceberg order whose first peak size was executed. The last row of the panel reports the average size of limit orders. On average, the limit order size is 1,000 shares whereas the average peak size of an iceberg order is 2,600 shares

⁴There are a few instances of marketable iceberg orders in our sample and we exclude them as well as all market orders or marketable limit orders when computing these figures.

implying that the visible part of an iceberg order is between two and three times the size of a typical limit order. The total size of an iceberg order is, on average, between seven and eight peak sizes which partly reflects the fact that there is clustering at even multiples such as five or ten times the peak size. The ratio of executed shares to peak size is between four and five, which implies that conditional on an iceberg order's first peak being executed its visible volume will be replenished almost four times and approximately 80% of the executed shares originate from the order's initially hidden volume.

The last row of the table reports the average ratio of the iceberg orders' duration to the limit order's duration and shows that iceberg orders stay in the order book for approximately seven times as long as limit orders. The longer duration of iceberg order's provides at least a partial explanation for the iceberg order's greater relative share of share executed and it also implies that the probability that a randomly selected order book has an iceberg order will exceed their share of orders submitted.

3 A Model Framework

Our model framework allows us to quantify the extent to which market participants can detect iceberg orders using publicly available information and the impact that such knowledge has on their order submission strategies. We use our framework to assess how the predicted and actual presence of iceberg orders change the surplus generated from trading and the division of the surplus between different market participants.

3.1 Market Participants

There are three types of market participants. At time t a liquidity demander, who may be informed, arrives and submits a market order that is matched with orders in the limit order book producing a trade. The limit order book contains limit orders submitted by liquidity suppliers and it may contain iceberg orders submitted by large traders. All market participants agree on a fundamental value of X for the stock at time t ; X may be interpreted as the liquidity providers' time t expectation of the liquidation value of the stock. All market participants observe the history of trades and order

book updates.

A liquidity demander, who may be informed about the future value of the stock, submits a market order at time t . Denote the size of the time t market order by m and the direction of the market order by d ; $d = +1$ for buy and $d = -1$ for sell market orders.

Liquidity suppliers submit limit orders to the order book and earn a surplus that depends on the value of the stock conditional on their limit orders executing. We assume that liquidity suppliers have an interest in detecting iceberg orders because the presence of iceberg orders may have both a direct and indirect effect on the surplus earned by liquidity suppliers.

Before time t , a large trader may arrive and submit a new or cancel an outstanding iceberg order. An iceberg order consists of a visible peak size and multiple hidden peaks. Large traders, like the liquidity suppliers, care about the surplus they earn from trading which depends on the stock value conditional on execution and the probability of execution.

3.2 The Order Book

The ask sides of the limit order book at time t are characterized by a series of quotes, $p_1^{ask}, p_2^{ask}, \dots, p_K^{ask}$, with the index starting from the best ask quote. The total visible volume offered at the k th best ask quote is denoted by q_k^{ask} . The cumulative visible volume offered at all quotes with equal or higher priority to the k th best quote, p_k^{ask} , is denoted by Q_k^{ask} , and is determined as, $Q_k^{ask} = \sum_{i \leq k} q_i^{ask}$. Accordingly the hidden volumes are denoted by \hat{q}_k^{ask} and the cumulative hidden volumes by \hat{Q}_k^{ask} . The indicator, I^{ask} , takes on a value of one, if there is hidden depth at the best ask quotes, and zero otherwise. The bid side has analogous notation with a different superscript.

3.3 Information

All market participants observe the visible limit order book. Liquidity demanders and suppliers unlike the large traders do not observe directly any hidden depth, \hat{q}^{ask} or \hat{q}^{bid} nor do they observe indicators for non-zero hidden depth, I^{ask} and I^{bid} . Instead they use recent price and order book dynamics to infer the existence of any iceberg orders. Below we describe the algorithm that the liquidity suppliers and demanders use to infer the existence of iceberg orders in our framework.

3.4 Detecting Iceberg Orders

We construct an iceberg detection algorithm that uses the order book dynamics to make predictions of whether or not the bid or the ask side has an iceberg order. When a trade occurs that involves the execution of volume in excess of the visible volume at given price level, it strongly suggests the presence of an iceberg order at that price level. The ratio between total size to peak size in Table 2 implies that after the replenishing of one peak, additional hidden depth is typically to be expected. This implies that one can detect iceberg orders by comparing the recent history of transaction prices and volume with the transition of the visible order book.

Our algorithm sets an iceberg-indicator, \hat{I} , to one every time new depth is added at a price consistent with the replenishment of an iceberg order. The algorithm resets the indicator for that price to zero only when an event occurs that could not have occurred had the iceberg order remained at that price. The algorithm stores a specific indicator and the expected volume until the next replenishment for each price level. The indicator remains unchanged unless a predicted replenishment fails to occur. Section A2 in the appendix provides a detailed example of how the algorithm works. The detection algorithm makes both type I errors setting $\hat{I} = 1$ when $I = 0$ or type II errors leaving $\hat{I} = 0$ when $I = 1$. Table 3 summarizes the possible combinations of predictions made by the algorithm, the actual iceberg states, and the terminology that we use.

Below we focus on the predictions that the algorithm makes for the best bid and ask quotes for a given order book. We also make a symmetry assumption for the bid and ask sides. We use the indicators \hat{I}^{own} and \hat{I}^{opp} to differentiate the cases in which an iceberg order is predicted at the ‘own’ side and ‘opposite’ side viewed from the perspective of a given limit order in the order book with I^{own} and I^{opp} denoting the true state.

3.5 Price Impact and Order Flow

We define a price impact function that builds on the The price impact of a market order is determined by its size and direction and it depends on the state of the order book. The change in the

fundamental value between t and $t + \tau$ is determined by:

$$X_{+\tau} = X + \mu + (\alpha + \beta m)d + (\alpha^{own} + \beta^{own} m)\hat{I}^{own}d + (\alpha^{opp} + \beta^{opp} m)\hat{I}^{opp}d + \epsilon_{+\tau}, \quad (1)$$

in which, α and β are the parameters for the baseline price impact, α^{own} and β^{own} determine the differential price impact when an iceberg order is predicted at the side of the book hit by the market order, and α^{opp} and β^{opp} determine the corresponding differential price impact when an iceberg is predicted at the side of the order book not hit by the market order.

We assume the market order size is exponentially distributed with a mean of λ in the baseline case, and $\lambda + \lambda^{own}$ when $\hat{I}^{own} = 1$, and $\lambda + \lambda^{opp}$ when $\hat{I}^{opp} = 1$. The expected hidden depth is denoted by η , $\eta = E[\hat{q}|\hat{I} = 1]$. The probability of a buy market order conditional on a market order is denoted by ϕ and it also depends on the iceberg state with $\phi + \phi^{own}$ when $\hat{I}^{own} = 1$ and $\phi + \phi^{opp}$ when $\hat{I}^{opp} = 1$.

3.6 Liquidity Supplier Surplus

We are interested in computing the expected surplus accruing to liquidity suppliers conditional on the public information available to them at time t . In order to do this we must take into account the probability of buy versus sell market orders conditional on a market order arriving, the distribution of market orders, and the expected price impact of the market order. Let h be a state vector that summarizes the current iceberg status conditional on the public information with $h = [1 \ \hat{I}^{own} \ \hat{I}^{opp}]$. Using this notation the price impact function in Equation 1 can be rewritten as:

$$X_{+\tau} = X + \mu + (\bar{\alpha}h)d + (\bar{\beta}h)(d \cdot m) + \epsilon_{+\tau},$$

with $\bar{\alpha} = [\alpha \ \alpha^{own} \ \alpha^{opp}]$ and $\bar{\beta} = [\beta \ \beta^{own} \ \beta^{opp}]$. The expected revision in the fundamental value is determined by the upper tail expectations at ask quote k :

$$X_{+\tau}(Q^k, d; \theta) = X + \mu + \left(\alpha h + (\beta h) \left(E[m|m \geq Q^k, h] + \hat{I}^{own} \eta \right) \right) d, \quad (2)$$

with θ denoting the vector of parameters. The expected surplus to all units take into account the probability of execution and the corresponding tail expectations. Let $\pi_k(q)$ denote the expected surplus for the q th unit at quote k . Using Equation (??) we can write the surplus $\pi_k(q)$ for the ask side as

$$\pi_k(q) = \int_q^\infty [p_k - x_{+\tau}(m, d = 1; \theta)] Pr(d = 1; h) f(m; h, \lambda) dm \quad (3)$$

where $Pr(d = 1; h)$ denotes the state dependent probability of a buy order and $f(m; h, \lambda)$ the probability distribution of the size of market orders with parameter vector λ . An equivalent formula is used for the bid side.

Denote the aggregate expected surplus for the visible volume at quote k by Π_k . We obtain it by integrating $\pi_k(q)$ for the visible volume at quote k

$$\Pi_k = \int_{Q_{k-1} + \hat{Q}_{k-1}}^{Q_k + \hat{Q}_{k-1}} \pi_k(q) dq \quad (4)$$

where we use the definition of $Q_0 = 0$ and $\hat{Q}_0 = 0$.

The aggregate expected surplus for the hidden volume at quote k is $\hat{\Pi}_k$ and calculated by

$$\hat{\Pi}^k = \int_{Q_k + \hat{Q}_{k-1}}^{Q_k + \hat{Q}_k} \pi_k(q) dq \quad (5)$$

The total liquidity provider surplus is obtained by adding across quotes of both sides of the order book. By conditioning on the state or a subset of order book levels we can calculate various components of the total surplus.

4 Empirical Results

We estimate the model parameters using GMM in two stages. In the second-stage estimation we use a Newey-West type weighting matrix with a Bartlett kernel with 10 lags. Table 4 summarizes the parameters and the moment conditions of our model. We estimate two specifications of our

model. The first specification, which we refer to as the public information specification, conditions only on the signals generated by the iceberg detection algorithm. The second specification, which we refer to as the full information specification, conditions on the true iceberg state in addition to the signals generated by the iceberg detection algorithm. One interpretation of the specifications is that the former specification conditions on the information of the liquidity suppliers and demanders whereas the latter conditions on the large traders' information. We start by presenting statistics on the performance of the iceberg detection algorithm.

4.1 Performance of the Iceberg Detection Algorithm

The iceberg detection algorithm generates predictions of whether or not a given order book contains iceberg orders based on publicly available data. Table 5 reports the average frequencies of the iceberg signals generated by the algorithm and the true iceberg states.⁵ The frequencies reported in the table are the cross-sectional averages with cross-sectional standard errors reported in parentheses. The diagonal entries are the percentages of correct predictions; for 85 to 90% of the observations the algorithm correctly detects no iceberg order, and for 6 to 9% of the observations the algorithm correctly detects an iceberg. The top right entry in each panel is the percentage of observations for which the algorithm falsely detects an iceberg. These instances are relatively rare with the average frequency being around 1%. The bottom left entry in each panel is the percentage of observations for which the algorithm fails to detect an iceberg order. Between 3 and 5% of the observations fall in this category. The algorithm's relative weakness is the instances of iceberg orders that it fails to detect. These percentages suggest that iceberg orders may succeed in remaining hidden, albeit perhaps only for a limited amount of time. Across the three categories the ratio of 'not detected' to the percentage of iceberg states is approximately one-third implying that one might expect an iceberg to be detected in three trades.

The frequencies reported in Table 5 are for the ask side only. The results for the bid side are very similar. Furthermore, the correlation between icebergs on the bid and the ask side is low so the frequency of an iceberg signal for at least one side of the book is approximately twice the

⁵Our algorithm for detecting iceberg orders and our results are comparable to those of De Winne and D'Hondt (2007).

percentages in the bottom right of each panel.

We cannot determine to what extent the predictions from our algorithm closely approximates the predictions of the market participants. Conversations with market participants suggest that it is reasonable to assume that active participants are able to collect this type of information. It may also be the case that at least some market participants apply algorithms that generate even more accurate predictions. For example, we have not used information in order sizes or any other characteristics that tends to be different for iceberg orders.⁶

4.2 Model Parameter Estimates

Table 8 reports the model parameter estimates for the public information case that uses the signals generated by the iceberg detection algorithm. The parameter estimates for the price impact function are reported in Panel A. The baseline price impact is determined by the α and β parameters. The estimates of the fixed impact of a market order range from 1.8 to 3.2 basis points. The estimates of the variable impact ranges from 0.4 to 1.3 basis points. The alpha and beta parameters with an *own* superscript capture the change in the price impact function when an iceberg is detected on the same side of the order book. The estimates for both α and β are negative and imply that the presence of an iceberg order on the same side of the order book tends to dampen the price impact. The fixed component is reduced by 0.6 to 1.2 basis points and the variable component is reduced by 0.2 to 0.4 basis points. Conversely, when the iceberg order is on the side opposite to the one hit by the market order, the *opp* case, the fixed component is 0.3 to 1.8 basis points higher but the variable component, β , is negative as in the *own* case albeit the magnitude of the reduction is smaller ranging from zero to 0.2 basis points.

Panel B of Table 8 reports the estimates for the market order size. The estimates for the baseline case, the λ s, are slightly below one suggesting that, on average, market orders are slightly below average size when no iceberg orders are predicted. The λ^{own} and λ^{opp} are both positive with the λ^{own} estimates around 0.3 and the λ^{opp} estimates around 0.1, which translates into an increase in the expected market order size by 10 to 30% when an iceberg order is predicted. One interpretation of

⁶We have also explored whether there is any evidence that immediate-or-cancel orders may be used to detect iceberg orders but while such orders are used quite frequently the connection to iceberg states is rather weak.

this finding is that large traders are successful in timing their iceberg order submissions to coincide with larger expected market orders.

Panel C of Table 8 reports the estimates for the expected hidden depth, η , and panel D reports the estimates of the probability of a buy market order, ϕ . The η estimates are between 7 and 8 implying that conditional on the market participants predicting an iceberg order, they expect the hidden depth to be equal to seven or eight average sized market orders. Using the figures from Table 2 for an average peak size of two to three units and a typical number of total peaks of seven or eight the above figures imply that the market participants rationally expect approximately one-half or somewhat less to be left of an iceberg order conditional on detection. The baseline estimates of ϕ are one-half as one would expect. The ϕ^{own} estimates are positive implying that when an iceberg order is predicted the side with the iceberg order is more likely to be hit by a market order; iceberg orders attract market orders. Given the ϕ is the probability of a buy market order conditional on a market order, the estimates of approximately 0.2 for ϕ^{own} imply that when an iceberg order is predicted on, say, the bid side the probability of a sell market order jumps to 0.7 and the probability of a buy market order falls to 0.3. The skewness in the market order flow combined with the asymmetric effect on α suggest that there may be a predictable drift in the mid-quote when icebergs are predicted. We will explore that issue below.

4.2.1 Full Information Model

Table 9 reports the model parameter estimates for the full information case that uses both the signals generated by the iceberg detection algorithm and the true iceberg states. The large information set requires a larger set of parameters. The baseline case corresponds to observations for which the algorithm correctly predicted no iceberg. The two iceberg scenarios of the public information case split into three sub-cases depending on whether the algorithm made a correct prediction ($\hat{I}^{own} \times I^{own}$ and $\hat{I}^{opp} \times I^{opp}$), failed to detect the iceberg ($(1 - \hat{I}^{own}) \times I^{own}$ and $(1 - \hat{I}^{opp}) \times I^{opp}$), or generated a false detection signal ($\hat{I}^{own} \times (1 - I^{own})$ and $\hat{I}^{opp} \times (1 - I^{opp})$). We use a subscript *det* for the correctly detected, a *fls* subscript for the falsely detected, and a subscript *nod* for the not detected iceberg orders.

Panel A reports the parameter estimates for the price impact function. The estimates for the baseline case, α and β , are very close to the estimates reported in Table 8. The α estimates for the false detection case have the opposite signs of the corresponding α estimates for the detected and not detected case consistent with a false detection producing a reversal. The estimates for $\alpha_n^{own\ od}$ and α_{nod}^{opp} are greater in magnitude than the corresponding α_{det}^{own} and α_{nod}^{opp} estimates consistent with a delayed reaction in cases with not detected iceberg orders.

The estimates of β for the own side case are all close to each other lending further support to the interpretation that periods in which iceberg are submitted or expected are associated with smaller price impact consistent with less asymmetric information.

Panel B reports the estimates for the market order size, λ . The estimates for the baseline case are somewhat smaller than in the public information case. The λ estimates for the false detection case do not differ from the baseline whereas the estimates for the not detected case are positive and economically large implying an increase in the average market order size by 40% or more. One interpretation of the increase for both the *det* and *nod* case and the lack of a change for the *fls* case is that the large traders submit iceberg orders when larger sized market orders are anticipated.

The estimates of η , the expected hidden depth, are larger for the *nod* case ranging from 11 to 12 units compared to 8 to 10 for the *det* case consistent with the not detected iceberg orders are orders that were submitted more recently.

The estimates of ϕ show that liquidity demanders react to the detected iceberg orders. The estimates of ϕ_{nod} are positive but less than 0.05 suggesting a substantially smaller skew. The estimates of both ϕ_{det} and ϕ_{fls} , however, are both positive and economically large with all point estimates between 0.15 and 0.2. The positive estimates for ϕ_{fls} show that conditional on the algorithm detecting an iceberg order the expected market order flow shifts from the fifty-fifty baseline case to approximately 0.65 vs. 0.35 favoring the side of the perceived iceberg order. This type of shift in the market order flow is evidence that the liquidity demanders are reacting to a perceived iceberg order.

4.3 Liquidity Provider Surplus

We compute different surplus estimates for liquidity suppliers based on the results in section 3.6. Additional details on the surplus calculations are provided in Appendix A1. Table 10 reports the average estimates of the expected liquidity provider surplus. Panel A reports the average expected surplus for all states conditional on the algorithm predicting (i) no iceberg (baseline) and (ii) at least one iceberg order at one of the best quotes. Standard errors are reported in parentheses and account for sampling error in the second-stage and the first-stage estimation error for the model parameters. The standard errors are computed under the assumption that the first- and second-stage errors are independent. The estimates of the overall expected surplus to liquidity suppliers, reported on the last row of the panel, are between 0.16 and 1.1 basis points and correspond to approximately 25 to 30% of the quoted half-spread in Table 1. The two rows above display the expected surplus in the baseline case and in the case of iceberg orders. The surplus estimates for the iceberg case are negative albeit statistically significant only for the small category and in general relatively close to zero.

Panel B of Table 10 reports the expected surplus figures based on the full information model and correspondingly there are four different states; baseline case, correctly detected iceberg, not detected iceberg, and falsely detected iceberg. The baseline estimates are close to those for the baseline case in Panel A. But the largest surplus estimates are for the case of not detected iceberg orders with the surplus increasing by 70 to 80% for the medium and small categories and by 180% for the large category. For the detected iceberg case the surplus estimates are positive but only for the medium category is the estimate statistically significant. The estimates for the falsely detected case are less precise and more mixed with only the medium category showing a positive expected surplus. Overall, the full information estimates that best capture the perspective of the large traders or as often is the case the large trader who obviously is aware of his iceberg order demonstrate that the greatest advantage is obtained when that iceberg order remains undetected. This is consistent with the idea that when liquidity suppliers detect an iceberg order they tend to adjust their order submission strategies and push down the surplus to a competitive level.

4.4 Surplus Breakdown-Public Information Case

Table 11 reports a breakdown of the expected surplus estimates for the public information specification. The top two rows report the baseline expected surplus by the order book position; the expected surplus for the best quote and the expected surplus for the second through the fourth best quotes. The expected surplus is negative, -0.2 to -0.6 basis points, for the best quote, consistent with the marginal order being set by a trader with some intrinsic reason for trading. Given that the half-spread is between 2 and 4.5 basis points it is plausible that a trader who will trade for sure is willing to forego immediate execution to obtain a lower transaction cost. The second through fourth best levels enjoy a positive surplus that ranges from 0.9 to 1.9 basis points.

The bottom seven rows of the table break down the expected surplus condition on an iceberg signal by order book position, the side of the iceberg order, and between displayed and hidden depth. The third and sixth rows show that the pattern observed in the baseline case for the best versus the second through fourth best quote levels persists with the best level having a negative expected surplus and the second through fourth enjoying a positive, but smaller, expected surplus.

When the surplus is split according to which side the iceberg order is on we get surplus figures that correspond to one side of the order book. For the best quote level being on the same side as the iceberg order results in a more favorable expected surplus than being on the opposite side of the iceberg. The iceberg states attract market orders increasing the execution probability while simultaneously reducing the price impact. Both effects raise the expected surplus and since the visible depth at the best quote is at the front of the queue it is unaffected by the large hidden depth. For the second through fourth best levels the effect is reversed with the *opp* surplus being higher than the *own* surplus which is zero. The zero expected surplus is primarily a consequence of the added depth ahead of the second through fourth levels. Finally, the hidden depth at the best quote expects a small negative surplus. Below we carry out the same breakdown for the full information specification of the model.

4.5 Surplus Breakdown-Full Information Case

Table 12 reports a breakdown of the expected surplus estimates for the full information specification. Panel A confirms that the baseline case generates surplus estimates that are close to those of the public information case.

Panel B reports the expected surplus estimates for observations with non-detected iceberg orders. The expected surplus for the best quote is split between the limit orders and the iceberg order's displayed peak volume and hidden volume. The iceberg order's expected surplus is between 1.4 and 3.2 basis points with approximately 70% of the surplus accruing to the visible peak volume. The own side of iceberg estimates for the limit orders at the best quotes and the displayed volume at the second through fourth best level are between 0.9 and 1.1 basis points implying that the liquidity suppliers and the large trader split the positive surplus in the not detected states. The greater surplus is consistent with less competition among the liquidity suppliers when the iceberg order is likely to be undetected.

Panel C reports the expected surplus estimates for the detected iceberg states. The expected surplus accruing to the iceberg orders is approximately 0.2 b.p. across all three categories. This positive surplus estimate together with the small positive estimate for the limit orders at the best quote (the own side case) show that even when the iceberg order is detected the side of the iceberg order captures a positive surplus that reflects a combination of higher execution probability and lower price impact. Even the second through fourth best level share in this positive surplus with estimates between 0.1 and 0.3 b.p.

Panel D reports the expected surplus estimates for observations with falsely detected iceberg orders. The estimates for the best quote are between -0.7 and -1.8 b.p. and the ordering between the *own* and *opp* cases is reversed with the own side estimates being the most negative. This is consistent with the fact that the false detection attracts market orders and raises the execution probability but the anticipated lower price impact does not materialize. Overall the breakdown shows that the iceberg order enjoys a comparative advantage when it is or is likely to be undetected. In general, liquidity suppliers who supply liquidity alongside the iceberg orders earn a higher expected surplus which reflects the two beneficial effects of the iceberg order from their perspective; iceberg orders

attract market order that tend to be larger and iceberg order dampen the price impact.

4.6 Interpreting the Evidence

4.6.1 The Order Book and Predicted Iceberg Orders

We use the signals generated by the iceberg detection algorithm to determine if the shape of the order book is different when the algorithm detects iceberg orders. Table 6 reports the mean of the spreads and the visible depths for baseline case of no iceberg and the mean change conditional on a signal of an iceberg order. Panel A of Table 6 reports that the mean bid-ask spread and the mean spread between the best and the second best quotes varies between approximately 3 and 9 basis points. Across all three categories, the bid-ask spread is 15 to 20% tighter when the algorithm detects an iceberg order. An iceberg signal has an asymmetric effect on the spread between the best and the second best quotes. On the side of the anticipated iceberg order the spread widens by 0.7 to 1.6 basis points. But this widening is counteracted by a tightening of the spread when the iceberg order is on the opposite side of 0.4 to 0.6 basis points. A typical iceberg scenario has a single iceberg order which implies that the change in the spreads across the two best levels would be the sum of the three estimates. For all three groups the sum of the changes is negative and it ranges from 0.6 to 0.9 basis points.

Panel B of Table 6 reports the mean visible depth at the best and second best quotes for the baseline case and the mean change conditional on an iceberg signal. We normalize depth by the average market order size. The visible depth at the best quote increases by 0.2 to 0.6 units when an iceberg is detected on the same side of the book and by 0.2 to 0.4 units when the iceberg is on the opposite side of the book. The corresponding changes for the second best quote are -0.2 for the own side and 0.1 to 0.3 for the opposite side. The net effect is that depth is higher for order books when an iceberg is predicted with the percentage change ranging from about 7% for the large category to over 20% for the small category.

Overall the results reported in Table 6 shows that predicted iceberg states are associated with greater visible liquidity. The total liquidity is by construction strictly greater than the visible liquidity when iceberg orders are present so that both visible and total liquidity is greater when

iceberg orders are predicted. This does correlation does not imply causation. It is possible that both large traders and liquidity suppliers change their order submissions in response to the expected liquidity demand.

4.6.2 Price Dynamics and Iceberg Orders

When the iceberg algorithm generates a signal we would expect market participants to adjust their strategies to incorporate any information conveyed by the presence of an iceberg order. But, what happens in situations in which the iceberg detection algorithm generates a false detection or fails to detect an iceberg order. Almost by construction the true state will be revealed to market participants so it is possible to quantify how market participants react in these scenarios by conditioning on both the ex ante signal generated by the detection algorithm and the true state know to the econometrician (and eventually learned by the market participants).

Panel A of Table 7 reports the expected mid-quote change conditional on the iceberg signal. Panel B reports the corresponding expected changes conditional on the iceberg signal and the true state. The cross-sectional standard errors of the estimates are reported in parentheses. The estimates in Panel A imply that the expected change in the mid-quote conditional on a signal of an iceberg order is economically and statistically close to zero consistent with informational efficiency.

Panel B shows that when the market is likely to falsely detect an iceberg or fail to detect an iceberg a reaction or a reversal is expected. The second row reports the expected change conditional on an iceberg that the algorithm fails to detect. The estimates range from -3 to -7 basis points which means that if there was an iceberg on the bid side that the algorithm failed to detect we expect to see the mid-quote increase by 3 to 7 basis points over the net 30 trades. Conversely, when the algorithm incorrectly generates a signal suggesting that there is an iceberg on the bid side when there is none we expect to see a reversal that is between 2.5 and 3 basis points. The expected mid-quote changes in the scenarios in which the algorithm makes a mistake suggests that there is information in the knowledge of iceberg orders and that information is consistent with asymmetric information or with some price pressure.

Overall the results reported in Table 7 are consistent with iceberg orders conveying some infor-

mation to the market.

5 Conclusions

Our results show that hidden liquidity changes the trading behavior of all market participant in several dimensions. The effects depend strongly on the state of detection by the other market participants. Undetected hidden liquidity apparently carry information about the future price movement of the stock. In that respect the iceberg order can be employed by an informed, but patient trader, instead of a market order which would require them to pay the spread, which would exceed the value of their information.

After detection the hidden liquidity has no further predictive power. Those are rather the period associated with greater overall liquidity and increased trading suggests that these are periods in which more gains from trade are realized. One interpretation of these findings is that market participants view the later stage of iceberg orders as positive shock to liquidity.

The similar amount of surplus that accrue from the expected executions at the best quote, show that the other liquidity suppliers take the effect of the detected hidden liquidity into account. They still have to forego their surplus of the quotes deeper in the book which they usually gain. The surplus on the order book side of the undetected liquidity is by a large fraction earned by the iceberg submitter and paid by mainly the suppliers of the opposite side and to a smaller fraction by the liquidity demanders.

A limitation of our approach is that we take as given the arrival and duration of iceberg orders. The alternatives for a trader who submits the iceberg order may include trading off the exchange or splitting up his order into smaller orders that are submitted to the order book over time. Careful modeling of these trade-offs could yield new insights about the economics of hidden liquidity and the trade-offs between transparency and liquidity. Among other things it may allow us to more definitively determine which of the above interpretations is closer to the truth. The study by Bessembinder and Venkataraman (2004) of trading at Euronext demonstrates that both iceberg orders and active trading outside the limit order book coexist.

References

- Aitken, M. J., H. Berkman, and D. Mak, 2001, “The Use of Undisclosed Limit Orders on the Australian Stock Exchange,” *Journal of Banking & Finance*, 25(8), 1589–1603.
- Anand, A., and D. G. Weaver, 2004, “Can Order Exposure Be Mandated?,” *Journal of Financial Markets*, 7(4), 405–426.
- Baruch, S., 2005, “Who Benefits from an Open Limit Order Book?,” *Journal of Business*, 78(4), 1267–1306.
- Bessembinder, H., M. Panayides, and K. Venkataraman, 2008, “Hidden Liquidity: An Analysis of Order Exposure Strategies in Electronic Stock Markets,” working paper, University of Utah.
- Bessembinder, H., and K. Venkataraman, 2004, “Does an Electronic Stock Exchange Need an Upstairs Market?,” *Journal of Financial Economics*, 73(1), 3–36.
- Biais, B., 1993, “Price Formation and Equilibrium Liquidity in Fragmented and Centralized Markets,” *Journal of Finance*, 48, 157–185.
- Biais, B., L. Glosten, and C. Spatt, 2005, “Market Microstructure: A Survey of Microfoundations, Empirical Results, and Policy Implications,” *Journal of Financial Markets*, 8, 217–264.
- Bloomfield, R., and M. O’Hara, 1999, “Market Transparency: Who Wins and Who Loses?,” *Review of Financial Studies*, 12(1), 5–35.
- , 2000, “Can Transparent Markets Survive?,” *Journal of Financial Economics*, 55, 425–459.
- Boehmer, E., G. Saar, and L. Yu, 2005, “Lifting the Veil: An Analysis of Pre-Trade Transparency at the NYSE,” *Journal of Finance*, 60(1), 783–816.
- Boni, L., and C. Leach, 2004, “Expandable Limit Order Markets,” *Journal of Financial Markets*, 7, 145–185.
- Buti, S., and B. Rindi, 2008, “Hidden Orders and Optimal Order Submission Strategies in a Dynamic Limit Order Markets,” working paper, Toronto University.

- De Winne, R., and C. D'Hondt, 2007, "Hide-and-Seek in the Market: Placing and Detecting Hidden Orders," *Review of Finance*, 11, 663–692.
- Esser, A., and B. Mönch, 2007, "The Navigation of an Iceberg: The Optimal Use of Hidden Orders," *Finance Research Letters*, 4, 68–81.
- Fleming, M. J., and B. Mizrach, 2008, "The Microstructure of a U.S. Treasury ECN: The BrokerTec Platform," working paper, Federal Reserve Bank of New York.
- Flood, M. D., R. Huisman, K. G. Koedijk, and R. J. Mahieu, 1999, "Quote Disclosure and Price Discovery in Multiple-Dealer Financial Markets," *Review of Financial Studies*, 12, 37–59.
- Foucault, T., S. Moinas, and E. Theissen, 2007, "Does Anonymity Matter in Electronic Limit Order Markets?," *Review of Financial Studies*, 20, 1705–1747.
- Harris, L., 1996, "Does a Large Minimum Price Variation Encourage Order Exposure?," working paper, University of Southern California.
- Hendershott, T., and C. M. Jones, 2005, "Island Goes Dark: Transparency, Fragmentation, and Regulation," *Review of Financial Studies*, 18, 743–793.
- Labys, W. P., 2001, "Essays on microstructure and the use of information in limit order markets," Ph.D. thesis, University of Pennsylvania (<http://repository.upenn.edu/dissertations>).
- Madhavan, A., 1995, "Consolidation, Fragmentation, and the Disclosure of Trading Information," *Review of Financial Studies*, 8, 579–603.
- , 1996, "Security Prices and Market Transparency," *Journal of Financial Intermediation*, 5, 255–283.
- , 2000, "Market Microstructure: A Survey," *Journal of Financial Markets*, 3, 205–258.
- Madhavan, A., D. Porter, and D. Weaver, 2005, "Should Securities Markets Be Transparent?," *Journal of Financial Markets*, 8, 266–288.

Moinas, S., 2007, “Hidden Limit Orders and Liquidity in Limit Order Markets,” working paper, University of Toulouse I.

O’Hara, M., 1995, *Market Microstructure Theory*. Blackwell Publishers, Cambridge, MA, 1st edn.

Pagano, M., and A. Röell, 1996, “Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading,” *Journal of Finance*, 51, 579–611.

Pardo, A., and R. Pascual, 2007, “On the Hidden Side of Liquidity,” working paper, <http://ssrn.com/abstract=459000>.

Tuttle, L., 2006, “Hidden Orders, Trading Costs and Information,” working paper, University of Kansas.

Table 1: **Sample Stocks**

	Ticker	MCAP [€b]	Trad. Vol. [€b]	Mid-Quote [€]	Trades [1000]	Trade Size [1000 shrs]	Bid-Ask Spread [b.p.]	Spread [cents]
Large	ALV	33.8	18.6	99.5	289.1	0.6	4.6	4.5
	DBK	38.2	19.8	67.6	253.4	1.2	4.0	2.7
	DCX	30.3	12.0	36.3	211.6	1.6	5.2	1.9
	DTE	34.9	22.4	15.6	284.4	5.0	7.1	1.1
	EOA	33.8	10.3	52.5	183.7	1.1	4.5	2.4
	MUV2	16.4	13.3	93.6	218.9	0.6	4.6	4.3
	SAP	27.4	11.8	131.0	179.4	0.5	4.6	6.0
	SIE	52.9	20.6	63.7	282.5	1.1	3.9	2.5
	Mean	33.5	16.1	70.0	237.9	1.5	4.8	3.2
Medium	BAS	25.4	8.0	43.1	165.1	1.1	4.7	2.0
	BAY	15.9	5.7	22.8	153.5	1.6	7.1	1.6
	BMW	12.2	5.6	34.7	134.9	1.2	5.6	1.9
	HVM	6.6	6.3	18.5	123.7	2.8	9.2	1.7
	IFX	4.8	9.4	11.6	179.0	4.5	10.0	1.2
	RWE	12.7	6.3	33.9	148.0	1.2	5.9	2.0
	VOW	9.7	6.7	39.2	162.8	1.0	5.2	2.0
	Mean	12.5	6.8	29.1	152.4	1.9	6.8	1.8
Small	ADS	4.1	2.0	92.5	62.6	0.4	6.7	6.1
	ALT	3.3	2.0	48.8	69.9	0.6	7.4	3.6
	CBK	7.6	3.4	15.3	92.7	2.4	9.6	1.5
	CONT	4.1	1.6	31.5	64.0	0.8	8.8	2.8
	DB1	4.8	2.3	47.0	62.8	0.8	7.0	3.3
	DPW	6.8	2.8	18.2	84.1	1.8	9.3	1.7
	FME	1.9	0.8	53.8	39.7	0.4	9.4	5.0
	HEN3	3.7	1.2	65.8	44.9	0.4	7.3	4.8
	LHA	4.5	2.8	14.1	86.5	2.3	10.6	1.5
	LIN	3.4	1.4	43.6	57.3	0.6	7.5	3.3
	MAN	2.4	1.8	27.7	67.6	0.9	9.1	2.5
	MEO	5.0	2.5	34.8	79.0	0.9	8.4	2.9
	SCH	7.1	3.3	41.0	97.4	0.8	6.7	2.7
	TKA	6.4	2.4	15.8	80.7	1.9	10.7	1.7
TUI	2.0	1.7	18.7	67.8	1.3	11.5	2.1	
Mean	4.5	2.1	37.9	70.5	1.1	8.7	3.0	
All	Mean	14.1	7.0	44.4	134.2	1.4	7.2	2.8

Table 1 reports the market capitalization, the trading volume, the average mid-quote, the total number of trades, the average trade size (1000 shares), and the average relative (basis points) and absolute (euro cents) bid-ask spreads for the sample stocks. The market capitalization is calculated using a free-float methodology. It is measured in billions of euros as of December 31st, 2003. All other figures are computed for the sample period January 2nd to March 31st, 2004. The categories, large, medium, and small are formed based on trading activity and used for all results below.

Table 2: **Iceberg Orders**

	Large	Medium	Small	All
Panel A: Iceberg Orders' Share of Non-Marketable Orders [%]				
Percent of shares submitted	6.8 (0.4)	11.3 (1.9)	9.8 (0.8)	9.3 (0.7)
Percent of shares executed	15.3 (1.1)	19.9 (1.4)	14.6 (1.3)	16.0 (0.9)
Panel B: Order Size				
Iceberg Orders				
Peak Size [1000 shares]	3.1 (1.1)	3.2 (0.7)	2.0 (0.2)	2.6 (0.4)
Total/Peak	7.3 (0.4)	7.1 (0.3)	7.9 (0.2)	7.6 (0.2)
Executed/Peak	4.8 (0.2)	4.9 (0.1)	4.4 (0.1)	4.6 (0.1)
Limit Orders [1000 shares]	1.4 (0.6)	1.3 (0.3)	0.7 (0.1)	1.0 (0.2)
Panel C: Median Distance between Order Price and Best Quote [basis points]				
Iceberg Orders	3.1 (0.5)	4.2 (0.9)	3.6 (0.5)	3.6 (0.3)
Limit Orders	4.0 (0.4)	4.8 (0.7)	3.5 (0.5)	3.9 (0.3)
Panel D: Order Duration to Execution or Cancellations				
Iceberg/Limit Order Duration	7.2 (0.7)	5.8 (0.6)	7.4 (0.4)	7.0 (0.3)

Table 2 reports descriptive statistics for iceberg orders and limit orders. Panel A reports the percentage of submitted and executed (non-marketable) orders that are iceberg orders. Panel B reports the average iceberg peak size (1000 shares), the average ratio of total to peak size, and the ratio of executed shares to peak size. The last row of Panel B reports the average limit order size (1000 shares). Panel C reports the average of the median distance between the order price and the best quote on the same side of the order book for iceberg orders and limit orders. Panel D reports the average ratio of the iceberg and limit order durations. Durations are computed as the time from the submission of an iceberg or limit order to the order's full execution or cancellation. The cross-sectional standard errors are reported in parentheses.

Table 3: Iceberg Detection Algorithm

True State	Algorithm Detects	
	No Iceberg $\hat{I} = 0$	Iceberg $\hat{I} = 1$
No Iceberg $I = 0$	Baseline Do not reject H_0	False Detection Type I Error $(1 - I) \times \hat{I} = I_{fls}$
Iceberg $I = 1$	Not Detected Type II Error $I \times (1 - \hat{I}) = I_{nod}$	Detected H_0 rejected $I \times \hat{I} = I_{det}$

Table 3 shows the four possible combinations of signals generated by the iceberg detection algorithm and the true state. We refer to a correct signal of no iceberg as the ‘baseline’ case, a correct signal of an iceberg as ‘detected,’ an incorrect signal of an iceberg when there is none as ‘false detection,’ and a combination of no signal of an iceberg or a no iceberg signal when an iceberg is present as a ‘not detected’ case. The off-diagonal cells correspond to Type I and Type II errors taking the null hypothesis to be ‘no iceberg.’ We use indicators with a hat, \hat{I} to indicate signals generated by the iceberg detection algorithm and indicators without a hat, I , for the true state of the book.

Table 4: **Model Summary**

PRICE IMPACT CONDITIONS		
E	$\begin{bmatrix} X_{+\tau} - X - \mu - (\bar{\alpha}h) d - (\bar{\beta}h) (d \cdot m) \\ (X_{+\tau} - X - \mu - (\bar{\alpha}h) d - (\bar{\beta}h) (d \cdot m)) d \\ (X_{+\tau} - X - \mu - (\bar{\alpha}h) d - (\bar{\beta}h) (d \cdot m)) (d \cdot m) \end{bmatrix}$	$\begin{bmatrix} \\ \\ \end{bmatrix} \otimes h = 0$
ORDER FLOW CONDITIONS		
$E [(m - \bar{\lambda}h) \otimes h] = 0$		
HIDDEN DEPTH CONDITIONS		
$E \begin{bmatrix} (\hat{q}^{ask} - \eta \hat{I}^{ask}) \hat{I}^{ask} \\ (\hat{q}^{bid} - \eta \hat{I}^{bid}) \hat{I}^{bid} \end{bmatrix} = 0$		
PROBABILITY OF BUY MARKET ORDER		
$E [(d^+ - \bar{\phi}h^{bid}) \otimes h^{bid}] = 0$		
MODEL VARIABLES		
X	$X_{+\tau}$	mid-quote - current and at $t + \tau$
\hat{q}^{bid}	\hat{q}^{ask}	hidden depth at 1st quotes
d	m	signed market order (buy: $d = +1$, sell: $d = -1$), and normalized market order size
I^{ask}	I^{bid}	Indicators for true iceberg orders
\hat{I}^{ask}	\hat{I}^{bid}	Indicators for iceberg orders predicted by algorithm
AUXILIARY VARIABLES		
$h^{bid} = [1 \ \hat{I}^{bid} \ \hat{I}^{ask}]'$		iceberg states - bid side view
$h^{ask} = [1 \ \hat{I}^{ask} \ \hat{I}^{bid}]'$		iceberg states - ask side view
h		iceberg states - interacted with realized market order
$h = h^{bid}$ if $d_t = -1$ and $h = h^{ask}$ if $d_t = +1$		
MODEL PARAMETERS		
μ		drift of the share price
$\bar{\alpha} = [\alpha \ \alpha^{own} \ \alpha^{opp}]$		fixed component of the price impact
$\bar{\beta} = [\beta \ \beta^{own} \ \beta^{opp}]$		variable component of the price impact
$\bar{\lambda} = [\lambda \ \lambda^{own} \ \lambda^{opp}]$		mean market order size
η		mean hidden depth
$\bar{\phi} = [\phi \ \phi^{own} \ -\phi^{own}]$		probability of buy market order conditional on market order

Table 5: Performance of the Iceberg Detection Algorithm (Ask Side)

True State	Large			Medium			Small		
	Prediction			Prediction			Prediction		
	No Iceberg	Iceberg	Sum	No Iceberg	Iceberg	Sum	No Iceberg	Iceberg	Sum
	$\hat{I}^{ask} = 0$	$\hat{I}^{ask} = 1$		$\hat{I}^{ask} = 0$	$\hat{I}^{ask} = 1$		$\hat{I}^{ask} = 0$	$\hat{I}^{ask} = 1$	
No Iceb.	90.0	1.0	91.0	84.7	1.5	86.2	89.9	0.8	90.7
$I^{ask} = 0$	(1.2)	(0.2)	(1.0)	(2.1)	(0.2)	(1.9)	(1.4)	(0.1)	(1.1)
Iceb.	2.9	6.1	9.0	4.6	9.2	13.8	3.1	6.1	9.3
$I^{ask} = 1$	(0.5)	(0.5)	(1.0)	(0.8)	(1.1)	(1.9)	(0.4)	(0.8)	(1.3)
Sum	92.9	7.1	.	89.3	10.7	.	93.0	7.0	.
	(0.7)	(0.7)	.	(1.3)	(1.3)	.	(0.9)	(0.8)	.

Table 5 reports, for the ask side of the book, the cross-sectional averages for the predictions (no iceberg versus iceberg) generated by the iceberg detection algorithm across the percentages of the true iceberg state (no iceberg versus iceberg). The columns of each 2×2 matrix correspond to the algorithm's predictions and the rows correspond to the true states. Entries on the diagonal correspond to correct predictions. The cross-sectional standard errors of the means are reported in parentheses below each mean. The results for the bid side are similar with the largest difference being of the order of less than one-half percentage points.

Table 6: **Limit Order Books and Predicted Icebergs**

	Large		Medium		Small	
Panel A: Spreads						
Bid-Ask Spread						
No Iceberg	4.91	(0.35)	6.98	(0.79)	8.76	(0.43)
Iceberg	-0.97	(0.13)	-1.00	(0.08)	-1.85	(0.16)
2nd Best Quote - Best Quote						
No Iceberg	3.26	(0.48)	5.27	(0.78)	5.75	(0.42)
Iceberg Own Side	0.79	(0.19)	0.66	(0.15)	1.57	(0.18)
Iceberg Opposite Side	-0.37	(0.05)	-0.37	(0.04)	-0.64	(0.05)
Panel B: Depths						
Depth at Best Quote						
No Iceberg	2.20	(0.54)	2.03	(0.23)	1.65	(0.07)
Iceberg Own Side	0.28	(0.08)	0.15	(0.10)	0.55	(0.13)
Iceberg Opposite Side	0.22	(0.10)	0.36	(0.08)	0.25	(0.03)
Depth at 2nd Best Quote						
No Iceberg	2.93	(1.08)	2.56	(0.44)	1.85	(0.12)
Iceberg Own Side	-0.23	(0.06)	-0.20	(0.13)	-0.23	(0.04)
Iceberg Opposite Side	0.07	(0.03)	0.31	(0.14)	0.15	(0.04)

Table 6 reports the cross-sectional average of the conditional means of (i) the bid-ask spread, (ii) the spread between the best and second best quotes, (iii) the visible depth at the best quote, and (iv) the visible depth at the second best quote. The first two rows report the mean bid-ask spread conditional on a ‘No Iceberg’ or an ‘Iceberg’ signal from the detection algorithm. The next three rows split the means conditional on an iceberg signal into those signaling an iceberg on the ‘own’ or the ‘opposite’ side of the order book. Cross-sectional standard errors are reported in parentheses. Spreads are measured in basis points. The depths have been normalized for each stock by the stock-specific average market order size so that the depth is measured in units of averaged-sized market orders.

Table 7: **Expected Mid-quote Changes and Iceberg Orders**

	Large		Medium		Small	
Panel A: Ex Ante (predicted)						
$\Delta m_{q+\tau} = c + a \left(\hat{I}^{ask} - \hat{I}^{bid} \right) + \epsilon_{+\tau}$.						
Intercept	-0.14	(0.03)	-0.27	(0.10)	-0.21	(0.09)
Predicted Iceberg (alg)	0.23	(0.09)	0.35	(0.14)	-0.42	(0.54)
Panel B: Ex Post (predicted + true state)						
$\Delta m_{q+\tau} = c + a_{nod} (I_{nod}^{ask} - I_{nod}^{bid}) + a_{det} (I_{det}^{ask} - I_{det}^{bid}) + a_{fls} (I_{fls}^{ask} - I_{fls}^{bid}) + \epsilon_{+\tau}$.						
Intercept	-0.15	(0.03)	-0.26	(0.10)	-0.21	(0.10)
Not detected (nod)	-2.96	(0.27)	-3.99	(0.19)	-7.39	(0.58)
Detected (det)	-0.18	(0.12)	-0.23	(0.18)	-1.03	(0.48)
False detection (fls)	2.56	(0.36)	2.85	(0.24)	2.89	(0.95)

Table 7 reports the mean of the mid-quote over a 30 trade horizon conditional on different information sets. Panel A reports the means conditional on the information generated by the iceberg detection algorithm; the algorithm either generates an iceberg or a no iceberg signal. The indicators I_{nod}^{ask} , I_{det}^{ask} , and I_{fls}^{ask} are defined as follows: $I_{nod}^{ask} = (1 - \hat{I}^{ask}) \times I^{ask}$; $I_{det}^{ask} = \hat{I}^{ask} \times I^{ask}$; and $I_{fls}^{ask} = \hat{I}^{ask} \times (1 - I^{ask})$, with the analogous definitions for the bid side. A negative coefficient on $-I_{nod}^{bid}$ implies an expected positive change in the mid-quote over the next 30 trades. Conversely, a positive coefficient on $-I_{fls}^{bid}$ implies an expected negative change in the mid-quote over the next 30 trades. Cross-sectional standard errors are reported in parentheses. Mid-quote changes are measured in basis points.

Table 8: **Model Parameter Estimates (Public Information)**

	Large		Medium		Small	
Panel A: Price Impact Function						
α	1.81	(0.02)	2.51	(0.03)	3.17	(0.04)
α^{own}	-0.65	(0.07)	-0.59	(0.11)	-1.21	(0.18)
α^{opp}	0.33	(0.07)	0.82	(0.12)	1.32	(0.20)
β	0.41	(0.01)	0.58	(0.02)	0.90	(0.02)
β^{own}	-0.18	(0.02)	-0.21	(0.03)	-0.37	(0.06)
β^{opp}	-0.02	(0.02)	-0.15	(0.04)	-0.22	(0.07)
μ	-0.15	(0.03)	-0.30	(0.05)	-0.21	(0.07)
Panel B: Market Order Size						
λ	0.97	(0.00)	0.95	(0.00)	0.96	(0.00)
λ^{own}	0.29	(0.01)	0.27	(0.01)	0.36	(0.01)
λ^{opp}	0.07	(0.01)	0.15	(0.01)	0.13	(0.01)
Panel C: Hidden Depth Size						
η	8.42	(0.06)	7.18	(0.05)	8.29	(0.07)
Panel D: Probability of Buy Market Order						
ϕ	0.50	(0.00)	0.50	(0.00)	0.50	(0.00)
ϕ^{own}	0.19	(0.00)	0.17	(0.00)	0.17	(0.00)

Table 8 reports the model parameter estimates for the large, medium, and small categories. Table 4 summarize the model equations and the parameters. The model parameters are estimated using GMM using a Newey-West 10-lag weighting matrix in the second stage. Standard errors are reported in parentheses. The symmetry assumption implies that, ϕ^{opp} , the probability of a buy market order conditional on a market order and an iceberg on the ask side, is simply $\phi - \phi^{own}$.

Table 9: Model Parameter Estimates (Full Information)

	Large		Medium		Small	
Panel A: Price Impact Function						
α	1.84	(0.02)	2.54	(0.03)	3.17	(0.04)
α_{nod}^{own}	-3.36	(0.08)	-4.63	(0.14)	-7.12	(0.22)
α_{det}^{own}	-1.14	(0.07)	-1.32	(0.11)	-2.07	(0.19)
α_{fls}^{own}	1.51	(0.13)	2.48	(0.21)	2.96	(0.38)
α_{nod}^{opp}	2.48	(0.09)	4.02	(0.16)	6.54	(0.24)
α_{det}^{opp}	0.69	(0.08)	1.35	(0.13)	1.98	(0.21)
α_{fls}^{opp}	-1.10	(0.17)	-1.08	(0.30)	-1.74	(0.53)
β	0.45	(0.01)	0.64	(0.02)	0.97	(0.03)
β_{nod}^{own}	-0.19	(0.02)	-0.14	(0.04)	-0.32	(0.07)
β_{det}^{own}	-0.19	(0.02)	-0.23	(0.03)	-0.41	(0.06)
β_{fls}^{own}	-0.20	(0.05)	-0.30	(0.08)	-0.25	(0.21)
β_{nod}^{opp}	-0.12	(0.03)	-0.11	(0.05)	-0.26	(0.09)
β_{det}^{opp}	-0.05	(0.03)	-0.16	(0.04)	-0.22	(0.08)
β_{fls}^{opp}	0.06	(0.06)	-0.15	(0.09)	-0.17	(0.17)
μ	-0.16	(0.03)	-0.29	(0.05)	-0.21	(0.07)
Panel B: Market Order Size						
λ	0.95	(0.00)	0.92	(0.00)	0.93	(0.00)
λ_{nod}^{own}	0.40	(0.01)	0.44	(0.01)	0.43	(0.01)
λ_{det}^{own}	0.35	(0.01)	0.33	(0.01)	0.42	(0.01)
λ_{fls}^{own}	0.04	(0.01)	0.02	(0.01)	0.06	(0.01)
λ_{nod}^{opp}	0.11	(0.01)	0.18	(0.01)	0.21	(0.01)
λ_{det}^{opp}	0.09	(0.01)	0.17	(0.01)	0.15	(0.01)
λ_{fls}^{opp}	0.02	(0.02)	0.07	(0.02)	0.10	(0.02)
Panel C: Size of Hidden Depth						
η_{nod}	12.35	(0.09)	10.71	(0.07)	12.02	(0.08)
η_{det}	9.86	(0.07)	8.30	(0.05)	9.42	(0.07)
Panel D: Probability of Buy Market Order						
ϕ	0.50	(0.00)	0.50	(0.00)	0.50	(0.00)
ϕ_{nod}	0.05	(0.00)	0.04	(0.00)	0.03	(0.00)
ϕ_{det}	0.20	(0.00)	0.18	(0.00)	0.18	(0.00)
ϕ_{fls}	0.17	(0.00)	0.15	(0.00)	0.16	(0.00)

Table 9 reports the model parameter estimates for the large, medium, and small categories for the full information specification. Superscripts *own* and *opp* denote an iceberg on the opposite or own side of the book from the perspective. The subscripts *det*, *nod*, and *fls* denote correctly detected, not detected, and falsely detected iceberg states. Table 4 summarize the model equations and the parameters. The model parameters are estimated using GMM using a Newey-West 10-lag weighting matrix in the second stage. Standard errors are reported in parentheses.

Table 10: **Liquidity Supplier Surplus**

	Large		Medium		Small	
Panel A: Public Information Model						
Baseline ($\hat{I} = 0$)	0.70	(0.02)	1.32	(0.04)	1.36	(0.06)
Iceberg ($\hat{I} = 1$)	-0.04	(0.06)	0.05	(0.09)	-0.36	(0.18)
All States	0.60	(0.02)	1.06	(0.04)	1.10	(0.06)
Panel B: Full Information Model						
Baseline ($I = 0 \wedge \hat{I} = 0$)	0.57	(0.02)	1.09	(0.04)	1.15	(0.06)
Not Detected ($I = 1 \wedge \hat{I} = 0$)	1.58	(0.07)	1.82	(0.12)	2.09	(0.20)
Detected ($I = 1 \wedge \hat{I} = 1$)	0.22	(0.07)	0.46	(0.10)	0.24	(0.19)
False Detection ($I = 0 \wedge \hat{I} = 1$)	-0.02	(0.10)	0.59	(0.16)	-0.13	(0.35)
All States	0.58	(0.02)	1.02	(0.04)	1.05	(0.06)

Table 10 reports the estimated expected surplus accruing to liquidity suppliers. The expected surplus is computed using the parameter estimates for the state-dependent market order distribution and price impact function and the empirical frequencies of buy versus sell orders in each state. The surplus is calculated by integrating the observed price schedule minus the price impact function with respect to the distribution of market order quantities. Panel A reports the surplus for the states of the public information model, panel B does the same for those of the full information model. In both cases surplus is calculated for the order book up to 4th order book quote. The calculation of panel A assumes expected hidden volumes at observations for which the algorithm indicates iceberg orders $\hat{I} = 1$, for panel B includes the actual hidden volumes. Standard errors in parenthesis take into account first-stage estimation error for the model parameters.

Table 11: Surplus Breakdown - Public Information

	Large		Medium		Small	
<hr/> <hr/> Baseline ($\hat{I} = 0$) <hr/>						
Displayed Volume at Best Quote	-0.20	(0.01)	-0.30	(0.02)	-0.57	(0.03)
Displayed Volume at 2nd-4th Quotes	0.90	(0.01)	1.60	(0.02)	1.91	(0.02)
<hr/> <hr/> Iceberg ($\hat{I} = 1$) <hr/>						
Displayed Volume at Best Quote	-0.09	(0.04)	-0.28	(0.05)	-0.60	(0.11)
- own Side of Iceberg	0.14	(0.04)	0.07	(0.06)	0.01	(0.11)
- opposite Side of Iceberg	-0.23	(0.02)	-0.35	(0.03)	-0.60	(0.05)
Displayed Volume at 2nd-4th Quotes	0.16	(0.01)	0.48	(0.02)	0.48	(0.03)
- own Side of Iceberg	0.00	(0.00)	0.00	(0.00)	0.00	(0.00)
- opposite Side of Iceberg	0.16	(0.01)	0.48	(0.02)	0.48	(0.03)
(Predicted) Hidden Volume at Best Quote	-0.11	(0.02)	-0.16	(0.04)	-0.25	(0.07)

Table 11 reports the estimated expected surplus in different parts of the order book for the public information model. The upper part shows the contribution of the best and second to fourth quotes at observations where the algorithm has not detected iceberg orders. The bottom part does the same for observations where the algorithm indicated iceberg orders. It further distinguishes between the sides of the order book relative to the assumed iceberg order. The last row is the surplus of the predicted hidden volume given by the model parameter η .

Table 12 on the next page reports the expected surplus for the full information model. For the baseline case the contribution of the different order book levels are shown. For the not-detected and the detected iceberg order it details the expected surplus at the best quote for submitters of limit orders and iceberg orders respectively. For limit orders the surplus is given relative to the iceberg order, whereas for the iceberg orders it differentiates between the displayed peak volume and the hidden volume. For a false detection the surplus is divided into liquidity submitted on the side of the (falsely) predicted iceberg and that on the opposite side.

For both tables the surplus is calculated by integrating the observed price schedule minus the price impact function with respect to the distribution of market order quantities. Standard errors in parenthesis take into account first-stage estimation error for the model parameters.

Table 12: Surplus Breakdown - Full Information

	Large		Medium		Small	
A. Baseline ($I = 0 \wedge \hat{I} = 0$)						
Displayed at Best Quote	-0.27	(0.01)	-0.38	(0.02)	-0.64	(0.03)
Displayed at 2nd-4th Quotes	0.84	(0.01)	1.47	(0.02)	1.78	(0.02)
B. Not Detected ($I = 1 \wedge \hat{I} = 0$)						
Limit Orders at Best Quote	0.01	(0.03)	-0.46	(0.05)	-1.25	(0.09)
- own side of iceberg	0.83	(0.02)	0.94	(0.03)	1.01	(0.04)
- opposite side of iceberg	-0.81	(0.03)	-1.40	(0.05)	-2.26	(0.09)
Iceberg Orders at Best Quote	1.41	(0.05)	1.82	(0.08)	3.15	(0.13)
- Displayed Peak Volume	1.03	(0.03)	1.24	(0.04)	2.18	(0.08)
- Hidden Volume	0.39	(0.02)	0.58	(0.03)	0.98	(0.06)
Displayed at 2nd-4th Quotes	0.16	(0.02)	0.46	(0.03)	0.18	(0.07)
- own side	0.04	(0.00)	0.12	(0.00)	0.10	(0.00)
- opposite side	0.12	(0.02)	0.34	(0.03)	0.08	(0.07)
C. Detected ($I = 1 \wedge \hat{I} = 1$)						
Limit Orders at Best Quote	-0.20	(0.02)	-0.37	(0.03)	-0.67	(0.05)
- own side of iceberg	0.13	(0.01)	0.14	(0.02)	0.11	(0.03)
- opposite side of iceberg	-0.32	(0.02)	-0.51	(0.03)	-0.79	(0.05)
Iceberg Orders at Best Quote	0.17	(0.06)	0.20	(0.08)	0.24	(0.16)
- Displayed Peak Volume	0.20	(0.03)	0.22	(0.05)	0.29	(0.09)
- Hidden Volume	-0.03	(0.02)	-0.02	(0.03)	-0.05	(0.07)
Displayed at 2nd-4th Quotes	0.24	(0.01)	0.62	(0.02)	0.66	(0.03)
- own side of iceberg	0.13	(0.00)	0.25	(0.01)	0.30	(0.01)
- opposite side of iceberg	0.11	(0.01)	0.38	(0.02)	0.37	(0.03)
D. False Detection ($I = 0 \wedge \hat{I} = 1$)						
Displayed at Best Quote	-0.72	(0.07)	-0.89	(0.10)	-1.83	(0.23)
- own side of (predicted) iceberg	-0.79	(0.07)	-1.05	(0.09)	-1.88	(0.21)
- opposite side of (predicted) iceberg	0.07	(0.04)	0.16	(0.08)	0.05	(0.13)
Displayed at 2nd-4th Quotes	0.70	(0.03)	1.47	(0.06)	1.68	(0.13)
- own side of (predicted) iceberg	0.44	(0.03)	0.82	(0.05)	0.93	(0.11)
- opposite side of (predicted) iceberg	0.26	(0.02)	0.65	(0.04)	0.74	(0.07)

For details refer to the previous page under table 11.

Appendix

A1 Surplus Calculation

The formulas presented in sections 3.6 are independent of the choice of the market order size distribution. Below we provide details about the surplus calculation given the choice of an exponential distribution.

The expected surplus, $\pi_k(q)$, of a liquidity provider for unit q in equation (3) is

$$\pi_k(q) = Pr(d = 1; h) [p_k - X - \mu - (\alpha h) - (\beta h)(q + \lambda h)] \left[1 - e^{-\frac{q}{\lambda h}}\right] \quad (\text{A } 1)$$

The aggregate expected surplus of a quote, used in equations (4) and (5), in the case of the exponential distribution is calculated by

$$\begin{aligned} \Pi = Pr(d = 1; h)(\lambda h) & \left\{ [p_k - X - \mu - (\alpha h) - (\beta h)(\lambda h)] \left[e^{-\frac{Q_s}{\lambda h}} - e^{-\frac{Q_e}{\lambda h}} \right] \right. \\ & \left. + (\beta h) \left[Q_s \left(e^{-\frac{Q_s}{\lambda h}} \right) - Q_e \left(e^{-\frac{Q_e}{\lambda h}} \right) \right] \right\} \quad (\text{A } 2) \end{aligned}$$

where the quote starts at Q_s and ends at Q_e . The two formulas above are valid for the ask side, similar equations can be derived for the bid side.

A two stage approach is applied for the standard errors of the surplus expressions. The first stage considers the uncertainty of the estimated parameters of the GMM model via the delta method. The second stage errors are those of the sample average used to calculate the surplus expressions itself. The two errors are combined assuming independence. Denote S as one of the surplus calculations, then an estimator of the variance of S , $\hat{\sigma}_S^2$, is derived by the following expression

$$\hat{\sigma}_S^2 = \frac{\partial S}{\partial \theta'} \hat{\mathbf{V}}(\hat{\theta}) \frac{\partial S}{\partial \theta} + s_S^2 \quad (\text{A } 3)$$

where $\hat{\mathbf{V}}(\theta)$ is the estimated variance matrix of the model parameters and s_S^2 is the sampling variance of the surplus calculation. The gradient vector of the surplus expression to the parameters $\frac{\partial S}{\partial \theta}$ is calculated numerically.

A2 Detection Algorithm: An Example

The following example illustrates the detection algorithm introduced in Section 3.4. The algorithm maintains and updates a detection vector with the following four elements for all price levels in the order book: price; detection flag; visible volume; volume until next replenishment. In the example, the best bid is initially 9.70 and there are no iceberg orders at this price level, i.e., the detection flag is zero. The first two elements (price and detection flag) are listed in the right-side margin.

- (9.70, 0)
1. A buy iceberg order is submitted at 9.77 with a total size of 9000, a peak size of 1000, implying a hidden volume of 8000 and a visible volume of 1,000 shares. The detection flag is currently zero. Detection vector: [9.77, 0, 1000, 0]. (9.77, 0)
 2. A sell market order for 1200 shares is submitted. The iceberg order's first peak size of 1000 shares is completely executed and another 200 share are automatically executed from the iceberg order's second peak size of 1000 shares. The remaining 800 shares of the iceberg order's second peak size are displayed in the book. The algorithm sets the detection flag to one and sets the volume until next replenishment to 800. Detection vector: [9.77, 1, 800, 800]. (9.77, 1)
 3. A buy limit order is submitted at 9.79 for 200 shares. The flag for 9.79 is zero. Detection vector: (9.79, 0) (9.79, 0, 200, 0). There is no change at 9.77 . (9.77, 1)
 4. A sell market order for 500 shares is submitted. 200 shares are executed at 9.79, and 300 shares are executed against the second peak size of the iceberg order at 9.77. The detection flag at 9.77 remains at one, and the volume until next replenishment is revised to 500. Detection vector: [9.77, 1, 500,500]. (9.77, 1)
 5. A buy limit order is submitted at 9.77 for 5000 shares. The flag at 9.77 remains at one and volume until replenishment is unchanged because the new limit order is behind the visible 500 share of the iceberg order's second peak size. Detection vector: [9.77, 1, 5500, 500]. (9.77, 1)
 6. The iceberg order at 9.77 is cancelled. The detection flag remains at one.⁷ Detection vector: [9.77, 1, 5000, 500]. (9.77, 1)
 7. A sell market order is submitted for 600 shares. The flag is reset to zero as the volume until next replenishment is exceeded without the expected replenishment. Detection vector: [9.77, 0, 4500, 0]. (9.77, 0)

⁷In principle, the algorithm could make use of the fact that the remaining volume of 5,000 at 9.77 exactly matches the size of the previously submitted limit order to infer that it is very likely that the drop in the visible volume at 9.77 was caused by the iceberg being cancelled. However, if the limit order in question was for 500 shares instead of 5,000 shares, it would be a fifty-fifty chance that the cancellation was due to the iceberg order. Our algorithm has not been optimized with respect to these scenarios so it is possible that its performance could be enhanced.