OPTIMAL EXECUTION HORIZON

David Easley[1]
dae3@cornell.edu

Marcos M. López de Prado[2]
marcos.lopezdeprado@tudor.com

Maureen O'Hara[3]
mo19@cornell.edu

v1.27 – August 2, 2012

ABSTRACT

Execution traders know that market impact greatly depends on whether their orders lean with or against the market. We introduce the OEH model, which incorporates this fact when determining the optimal trading horizon for an order, an input required by many sophisticated execution strategies. From a theoretical perspective, OEH explains why market participants may rationally "dump" their orders in an increasingly illiquid market. OEH is shown to perform better than participation rate schemes and VWAP strategies. We argue that trade side and order imbalance are key variables needed for modeling market impact functions, and their dismissal may be the reason behind the apparent disagreement in the literature regarding the functional form of the market impact function. Our backtests suggest that OEH contributes substantial "execution alpha" for a wide variety of futures contracts. An implementation of OEH is provided in Python language.

Keywords: Liquidity, flow toxicity, broker, VWAP, market microstructure, adverse selection, probability of informed trading, VPIN, OEH.

JEL codes: C02, D52, D53, G14, G23.

We thank Tudor Investment Corporation, Robert Almgren, Riccardo Rebonato, Myck Schwetz, Brian Hurst, Hitesh Mittal, David Bailey, David Leinweber, John Wu, the CIFT group at the Lawrence Berkeley National Laboratory and participants at the Workshop on Ultra-High Frequency Econometrics, Market Liquidity and Market Microstructure for providing useful comments. We are grateful to Sergey Kosyakov for research assistance.

* The authors have applied for a patent on 'VPIN' and have a financial interest in it.

[1] Scarborough Professor of Social Science, Department of Economics, Cornell University.
[2] Head of Global Quantitative Research & HFFT at Tudor Investment Corp.; Research Affiliate at CIFT, Lawrence Berkeley National Laboratory.
[3] Purcell Professor of Finance, Johnson Graduate School of Management, Cornell University.

1. INTRODUCTION

Optimal execution strategies compute a trajectory that minimizes the shortfall cost of acquiring or disposing of a position in an asset. Well-known contributions to this subject are Perold [1998], Bertsimas and Lo [1998], Almgren and Chriss [2000], and Kissell and Glantz [2003], to cite only a few. These strategies provide an abstract and general framework to model the costs that trading imposes on a particular investor. One drawback of this abstraction is that it does not explicitly define how market impact arises from a trade's perturbation of the liquidity provision process. This paper attempts to cover that gap in the market impact literature, offering an answer to some of its long-standing questions. For example, execution traders know that market impact greatly depends on whether their orders lean with or against the market. The execution strategy literature, however, has not yet incorporated the two variables associated with this phenomenon: order side and market imbalance.

The PIN theory (Easley et al. [1996]) shows that market makers adjust the range at which they are willing to provide liquidity based on their estimates of the probability of being adversely selected by informed traders. Easley, López de Prado and O'Hara [2012] show that, in high frequency markets, this probability can be accurately approximated as a function of the absolute order imbalance (absolute value of buy volume minus sell volume). Suppose that we are interested in selling a large amount of E-mini S&P500 futures in a market that is imbalanced towards sells. Because our order is leaning with previous orders, it reinforces market makers' fears that they are being adversely selected, and that their current (long) inventory will be harder to liquidate without incurring a loss. As a result, market makers will further widen the range at which they are willing to provide liquidity, increasing our order's market impact. Alternatively, if our order were on the buy side, market makers would narrow their trading range as they believe that their chances of turning over their inventory improve, in which case we would experience a lower market impact than with a sell order. Thus, order imbalance and market maker behavior set the stage for understanding how orders fare in terms of execution costs. Our goal in this paper is to apply this theory in the context of market impact and execution strategies.

Besides reducing transaction costs, there are a number of reasons why traders care about not increasing the order imbalance. First, Federal Regulation and Exchange Circulars limit a trader's ability to disrupt or manipulate market activity.[4] A trading strategy that disrupts the market can bring fines, restrictions and sanctions.[5] Second, traders often will have to come back to the market shortly after completing the initial trade. If a trader got great fills in the previous trade at the expense of eroding liquidity, the previous trade's gains may transform into losses on the successive trades.[6] Third, the position acquired will be marked-to-market, so post-trade liquidity conditions will be reflected in the unrealized P&L of the position.[7] Thus, it would be useful to

[4] We could envision a future in which regulators and exchanges limit the assets under management of an investment firm based on that market participant's technology as it relates to the disruption of the liquidity provision process.

[5] See, for example, the "stupid algo" rules recently introduced by the Deutsche Börse (discussed in *"Superfast traders feel heat as Bourses act"*, Financial Times, March 5, 2012). The CFTC and SEC have also recently introduced explicit rules relating to algorithmic impact on markets.

[6] Even if there are no successive trades, leaving a footprint leaks information that can be recovered by competitors and used against that trader in futures occasions. See Easley et al. [2012c] for examples.

[7] If a buyer pushes the mid-price up at the expense of draining liquidity, she may find that the liquidation value of that position implies a loss.

determine the amount of volume needed to "conceal" a trade so that it leaves a minimum footprint on the trading range.

Many optimal execution strategies make the assumption that liquidity cost per share traded is a linear function of trading rate or of the size of the block to be traded. This seems an unrealistic assumption and it is the motivation for Almgren's [2003] study of optimal execution. The model we introduce here also incorporates a nonlinear response in trading cost to size of the trade. We are not aware of other methodologies explicitly taking into account the presence of asymmetric information in determining the optimal execution horizon. In this context, we find that the side of the trade is as important as its size.

Our **O**ptimal **E**xecution **H**orizon model, which for brevity we denote OEH henceforth, is related to a growing number of recent studies concerned with execution in the context of high frequency markets and tactical liquidity provision. Schied and Schöneborn [2009] point out that the speed by which the remaining asset position is sold can be decreasing in the size of the position, but increasing in the liquidity price impact. Gatheral and Schied [2011] extend the Almgren-Chriss framework to an alternative choice of the risk criterion. Forsyth [2011] formulates the trading problem as an optimal stochastic control problem, where the objective is to maximize the mean-variance tradeoff as measured at the initial time. Bayraktar and Ludkovski [2011] propose an optimal trade execution scheme for dark pools.

OEH does not replace or supersede minimum market impact strategies. On the contrary, it is complementary to them. OEH does not address the question of how to slice the orders (create a "trading schedule"), which has been studied by Hasbrouck and Schwartz [1988], Berkowitz et al. [1988], Grinold and Kahn [1999], Konishi and Makimoto [2001] among others previously cited. Our main concern is with understanding how transaction costs are derived from the impact that a trade has on liquidity providers. There are some connections, however, between OEH and the previous studies. Like earlier models, OEH minimizes the impact on liquidity subject to a timing risk. From the standard VWAP (Madhavan [2002]) to sophisticated nonlinear approaches (Almgren [2003]; Dai and Zhong [2012]), many execution strategies require as an input the *trading horizon*, which is typically exogenous to those models. One of our contributions is to provide a framework for determining this critical input variable from a new perspective: *Informational leakage*. In doing so, we stress the importance of modeling the order side and its asymmetric impact on the order imbalance.

The current paper is organized as follows. Section 2 reviews how order imbalance and trading range are related. Section 3 explains why our trading actions leave a footprint on the market makers' trading range. Section 4 incorporates timing risk and risk aversion to our analysis. Section 5 develops the Optimal Execution Horizon (OEH) algorithm that determines the optimal horizon over which to execute a trade. Section 6 presents three numerical examples, illustrating several stylized cases. Section 7 argues why the apparently irrational behavior of some market participants during the "*flash crash*" may be explained by OEH. Section 8 compares OEH's performance with that of trading schemes that target a volume participation rate. Section 9 discusses how our model relates to alternative functional forms of the market impact function. Section 10 provides backtests of the performance of EOH, and compares it with a VWAP execution strategy. Section 11 summarizes our conclusions. Appendices 1 provides a proof.

Appendix 2 offers a generalization of our approach. Appendix 3 contains an implementation of Appendices 1 and 2 in Python language. Appendix 4 provides a procedure for estimating market imbalance.

2. TRADING RANGE AND FLOW TOXICITY

We begin by summarizing a standard sequential-trade market microstructure approach to determining the trading range for an asset. In a series of papers, Easley et al. [1992a, 1992b, 1996] demonstrate how a microstructure model can be estimated for individual assets using trade data to determine the probability of information-based trading, PIN. This microstructure model views trading as a game between liquidity providers and traders (position takers) that is repeated over trading periods $i=1,\ldots,I$. At the beginning of each period, nature chooses whether an information event occurs. These events occur independently with probability α. If the information is good news, then informed traders know that by the end of the trading period the asset will be worth $\overline{S}$; and, if the information is bad news, that it will be worth $\underline{S}$, with $\overline{S} > \underline{S}$. Good news occurs with probability ($1-\delta$) and bad news occurs with probability δ. After an information event occurs or does not occur, trading for the period begins with traders arriving according to Poisson processes. During periods with an information event, orders from informed traders arrive at rate μ. These informed traders buy if they have seen good news, and sell if they have seen bad news. Every period, orders from uninformed buyers and uninformed sellers each arrive at rate ε.

Easley, Kiefer, O'Hara and Paperman [1996] argue that, for the natural case that $\delta = \frac{1}{2}$, the trading range at which market makers are willing to provide liquidity is

$$\Sigma = \frac{\alpha\mu}{\alpha\mu + 2\varepsilon}\left[\overline{S} - \underline{S}\right] \tag{1}$$

This trading range gives the market maker's targeted profit per portfolio turnover. The first term in Eq. (1) is the probability that a trade is information-based. It is known as PIN. Easley, Engle, O'Hara and Wu [2008] show that expected absolute trade imbalance can be used to approximate the numerator of PIN. They demonstrate that, letting V^S and V^B represent sell and buy volume, respectively, $E[|V^S - V^B|] \approx \alpha\mu$ for a sufficiently large μ. Easley, López de Prado and O'Hara [2011a, 2012a] argue that in volume-time space, PIN can be approximated as

$$PIN \equiv \frac{\alpha\mu}{\alpha\mu + 2\varepsilon} \approx VPIN \equiv \frac{E[|V^B - V^S|]}{V} = E[|OI|] \tag{2}$$

where we have grouped trades in equal volume buckets of size $V^B + V^S = V = \alpha\mu + 2\varepsilon$, and $OI = \frac{V^B - V^S}{V}$ represents the order imbalance within V.[8] This volume-time approximation of PIN, known as VPIN, has been found to be useful in a number of settings (see Easley, López de Prado and O'Hara [2012a] or Bethel et al. [2011] for example). In the next section we will show how VPIN's expectations play a role in modeling the liquidity component of an execution strategy.

[8] For a procedure that can be used to estimate OI, see Easley, López de Prado and O'Hara [2012b].

3. THE LIQUIDITY COMPONENT

For the reasons argued in the introduction, traders are mindful of the footprint that their actions leave on the trading range, Σ. According to Eq. (2) the trading range will change due to adjustments on VPIN expectations. In the next two subsections we discuss the impact that our trade has on the order imbalance, how that may affect VPIN expectations and consequently the transaction cost's liquidity component.

3.1. IMPACTING THE ORDER IMBALANCE

Suppose that we are an aggressive trader for *m* contracts to be traded in the next volume bucket.[9] For now, we set the bucket size at an exogenously given *V* (we explain in Section 5 how to optimize this bucket size). Our trade of *m* contracts will cause the bucket to fill faster than it would otherwise fill. Let us suppose that our trade pushes buys and sells from other traders out of the next bucket at equal rates. Given our knowledge of OI, we let $\widetilde{V^B}$ and $\widetilde{V^S}$ be our forecasts of buy and sell volume that would occur in the next bucket without our trade.[10] The expected order imbalance over *V* using our private information about our trade of *m* contracts can be represented by

$$\begin{aligned}\widehat{OI} &\equiv \frac{\widehat{V^B}-\widehat{V^S}}{V} = \frac{\frac{\widetilde{V^B}}{V}(V-|m|)-\frac{\widetilde{V^S}}{V}(V-|m|)+m}{V} = \\ &= \frac{\widetilde{V^B}-\widetilde{V^B}\frac{|m|}{V}-\widetilde{V^S}+\widetilde{V^S}\frac{|m|}{V}+m}{V} = \frac{\left(\widetilde{V^B}-\widetilde{V^S}\right)\left(1-\frac{|m|}{V}\right)+m}{V} \\ &= \underbrace{(2v^B-1)}_{OI}\left(1-\frac{|m|}{V}\right)+\frac{m}{V}\end{aligned} \tag{3}$$

where $V \geq m$, $v^B \equiv \frac{\widetilde{V^B}}{V}$ represents the forecasted fraction of buy volume in absence of our trade and $OI \equiv 2v^B - 1$ is the order imbalance in absence of our trade.[11] In the nomenclature above, $\widehat{OI}$ incorporates private information and OI only public information. If $\frac{|m|}{V} \approx 0$, $\widehat{OI} \approx OI$. Alternatively, in the extreme case of $\frac{|m|}{V} \approx 1$, we have $\left|\widehat{OI}\right| \approx 1$. Generally, the impact of our trade on $\widehat{OI}$ will depend on *m*, but also the OI that would otherwise occur in the next volume bucket. If, for instance, $V^B < V^S$, a trade of size $m = \frac{\left(\widetilde{V^S}-\widetilde{V^B}\right)V}{\left(\widetilde{V^S}-\widetilde{V^B}\right)+V}$ will make $\widehat{OI} = 0$.

3.2. INFORMATIONAL LEAKAGE ON THE LIQUIDITY COMPONENT

[9] Of course, *m* must be smaller than or equal to *V* in order for us to do it within one volume bucket. The meaning of "aggressive trader" is that we determine the timing of the trade (as opposed to being a "passive trader").

[10] Easley et al. [2012a] present evidence that order imbalance shows persistence over (volume) time. Easley et al. [2012d] present a forecasting model for v^B.

[11] For a procedure that can be used to estimate v^B, see Easley, López de Prado and O'Hara [2012b].

The previous section explained how trading m contracts interspersed with $V - m$ external volume displaces the expected order imbalance from OI to $\widehat{OI}$. Next, we discuss how that displacement triggers an updating of the market makers' expectations on the order imbalance, and thus VPIN.

Market makers adjust their estimates of VPIN as a result of the information leaked during the trading process. *Ceteris paribus*, if $|m|$ is relatively small, we expect it to convey little information to market makers. For example, if a small trade is executed in block but it is however large enough to fill a small bucket of size *V*, market makers may expect VPIN to remain around forecasted levels rather than jump to 1. We model the expected order imbalance (leaked to market makers) during execution, $\widetilde{OI}$, as a convex combination of two extreme outcomes: No leakage ($\varphi[|m|] \to 0$) and complete ($\varphi[|m|] \to 1$) informational leakage.

$$\widetilde{OI} \equiv \varphi[|m|] \underbrace{\left[(2v^B - 1)\left(1 - \frac{|m|}{V}\right) + \frac{m}{V}\right]}_{\widehat{OI}} + (1 - \varphi[|m|]) \underbrace{(2v^B - 1)}_{OI} \tag{4}$$

where $\varphi[.]$ is a monotonic increasing function in $|m|$ with values in the range (0,1). $\widetilde{OI}$ contains private information to the extent that it has been leaked by *m*. That occurs as a result of the trade's size ($|m|$). The role of $\varphi[.]$ is to determine the degree by which the effective order imbalance during our execution ($\widehat{OI}$) impacts the market markers' expectations on VPIN, and consequently leads to an adjustment of the range at which they provide liquidity, Σ.

It is critical to understand that the privately known order imbalance ($\widehat{OI}$) may differ from the one inferred by the market makers ($\widetilde{OI}$). The reason is, some private information may have been leaked by *m*, but not all. If $\varphi[|m|] \to 0$, there is no informational leakage, and $\widetilde{OI} \approx OI$, regardless of $\widehat{OI}$. In this case, it is as if *m* had not been traded and the effect on VPIN will be imperceptible. However, when the order is so large that $\varphi[|m|] \to 1$, the leak is complete and the market maker knows as much as the trader, $\widetilde{OI} \approx \widehat{OI}$. For instance, an order of 75,000 contracts may dramatically displace the expectation of VPIN, even if blended among a volume 10 times greater (see SEC-CFTC [2010] in connection with the Waddell & Reed order). This displacement of the expectation from OI to $\widetilde{OI}$ is the *footprint* left by *m* in the liquidity provision process. Different footprint specifications could be considered, but that would not conceptually alter the analysis and conclusions presented here.

For simplicity, we have assumed that $E[|OI|] = |OI|$, because this expectation is based on past, public information. Similarly, $E\left[\left|\varphi[|m|]\widehat{OI}\right|\right] = \varphi[|m|]\left|\widehat{OI}\right|$ because $\varphi[|m|]\widehat{OI}$ is precisely the portion of $\widehat{OI}$ that has become public due to *m*'s leakage. In this particular model we have not considered other sources of uncertainty in forming VPIN's expectation, thus $\widehat{VPIN} \equiv E\left[\left|\widetilde{OI}\right|\right] = \left|\widetilde{OI}\right|$.

4. THE TIMING RISK COMPONENT

Minimizing the footprint that trading has on Σ is an important factor when deciding the execution horizon. But it is not the only factor that matters to a trader. Minimal impact may require slicing

a large order into multiple small trades, resulting in a long delay in executing the intended trade. This delay involves a risk, called *timing risk*. To model this timing risk, we introduce a simple, standard model of how the midpoint of the spread evolves. It is useful to view the trading range, Σ, as being centered on the mid-price, *S*, which moves stochastically as (volume) time passes. We model this process as an exogenous arithmetic random walk

$$\Delta S \equiv \hat{\sigma}\sqrt{\frac{V}{V_\sigma}}\xi \tag{5}$$

where $\xi \sim N(0,1)$ i.i.d. and V_σ is the volume used to compute each mid-price change. So V/V_σ is the number of price changes recorded in a bucket of size *V*. The standard deviation, σ, of price changes is unknown, but it can be estimated from a sample of *n* equal-volume buckets as $\hat{\sigma}$, in which case $\frac{(n-1)\hat{\sigma}^2}{\sigma^2} \sim \chi^2_{n-1}$.

According to this specification, the market loss from a trade of size m can be probabilistically bounded as

$$P\left[Sgn(m)\Delta S > Z_\lambda \hat{\sigma}\sqrt{\frac{V}{V_\sigma}}\right] = 1 - \lambda \tag{6}$$

where $Sgn(m)$ is the sign of the trade *m*, λ is the probability with which we are willing to accept a loss greater than $Z_\lambda \hat{\sigma}\sqrt{\frac{V}{V_\sigma}}$, and Z_λ is the critical value from a Standard Normal distribution associated with $\lambda \in \left(0, \frac{1}{2}\right]$. λ can also be interpreted as a *risk aversion* parameter, as it modulates the relative importance that we give to timing risk.

The specification above abstracts from any direct impact that *m* may have on the midpoint price process. As is well known from microstructure research, the private information in our trade can introduce permanent effects on prices. In Appendix 2, we provide a generalization of this component taking into account the possibility that *m* leaks part of our private information into the mid-price.

5. OPTIMAL EXECUTION HORIZON

We have argued that the impact of our trade of *m* on Σ will depend on the size of *m* relative to *V*. In this section we take the intended trade, *m*, as fixed and determine the optimal trading volume *V* in which to hide our trading action *m* without incurring excessive timing risk. In order to compute this quantity, we define a *probabilistic loss* Π which incorporates the liquidity and timing risk components discussed earlier. The probabilistic loss models how different execution horizons will impact our per unit portfolio valuation after the trade's completion This is not an implementation cost, as that is contingent upon the execution model adopted to slice the parent order into child orders. OEH's goal is to determine the V^* that optimally conceals (in the sense of

minimizing Π for a given risk aversion λ) our order m, given the prevalent market state ($\varphi[|m|]$, v^B, $[\overline{S}-\underline{S}]$, V_σ, $\hat{\sigma}$), where

$$\Pi \equiv \underbrace{\left|\varphi[|m|]\left[(2v^B-1)\left(1-\frac{|m|}{V}\right)+\frac{m}{V}\right]+(1-\varphi[|m|])(2v^B-1)\right|[\overline{S}-\underline{S}]}_{liquidity\ component} \underbrace{-Z_\lambda\sqrt{\frac{V}{V_\sigma}}\hat{\sigma}}_{timing\ risk\ component} \tag{7}$$

with $Z_\lambda \leq 0$, thus the subtraction. The greater V, the smaller the impact on the order imbalance, but also the larger the possible change in the center of the trading range. Appendix 1 demonstrates the solution to minimizing Π, which can be implemented through the following algorithm (see Appendix 3 for an implementation in Python language):

1. If $(2v^B-1)|m| < m$, try $V_1 = \left(2\varphi[|m|][(2v^B-1)|m|-m][\overline{S}-\underline{S}]\frac{\sqrt{V_\sigma}}{Z_\lambda\hat{\sigma}}\right)^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_1, $\widetilde{OI}[V_1]$.
    a. If $\widetilde{OI}[V_1] > 0$ and $v^B \leq v^{B+}$, then $V^* = V_1$ is the solution.
    b. If $\widetilde{OI}[V_1] > 0$ and $v^B > v^{B+}$, then $V^* = |m|$ is the solution.
2. If $(2v^B-1)|m| > m$, try $V_2 = \left(2\varphi[|m|][m-(2v^B-1)|m|][\overline{S}-\underline{S}]\frac{\sqrt{V_\sigma}}{Z_\lambda\hat{\sigma}}\right)^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_2, $\widetilde{OI}[V_2]$.
    a. If $\widetilde{OI}[V_2] < 0$ and $v^B \geq v^{B-}$, then $V^* = V_2$ is the solution.
    b. If $\widetilde{OI}[V_2] < 0$ and $v^B < v^{B-}$, then $V^* = |m|$ is the solution.
3. If $(2v^B-1)|m| = m$, then $V_3 = |m|$ is the solution.
4. Else, try $V_4 = \varphi[|m|]\left(|m| - \frac{m}{2v^B-1}\right)$.
    a. If $v^B \geq v^{B=}$, then $V^* = V_4$ is the solution.
    b. If $v^B < v^{B=}$, then $V^* = |m|$ is the solution.

where

$$\widetilde{OI}[V^*] = \varphi[|m|]\left[\frac{m-(2v^B-1)|m|}{V^*}+2v^B-1\right]+(1-\varphi[|m|])(2v^B-1) \tag{8}$$

$$v^{B+} = \frac{\frac{|m|}{m}+1}{2}+\frac{Z_\lambda\hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\overline{S}-\underline{S}]\sqrt{V_\sigma}} \tag{9}$$

$$v^{B-} = \frac{\frac{|m|}{m}+1}{2}-\frac{Z_\lambda\hat{\sigma}\sqrt{|m|}}{4\varphi[|m|][\overline{S}-\underline{S}]\sqrt{V_\sigma}} \tag{10}$$

$$v^{B=} = \frac{1}{2}\left(\frac{|m|\varphi[|m|]}{m(\varphi[|m|]-1)} + 1\right) \tag{11}$$

$\widetilde{OI} \in [-1,1]$ is the signed order imbalanced as a result of the informational leak that comes with trading m. v^{B-}, v^{B+} and $v^{B=}$ set the boundaries for v^B in order to meet the constraint that $V^* \geq |m|$. The liquidity component is nonlinear in m and takes into account the side of the trade (leaning against or with the market), which we illustrate with numerical examples in Section 6.

Note that V^* has been optimized to minimize Π, and it will generally differ from the V used to compute VPIN. We reiterate the point made earlier that various uses of VPIN require different calibration procedures. In this paper, we propose a method for determining the V^* that minimizes the probabilistic loss Π, not the V that maximizes our forecasting power of v^B or short-term toxicity-induced volatility.

In practice, V^* could be estimated once before submitting the first child order and re-estimated again before submitting every new slice. By doing so, we incorporate the market's feedback into the model: If our initial slices have a greater (or smaller) impact than expected on the liquidity provision, we could adjust in real-time. That is, we assume that the values observed will remain constant through the end of the liquidation period, and determine the statically optimal strategy using those values. As the input parameters change, we could recompute the stationary solution. Almgren's [2009] shows that this *"rolling horizon"* approach, although not dynamically optimal, provides reasonably good solutions to a related problem.

6. NUMERICAL EXAMPLES

In this section, we develop numerical examples using the following state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] \to 1$. We want to find the optimal horizon to buy our desired amount of 1,000, recognizing that the optimal strategy must take account of both liquidity costs and timing risk costs. This strategy will depend, in part, on the imbalance in the market, which we capture as the fraction of buy volume, v^B. Three scenarios appear relevant: $v^B < \frac{1}{2}$, $v^B = \frac{1}{2}$ and $v^B > \frac{1}{2}$. We limit our attention to $m > 0$, with the understanding that a symmetric outcome would be attained with a sell order $\widetilde{m} < 0$ and $\widetilde{v^B} = 1 - v^B$ (see Figure 5).

[FIGURE 1 HERE]

Figure 1(a) displays the optimal volume horizons (V^*) for various values of the fraction of buy volume, v^B. As is apparent, the optimal trading horizon differs dramatically with imbalance in the market. The reason for this is illustrated in Figure 1(b), which demonstrates how the probabilistic loss of trading 1,000 shares is also a function of the buy imbalance v^B. The probabilistic loss is the sum of the loss from the liquidity component and the loss from the timing

component. Notice that with $m > 0$, the liquidity component does not contribute to Π until shortly before the market is balanced ($v^B = \frac{1}{2}$).[12]

6.1. SCENARIO I: $v^B = 0.4$

In this scenario we want to buy 1,000 contracts from a market that we believe to be slightly tilted towards sales (projected buys are only 40% of the total volume). If we plot the values of $\Pi[V,.]$ for different levels of *V*, we obtain Figure 2.

[FIGURE 2 HERE]

In this case $\widetilde{OI}[V^*]$=*0* at the optimum trading horizon of $V^* = 6{,}000$ contracts (of which 5,000 come from other market participants). The optimum occurs at the inflexion point where the liquidity component function changes from convex decreasing to concave increasing. This may seem counter-intuitive as it is natural to expect the liquidity component to be decreasing in *V*. But this reasoning misses the important role played by market imbalance. Because we are buying in a selling market, there is a V^* at which we are narrowing Σ to the minimum possible and our liquidity cost is zero. Once we pass that liquidity-component optimal V^*, the trading range Σ necessarily widens again. The only way this can happen is with an increasing concave section in the liquidity component function, which explains the appearance of the inflexion point.

If we trade the desired 1,000 contracts within less than 6,000 total contracts traded our loss Π will increase because Σ increases more rapidly than the timing component of loss declines. If, on the other hand, we trade those 1,000 while more than 6,000 contracts are traded, our loss Π will increase because we will be taking on both excessive liquidity and timing costs.

One question to consider is, why does OEH allow the purchase to occur in the presence of selling flow which may result in lower future prices? Or put differently, why is not optimal to have a larger, possibly infinite execution horizon for a buy order, as long as $v^B < 0.5$? After all, prices may go lower as a result of the selling pressure, and that would give the trader a chance to buy at a better level. The reason is, this is an execution model, not an investment strategy. *The portfolio manager has decided that the trade must occur as soon as liquidity conditions allow it.* It is not the prerogative of OEH to speculate on the appropriateness of the trader's decision, which may be motivated for a variety of reasons (she holds private information, it is part of her asset management mandate, she must obey a stop loss or abide by a risk limit, a release is about to occur, etc.). The role of the OEH model is merely to determine the execution horizon that minimizes the informational leakage.

6.2. SCENARIO II: $v^B = 0.5$

In this scenario, the market is expected to be balanced (buys are 50% of the total volume). If we plot the values of $\Pi[V,.]$ for different levels of *V*, we obtain Figure 3.

[FIGURE 3 HERE]

[12] The liquidity component is positive at $v^B < \frac{1}{2}$ as the order m=1,000 makes the market unbalanced toward buys for $v^B < \frac{1}{2}$.

The optimum now occurs at $V^* = 11{,}392$ contracts, with a value for $\widetilde{OI}[V^*] = 0.088$. The model recognizes that a larger volume horizon is needed to place a buy order in an already balanced market than in a market leaning against our order (note the change in shape of the liquidity component). In this scenario, our order does not narrow Σ abruptly (as in Scenario I), and the only way to reduce Π is by substantially increasing V (i.e., disguising our order among greater market volume). The liquidity component function is now convex decreasing, without an inflexion point, because the market is not leaning against us. But the optimal V^* is still limited because of greater timing risk with increasing V.

6.3. SCENARIO III: $v^B = 0.6$

The market is now expected to be tilted toward buys, which represent 60% of the total volume. If we plot the values of $\Pi[V, .]$ for different levels of V, we obtain Figure 4.

[FIGURE 4 HERE]

The optimum now occurs at $V^* = 9{,}817$ contracts, with a value for $\widetilde{OI}[V^*] = 0.2816$. Two forces contribute to this outcome: On one hand, we are leaning with the market, which means that we are competing for liquidity ($v^B > \frac{1}{2}$), and we need a larger volume horizon than in Scenario I. On the other hand, the gains from narrowing Σ are offset by the additional timing risk, and Π eventually cannot be improved further. Note that this convex, decreasing liquidity component function asymptotically converges to the same level as the concave, increasing liquidity component function in Scenario I, due to the same absolute imbalance (VPIN) of both scenarios. The equilibrium between these two forces is reached at 9,817 contracts, a volume horizon between those obtained in Scenarios I and II.

7. EXECUTION HORIZONS AND THE OSCILLATORY NATURE OF PRICES

Figure 1 illustrates how traders behave as predicted toxicity increases. For buyers in a market tilted toward buys, the execution horizon is a decreasing function of v^B. For sellers in a market tilted toward sells, the execution horizon is also a decreasing function. Figure 5 depicts the optimal horizon for a sell order. Note that, in this case, v^S decreases as we move from left to right in the graph. For both buyers and sellers, therefore, their optimal trading horizon will be influenced by the toxicity expected in the market.

[FIGURE 5 HERE]

We next consider the combined impact of alternative trade sizes, sides and v^B. For simplicity, we assume that $\varphi[|m|]$ is linear in $|m|$, with $\varphi[|m|] \equiv \frac{m}{10^3}$ for $m \in (0, 10^3)$. The largest execution horizons occur in a perfectly balanced market, due to the inherent difficulty of concealing trading intentions. Again we observe in Figure 6 that *leaning against the market* (selling in a buying market or buying in a selling market) allows for shorter execution horizons, thanks to the possibility of achieving zero liquidity cost. *Leaning with the market* (selling in a selling market or buying in a buying market) leads to shorter execution horizons than those in a balanced market, but longer than when the trade leans against the market.

[FIGURE 6 HERE]

Now consider a recent event in which toxicity played a significant role. The CFTC-SEC examination of the "flash crash" (see CFTC-SEC [2010]) indicates that, during the crash, market participants dumped orders on the market. Why would anyone reduce their execution horizon in the midst of a liquidity crisis? Perhaps the most cited example is the Waddell & Reed order to sell 75,000 E-mini S&P500 contracts. It seems at first unreasonable to execute large orders in an increasingly illiquid market. The model introduced in this paper provides a possible explanation for this behavior, and illustrates how it contributes to the oscillatory nature of prices.[13]

Suppose that v^S increases to a level sufficient to have a measurable impact on PIN. This prompts sellers to shorten their execution horizons in volume-time (Scenario III, left side of Figure 5) because small gains from the liquidity component come at the expense of substantial timing risks. As several sellers compete for the same gradually scarcer liquidity, an increase in the volume traded per unit of time will likely occur, which is consistent with the observation that market imbalance accelerates the rate of trading (i.e., greater volume occurs during liquidity crises, like the "flash crash"). Prices are then pushed to lower levels, at which point buyers will have even shorter execution horizons than the sellers (Scenario I, left side of Figure 2). This is caused by the convex increasing section of the optimum for buyers (Figure 1(a)), compared to the concave decreasing section of the optimum for sellers (Figure 5). The increasing activity of buyers causes prices to recover some of the lost ground, v^S returns to normal levels, and execution horizons expand (Scenario II). The outcome is an oscillatory price behavior induced solely by timing reactions of traders to an initial market imbalance.[14]

8. VOLUME PARTICIPATION STRATEGIES

Many money managers conceal large orders among market volume by targeting a certain participation rate. Part of OEH's contribution is to show that, to determine an optimal participation rate, order side and order imbalance should also be taken into account. In this section, we will illustrate how a typical volume participation scheme performs in relative terms to an optimal execution horizon strategy.

Suppose we apply the same parameter values used in Section 6: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$. For simplicity, consider a $\varphi[|m|]$ linear in $|m|$, with $\varphi[|m|] \equiv \frac{m}{10^4}$ for $m \in (0, 10^4)$. Figure 7 compares the probabilistic loss from OEH and a scheme that participates in 5% of the volume, for various buy order sizes when $v^B = \frac{1}{2}$ (i.e. when the markets is balanced between buying and selling). What is apparent is that volume participation results in a far mostly costly trading outcome. For a desired trade size of 2000 contracts, the Volume Participation probabilistic loss is almost 50% greater than that of the OEH algorithm.

[13] We are not claiming that the Waddell & Reed order was executed optimally. Our argument instead is that trading faster in toxic markets is not necessarily irrational.

[14] Note that in this discussion we are taking the existence of buyers and sellers as given. The effects discussed in the text arise only from optimal trading horizons. If the impacts on prices also induce changes in desired trades, the effects on market conditions may be exacerbated or moderated.

[FIGURE 7 HERE]

The divergent behavior of these algorithms is also influenced by market imbalance. Figures 8 and 9 present the equivalent results for $v^B = 0.4$ and $v^B = 0.6$ respectively. OEH's outperformance is particularly noticeable in those cases when the order leans against the market (i.e. buying in a selling market).

[FIGURE 8 HERE]

[FIGURE 9 HERE]

9. THE SQUARE ROOT RULE

Loeb [1983] was the first to present empirical evidence that market impact is a square root function of the order size. Grinold and Kahn [1999] justified this observation through an inventory risk model: Given a proposed trade size $|m|$, the estimated time before a sufficient number of trades appears in the market to clear out the liquidity supplier's net inventory is a linear function of $\frac{|m|}{\bar{V}}$, where $\bar{V}$ is the average daily volume. Because prices are assumed to follow an arithmetic random walk, the transaction cost incorporates an inventory risk term of the form $\hat{\sigma}\sqrt{\frac{|m|}{\bar{V}}}$.

Over the last three decades, studies have variably argued in favor of linear (Breen, Hodrick and Korajczyk [2002], Kissell and Glantz [2003]), square-root (Barra [1997]) or power-law (Lillo, Farmer and Mantegna [2003]) market impact functions. A possible explanation for these discrepancies is that these studies did not control for two critical variables affecting transaction costs: order side and its relation to order imbalance. Consider once again the same parameter values used in Section 6: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$. Suppose that $\varphi[|m|]$ is linear in $|m|$, with $\varphi[|m|] \equiv \frac{m}{10^4}$ for $m \in (0, 10^4)$.

Figure 10 plots the probabilistic loss Π that results from executing a buy order of size m at optimal horizons $V^*(m)$ given $v^B = \frac{1}{2}$. A power function fits Π almost perfectly, with a power coefficient very close to the $\frac{3}{5}$ reported by Almgren, Thum, Hauptmann and Li [2005]. However, Figure 11 shows that, if $v^B = 0.4$, Π has a linear end and is below the levels predicted by the square root. Finally, when $v^B = 0.6$, Figure 12 displays Π values greater than in the other cases, as the order now competes for liquidity. These conclusions depend on our assumption of a linear $\varphi[|m|]$, but they can be generalized to other functional forms. For example, Figure 13 shows that, for $\varphi[|m|] \propto \sqrt{m}$ and $v^B = \frac{1}{2}$, Π fits the square root perfectly, as originally reported by Loeb [1983].

[FIGURE 10 HERE]

[FIGURE 11 HERE]

[FIGURE 12 HERE]

[FIGURE 13 HERE]

10. EMPIRICAL ANALYSIS

In the previous sections we presented a theory showing that OEH optimally uses order imbalance to determine the amount of volume needed to disguise a trade. We have also shown that OEH provides an explanation for the apparently contradictive views on the functional form of the market impact function. In this section we study the empirical performance of OEH relative to a VWAP strategy.

We consider an informed trader who wishes to trade m units of a particular futures contract. To keep the discussion as general as possible, we will discuss two scenarios. In the first scenario, the trader has information about the sign of the price change over the next volume bucket, but not its magnitude. In this case, we model the desired trade by $m = Sgn(\Delta P_\tau)q$, where q is a standard trade size determined by the trader. In the second scenario, the trader has information about the sign as well as the magnitude of the price move, and in this case we model the desired trade by $m = f(\Delta P_\tau)q$. As in Section 5, we can analyze the performance of an execution strategy in terms of two components: Liquidity and timing costs. The *timing* cost can be directly observed, as measured by the price change with respect to the fill price, excluding the impact that the order has on *liquidity*. For example, in the first scenario the trader may know that prices will go up, but she may not know the magnitude of the price increase, so for a sufficiently large m the expected timing gain may turn into a loss due to the liquidity cost.

When the trader executes using OEH, she combines her imperfect information with market estimates of σ and v^B over the next bucket. In this particular exercise, the latter is based on the procedure described in Appendix 4, and the former is the standard deviation of price changes over a given sample length. Then, for some $\varphi[|m|]$, λ, and a long-run volatility $\bar{\sigma}$, assuming $[\overline{S} - \underline{S}] = -Z_\lambda \bar{\sigma}$,[15] she can compute the optimal execution horizon V^* over which to trade m contracts. Let's denote the average fill price over the horizon V^*, in absence of her trade, by $\bar{P}_{OEH}$. The realized timing component is $(P_\tau - \bar{P}_{OEH,\tau})m$. The liquidity component can be estimated as $|\widetilde{OI}_{OEH,\tau}|[\overline{S} - \underline{S}]|m|$, following Eqs. (4) and (7). Thus, the total profit during volume bucket τ is

$$PL_{OEH,\tau} = \underbrace{-|\widetilde{OI}_{OEH,\tau}|[\overline{S} - \underline{S}]|m|}_{PL^L_{OEH,\tau}} + \underbrace{(P_\tau - \bar{P}_{OEH,\tau})m}_{PL^T_{OEH,\tau}} \tag{12}$$

[15] Other specifications could be considered to model the maximum trading range at which market makers are willing to provide liquidity. Here, we express that number as a function of the long term volatility, multiplied by a market makers' risk aversion factor. That factor may not coincide with Z_λ, however we see an advantage in keeping the model as parsimonious as possible, rather than introducing a new variable.

where we have expressed the profit in terms of its liquidity ($PL^L_{OEH,\tau}$) and timing ($PL^T_{OEH,\tau}$) components. Similarly, if the trader executes through a VWAP with fixed horizon, the total profit during volume bucket τ is

$$PL_{VWAP,\tau} = \underbrace{-\left|\widetilde{OI}_{VWAP,\tau}\right|\left[\overline{S} - \underline{S}\right]|m|}_{PL^L_{VWAP,\tau}} + \underbrace{\left(P_\tau - \bar{P}_{VWAP,\tau}\right)m}_{PL^T_{VWAP,\tau}} \tag{13}$$

Based on the previous equations, we can compute the relative outperformance of OEH over VWAP in terms of its information ratio,

$$IR \quad = \frac{E\left[PL_{OEH,\tau} - PL_{VWAP,\tau}\right]}{\sigma\left[PL_{OEH,\tau} - PL_{VWAP,\tau}\right]}\sqrt{n} \tag{14}$$

where $\sqrt{n}$ is the annualization factor, and *n* the number of independent trades per year.

Our goal is to evaluate the performance of OEH relative to a VWAP benchmark, for a trader that needs to execute a large order on a daily basis (*n=260*). To compute this performance, we estimate v^B using the methodology discussed in Easley, López de Prado and O'Hara [2012b], on one volume bucket per day, where each volume bucket is composed of 25 volume bars. σ_τ is estimated for each bucket looking back enough trades to have a combined volume of 5 times the ADV (an average of one week of trading activity). We then compute the optimal execution horizons for trades of size $q = \left\{\frac{ADV}{100}, \frac{ADV}{20}, \frac{ADV}{10}\right\}$, with a risk aversion $\lambda = 0.05$ and an informational leakage function $\varphi[|m|] = min\left(\sqrt{\frac{|m|}{ADV}}, .999\right)$, where ADV is the average daily volume.[16]

[TABLE 1 HERE]

Table 1 summarizes the data used in our calculations. We have selected these products because they encompass a wide variety of asset classes and liquidity conditions. For example, E-Mini S&P500 Futures are traded in the Chicago Mercantile Exchange, it is an equity index product, our sample contains 476,676,009 transactions recorded between January 1st 2007 and July 26th 2012, rolls occur 12 days prior to expiration date, and the ADV for that period has been 1,964,844.89 contracts.

[TABLES 2-4 HERE]

Tables 2, 3 and 4 report the outperformance of OEH over VWAP for trading sizes equivalent to 1%, 5% and 10% of ADV respectively, when the trader has information about the *sign* of the price move over the next volume bucket (not the size of the move). To interpret these results, suppose that a trader has information that the price of the E-Mini S&P500 futures will increase

[16] These are rather general values, and alternative ones could be adopted, depending on the user's specific objective. For example, λ could have been calibrated in order to maximize OEH's risk-adjusted performance over VWAP.

over the next volume bucket, and for that reason she wishes to acquire a position equivalent to 1% of E-Mini S&P500 futures' ADV, i.e. buy 19,648 contracts. We compute the liquidity and timing components for that trade using VWAP, and evaluate by how much OEH beats VWAP on that same trade in dollar terms. After repeating that calculation for each of the 1,450 volume buckets, we can estimate OEH's performance over VWAP. For instance, Table 2 reports that the maximum profit attainable with this information is 12.1428 index points on average, if execution were instantaneous and costless.[17] Out of that, OEH was able to capture 10.5104 index points on average, or 86.56% of the average maximum profit. This represents an outperformance of 4.3262 index points over VWAP's performance on the same trade, or 35.63% of the average maximum profit, with an information ratio of 10.04. As expected, OEH's IR edge over VWAP decays as we approach very large daily trades. For example, the information ratio associated with an OEH trade for 196,484 E-Mini S&P500 futures contracts (10% of its ADV) is 2.59. This occurs because for trades of that extremely large size, it becomes increasingly difficult to conceal the trader's intentions. Although the information ratio is smaller, the average dollar amount of savings is greater (0.9577 points per contract), because the trade is for a size 10 times larger.

[TABLES 5-7 HERE]

Tables 5, 6 and 7 present the results of applying our methodology when the trader has information about the *side and size* of the price change, for an amount equivalent to 1%, 5% and 10% of ADV respectively. OEH's edge over VWAP also decays as we approach very large daily trades, however the extent of this performance decay is much less pronounced when the trader holds size as well as sign information with regards to price changes.

11. CONCLUSIONS

The choice of execution horizon is a critical input required by many optimal execution strategies. These strategies attempt to minimize the trading costs associated with a particular order. They do not typically address the footprint that those actions leave in the liquidity provision process. In particular, most execution models do not incorporate information regarding the order's side, without which it is not possible to understand the asymmetric impact that the order will have on the liquidity provision process.

In this paper, we introduce the OEH model, which builds on asymmetric information market microstructure theory to determine the optimal execution horizon. OEH allows existing optimal execution models to minimize both an order's trading costs and its footprint in the market. In a high frequency world, this latter ability takes on increased importance. OEH is shown to perform better than schemes that target a participation rate. Our model also provides an explanation of the apparent disagreement in the literature regarding the functional form of the market impact function as a result of not controlling for the order side and its relation to the order imbalance. Overall, our analysis suggests a new way to trade in the high frequency environment that characterizes current market structure.

Our empirical study shows that EOH allows traders to achieve greater profits on their information, as compared to VWAP. If the trader's information is right, OEH will allow her to

[17] The value of an index point is USD50 in the case of the E-Mini S&P500 futures.

capture greater profits on that trade. If her information is inaccurate, OEH will deliver smaller losses than VWAP. OEH is not an investment strategy on its own, but delivers substantial "execution alpha" by boosting the performance of "investment alpha".

APPENDIX

A.1. COMPUTATION OF THE OPTIMAL EXECUTION HORIZON

We need to solve the optimization problem

$$\begin{aligned}\min_{V} \quad & \Pi\left[V, v^B, m, \left[\overline{S} - \underline{S}\right], \lambda, \hat{\sigma}, V_\sigma\right] \\ & \equiv \left|\varphi[|m|]\left[(2v^B - 1)\left(1 - \frac{|m|}{V}\right) + \frac{m}{V}\right]\right. \\ & \left. + (1 - \varphi[|m|])(2v^B - 1)\right| \left[\overline{S} - \underline{S}\right] - Z_\lambda \sqrt{\frac{V}{V_\sigma}}\hat{\sigma}\end{aligned} \tag{15}$$

subject to $V \geq |m|$.

Let's denote

$$\begin{aligned}\widetilde{OI} &\equiv \varphi[|m|]\left[(2v^B - 1)\left(1 - \frac{|m|}{V}\right) + \frac{m}{V}\right] + \underbrace{(1 - \varphi[|m|])(2v^B - 1)}_{c} \\ &= \varphi[|m|]\left[\frac{m - (2v^B - 1)|m|}{V} + 2v^B - 1\right] + c\end{aligned} \tag{16}$$

Observe that, if $V \geq |m|$, then $-1 \leq \widetilde{OI} \leq 1$. The objective function contains the absolute value of $\widetilde{OI}$ which depends on the choice variable V and it is not continuously differentiable around $\widetilde{OI} = 0$. Accordingly, we solve the problem in cases based on the sign of $\widetilde{OI}$ at an optimum.

A.1.1. CASE 1: Suppose that $\widetilde{OI} > 0$. Then

$$\frac{\partial\left|\widetilde{OI}\right|}{\partial V} = \varphi[|m|]\frac{(2v^B - 1)|m| - m}{V^2}, \text{ and} \tag{17}$$

$$\frac{\partial \Pi(V)}{\partial V} = \varphi[|m|]\frac{(2v^B - 1)|m| - m}{V^2}\left[\overline{S} - \underline{S}\right] - \frac{Z_\lambda \hat{\sigma}}{2\sqrt{V_\sigma V}}. \tag{18}$$

Note that the second term in (18) is positive, and so an interior solution can occur only if the first term in (18) is negative. We first suppose that the solution to the problem does not occur at the constraint $V \geq |m|$.

Using the first order necessary condition for an interior solution that $\frac{\partial \Pi(V)}{\partial V} = 0$ and multiplying by $\sqrt{V}$, we obtain

$$\varphi[|m|][(2v^B - 1)|m| - m]\left[\overline{S} - \underline{S}\right]V^{-3/2} = \frac{Z_\lambda \hat{\sigma}}{2\sqrt{V_\sigma}}. \tag{19}$$

Thus,

$$V^* = \left(2\varphi[|m|][(2v^B - 1)|m| - m]\left[\overline{S} - \underline{S}\right]\frac{\sqrt{V_\sigma}}{Z_\lambda \hat{\sigma}}\right)^{2/3}. \tag{20}$$

The V^* given by Eq. (17) is a candidate for an interior solution if $\widetilde{OI} > 0$ evaluated at V^* and $V^* \geq |m|$. It is straightforward to show that $V^* \geq |m|$ if $v^B \leq v^{B+}$ where

$$v^{B+} = \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|]\left[\overline{S} - \underline{S}\right]\sqrt{V_\sigma}} + \frac{\frac{|m|}{m} + 1}{2}. \tag{21}$$

Alternatively, if $v^B > v^{B+}$ then the candidate for a solution in this region is $|m|$. Note that the condition $v^B \leq v^{B+}$ implies that m is negative which in turn implies that the first term in (18) is positive and so the V^* given by Eq. (17) is well defined.

In this region the second derivative of $\Pi(V)$ is

$$\frac{\partial^2 \Pi(V)}{\partial V^2} = -2\varphi[|m|]\frac{(2v^B - 1)|m| - m}{V^3}\left[\overline{S} - \underline{S}\right] + \frac{Z_\lambda \hat{\sigma}}{4V\sqrt{V_\sigma V}}. \tag{22}$$

It can be shown that $\frac{\partial^2 \Pi(V)}{\partial V^2} > 0$ if $v^B \leq v^{B+}$.

A.1.2. CASE 2: Suppose that $\widetilde{OI} < 0$. Then

$$\frac{\partial\left|\widetilde{OI}\right|}{\partial V} = \varphi[|m|]\frac{m - (2v^B - 1)|m|}{V^2}. \tag{23}$$

Using an argument similar to that in Case 1, we see that in this region the candidate for an interior solution is

$$V^* = \left(2\varphi[|m|][m - (2v^B - 1)|m|]\left[\overline{S} - \underline{S}\right]\frac{\sqrt{V_\sigma}}{Z_\lambda \hat{\sigma}}\right)^{2/3}. \tag{24}$$

The V^* given by Eq. (21) is a candidate for a solution if $\widetilde{OI} < 0$ evaluated at V^* and $V^* \geq |m|$. Note that $V^* \geq |m|$ if $v^B \geq v^{B-}$ where

$$v^{B-} = \frac{\frac{|m|}{m} + 1}{2} - \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|]\left[\overline{S} - \underline{S}\right]\sqrt{V_\sigma}}. \tag{25}$$

Note that the condition $v^B \geq v^{B-}$ implies that m is negative which in turn implies that the V^* given by Eq. (24) is well defined.

Alternatively, if $v^B < v^{B-}$ then the candidate for a solution in this region is $|m|$.

In this region the second derivative of $\Pi(V)$ is

$$\frac{\partial^2 \Pi(V)}{\partial V^2} = -2\varphi[|m|]\frac{m-(2v^B-1)|m|}{V^3}\left[\overline{S}-\underline{S}\right] + \frac{Z_\lambda \hat{\sigma}}{4V\sqrt{V_\sigma V}}. \tag{26}$$

It can be shown that $\frac{\partial^2 \Pi(V)}{\partial V^2} > 0$ if $v^B \geq v^{B-}$.

A.1.3. CASE 3: Suppose that $\widetilde{OI} = 0$. In this case $\frac{\partial \Pi(V)}{\partial V}$ does not exist**,** and the optimum cannot be computed using calculus. However, we can still compute V^*, because in this case

$$\varphi[|m|]\left[\frac{m-(2v^B-1)|m|}{V} + (2v^B-1)\right] + (1-\varphi[|m|])(2v^B-1) = 0, \tag{27}$$

and thus

$$V^* = \varphi[|m|]\left(|m| - \frac{m}{2v^B-1}\right). \tag{28}$$

It is straightforward to show that $V^* \geq |m|$ if $v^B \geq v^{B=}$ where

$$v^{B=} = \frac{1}{2}\left(\frac{|m|\varphi[|m|]}{m(\varphi[|m|]-1)} + 1\right). \tag{29}$$

A.2. INFORMATIONAL LEAKAGE ON THE MID-PRICE

The literature devoted to optimal execution typically differentiates between the temporary and the permanent impact that an order *m* has on mid-prices (e.g., Stoll [1989]). The temporary impact arises from the actual order slicing (also called *trajectory*) within the execution horizon, while the permanent impact is a result of the information leaked by the order flow. Because our model is not concerned with the trajectory computation, we do not take into consideration the temporary impact of the actual order slicing. With regard to the permanent impact, Almgren and Chriss [2000] and Almgren [2003] propose a linear permanent impact function that depends on $|m|$.

Let us suppose that, as in the above referred models, the permanent impact function depends on $|m|$, but that unlike those models it is also a function of our volume participation rate, $\frac{|m|}{V^*}$, and whether we are leaning with or against the market. As we have noted, this is not the approach chosen by most execution strategies, but we believe it is worth considering because for large enough trades our orders will have a noticeable impact. For example, in the extreme case that

$\frac{|m|}{V^*} \approx 1$, and a sufficiently large $|m|$, our parent trade will not go unnoticed, no matter how well we slice it (especially if we are leaning with the market). We would expect leakage to be greatest when informed traders dominate market activity over a considerable horizon. Under this conjecture, we could conceive a permanent impact function that indeed depends on factors other than $|m|$.

In fact, the specification presented earlier in the paper already modeled an informational leakage process. In the text, we only considered the effect on the trading range, but there is no reason why we could not assume that the same way our parent trade leaves a footprint in the liquidity component, it also leaves a footprint in the mid-price. The probabilistic loss function could then take the form

$$\Pi \equiv \left|\varphi[|m|]\left[(2v^B-1)\left(1-\frac{|m|}{V}\right)+\frac{m}{V}\right]+(1-\varphi[|m|])(2v^B-1)\right|\left(\left[\overline{S}-\underline{S}\right] + |m|k\right) - Z_\lambda \hat{\sigma}\sqrt{\frac{V}{V_\sigma}} \tag{30}$$

for $k \geq 0$. If $k = 0$, the probabilistic loss reduces to our original specification. If $k > 0$, there is a permanent impact on prices that is linear in the informational leakage and the parent order's size. In this way we are taking into account not only the size of the order, but also its side, as we should expect a greater permanent impact when we compete with the market for scarce liquidity. The optimal execution horizon can then be computed with the following algorithm:

1. If $(2v^B-1)|m| < m$, try $V_1 = \left(2\varphi[|m|][(2v^B-1)|m|-m]\left[\overline{S}-\underline{S}+|m|k\right]\frac{\sqrt{V_\sigma}}{Z_\lambda\hat{\sigma}}\right)^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_1, $\widetilde{OI}[V_1]$.
    a. If $\widetilde{OI}[V_1] > 0$ and $v^B \leq v^{B+}$, then $V^* = V_1$ is the solution.
    b. If $\widetilde{OI}[V_1] > 0$ and $v^B > v^{B+}$, then $V^* = |m|$ is the solution.
2. If $(2v^B-1)|m| > m$, try $V_2 = \left(2\varphi[|m|][m-(2v^B-1)|m|]\left[\overline{S}-\underline{S}+|m|k\right]\frac{\sqrt{V_\sigma}}{Z_\lambda\hat{\sigma}}\right)^{2/3}$ and compute the value of $\widetilde{OI}$ associated with V_2, $\widetilde{OI}[V_2]$.
    a. If $\widetilde{OI}[V_2] < 0$ and $v^B \geq v^{B-}$, then $V^* = V_2$ is the solution.
    b. If $\widetilde{OI}[V_2] < 0$ and $v^B < v^{B-}$, then $V^* = |m|$ is the solution.
3. If $(2v^B-1)|m| = m$, then $V_3 = |m|$ is the solution.
4. Else, try $V_4 = \varphi[|m|]\left(|m| - \frac{m}{2v^B-1}\right)$.
    a. If $v^B \geq v^{B=}$, then $V^* = V_4$ is the solution.
    b. If $v^B < v^{B=}$, then $V^* = |m|$ is the solution.

with

$$\widetilde{OI}[V^*] \equiv \varphi[|m|]\left[\frac{m-(2v^B-1)|m|}{V^*}+2v^B-1\right]+(1-\varphi[|m|])(2v^B-1) \tag{31}$$

$$v^{B+} = \frac{\frac{|m|}{m}+1}{2} + \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|]\left[\overline{S} - \underline{S} + |m|k\right]\sqrt{V_\sigma}} \tag{32}$$

$$v^{B-} = \frac{\frac{|m|}{m}+1}{2} - \frac{Z_\lambda \hat{\sigma}\sqrt{|m|}}{4\varphi[|m|]\left[\overline{S} - \underline{S} + |m|k\right]\sqrt{V_\sigma}} \tag{33}$$

$$v^{B=} = \frac{1}{2}\left(\frac{|m|\varphi[|m|]}{m(\varphi[|m|]-1)} + 1\right) \tag{34}$$

A.3. ALGORITHM IMPLEMENTATION

The following is an implementation in Python of the algorithm described in this paper. Set the values given in the "Parameters" section according to your particular problem. The *k* parameter is optional. If *k* is not provided (or given a value *k=0*), the procedure described in Appendix 1 is followed. Otherwise, the specification discussed in Appendix 2 is applied. If you run this code with the parameters indicated below, you should get $V^* = 6{,}000$ (the solution in Section 6.1).

```
#!/usr/bin/env python
# By MLdP on 20120316 <ml863@cornell.edu>
# It computes the Optimal Execution Horizon

#-------------------------------------------------------
# PARAMETERS
sigma=1000
volSigma=10000
S_S=10000
zLambda=-1.644853627 #CDF(0.05) from the Std Normal dist
vB=0.4
phi=1
m=1000

#-------------------------------------------------------
def signum(int):
    if(int < 0):return -1
    elif(int > 0):return 1
    else:return 0

#-------------------------------------------------------
def getOI(v,m,phi,vB,sigma,volSigma):
    return phi*(float(m-(2*vB-1)*abs(m))/v+2*vB-1)+ \
        (1-phi)*(2*vB-1)

#-------------------------------------------------------
def getBounds(m,phi,vB,sigma,volSigma,S_S,zLambda,k=0):
    vB_l=float(signum(m)+1)/2-zLambda*sigma*abs(m)**0.5/ \
        float(4*phi*(S_S+abs(m)*k)*volSigma**0.5)
    vB_u=float(signum(m)+1)/2+zLambda*sigma*abs(m)**0.5/ \
```

```
        float(4*phi*(S_S+abs(m)*k)*volSigma**0.5)
    vB_z=(signum(m)*phi/float(phi-1)+1)/2.
    return vB_l,vB_u,vB_z

#-------------------------------------------------------
def minFoot(m,phi,vB,sigma,volSigma,S_S,zLambda,k=0):
    # compute vB boundaries:
    vB_l,vB_u,vB_z=getBounds(m,phi,vB,sigma,volSigma,S_S, \
        zLambda,k)

    # try alternatives
    if (2*vB-1)*abs(m)<m:
        v1=(2*phi*((2*vB-1)*abs(m)-m)*S_S*volSigma**0.5/ \
            float(zLambda*sigma))**(2./3)
        oi=getOI(v1,m,phi,vB,sigma,volSigma)
        if oi>0:
            if vB<=vB_u: return v1
            if vB>vB_u: return abs(m)
    elif (2*vB-1)*abs(m)>m:
        v2=(2*phi*(m-(2*vB-1)*abs(m))*S_S*volSigma**0.5/ \
            float(zLambda*sigma))**(2./3)
        oi=getOI(v2,m,phi,vB,sigma,volSigma)
        if oi<0:
            if vB>=vB_l: return v2
            if vB<vB_l: return abs(m)
    elif (2*vB-1)*abs(m)==m: return abs(m)
    if vB>=vB_z: return phi*(abs(m)-m/float(2*vB-1))
    if vB<vB_z: return abs(m)
 #-------------------------------------------------------
def main():
    print minFoot(m,phi,vB,sigma,volSigma,S_S,zLambda)

#-------------------------------------------------------
if __name__ == '__main__': main()
```

A.4. EXPECTED ORDER IMBALANCE

Given a sample of L buckets, we fit a forecasting regression model of the form

$$Ln(v_{\tau+1}^B) = \beta_0 + \beta_1 Ln(v_\tau^B) + \varepsilon_\tau \tag{35}$$

and because $v_\tau^B \in [0,1]$ we must obtain that $|\beta_1| < 1$.[18] The expected value of the order imbalance at τ over $\tau + 1$ is

$$E_\tau[OI_{\tau+1}] = 2e^{E_\tau[Ln(v_{\tau+1}^B)]} - 1 \tag{36}$$

where

[18] When fitting this regression, observations pairs $(v_{\tau+1}^B, v_\tau^B)$ where either $v_{\tau+1}^B = 0$ or $v_\tau^B = 0$ are removed.

$$E_\tau[Ln(v_{\tau+1}^B)] = \hat{\beta}_0 + \hat{\beta}_1 Ln(v_\tau^B) \tag{37}$$

In general we would expect roughly balanced markets, and so $v_\tau^B = \frac{1}{2} \Longrightarrow E_\tau[v_{\tau+1}^B] = \frac{1}{2}$. This means that the regression should cross through the equilibrium point $(Ln(v_\tau^B), E_\tau[Ln(v_{\tau+1}^B)]) = \left(Ln\left(\frac{1}{2}\right), Ln\left(\frac{1}{2}\right)\right)$, rather than having an equilibrium is systemically disrupted by some arbitrary intercept value $\hat{\beta}_0$. We impose that condition as

$$Ln(v_{\tau+1}^B) - Ln\left(\frac{1}{2}\right) = \beta_1^* \left[Ln(v_\tau^B) - Ln\left(\frac{1}{2}\right)\right] + \varepsilon_\tau \tag{38}$$

and we use,

$$Ln(v_{\tau+1}^B) = \beta_0^* + \beta_1^* Ln(v_\tau^B) + \varepsilon_\tau \tag{39}$$

where $\hat{\beta}_1^* = \frac{\sum_{l=1}^{L}\left(Ln(v_{\tau+1}^B)Ln(v_\tau^B) + \left[Ln\left(\frac{1}{2}\right)\right]^2 - Ln\left(\frac{1}{2}\right)\left(Ln(v_{\tau+1}^B) + Ln(v_\tau^B)\right)\right)}{\sum_{l=1}^{L}\left(\left[Ln(v_\tau^B)\right]^2 + \left[Ln\left(\frac{1}{2}\right)\right]^2 - 2Ln\left(\frac{1}{2}\right)Ln(v_\tau^B)\right)}$ and $\hat{\beta}_0^* = Ln\left(\frac{1}{2}\right)\left(1 - \hat{\beta}_1^*\right)$.

FIGURES

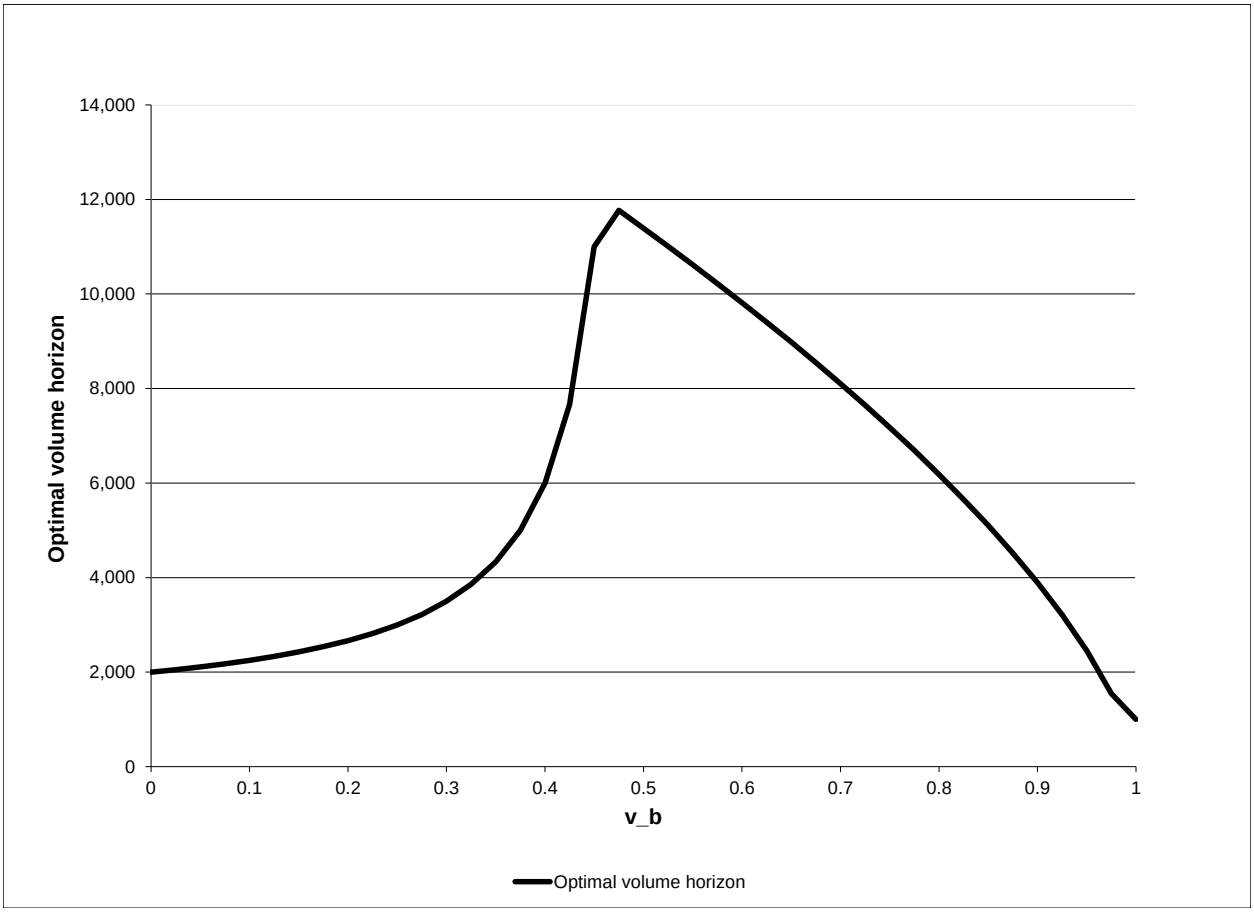

Figure 1(a) – Optimal V for various v^B on a buy order

This Figure demonstrates how the optimal trading horizon for a buy order depends upon the expected fraction of buy orders in the market. When all orders are buys, v^B is 1, while if all orders are sells then v^B is 0. The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\widehat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

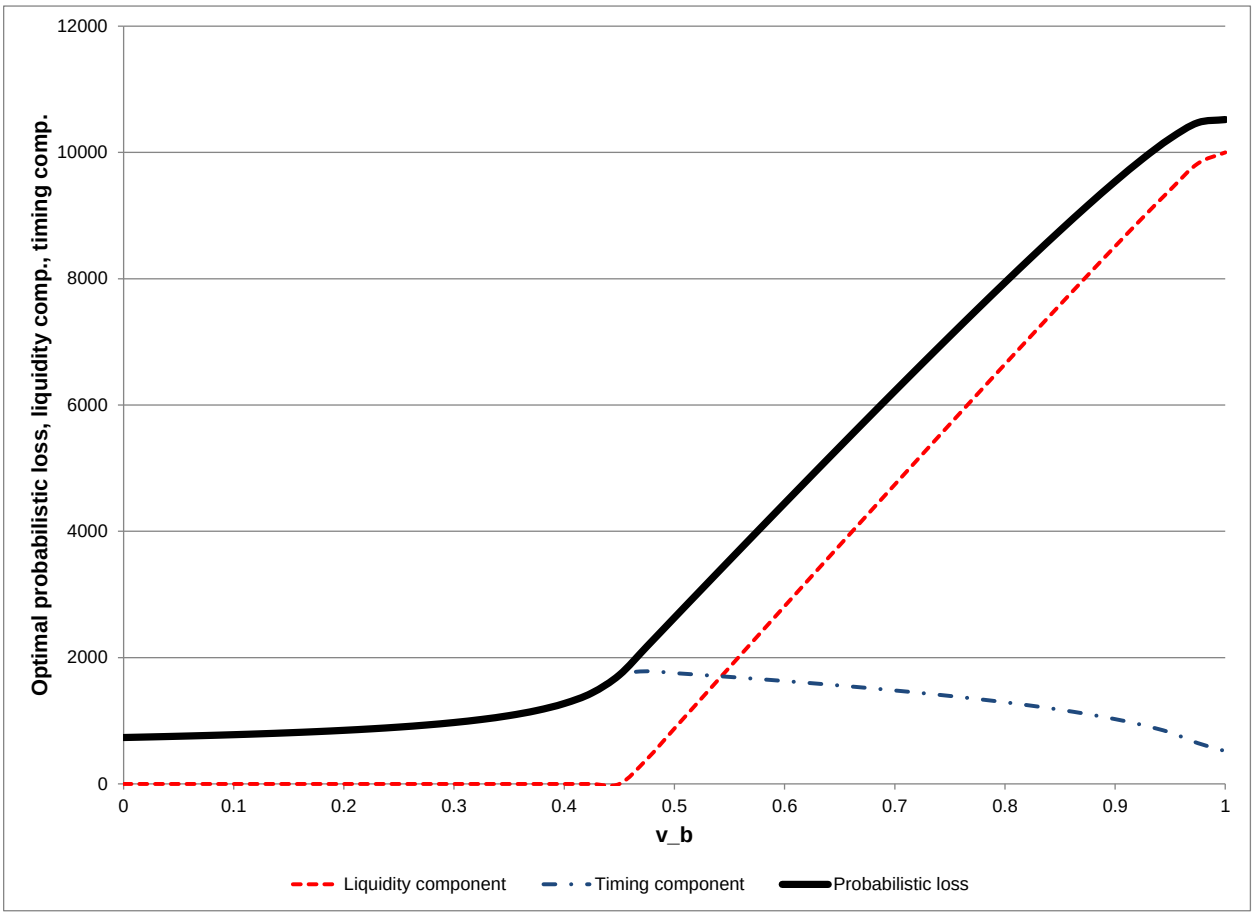

Figure 1(b) – Optimal probabilistic loss and its components for various v^B on a buy order

This Figure shows the expected total trading loss for a buy order and its relation to market imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The expected market imbalance is given by the expected fraction of buy orders in the market. When all orders are buys, v^B is 1, while if all orders are sells v^B is 0. The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

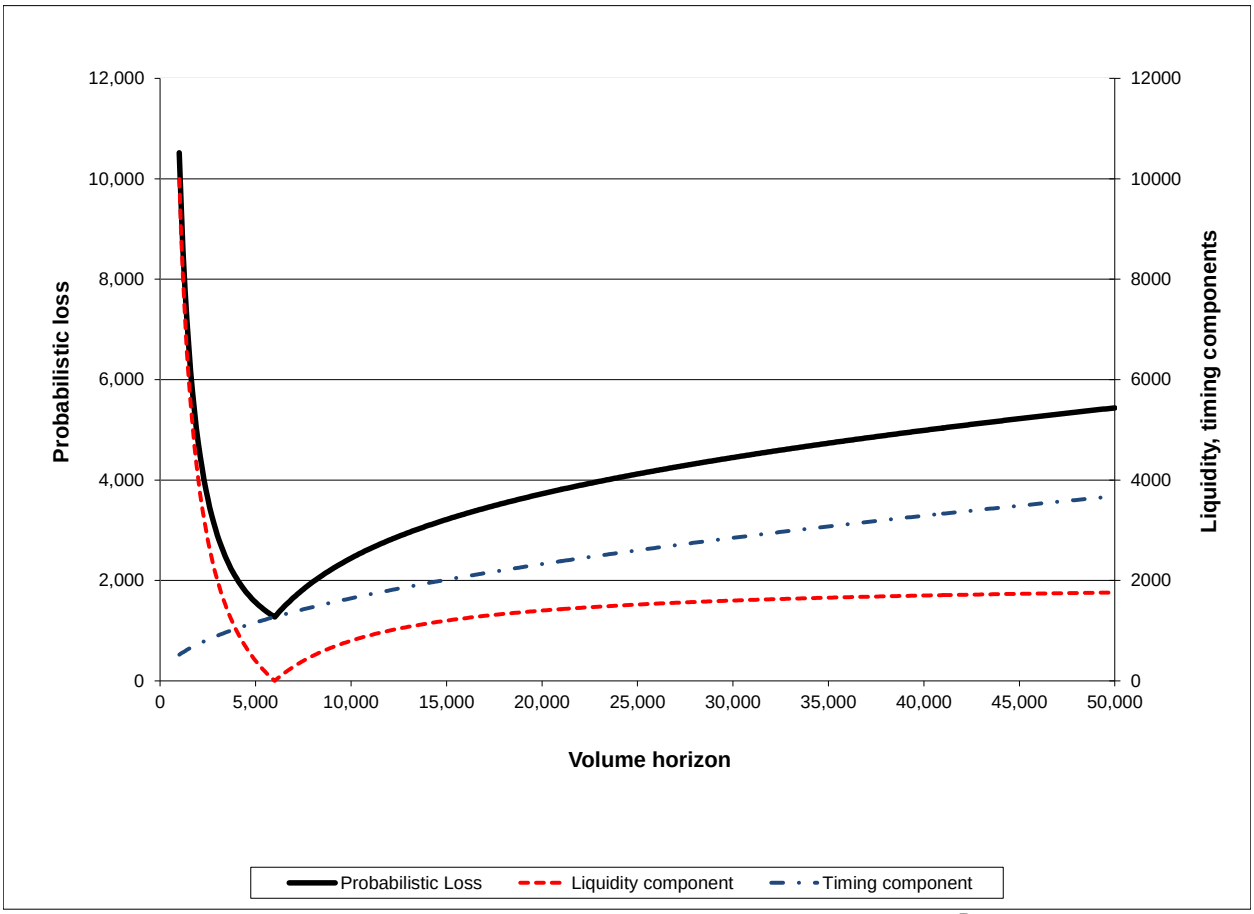

Figure 2 - $\Pi(V,.)$ for different volume horizons (V), with $v^B = 0.4$

This Figure demonstrates the expected total trading loss for a buy order and its relation to the volume horizon when the market is expected to have more selling activity. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

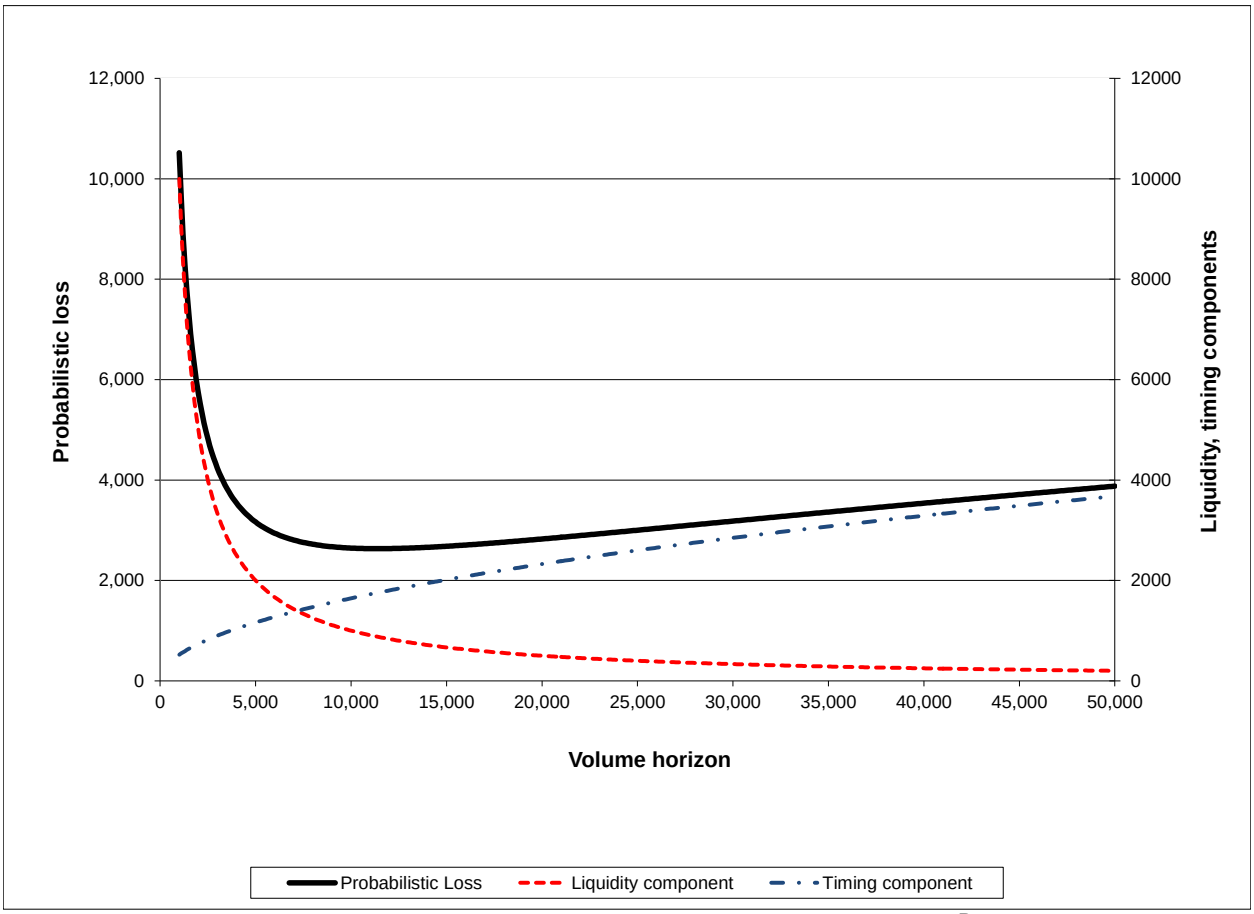

Figure 3 - $\Pi(V,.)$ for different volume horizons (V), with $v^B = 0.5$

This Figure shows the expected total trading loss for a buy order and its relation to the volume horizon when the market is expected to be in a balanced state. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

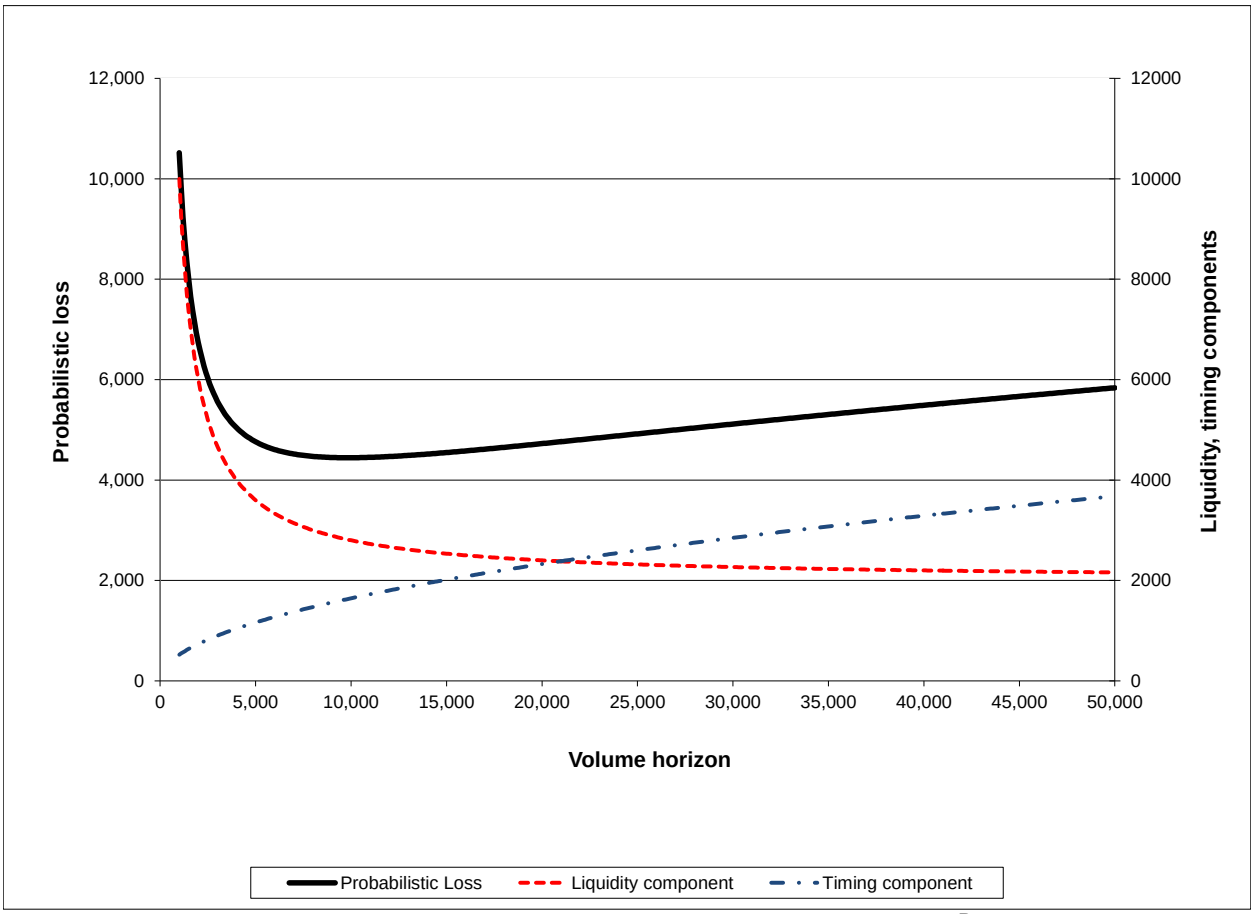

Figure 4 - $\Pi(V,.)$ for different volume horizons (V), with $v^B = 0.6$

This Figure demonstrates the expected total trading loss for a buy order and its relation to the volume horizon when the market is expected to have a greater buy imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

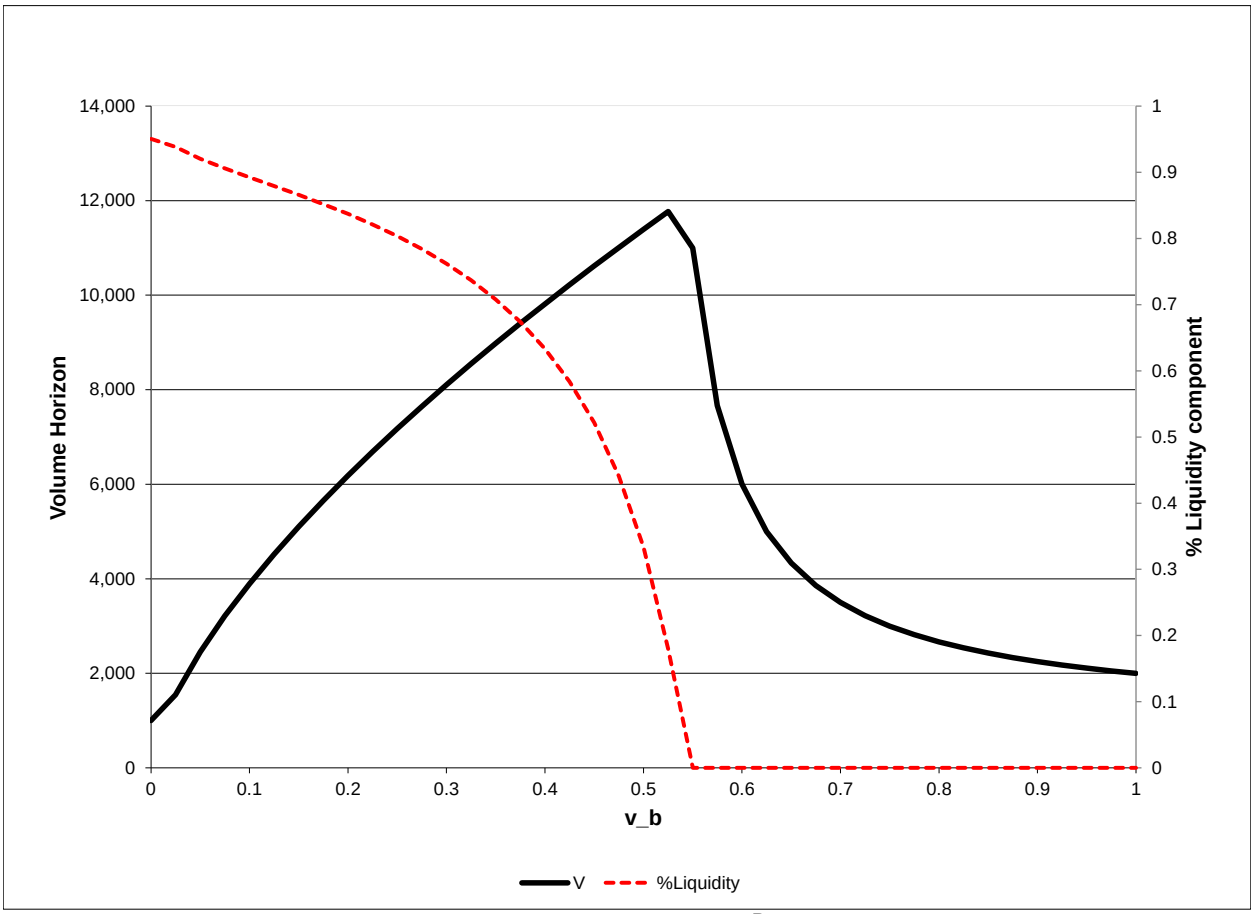

Figure 5 – Optimal V for various v^B on a sell order

This Figure shows how the optimal trading horizon for a sell order depends upon the expected order imbalance in the market. When all orders are buys (sells), v^B is 1, while if all orders are sells v^B is 0. The optimal volume horizon is defined over shares or contracts. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|] = 1$.

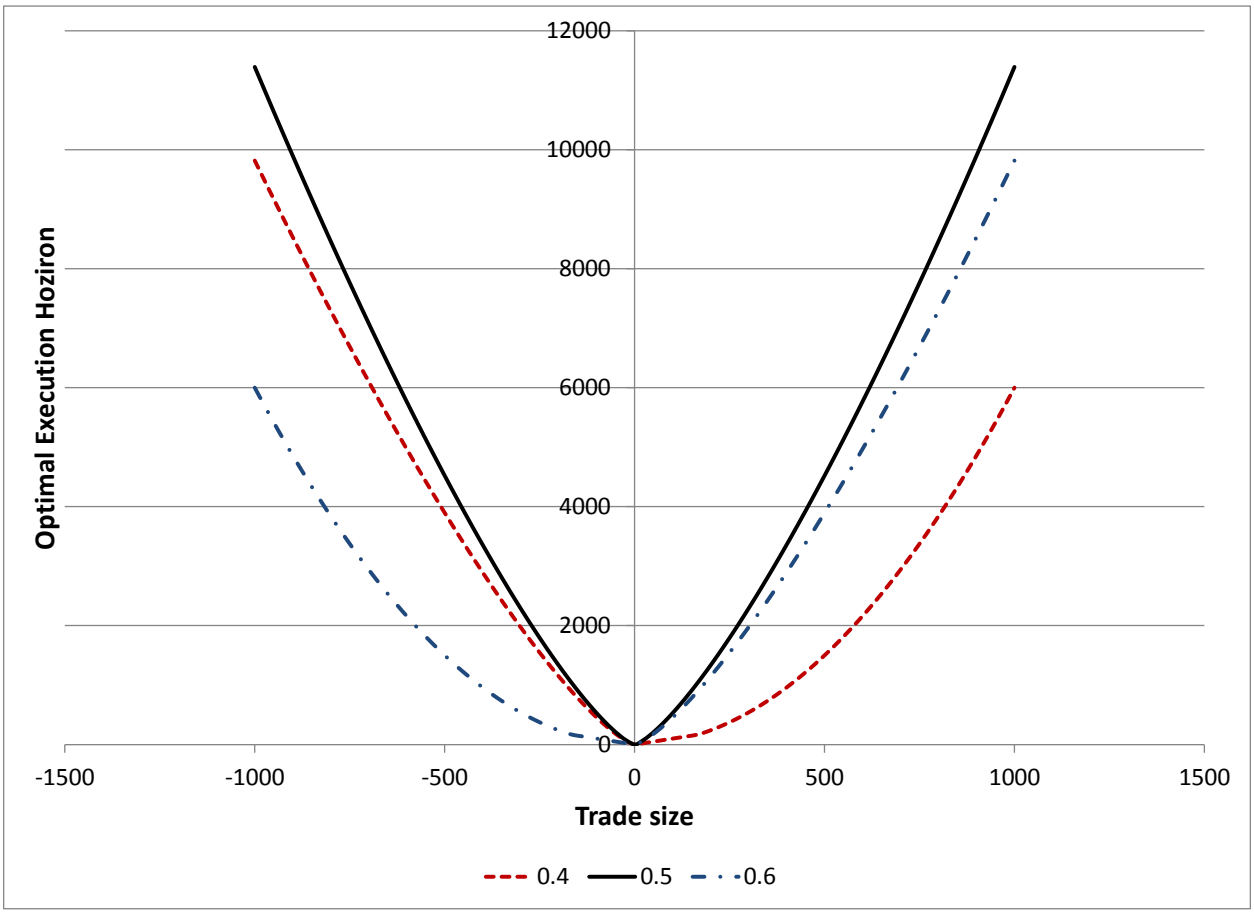

Figure 6 – Optimal execution horizons for various order imbalances and trade sizes/sides

Combining alternative trade sizes and sides with our three scenarios ($v^B = 0.4, v^B = \frac{1}{2}, v^B = 0.6$) results in the optimal execution horizons displayed in the figure above. $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $m = 1{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

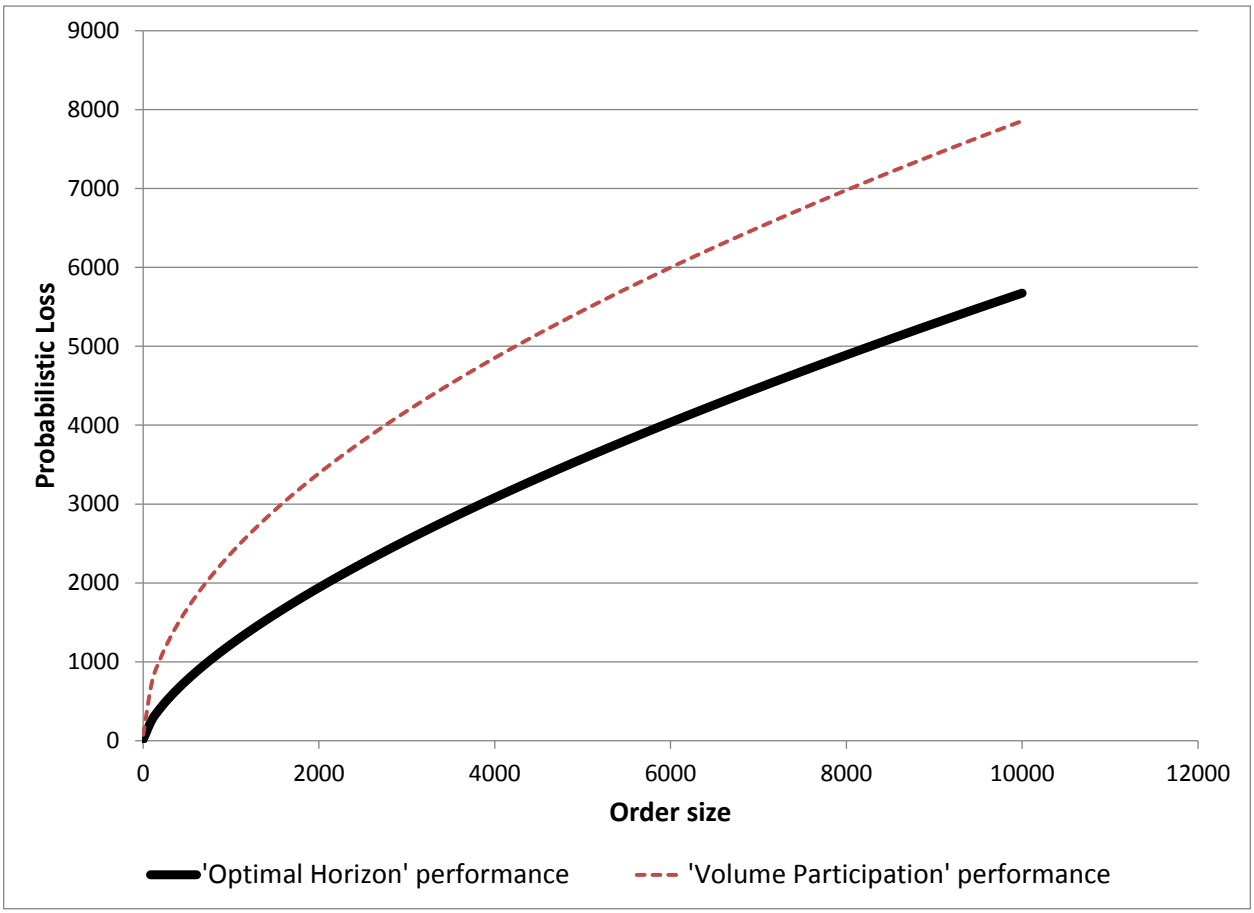

Figure 7 – OEH's performance vs. a Volume Participation strategy when $v^B = \frac{1}{2}$

This Figure shows the expected total trading loss for a buy order arising from either a Volume participation strategy or the OEH strategy when the market is expected to be in a balanced state. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The volume participation strategy is assumed to participate in 5% of the volume. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

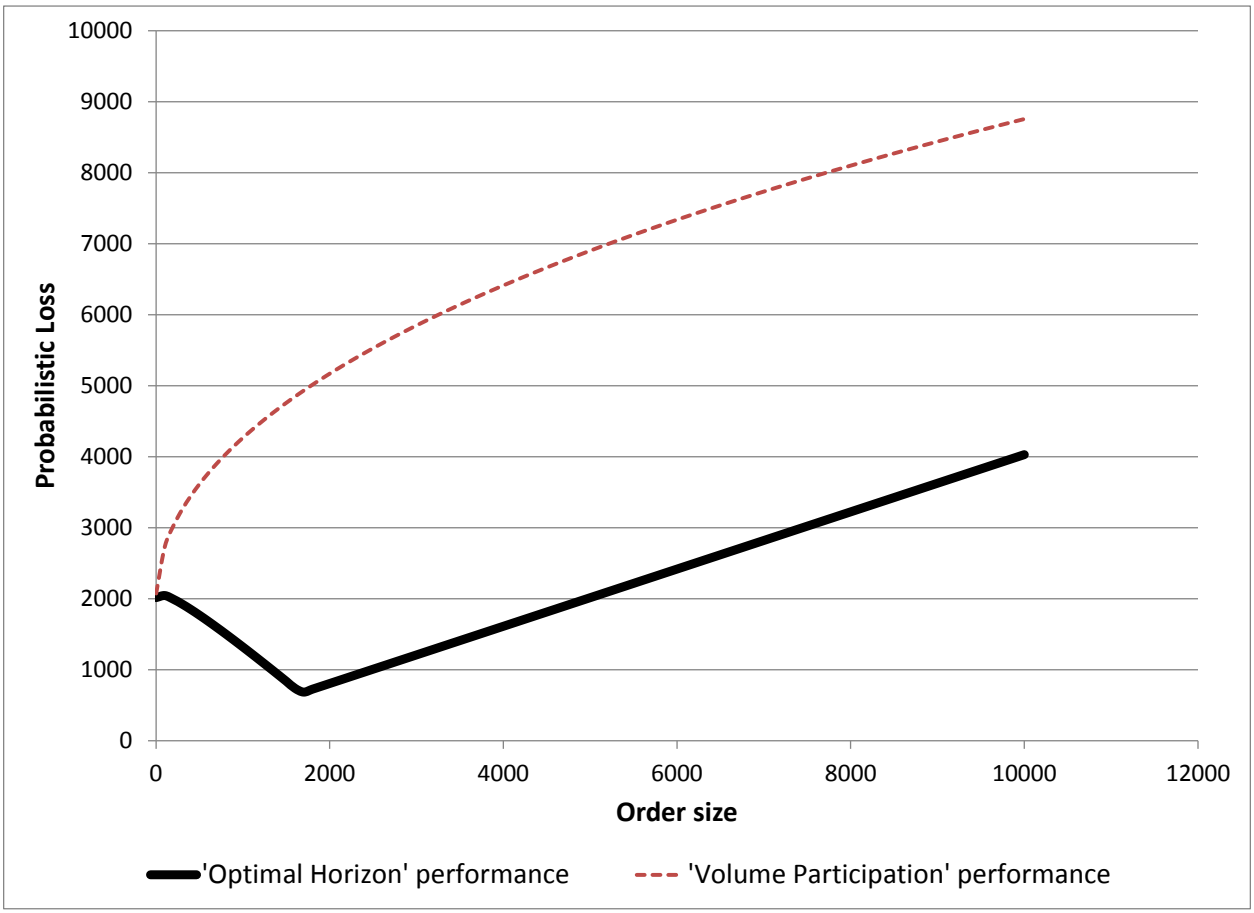

Figure 8 – OEH's performance vs. a Volume Participation strategy when $v^B = 0.4$

This Figure shows the expected total trading loss for a buy order arising from either a Volume participation strategy or the OEH strategy when the market is expected to have a greater sell imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The volume participation strategy is assumed to participate in 5% of the volume. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

Figure 9 – OEH's performance vs. a Volume Participation strategy when $v^B = 0.6$

This Figure shows the expected total trading loss for a buy order arising from either a Volume participation strategy or the OEH strategy when the market is expected to have a greater buy imbalance. The total loss is a function of the liquidity component (the cost from trading immediately) and a timing component (the cost from delaying trading). The optimal volume horizon is defined over shares or contracts. The volume participation strategy is assumed to participate in 5% of the volume. The figure is drawn for state variables: $\hat{\sigma} = 1{,}000$, $V_\sigma = 10{,}000$, $[\overline{S} - \underline{S}] = 10{,}000$, $\lambda = 0.05$ and $\varphi[|m|]$ linear.

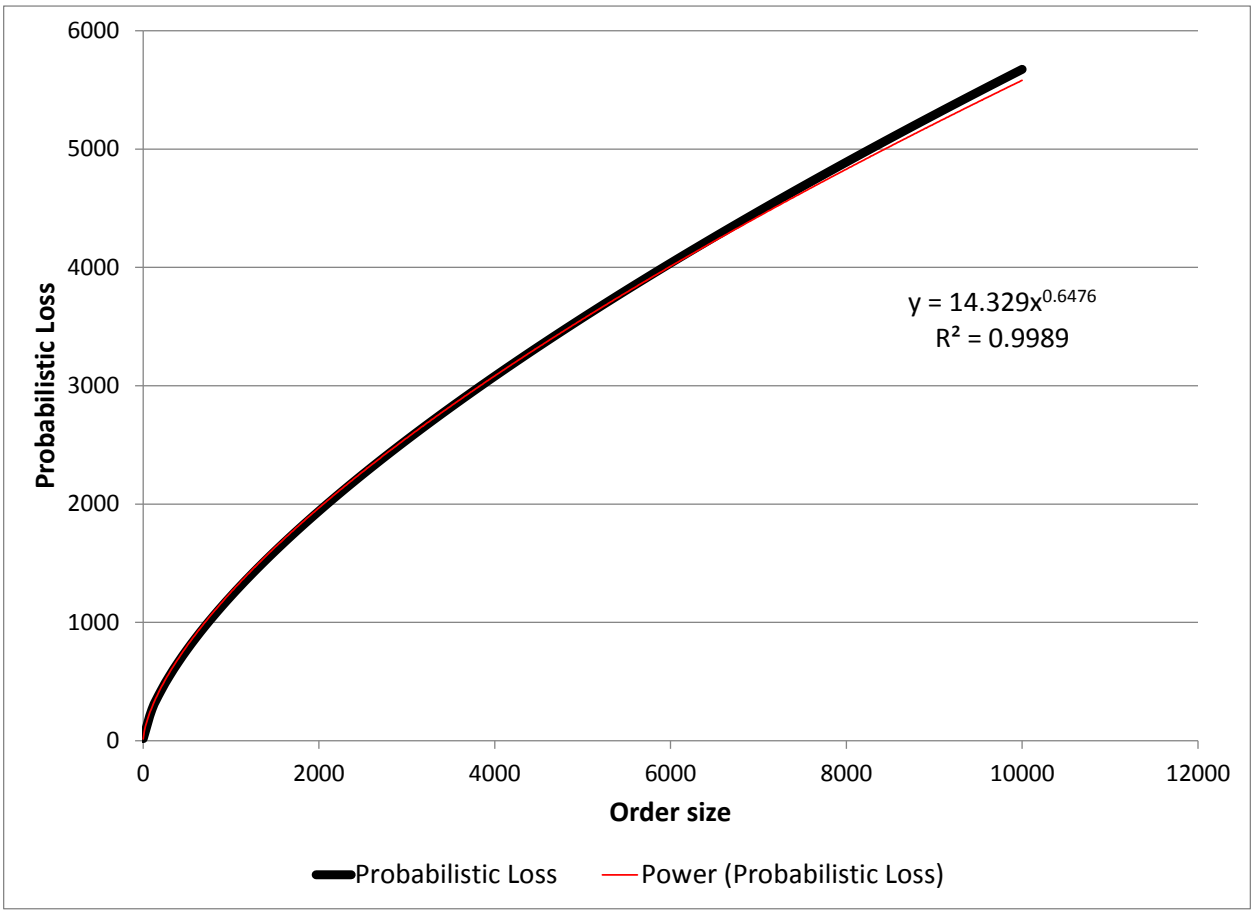

Figure 10 – Probabilistic loss under $v^B = \frac{1}{2}$ *and* $\varphi[|m|]$ *linear*

When order flow is balanced, the probabilistic loss follows a functional form close to the square root.

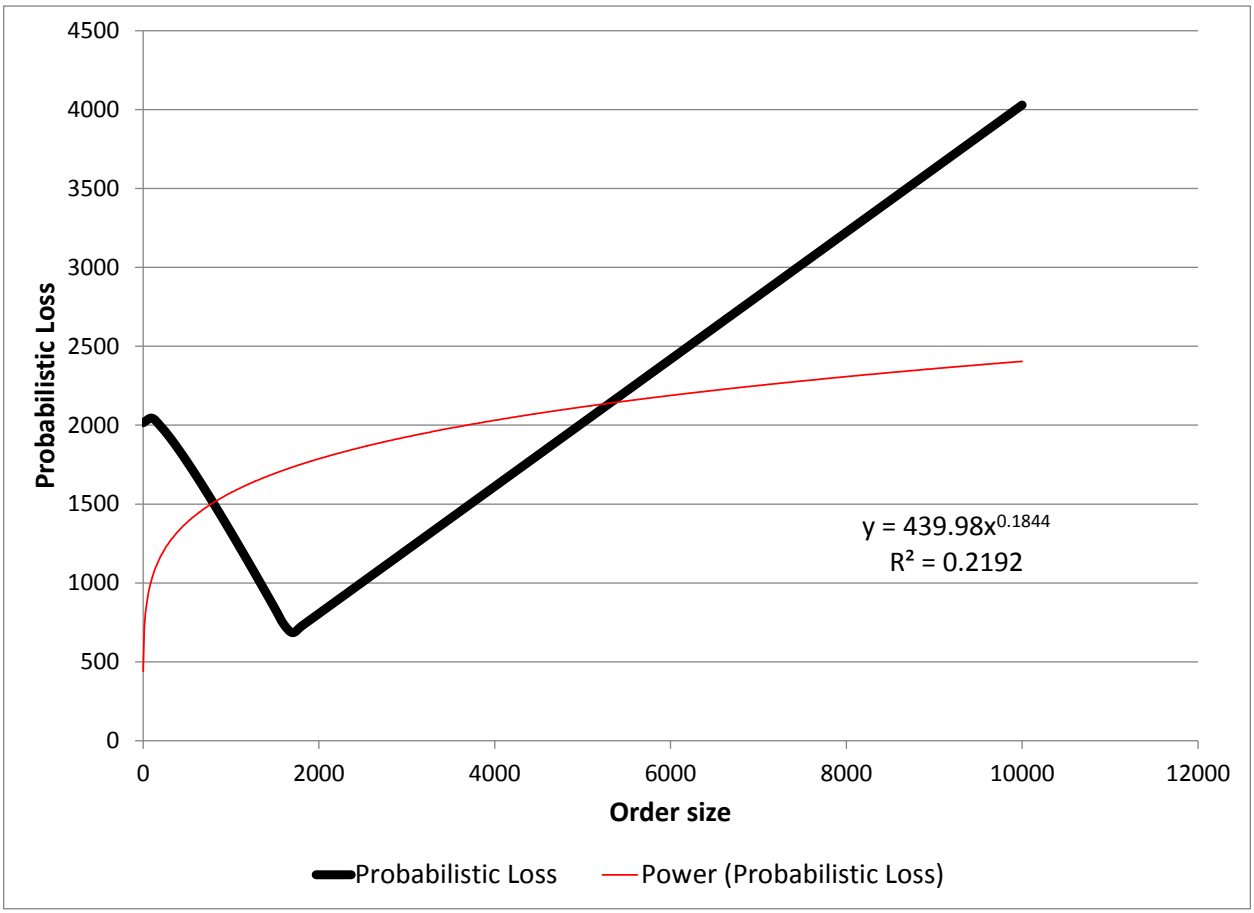

Figure 11 – Probabilistic loss under $v^B = 0.4$ *and* $\varphi[|m|]$ *linear*

When order flow is leaning against the market, the probabilistic loss has a piecewise linear functional form.

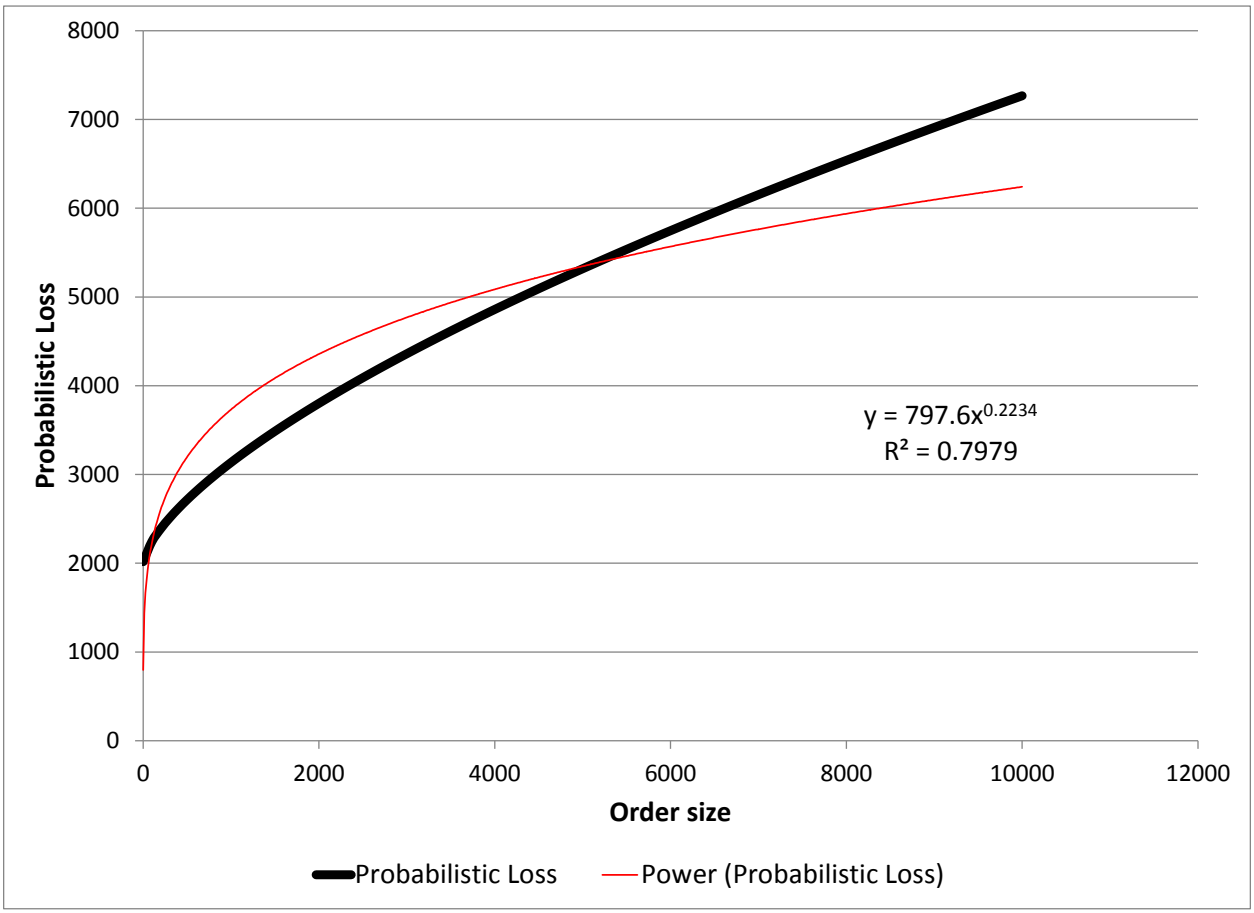

Figure 12 – Probabilistic loss under $v^B = 0.6$ *and* $\varphi[|m|]$ *linear*

When the order is large and competing for liquidity, the probabilistic loss is greater than predicted by the square root.

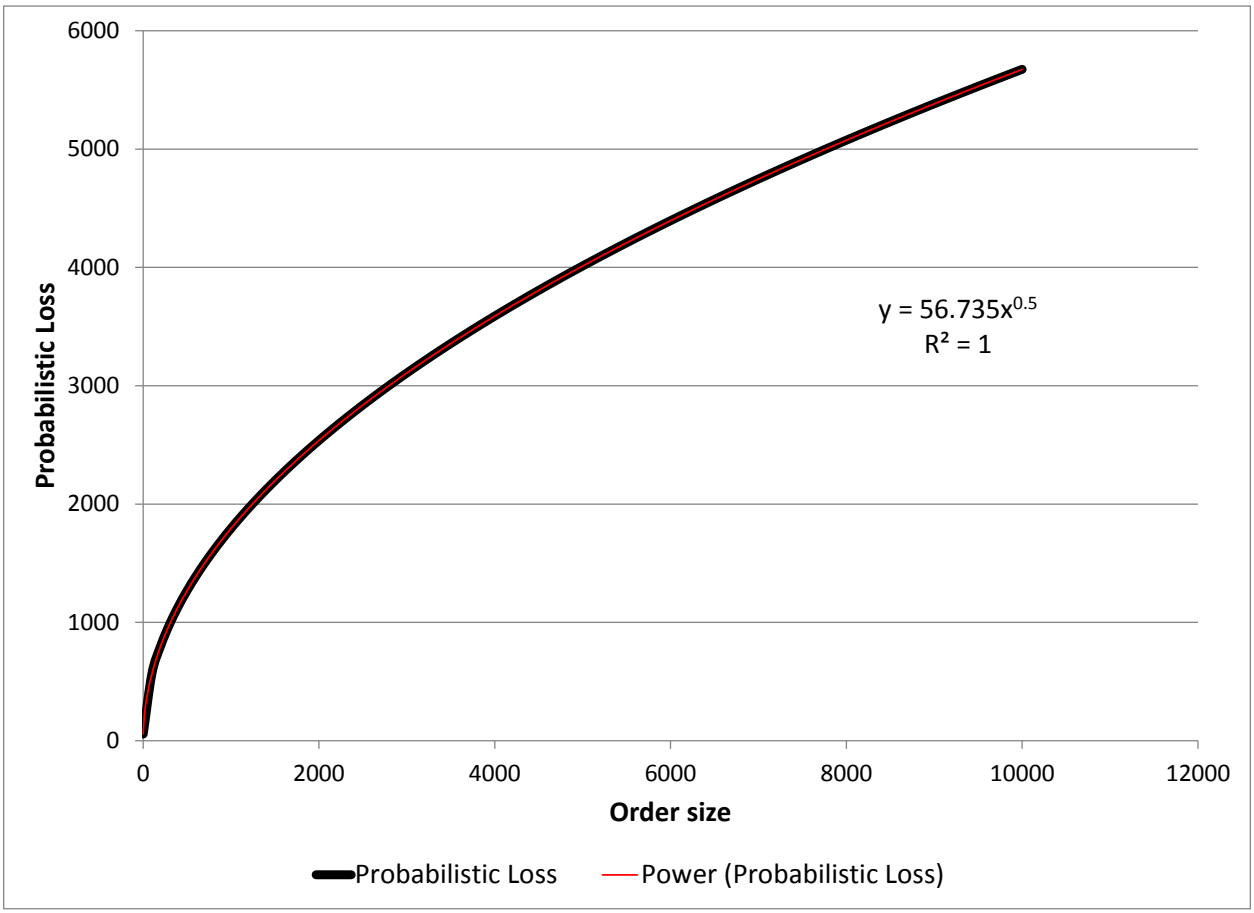

Figure 13 – Probabilistic loss under $v^B = \frac{1}{2}$ *and* $\varphi[|m|] \propto \sqrt{m}$

The probabilistic loss exactly fits the square root when $v^B = \frac{1}{2}$ and $\varphi[|m|] \propto \sqrt{m}$.

Futures Contract	Exchange	Group	Start	End	Roll	Records	ADV
E-Mini S&P500	CME	Equity	1/1/2007	7/26/2012	12	476,676,009	1,964,844.89
T-Note	CBOT	Rates	1/1/2007	7/26/2012	28	95,091,010	921,056.33
EUR/USD	CME	FX	1/1/2007	7/26/2012	10	188,197,121	233,201.17
WTI Crude Oil	NYMEX	Energy	1/1/2007	7/26/2012	19	164,619,912	194,902.36
Gold	COMEX	Metals	1/1/2007	7/26/2012	27	62,672,073	81,854.96
Corn	CBOT	Softs	1/1/2007	7/26/2012	20	41,833,299	73,860.53
Natural Gas	NYMEX	Energy	1/1/2007	7/26/2012	Volume	50,575,494	61,685.78
Lean Hogs	CME	Meat	1/1/2007	7/26/2012	24	5,499,602	6,544.67
Cotton#2	ICE	Softs	1/1/2007	7/26/2012	20	4,494,294	6,171.32

Table 1 – Description of the data series used in the numerical examples

Futures Contract	Information	Trade Size	Max Profit	OEH Profit (Pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign	0.01*ADV	12.1428	10.5104	4.3262	35.63%	10.04
T-Note	Sign	0.01*ADV	0.3966	0.3441	0.1322	33.33%	9.18
EUR/USD	Sign	0.01*ADV	0.0074	0.0064	0.0028	37.28%	10.62
WTI Crude Oil	Sign	0.01*ADV	1.3913	1.1949	0.4582	32.93%	10.02
Gold	Sign	0.01*ADV	9.4932	8.1780	3.2875	34.63%	9.68
Corn	Sign	0.01*ADV	8.4173	7.2806	3.1640	37.59%	9.67
Natural Gas	Sign	0.01*ADV	0.1098	0.0945	0.0409	37.26%	9.50
Lean Hogs	Sign	0.01*ADV	0.7451	0.6334	0.2613	35.07%	10.51
Cotton#2	Sign	0.01*ADV	1.3211	1.1358	0.4675	35.38%	7.66

Table 2 – EOH's outperformance over VWAP for trades equivalent to 1% of ADV and information regarding the side of the price move over the next bucket

Futures Contract	Information	Trade Size	Max Profit	OEH Profit (Pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign	0.05*ADV	12.1428	8.2047	2.2770	18.75%	5.63
T-Note	Sign	0.05*ADV	0.3966	0.2682	0.0649	16.37%	4.91
EUR/USD	Sign	0.05*ADV	0.0074	0.0051	0.0015	20.78%	6.51
WTI Crude Oil	Sign	0.05*ADV	1.3913	0.9275	0.2217	15.94%	5.33
Gold	Sign	0.05*ADV	9.4932	6.3202	1.6253	17.12%	5.15
Corn	Sign	0.05*ADV	8.4173	5.6471	1.7222	20.46%	5.73
Natural Gas	Sign	0.05*ADV	0.1098	0.0731	0.0221	20.14%	5.68
Lean Hogs	Sign	0.05*ADV	0.7451	0.4820	0.1230	16.51%	5.32
Cotton#2	Sign	0.05*ADV	1.3211	0.8776	0.2372	17.96%	4.24

Table 3 – EOH's outperformance over VWAP for trades equivalent to 5% of ADV and information regarding the side of the price move over the next bucket

Futures Contract	Information	Trade Size	Max Profit	OEH Profit (Pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign	0.1*ADV	12.1428	6.3551	0.9577	7.89%	2.59
T-Note	Sign	0.1*ADV	0.3966	0.2065	0.0212	5.36%	1.82
EUR/USD	Sign	0.1*ADV	0.0074	0.0039	0.0007	9.18%	3.17
WTI Crude Oil	Sign	0.1*ADV	1.3913	0.7037	0.0600	4.31%	1.58
Gold	Sign	0.1*ADV	9.4932	4.7746	0.5125	5.40%	1.82
Corn	Sign	0.1*ADV	8.4173	4.2081	0.6715	7.98%	2.63
Natural Gas	Sign	0.1*ADV	0.1098	0.0549	0.0091	8.30%	2.57
Lean Hogs	Sign	0.1*ADV	0.7451	0.3580	0.0314	4.22%	1.46
Cotton#2	Sign	0.1*ADV	1.3211	0.6582	0.0869	6.58%	1.72

Table 4 – EOH's outperformance over VWAP for trades equivalent to 10% of ADV and information regarding the side of the price move over the next bucket

Futures Contract	Information	Trade Size	Max Profit	OEH Profit (Pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign, Size	0.01*ADV	15.8723	14.1671	6.4076	40.37%	8.52
T-Note	Sign, Size	0.01*ADV	0.5291	0.4721	0.1959	37.03%	6.74
EUR/USD	Sign, Size	0.01*ADV	0.0098	0.0087	0.0039	39.98%	8.74
WTI Crude Oil	Sign, Size	0.01*ADV	1.8682	1.6672	0.6830	36.56%	8.39
Gold	Sign, Size	0.01*ADV	12.5753	11.4060	4.7222	37.55%	6.96
Corn	Sign, Size	0.01*ADV	12.3966	11.0999	5.1200	41.30%	5.77
Natural Gas	Sign, Size	0.01*ADV	0.1380	0.1230	0.0566	40.98%	7.97
Lean Hogs	Sign, Size	0.01*ADV	0.8442	0.7552	0.3348	39.66%	7.92
Cotton#2	Sign, Size	0.01*ADV	1.7879	1.6020	0.7070	39.54%	6.52

Table 5 – EOH's outperformance over VWAP for trades equivalent to 1% of ADV and information regarding the side and size of the price move over the next bucket

Futures Contract	Information	Trade Size	Max Profit	OEH Profit (Pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign, Size	0.05*ADV	15.8723	11.4107	4.0384	25.44%	7.34
T-Note	Sign, Size	0.05*ADV	0.5291	0.3744	0.1120	21.17%	6.30
EUR/USD	Sign, Size	0.05*ADV	0.0098	0.0071	0.0025	25.40%	8.26
WTI Crude Oil	Sign, Size	0.05*ADV	1.8682	1.3453	0.4091	21.90%	7.25
Gold	Sign, Size	0.05*ADV	12.5753	9.3511	2.9978	23.84%	5.94
Corn	Sign, Size	0.05*ADV	12.3966	8.6294	3.0115	24.29%	6.28
Natural Gas	Sign, Size	0.05*ADV	0.1380	0.0988	0.0358	25.98%	6.91
Lean Hogs	Sign, Size	0.05*ADV	0.8442	0.6141	0.2119	25.11%	7.50
Cotton#2	Sign, Size	0.05*ADV	1.7879	1.2726	0.4329	24.21%	5.46

Table 6 – EOH's outperformance over VWAP for trades equivalent to 5% of ADV and information regarding the side and size of the price move over the next bucket

Futures Contract	Information	Trade Size	Max Profit	OEH Profit (Pts)	Outperf.(Pts)	Outperf.(%)	IR
E-Mini S&P500	Sign, Size	0.1*ADV	15.8723	8.8113	2.2197	13.98%	5.98
T-Note	Sign, Size	0.1*ADV	0.5291	0.2891	0.0546	10.32%	4.69
EUR/USD	Sign, Size	0.1*ADV	0.0098	0.0055	0.0014	14.79%	7.10
WTI Crude Oil	Sign, Size	0.1*ADV	1.8682	1.0389	0.2004	10.73%	5.16
Gold	Sign, Size	0.1*ADV	12.5753	6.7359	1.4143	11.25%	6.36
Corn	Sign, Size	0.1*ADV	12.3966	6.4158	1.5333	12.37%	5.39
Natural Gas	Sign, Size	0.1*ADV	0.1380	0.0768	0.0208	15.11%	5.26
Lean Hogs	Sign, Size	0.1*ADV	0.8442	0.4827	0.1199	14.21%	6.57
Cotton#2	Sign, Size	0.1*ADV	1.7879	0.9576	0.2466	13.79%	4.26

Table 7 – EOH's outperformance over VWAP for trades equivalent to 10% of ADV and information regarding the side and size of the price move over the next bucket

REFERENCES

- Almgren, R. (2009): *"Optimal Trading in a Dynamic Market"*, working paper.
- Almgren, R. (2003): *"Optimal Execution with Nonlinear impact functions and trading-enhanced risk"*, Applied Mathematical Finance (10), 1-18.
- Almgren, R. and N. Chriss (2000): *"Optimal execution of portfolio transactions"*, Journal of Risk, Winter, pp.5-39.
- Almgren, R., C. Thum, E. Hauptmann and H. Li (2005): *"Direct estimation of equity market impact"*, working paper.
- Barra (1997): *"Market Impact Model Handbook"*, Barra.
- Bayraktar, E., M. Ludkovski (2011): *"Optimal Trade Execution in Illiquid Markets"*, Mathematical Finance, Vol. 21(4), 681-701.
- Berkowitz, S., D. Logue and E. Noser (1988): *"The total cost of transactions on the NYSE"*, Journal of Finance, 41, 97-112.
- Bertsimas, D. and A. Lo (1998): *"Optimal control of execution costs"*, Journal of Financial Markets, 1, 1-50.
- Bethel, W., D. Leinweber, O. Ruebel, K. Wu (2011): *"Federal Market Information Technology in the Post Flash Crash Era: Roles for Supercomputing"*, Journal of Trading, Vol. 7, No. 2, pp. 9-25. SSRN: http://ssrn.com/abstract=1939522
- Breen, W., L. Hodrick and R. Korajczyk (2002): *"Predicting equity liquidity"*, Management Science 48(4), 470-483.
- CFTC-SEC (2010), *"Findings Regarding the Market Events of May 6, 2010"*.
- Dai, M. and Y. Zhang (2012), *"Optimal Selling/Buying Strategy with Reference to the Ultimate Average,"* Mathematical Finance, 22(1), 165-184.
- Easley, D. and M. O'Hara (1992b): *"Time and the process of security price adjustment"*, Journal of Finance, 47, 576-605.
- Easley, D., Kiefer, N., O'Hara, M. and J. Paperman (1996): *"Liquidity, Information, and Infrequently Traded Stocks"*, Journal of Finance, September.
- Easley, D., R. F. Engle, M. O'Hara and L. Wu. (2008) *"Time-Varying Arrival Rates of Informed and Uninformed Traders"*, Journal of Financial Econometrics.
- Easley, D., M. López de Prado and M. O'Hara (2011a): *"The Microstructure of the Flash Crash: Flow toxicity, liquidity crashes and the probability of Informed Trading"*, The Journal of Portfolio Management, Vol. 37, No. 2, pp. 118-128, Winter. http://ssrn.com/abstract=1695041 .
- Easley, D., M. López de Prado and M. O'Hara (2011b): *"The Exchange of Flow Toxicity"*, Journal of Trading, Spring 2011, 8-14. http://ssrn.com/abstract=1748633
- Easley, D., M. López de Prado and M. O'Hara (2012a): *"Flow toxicity and Liquidity in a High Frequency world"*, Review of Financial Studies, Vol. 25 (5), pp. 1457-1493. http://ssrn.com/abstract=1695596.
- Easley, D., M. López de Prado and M. O'Hara (2012b): *"Bulk Volume Classification"*, SSRN, working paper. http://ssrn.com/abstract=1989555
- Easley, D., M. López de Prado and M. O'Hara (2012c): *"The Volume Clock: Insights into the High Frequency Paradigm"*, The Journal of Portfolio Management, forthcoming (Fall, 2012). http://ssrn.com/abstract=2034858

- Easley, D., M. López de Prado and M. O'Hara (2012d): *"TIM and the Microeconomics of the Flash Crash"*. Working paper.
- Forsyth, P. (2011), *"A Hamilton Jacobi Bellman approach to optimal trade execution"*, Applied Numerical Mathematics, Vol. 61 (2), February, 241-265.
- Gatheral, J. and A. Schied (2011): *"Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework"*, International Journal of Theoretical and Applied Finance 14(3), 353-368.
- Grinold, R. and R. Kahn (1999): *"Active portfolio management",* McGraw-Hill, pp. 473-475.
- Hasbrouck, J. and R. Schwartz (1988): *"Liquidity and execution costs in equity markets"*, Journal of Portfolio Management, 14 (Spring), 10–16.
- Kissell, R. and M. Glantz (2003): *"Optimal trading strategies"*, American Management Association.
- Konishi, H. and N. Makimoto (2001): *"Optimal slice of a block trade"*, Journal of Risk 3(4).
- Lillo, F., J. Farmer and R. Mantegna (2003): *"Master curve for price impact function"*, Nature 421, 129-130.
- Madhavan, A. (2002): *"VWAP strategies"*, Transaction Performance, Spring.
- Perold, A. F. (1988): *"The implementation shortfall: Paper versus reality"*, The Journal of Portfolio Management 14, Spring, 4-9.
- Schied, A. and T. Schöneborn (2009): *"Risk Aversion and the Dynamics of Optimal Liquidation Strategies in Illiquid Markets"*, Finance and Stochastics, 13(2), 181–204.
- Stoll, H. (1989): *"Inferring the components of the bid-ask spread: Theory and empirical tests"*, Journal of Finance, 44, 115-134.