Finance Research Letters
Volume 4, Issue 2, June 2007, Pages 68-81

The navigation of an iceberg: The optimal use of hidden orders

Angelika Esser [a] [b], Burkart Mönch [a] [b]

https://doi.org/10.1016/j.frl.2006.12.003

Abstract

Market participants with large orders to execute are often reluctant to expose these to an open order book in their entirety in order to avoid a potential adverse market impact. Therefore, investors often split large orders into smaller tranches. Iceberg orders facilitate these trading practices by executing such business automatically in the order book. This article analyzes the rationale for the use of iceberg orders by assessing the costs and benefits of this trading instrument.

Introduction

Electronic trading platforms often provide market participants with information on an anonymous open order book during continuous trading in real time. Usually the limits, the accumulated order volumes of each limit, and the number of orders in the book at each limit are displayed, so that traders can assess the altering order flow and the market liquidity.

What does an open order book imply for institutional investors who want to submit limit orders, the total volume of which is large relative to others in the market? No doubt, exposing large limit orders in an open order book may reveal the investor's motives for trading and may raise suspicion that the originator of the large order has access to private information about the true value of the security under consideration. Consequently, other market participants change their own order submission strategy, which in turn lowers the probability that the large order will be executed at the prespecified limit. The institutional investor then has to choose a less favorable limit in order to increase the probability of execution and, thus, suffers losses from the *indirect* adverse price impact of the large exposure in the order book. A possible solution is not to submit one large limit order but to split the order into several smaller limit orders, which are submitted over time. For this reason many electronic trading platforms introduced so-called iceberg orders. Euronext, the Toronto Stock Exchange, the London Stock Exchange (with its order driven services SETS, SETSmm, and IOB), and XETRA are just some examples. Iceberg orders allow market participants to submit orders with only a certain portion of the order publicly disclosed. The metaphor alludes to the fact that, in nature, an iceberg's biggest part floats unobservable under the water. Only one-ninth of the mass of ice is seen above the water surface.

An iceberg order is specified by its mandatory limit, its overall volume, and a peak volume. The peak is the visible part of the iceberg order and is introduced into the order book with the original time stamp of the iceberg order according to price/time priority. As soon as the disclosed volume of an iceberg order has received a complete fill and a hidden volume is still available, a new peak is entered into the book with a new time stamp. The new peak behaves in an identical manner to a conventional limit order. From this point of view, a pure limit order is basically a special case of an iceberg order where the peak volume coincides with the total order volume.

However, it is important to note that iceberg orders exhibit a less favorable time priority compared with pure limit orders. After the peak of an iceberg order is completely matched, all visible limit orders at the same limit that were entered before the new peak is displayed take priority, i.e. they must have received a complete fill before the new peak comes to execution.

When submitting an iceberg order to the market, several issues have to be considered. Imagine, for example, that the management of a mutual fund has to close a large position in a single stock within one trading day. Using an iceberg order with only a small peak size allows it to minimize the adverse informational impact of disclosing the actual order volume. However, the smaller the peak size, the less favorable the time priority of the overall order. Thus, choosing a peak size that is too small seems suboptimal. Such a strategy would significantly lengthen the time to complete execution or would make a complete fill unlikely. Moreover, the fund managers have to choose a reasonable limit for the order. If the limit is too low, one may miss some trade opportunities, i.e. one would give away the chance to participate in raising stock prices. Otherwise, if the limit is too ambitious, the order is unlikely to receive a complete fill.

We model this tradeoff analytically in a continuous time setup where a large position in a single stock is to be liquidated within a finite trading window.[1] The model yields the optimal peak size and the optimal order limit by maximizing the expected payoff of the liquidation strategy under certain assumptions concerning the execution risk of the iceberg order. Note that a pure limit order would also be an admissible solution to our optimization problem.

Unless an iceberg sell order is immediately executable, i.e. the limit is so low that it is actually a market sell order, the probability of receiving a complete fill within a finite time horizon is strictly smaller than one. In principle, at least two alternative approaches would be able to incorporate execution risk into a liquidation model.

First, one may assume that the investor is forced to trade the remaining shares with a market order if the iceberg order fails to receive complete execution. Consequently, the investor has to bear a liquidity discount, so that he or she gets penalized for every share that could not be sold via the iceberg order. However, in our opinion such a rigorous assumption may not always be justified in practice, especially if the remaining order volume under consideration cannot be absorbed by the market without a significant price change.

In this case, investors typically follow a dynamic strategy, i.e. they review their orders frequently and adjust them if the market moves away from the prespecified order limit. For this reason, we propose a different approach that considers the execution probability as a boundary condition, i.e. only those combinations of peak size and limit are admissible that assure a certain execution probability within a prespecified time horizon. Compared with the first concept, this approach is rather flexible and does not require any assumption concerning the liquidation of the unexecuted part of the iceberg order. To get a flavor of the concept, imagine, for example, an investor who wants to liquidate a large position, say, within one week. At the end of each trading day, the investor inspects the state of the iceberg order and, if necessary, adjusts the limit or the peak size to reach the target.[2] Our approach can assist the investor in this procedure. It deals with the optimal combination of order limit and peak size that maximizes the expected liquidation revenues in the case of complete execution, given that the probability to receive a complete fill exceeds a certain level, for example 40% within one trading day. If the order remains partially or completely unexecuted by the end of the first day, the investor may wish to rerun the optimization on the second day and, thereby, increase the execution probability, let us say to 60% and so on. If a substantial part of the order is still unexecuted on the last day of the week, the investor will probably choose a minimum execution probability that is close to one.

The technical design of the model can be summarized as follows: During continuous trading, a transaction takes place if an order becomes executable against orders on the other side of the book. Thus, for an iceberg sell order that is stored on the ask side of the book, the dynamics of the best bid price are of special interest. We model the best bid price as a stochastic process in continuous time and assume a constant best bid size. If the stochastic process hits the limit of the iceberg order, a transaction is executed and the stochastic process jumps back to the next lower limit. Whether the peak of the iceberg order or another sell order at the same limit is processed at this event depends on the relative time priority of the orders. If new orders with the same limit as the iceberg order are submitted continuously to the book, the time priority of the iceberg order deteriorates compared with a pure limit order. The smaller the peak size of the iceberg order, the more often the limit must be hit such that the iceberg order receives a complete fill.

On the other hand, a smaller peak size lowers the adverse informational impact of showing the actual order volume in an open book. We model the drift of the stock price process as a function of the visible order imbalance. When the peak size of an iceberg order enters the book, the visible order imbalance changes. We exemplify empirically, using order book data, that current variations in the visible order book imbalance are positively correlated with future returns. Thus, the higher the peak size of an iceberg sell order, the smaller the order imbalance and the smaller the expected returns in the next time intervals. Consequently, a higher peak size results in a smaller probability that the stock price process will reach the prespecified limit within the given time horizon.

In total, one can observe two opposite effects if the peak size of an iceberg sell order is reduced in our model:

- The drift of the stochastic process is reduced to a smaller extent when the order enters the book.
- The number of times the limit must be hit in order to process the iceberg order completely increases.

While the first effect is beneficial for the originator of the iceberg order, the latter is not. Our framework weights these effects and identifies the optimal combination of peak size and order limit.

The rest of the paper is organized as follows: Section 2 briefly reviews the related literature. Section 3 introduces the theoretical setup. In Section 4, we explicitly model the drift as a function of the order imbalance. Our approach to determine the optimal combination of order limit and peak size is calibrated with a clinical order book data sample in Section 5. The paper concludes in Section 6 with a brief summary and a discussion of issues for further research.

Section snippets

Related literature

To our knowledge, the model discussed in Moinas (2005) is the only theoretical approach that considers iceberg orders as trading instruments. In this model, liquidity suppliers are potentially informed and can submit limit or iceberg orders. The informed traders try to mimic the uninformed investors in order not to be detected. Depending on the degree of adverse selection, the informed investors use iceberg orders to trade large volumes without signaling their presence.

Lo et al. (2002) propose…

General idea and dynamics

This section introduces the basic concepts and provides the motivation for the assumptions that are made. Assume that the large investor holds ϕ_0 shares that should be liquidated before time *T*. For this purpose, the trader submits an iceberg sell order that is stored on the ask side of the order book. The investor assigns a peak size ϕ_p and a limit $\overline{S}$ to the iceberg order. The latter is strictly higher than the initial best bid price S_0 such that the first proportion of the order is not…

Modeling of the drift component

Up to now, the drift of the best bid price has been assumed to be a constant. This section completes the theoretical framework by modeling explicitly the impact of the peak size on the drift following the intuition that the disclosure of large order volumes has an adverse market impact. For this purpose, we will model the drift μ_t as a function of the order imbalance B_t.

Similar to Brown et al. (1997), we define the imbalance B_t of the order book as the number of shares displayed on the bid side …

Dataset

XETRA is an order driven market where investors, by placing limit orders, establish prices at which other participants can buy or sell shares. A trade takes place whenever a counterpart order hits the quotes. At the time the data were collected, XETRA was open for trading from 9:00 a.m. to 8:00 p.m. Market participants can see all nonhidden entries on each side of the order book, but trading in XETRA is anonymous, i.e. market participants cannot identify the counterparts.

Based on event…

Summary and conclusion

This paper introduces a setup that enables the determination of the optimal combination of limit and peak size of an iceberg order, given a large position in a security that should be liquidated within a finite time horizon. The framework balances the direct advantage of a large peak size that leads to a better time priority of an iceberg order and the adverse informational impact of revealing large order volumes in an open order book. Furthermore, it assesses the tradeoff between the order…

Acknowledgements

This research was partially supported by the German Research Foundation (Deutsche Forschungsgemeinschaft). We are grateful to the Trading Surveillance Office of the Deutsche Börse AG for providing us a limit order book dataset. We are indebted to Nicole Branger and Christian Schlag for inspiration and insightful discussions. Special thanks for helpful comments go to an anonymous referee, Christoph Benkert, Matthias Birkner, Diego Garcia, Jeff Harris, Götz Kersting, Micong Klimes, Michael…

References (8)

A. Ranaldo
Order aggressiveness in limit order markets
Journal of Financial Markets (2004)

A.W. Lo *et al.*
Econometric models of limit-order execution
Journal of Financial Economics (2002)

P. Brown *et al.*
The interaction between order imbalance and stock price
Pacific-Basin Finance Journal (1997)

Cao, C., Hansch, O., Wang, X., 2004. The informational content of an open limit order book. Working paper, Penn State...

There are more references available in the full text version of this article.

Cited by (24)

Futures volatility forecasting based on big data analytics with incorporating an order imbalance effect
2022, International Review of Financial Analysis

Show abstract

The impact of reporting changes on hidden liquidity: Evidence from the Chicago stock exchange
2022, Global Finance Journal

Show abstract

Blockchain for Internet of Energy management: Review, solutions, and challenges
2020, Computer Communications

Show abstract

Optimal order execution using hidden orders

2018, Journal of Economic Dynamics and Control

> *Citation Excerpt :*
> …One strand of this literature considers similar order execution or exposure problems faced by traders in a static framework. For example, Esser and Mönch (2007) derived the optimal limit price and peak size for an iceberg order for a static liquidation strategy. Cebiroğlu and Horst (2015) studied a single-stage optimal order exposure problem (how many units to display or hide) which a trader faces, where the trader's exposed limit order affects incoming order flows.…

Show abstract

Pre-trade transparency and informed trading: Experimental evidence on undisclosed orders

2016, Journal of Financial Markets

> *Citation Excerpt :*
> …Before I describe the details of my experimental design, I briefly revisit the key predictions recent theories provide on the source of undisclosed orders and the impact of opacity on market quality. Liquidity traders submit undisclosed orders to reduce the price impact of large orders (Esser and Mönch, 2007). Submitting an undisclosed order with a fixed peak size is a dominant order-splitting strategy as opposed to one splitting large limit orders because, given the time priority rule in the market, the visible part of the undisclosed order is automatically executed.…

Show abstract

Optimal order display in limit order markets with liquidity competition

2015, Journal of Economic Dynamics and Control

> *Citation Excerpt :*
> …When passive (limit) orders are considered, the analysis is usually confined to dark pool orders which induce little or no market impact.5 An exception is Esser and Mönch (2007) who allow for market impact of passive orders on prices, but not liquidity competition. Complementing previous work on the market impact of limit orders of Hautsch and Huang (2012) we show that limit order placements primarily affect the supply side of liquidity through an increase of liquidity competition.…

Show abstract

View all citing articles on Scopus

Recommended articles (6)

Research article

Effects of lit and dark market fragmentation on liquidity

Journal of Financial Markets, Volume 35, 2017, pp. 1-20

Show abstract

Research article

Trading under market impact: Crossing networks interacting with dealer markets

Journal of Economic Dynamics and Control, Volume 100, 2019, pp. 131-151

Show abstract

Research article

An examination of the NYSE's retail liquidity program

The Quarterly Review of Economics and Finance, Volume 80, 2021, pp. 367-373

Show abstract

Research article

The microstructure of a U.S. Treasury ECN: The BrokerTec platform

Journal of Financial Markets, Volume 40, 2018, pp. 2-22

Show abstract

Research article

Limited investor attention and biased reactions to information: Evidence from the COVID-19 pandemic

Journal of Financial Markets, Volume 62, 2023, Article 100757

Show abstract

Research article

Closing auctions: Nasdaq versus NYSE

Journal of Financial Economics, Volume 143, Issue 3, 2022, pp. 1120-1139

Show abstract

View full text