Dynamic Order Execution Mechanisms

Aryan Ayyar[*]

This work is a thesis for three-credit project taken by the student at the Hong Kong Polytechnic University for the academic year 2024-25, semester two of the Master of Finance (Investment Management) program. Submitted to the Department of Accounting and Finance, The Hong Kong Polytechnic University for the consideration of the project assessment committee. The author also wishes to thank Te Feng Chen, the project supervisor for his mentorship. He also wishes to thank Tarun Chordia, Madhu Veeraraghavan, Linti Zhang, Li Jiang, Xiaoyang Ben, Yangyang Chen, Krishna Prasad, Increase Onwachukwu, and others for helpful comments and guidance. This work is intended as an appreciation for literature based on the author's reading and interpretation. He thanks Robert Almgren, Neil Chriss and Albert Kyle for their contribution to this field. All errors are the author's own and any discrepancies can be reported to aryan.ayyar@outlook.in. Comments and suggestions would be highly appreciated.

श्रीमत्सुन्दरजामातृमुनिमानसवासिने
सर्वलोकनिवासाय श्रीनिवासाय मङ्गलम्।

- श्री वेङ्कटेश्वर सुप्रभातम्

śrīmat-sundara-jāmātṛ-muni-mānasa-vāsine
sarva-loka-nivāsāya śrīnivāsāya maṅgalam

- *Śrī Vēṅkaṭēśvara Suprabhātam*

Abstract

The work investigates optimal execution strategies in financial markets under both deterministic and stochastic frameworks, focusing on minimizing transaction costs while managing informational asymmetries. We build upon foundational works such as Almgren and Chriss (2000) and Kyle (1985) bridging discrete-time optimization and continuous-time stochastic control to develop a unified approach for informed traders operating in markets with price impact and dynamic liquidity. Further, the thesis extends the canonical Kyle model by introducing a finite-horizon equilibrium with risk-averse agents and learning market makers, endogenizing the price impact function as a response to order flow informativeness. We also perform an empirical estimation of relevant parameters using data of close to 100 Million trades for over 8,000 stocks in the NYSE and NASDAQ for the year 2022 validating such a model's theoretical constructs and providing insight into the distribution and interdependence of price impacts and the key role they play in optimal execution of transactions.

1. Introduction

Liquidity is a multifaceted concept, varying in its interpretation across different contexts. Financial Economists and Mathematicians alike have been studying this for over four decades now. Liquidity (and often illiquidity) is in my opinion an elusive construct—recognizable in practice but challenging to precisely define (see Amihud, Mendelson, and Pedersen, 2012). From a fundamental asset pricing perspective, liquidity can be conceptualized as the present value of all future discounted transaction costs. There exist conflicting views but for one thing is certain – transaction costs can't be ignored. The consequences of this may be manyfold – First, liquidity directly relates to asset pricing and leads to a reduction in value of asset prices due to discounted transaction costs. Second, liquidity risk has been a subject of contention for a long time now, but the implications of this on trading and transaction costs have in my opinion well researched although a framework that unifies many conflicting viewpoints is yet to take a seat. Earlier, traders, when submitting their demands, would consider the information content of the price but not the price impact of their trade. To understand this, let's consider a hypothetical example of a hedge fund holding a position of 5 million shares in Tesla (TSLA) which is advised to relocate the shares due to an ongoing portfolio reconstruction. Such a relocation may also be made if there exists some information that the hedge fund possesses such that liquidating this large position shall make the insider (the hedge fund in our case) gain profits. But , with an average daily volume of 30 million shares, an abrupt liquidation of 5 million shares (approx. 16% of ATV) would have a significant impact on the price. There are two obvious options available to the fund – (a) sell everything now at a known price entailing a high price impact or (b) sell in equal size packets over a given time horizon to minimize price impact but at the same time holding significant risk of price change over time. A third alternative may be as follows - split orders proportionately to the intraday volume by a VWAP algorithm, which may be better than a naïve execution strategy but may still have significant slippage. From a purely mathematical perspective, we have evidence to believe that all of these strategies may in fact be inefficient as such order execution mechanisms don't take into consideration either/both the implementation shortfall or variance of the price dynamics. By the means of such an essay, we intend to discuss a class of algorithms that deal with the optimal trading strategy of a rational trader to overcome such limitations.

2. Relevant Literature

The effect of liquidity costs on security prices and returns was first studied by Amihud and Mendelson (1986) who defined liquidity as a measure of trading costs through the mechanism of the bid-ask spread. Amihud (2002) proposes that expected market illiquidity has a significant impact on the returns of individual stocks. Utilizing an autoregressive model, his study introduces the concept of unexpected liquidity shocks, which exhibits a peculiar characteristic one would like to describe as follows - Illiquidity in a given year increases the anticipated illiquidity for the subsequent year, which is associated with higher expected stock returns or lower asset prices; however, an unexpected increase in illiquidity tends to result in a decline in stock returns, contrary to the effects observed with expected illiquidity. Hence, not all forms of liquidity follow a uniform logic. While expected illiquidity is factored into market dynamics and investor expectations, absorbing its costs through pricing mechanisms, unexpected liquidity shocks operate outside the realm of foresight, triggering reactions that markets are ill-prepared to accommodate.

Chordia, Roll, and Subrahmanyam (2001), along with Hasbrouck and Seppi (2001), argue that market-wide liquidity significantly influences individual stock liquidity, a phenomenon they refer to as liquidity commonality. Their findings suggest that a portion of stock-specific liquidity is driven by broader market liquidity conditions, implying that fluctuations in the average liquidity of the market are reflected in the liquidity levels of individual stocks. Liquidity and Asset Pricing have been subject to a rigorous study by Acharya and Pedersen (2005). The Acharya-Pedersen framework decomposes (while extending the standard CAPM to accommodate for transaction costs) the net beta into risks pertaining to (i) commonality in liquidity, (ii) return sensitivity to market liquidity and (iii) liquidity sensitivity to market returns. They also explore cross-sectional predictions of such a model using NYSE and AMEX stocks for a large sample period. In particular, (iii) seems to be of interest to the asset pricing literature. Brunnermeier and Pedersen (2009) highlight that market liquidity (which they refer to as the ease of trading) is tightly connected to a trader's funding liquidity, i.e. the availability of capital. When an initial liquidity shock occurs, traders face funding constraints which reduce their capacity to provide liquidity. This effect, in turn, increases volatility, as lower liquidity leads to larger price impacts and wider order imbalances. Within a general equilibrium framework, they show that uninformed financiers tightening margins during volatile periods triggers a destabilizing feedback loop.

With a series of models introduced, the first of which by Kyle (1985) took into cognizance the optimal strategy of a trader based on the risks of price impacts using a conjectured equilibrium approach. This created a new class of models effectively separating trade models into two categories – (a) Walrasian batch models and (b) Dealer or sequential trade models. Notable extension of the former class of models include that by Admati and Pfleiderer (1988), Kyle (1989) and Foster and Vishwanathan (1990). However, both these classes of models were studies for the means of understanding the mechanism of price formation in markets, although in my opinion, the optimal strategy of the trader was not an object of interest by itself. Over time, empirical research has consistently demonstrated that intraday liquidity patterns—encompassing volatility, trading volume, and order flow—exhibit a characteristic U-shaped pattern. That is, these variables tend to be elevated at market open, decline through midday, and subsequently rise again toward the market close (see, for example, Harris (1986), Jain and Joh (1988), and Kirchner and Schlag (1998)). Although it seems that such patterns were not fully explainable, partial explanations include that regarding reduced adverse selection costs (see Hasbrouck (1991) and Foster and Vishwanathan (1993)). Some studies such as that by Bessembinder (1994) and Amihud and Mendelson (1982) indicate increased overnight inventory holding costs as a possible explanation for the surge in intraday patterns before the end of the day. An innovation was introduced when Bertsimas and Lo (1998) proposed that a dynamic optimization approach be formally introduced to derive trading strategies that minimize the expected cost of executing trades over a given period. Such a framework demonstrated that, under a variety of conditions, an optimal trading trajectory could be identified to achieve cost efficiency. However, their model does not explicitly account for the volatility of order execution costs. They further establish that a naïve strategy of executing trades by selling a constant number of shares at equally spaced time intervals is optimal if and only if price impacts are both linear and permanent, and if stock prices follow an arithmetic random walk. Additionally, their study acknowledges the influence of exogenous factors, such as prevailing market conditions and initial holdings, on the optimal liquidation strategy.

A profound evolution in the theory of optimal execution emerged with the work of Almgren and Chriss (2000), who introduced a class of algorithms designed to balance execution efficiency with risk control. Their framework is rooted in the minimization of a utility function, wherein the objective is not merely to reduce transaction costs but to navigate the trade-off between expected revenue and the inherent uncertainty in execution better known today as the Implementation shortfall. This strategy works in a framework of maximizing the expected revenue of a trade (in other words minimizing execution costs) with a suitable penalty for the uncertainty of revenue. In such an endeavor, the Almgren-Chriss (henceforth, AC) framework comes up with an efficient frontier of optimal execution strategies, each for a given level of risk aversion. Much like the efficient frontier in portfolio theory, this approach developed the fundamental idea that any reduction in risk must necessarily come at the expense of cost efficiency, and vice versa. Thus, the problem of optimal execution is not mere cost minimization and henceforth, enters the realm of decision-making under uncertainty, where market participants must calibrate their strategies based on their individual risk tolerance and market conditions. Subsequent extensions of such a model have been proposed, notably by Hisata and Yamai (2000) and Dubil (2002), who introduce endogenous considerations regarding the final liquidation horizon. A common methodological feature among these studies is the reliance on a VaR framework to model execution risks, alongside the assumption of a constant trading velocity. Hisata and Yamai (2000) further contribute by providing closed-form solutions for market impact functions characterized by the square-root law and by developing a numerical framework for obtaining such solutions. Dubil (2002), in contrast, examines both linear and power-law impact functions, deriving specialized parameterizations for each case.

More recent models extend the liquidation framework to incorporate trading via limit orders. Among these, Gueant, Lehalle, and Tapia(2012) link the optimal trade execution schedule to the pricing strategy within the limit order book. Their framework employs a HJB equation-based approach, modeling the liquidation process as a control problem that accounts for both non-execution and price-based risks.The aim of this essay is multifold – First, we demonstrate an appreciation for existing models for they set among other things a very clear and baseline framework to work with. Second, we attempt to formulate an informed trader's optimal execution trajectory within two complementary frameworks - (i) as a constrained optimization problem in discrete time, following the Almgren–Chriss (2000) model extended as a stochastic control problem, generalizing the approach developed by Guéant (2016)., and (ii) as a competitive REE equilibrium with dealer learning. We show that in equilibrium, the process by which an informed trader trades are in fact heavily influenced (a) by his/her intentions and (b) whether or not the market maker knows of their intentions. We have in our humble attempt, tried to unify two frameworks that may in the first instance seem conflicting or contradicting but on a deeper level pertain to some interesting conclusions.

3. **Optimal Execution Models**

In this section, we begin by establishing a deterministic structure of the AC framework under various execution cost assumptions, before transitioning into its stochastic formulation. We consider a single informed trader holding a portfolio comprising N risky assets, where his position at time t could be represented by a grid $x_t \in \mathbb{R}^n$. Each component $(x_t)_i$ denotes the number of shares (or proportional wealth allocation) held in asset i at time t. Due to exogenous events, the trader seeks to optimally transition from such an initial portfolio $x_0 \in \mathbb{R}^n$ to a terminal target portfolio state $x^* \in \mathbb{R}^n$, typically chosen as $x^* = 0$ in order to represent a state of complete liquidation. It may now be obvious that naïve execution strategies discussed in the beginning of this essay be in fact not be the wisest choice to proceed with. Hence, execution process is now discretized over N periods with time step $\tau = \frac{T}{N}$, where $T > 0$ can be considered the total liquidation horizon. For the purpose of demonstration, we consider a single-asset case although a generalization to the multi-asset setting is definitely possible but as a case with similar outcomes, is not an article of discussion in this essay. Since we know that the informed trader now holds x_0 shares of the stock, we let x_k denote the remaining inventory at the beginning of interval $k \in \{1, \ldots, N\}$ where each interval is of a time-step τ. Defining the number of shares sold during interval $[t_{k-1}, t_k]$ as $n_k = x_{k-1} - x_k$, inventory trajectory needs to satisfy $x_t = X - \sum_{k=1}^{t} n_k$. Here, x_t represents the total amount of shares the trader is holding at a given time t. We now proceed to show the execution strategies obtained by Almgren and Chriss (2000), generalize these and then analyze whether such a strategy may be considered optimal from the insider's point of view. Let's take the security's price to evolve according to an arithmetic Brownian motion given by[1]

$$S_k = S_{k-1} + \sigma_s \, \tau^{0.5} \, \epsilon_k - \tau g\left(\frac{n_k}{T}\right)$$

where σ_s is the volatility parameter, ϵ_k is a draw from a set of standard normally distributed functions which depend on the average rate of trading $v = \frac{n_k}{T}$, which we shall soon also refer to the velocity of trade. In an ideal world, if all shares were sold at the price S_0 without any price impact or risk, the total trading revenue may be given as $\pi = XS_0$. However, considering a more realistic scenario, there are price impacts and other costs associated with such a liquidation process which in the horizon of the liquidation process can be both (a) substantial given the large position size and (b) interfere with the optimal liquidation process as such costs reduce the profits of the informed trader. Hence, we define the implementation shortfall as difference between the value of a portfolio in an idealized frictionless world and the actual value after trading has occurred, considering price impact and other trading costs. The new profit if each of these N buckets of shares are executed at a price S' may now be given as $S'n_k$. Hence, the implementation shortfall which is basically the difference between the real world and theoretical profits may be given as

$$I.S = XS_0 - \sum_{t=0}^{T} n_k S'_k$$

We now proceed to take notice of the execution price[2] which may be defined as

[1] Such an assumption is based on the standard practice of modelling security dynamics as an Arithmetic Brownian Motion compared to a Geometric Brownian Motion for relatively shorter horizons. An assumption like this has many advantages over modelling the price dynamics as a GBM

[2] One must take note of the fact the execution price of a security is not exactly the price of the security. Consider a case where the fundamental value of a security corresponds to the last traded price (which in our case is defined as the mid-point of the bid-ask spread). Even in such a case, if a trader sells (buys) a position of considerable price, the execution price may be determined by the depth of the market in the limit order book and in certain cases would have to be taken as a weighted average of the shares executed at each level of the LOB considering a case of only market

$$S_t' = S_{t-1} - h(v)$$

Notice that the execution price is dependent on the existing fundamental price (at time $t = t-1$) and the temporary price impact $h(.)$ which may or may not be linearly dependent on v .Hence the terms $\Sigma n_k S_k'$ could in fact be given by the relation

$$\sum_{k=0}^{N} n_k S_k' = XS_0 + \sum_{k=1}^{N} (\left(\sigma_s^{0.5}\epsilon_k - \tau g(v)\right) x_k - \sum_{k=1}^{N} n_k h(v)$$

The first term of the Implementation Shortfall (henceforth I.S) - XS_0 is a measure the initial profits that one might gain from executing a trade at the best price, but the successive terms now measure penalties for both (a) permanent and (b) temporary price impacts[3]. The price impact due to volatility may be expressed as $(\sigma_s{}^{0.5}\epsilon_k)$. Thus, an increase in volatility directly leads to an increase the implementation shortfall. This is a suitable penalty imposed for high volatility stocks. Also, the permanent impact of the trade represented by $\{\tau g(v)\}$ are non-transitory in nature and hence depend on the selling amount. These may be carried forward through trades and hence directly influence the fundamental value of a security[4]. Transitory price impacts which may be given by $\sum_{k=1}^{N} n_k h(v)$ are usually only dependent on time k and are not carried forward to the next trade. Such impacts which don't get carried forward are a major source of order book resilience which is a discussion of a growing body of literature. The trader's value function hence, is to come up with a liquidation strategy that chooses to optimally liquidate the position with a suitable penalty for the implementation shortfall. If the difference between the execution price and the fundamental price of the security is given as x , then it may be convenient to define $E(\zeta_t)$ given by[5]

$$E(\zeta_t) = \frac{1}{2}\gamma X_0^2 + \epsilon \sum_{k=1}^{N} |n_k| + \frac{\eta'}{\tau} \sum_{k=1}^{N} n_k^2$$

The key to solving this problem is the estimation of the variable $I.S = E(\zeta) + \lambda\, Var(\zeta)$. Hence, the trader's value function for the convex optimization becomes

$$\min_{\zeta_t} U(\zeta_t) = E(\zeta_t) + \lambda Var(\zeta_t)$$

$$U(\zeta_t) = \frac{1}{2}\gamma X^2 + \epsilon \sum_{k=1}^{N} |n_k| + \frac{\eta'}{\tau} \sum_{k=1}^{N} n_k^2 + \lambda(\sigma^2 \sum_{k=1}^{N} \tau\, \mathrm{x}_\mathrm{k}^2)$$

This is the optimization problem that the informed trader needs to solve. We now proceed with an analysis of different cases to be considered for such an optimization process.

---

orders for demonstrating our case. In a later analysis, it may also be possible to consider the case of limit orders (see for example , Gueant, Lehalle, and Tapia(2012))

[3]We also must remember that $v = \frac{n_k}{\tau}$ is the velocity of liquidation or the rate of trade execution. Such a note is provided to the reader for their reference.

[4] A detailed analysis of temporary and permanent price impacts has been done on more than a few occasions although the results of such an analysis are often difficult to observe and estimate (see for example Glosten and Harris (1988)) although, in my opinion a general intuition may be more helpful for understanding our case.

[5] $\eta' = \eta - \frac{1}{2}\gamma\tau$ serves as the more complex parameter for estimation.

3.1 Sequential Liquidation in Discrete Intervals

The first of these cases deals with batching orders into smaller packets. Consider we batch a total position of X shares into N equally distributed packets , each packets holding $n_k = \frac{X}{N}$ shares while at the same time also obtaining

$$x_k = (N-k)n_k = \frac{(N-k)X}{N}$$

as the number of shares liquidated in total by time $t = k$. We now use the expected value and variance of the optimal execution functions discussed previously to get

$$U(\zeta) = \frac{1}{2}\gamma X^2 + \epsilon X + \frac{\eta' X^2}{\tau} + \lambda\left(\frac{\sigma^2 T X^2 (N-1)(2N-1)}{6N^2}\right)$$

We may also notice that since ϵ is from a random normally distributed draw, if the informed trader ends up being unlucky, there may be a significant implementation shortfall associated with such a draw. The main argument here is that for an infinitely small basket size (implying an infinitely large number of baskets) there is a solution for the optimal portfolio liquidation which turns out to be finite and dependent on a random draw[6]. Hence, such a strategy seem look great on the outward but may not be among the wisest available options on the table.

3.2 Instantaneous Liquidation with No Price Uncertainty

Let us draw to another extreme version of this where we liquidate the entire portfolio in the first step itself with given whatever price is available. Hence, the revised parameters for $n_0 = X, n_k = 0$ and , $\zeta_k = 0$ give us the utility function

$$U(\zeta) = E(\zeta) = \epsilon X + \frac{\eta X^2}{\tau}, Var(\zeta) = 0$$

In such a scenario, we find the variance of the implementation shortfall to be negligible but on the other hand, the expected value for the utility function assumes a random shock dependent on the initial position X. If $\tau \ll \tau'$or in other words, we sell a large number of shares in a very short time span, it may have a significant impact on the shortfall. This now points out that for a given level of risk aversion, there could exist a unique optimal strategy for the execution of such a portfolio transaction. The essence of the AC framework shows that this is a typical constrained optimization problem which can be solved using Lagrange Multipliers. Specifically, the problem may be stated as - $\min_{\zeta}\left(\mathbb{E}(\zeta) + \lambda Var(\zeta)\right)$ Since for all $\lambda > 0$ there exists a unique solution $\zeta^*(\lambda)$ which minimizes the expected value of the implementation shortfall for both (a) a given level of risk aversion and (b) a certain amount of variance thus giving rise to an optimal portfolio trajectory which is in the form of an efficient frontier. To obtain a closed form solution, we shall have to set the partial derivatives of these parameters to zero (since $n_k^2 = x_k^2 + x_{k-1}^2 - 2x_k x_{k-1}$) .Hence, the first order condition is now given as

$$U'(\zeta) = 2\tau\left(\lambda\sigma^2\zeta_j - \frac{\eta'\left(\zeta_{j-1} - 2\zeta_j + \zeta_{j+1}\right)}{\tau^2}\right) + \epsilon = 0$$

[6] For smaller and smaller sizes of baskets, it may spread out to evaluating this expression for a very large number of baskets (note that a very large number of baskets implies a small size of basket). Hence it may be interesting to evaluate $\lim_{N\to\infty} U(x)$ which gives us the surprising (finite) solution $\lim_{N\to\infty} U(x) = \frac{1}{2}\gamma X^2 + \epsilon X \frac{\eta' X^2}{\tau} + \frac{\lambda\sigma^2 T X^2}{3}$

We now obtain a closed form solution[7] using hyperbolic sine and cosine functions given by

$$x_j = \frac{\sinh\left(\kappa(T - t_j)\right)}{\sinh(\kappa T)} X$$

$$n_j = \frac{2\sinh\left(\frac{1}{2}\kappa T\right)}{\sinh(\kappa T)} \cosh\left(\kappa\left(T - \left(j - \frac{1}{2}\right)\tau\right)\right)$$

3.3 **Empirical Results**

We proceed to estimate transitory and permanent price impacts based on the econometric approach of Glosten and Harris (1988) as adapted by Sadka (2006). We estimate price impacts for 8,549 stocks from the NYSE and NASDAQ for a single day chosen randomly[8] in the year 2022 using a sample of ~100 Million Trades. The estimation procedure is conducted at the individual stock level using intraday transaction data and proceeds in the following steps - (1) modeling price changes to isolate permanent and transitory effects, and (2) aggregating firm-level estimates to construct market-wide illiquidity measures. In the first step, trade direction is classified as buyer- or seller-initiated based on the Lee and Ready (1991) algorithm. Specifically, trades executed above the prevailing NBBO midpoint are classified as buyer-initiated, those below as seller-initiated, and trades at the midpoint are excluded. We employ consolidated trade and quote data from the TAQ database on WRDS, and follow Chordia, Roll, and Subrahmanyam (2000) in using the National Best Bid and Offer (NBBO) as the reference quote.

[Insert Fig. 1]

Our model characterizes price impacts as linear functional forms (see Glosten and Harris, 1988) and hence the price impact model assumes that transaction prices are subject to both permanent (informational) and transitory (non-informational) effects. Let $D_t \in \{+1, -1\}$ indicate trade direction at event time t, where +1 denotes a buyer-initiated trade and -1 a seller-initiated trade, with V_t the corresponding trade size. Let m_t denote the market maker's expectation of the security's fundamental value at time t. We specify the evolution of this value given as

$$m_t = m_{t-1} + C(D_t - \mathbb{E}_{t-1}[D_t]) + \lambda(D_t V_t - \mathbb{E}_{t-1}[D_t V_t] + y_t$$

where C is the fixed permanent component of price impact, λ is the variable permanent component (per share traded), and y_t is a public information shock. This formulation captures for the unexpected order flow, adjusting for autocorrelation using a five-lag AR process given as

$$D_t V_t = \alpha_0 + \sum_{j=1}^{5} \alpha_j D_{t-j} V_{t-j} + \epsilon_t$$

Assuming $\epsilon_t \sim N(0, \sigma_\epsilon^2)$, the expected trade direction can be given as

[7] If we take $\frac{\lambda\sigma^2}{\eta'}\zeta_j = \kappa^2\zeta_j$, the above expression simplifies into $\zeta_k = \frac{1}{\kappa^2\tau^2}(\zeta_{k-1} - 2\zeta_k + \zeta_{k+1})$

[8] The date chosen by us is September 1st, 2022 and we estimate the price impact for all stocks on this particular day of the year.

$$\mathbb{E}_{t-1}[D_t] = 1 - 2\phi\left(\frac{-\mathbb{E}_{t-1}[D_tV_t]}{\sigma_\epsilon}\right),$$

Where $\phi(.)$ denotes the cumulative standard normal distribution. Next, observed transaction prices p_t are assumed to reflect transitory costs ζ and l where ζ represents the fixed transitory price impact and l is the variable transitory component impact per share. Taking first differences and substituting the components, we estimate

$$\Delta p_t = C(D_t - \mathbb{E}_{t-1}[D_t]) + \lambda(D_tV_t - \mathbb{E}_{t-1}[D_tV_t]) + \zeta D_t + \ell\, D_tV_t + \eta_t$$

With η_t, a composite error term incorporating both pricing errors and microstructure noise. This regression is estimated for each stock using OLS, with corrections for serial correlation in the error term.

[Insert Table 1]

4. Stochastic Control in Dynamic Execution

Let us now proceed to consider a more generalized case of the previous section. Take an informed trader's initial position over the time interval [0, T] modelled by the process $q_t,\ \forall t \in [0, T]$ with the dynamics given by $q'(t) = dq_t = v_t dt$. Here, v_t represents the trading velocity which in our case is deterministic. The aim of this control problem is to find a solution that maximizes an objective function by altering the control variable. Hence, the number of shares available at each time period is modelled as a function of the trade velocity and time[9].

Now, we model the price of the stock as a stochastic process given by $dS_t = \sigma dW_t + k\, v_t dt$ such that the price (i.e. the midpoint of the spread) is a function of the trading velocity as well as a Weiner process with the parameter k given by the permanent price impact of the trade. It's important to note here that the actual execution cost may/may not be a linear function of only the trader's volume but could also be related with the existing market volume of other agents as well. Hence, the market volume process $V_t, \forall\, t \in [0, T]$ may be taken to be a deterministic continuous process which is bounded. The execution price may be directly dependent on the transaction volume and inversely dependent on the market volume, a argument that is not just logical but practical (see Gueant, 2016) . This results in the execution price being modelled by

$$S'_t = S_t + g\left(\frac{v_t}{V_t}\right), \qquad g(0) = 0$$

[9] . It may now be fair to discuss the three important conditions that need to be met for any control to be admissible- First, the control variable v_t must be progressively measurable. In other words, a decision made at time t is only possible with information available until till such a time. Second, such a process must also satisfy an unwinding constraint $s.t \int_0^T v_t dt = -q_0$. Such a condition may be logical as the trader must completely liquidate their position by the end of their trading endeavor. Since the initial holding is given as q_0, $q_t = q_0 + \int_0^T v_t dt$ and at time $t = T$, the condition of q_t at time $t = T$ may be zero shares.Last, process must also be bounded by the trading volume so it may not be possible to liquidate an unjustifiably large position of shares at any given point in time.

[Insert Fig. 2]

Consider that the informed trader executes a small position of $q'(t) = v_t$ at an execution price of S'_t, then his signed profit (negative for buying shares as one may pay to buy and get paid to sell) as a function of the trade velocity can be expressed as

$$dX_t = -S'_t dq = -S'_t v_t dt$$

For a total of N shares, the execution cost may be given by the function $L(N) = Ng(N), L(0) = 0$. In other words,

$$L: \ dX_t = -\left(S_t + g\left(\frac{v_t}{V_t}\right)\right) v_t dt$$

or

$$L: dX_t = -S_t v_t dt - V_t \, L\left(\frac{v_t}{V_t}\right) dt$$

For a more practical scenario, we may choose for L, a strictly convex power function somewhat like $L(x) = \eta|x|^{1+k}, k > 0$, but for the standard Almgren-Chriss framework, we choose $L(x) = \eta x^2$. To solve this, consider CARA function, a form that could be represented by $U = \mathbb{E}[-\exp(\gamma X_T)]$ where gamma is the absolute risk aversion coefficient. We now consider a case where the admissible control uses a deterministic strategy[10]. Now, to derive the optimal solution[11] for such an optimization problem, first we have the trader's deterministic profit given by $x = \int_0^T dX_t$

$$\int_0^T dX_t = X_0 - \int_0^T S_t v_t dt - \int_0^T L\, V\left(\frac{v_t}{V_t}\right) dt$$

$$X_T = X_0 + q_0 S_0 - \frac{k}{2} q_0^2 + \sigma \int_0^T q_t dW_t - \int_0^T L\, V\left(\frac{v_t}{V_t}\right) dt$$

Considering the trader's profits are normally distributed[12], the expected value and variance of the cashflows are given by

$$\mathbb{V}[X_T] = \sigma^2 \int_0^T q_t^2 dt$$

$$\mathbb{E}[X_T] = X_0 + q_0 S_0 - \frac{k q_0^2}{2} - \int_0^T V_t L\left(\frac{v_t}{V_t}\right) dt$$

What may be important here is to understand that the final term in the expected value corresponds to the dynamic order execution costs which in the AC framework are modelled by the transitory

---

[10] One may argue that the execution strategy which is optimal may not necessarily be deterministic but could be stochastic as well. We completely agree to such an argument but on a deeper examination, it may be possible to show that such a strategy is in fact deterministic as well optimal (see Gueant, 2016)

[11] One must remember the fact that $dS_t = \sigma dW_t + k\, v_t dt$ has the solution -

$$S_t = S_0 + \sigma W_t + k \int_0^t v_t dt = S_0 + \sigma W_t - q_t$$

Hence, we may apply the same in the integral to obtain our result.

[12] We again agree that this may be a heavy assumption but for the sake of simplicity and demonstration it may be important to consider such a assumption for the results obtained hold a strong intuition.

price impacts. Gueant (2016) shows that such a problem may be given by a Hamiltonian system shown by

$$\begin{cases} p'(t) = \gamma\sigma^2 q^*(t) \\ q^{*\prime}(t) = V_t H'(p(t)) \\ q^*(0) = q_0 \\ q^*(T) = 0 \end{cases}$$

In the case of quadratic execution costs, we consider a quadratic function $L(\rho) = \eta\rho^2$ then from the Hamiltonian system, we have $H(\rho) = \frac{\rho^2}{4\eta}$ [13] or the system reduces to the form of

$$\begin{cases} p'(t) = \gamma\sigma^2 q^*(t) \\ q^{*\prime}(t) = \frac{V_t}{2\eta} p(t) \\ q^*(0) = q_0 \\ q^*(T) = 0 \end{cases}$$

If $(V_t)_t$ is assumed as constant, (i.e. $V_t = V, \forall t \in [0, T]$) , then we obtain $q^*(t)$ as the classic result given by

$$q^*(t) = \frac{q_0 \sinh\left(\sqrt{\frac{\gamma\sigma^2 V}{2\eta}}(T-t)\right)}{\sinh\left(\sqrt{\frac{\gamma\sigma^2 V}{2\eta}}T\right)}$$

With the trade velocity expressed as

$$q^{*\prime}(t) = -\frac{q_0\sqrt{\frac{\gamma\sigma^2 V}{2\eta}}\cosh\left(\sqrt{\frac{\gamma\sigma^2 V}{2\eta}}(T-t)\right)}{\sinh\left(\sqrt{\frac{\gamma\sigma^2 V}{2\eta}}T\right)}$$

[Insert Fig. 3]

An important consideration in the context of optimal execution strategies is the convexity of the unwinding process, denoted as q^*, which implies that the liquidation process is initially fast but progressively decelerates. Within such a framework, the trader can afford to trade aggressively in the early stages of execution, as this reduces market impact when the position is large. However, as liquidation continues and the position diminishes, the cost of further trade becomes increasingly sensitive to risk aversion, as the trader approaches the end of the liquidation horizon and seeks to avoid drastic price deviations. As the remaining position decreases, the risk aversion component grows in importance, and the trader becomes more cautious. The cost of trading aggressively towards the end of the execution period rises due to the heightened volatility associated with a

---

[13] H is the Legendre-Fenchel transform of the function L given as $H(\rho) = \sup_{p} p\rho - L(\rho)$. It can be shown that since L is strictly convex, such a Hamiltonian system can be deemed to be as correct.

smaller position and greater market impact. Hence, the trader adjusts the execution schedule to slow down, minimizing the potential for significant price deviations in the final stages of the trade. However, we show in the subsequent section that it may be very important for the informed trader to consider the market maker's learning process[14] in case his strategy is purely speculative[15]. Hence, for all practical scenarios, if a market maker were to conjecture such a strategy used by an informed trader, she would set the price to increase the informed trader's price impact.

5. Optimal Liquidation under Asymmetric Information

We now proceed to show the audience how the Kyle (1985) framework fits in the results derived by us although the purpose of such analysis conducted was different from that intended to. We extend the canonical Kyle model to a finite horizon with risk-averse agents and time-varying impact. To show this, we now consider a discrete-time market with trading periods indexed by $t = 1, \dots, T$. There is a security with terminal asset value $v \sim N(\mu, \sigma_v^2)$ and three rational profit maximizing agents namely (i) informed trader who has private information about the fundamental value of the asset and intends to trade over a finite horizon T, (ii) noise traders who submit random orders $u_t \sim N(0, \sigma_u^2)$ and are modelled as i.i.d ; and (iii) a risk averse market maker. The market maker observes the total order flow $y_t = x_t + u_t$ consisting of both the informed and uninformed trader's demand schedules. The informed trader conjectures that the market maker follows a linear price adjustment rule given by $p_t = p_0 + \lambda_t y_t + \mu$. For each $t \in T$, the informed trader submits a demand schedule q_t and subsequently noise traders a schedule u_t while the market maker observes the total order flow. The aim of the informed trader in this model also remains the same – Maximize profits while minimizing transaction costs for we consider his intentions to be speculative in nature. He may also want to reduce the risk of variance between the execution price and the fundamental value. Hence, the informed trader's profit maximizing objective function is given as $\max_{x_t}[E[\pi_t] - \gamma Var(\pi_t)]$ where x_t is the number of shares liquidated (which in our case is for a sell program) at time t. The market maker now sets the price p_t as the conditional expectation of the fundamental value based on the information set observed at time t. Hence, $p_t = \mathbb{E}[v|y] = p_0 + \lambda_t y_t + \mu$. However, this time the market maker conjectures that the informed trader trades linearly in the divergence from the price from its fundamental value, i.e. his trading strategy is of the form $x_t = \beta_t(v_t - p_{t-1})$ where β_t is his trading intensity in period *t*. In such a given strategy, it may be possible to show that the price impact of the informed trade schedule may not be constant but one that increases with time and is in fact given by

$$\lambda_t = \frac{\beta_t \Sigma_{t|t-1}}{\beta^2 \Sigma_{t|t-1} + \sigma_u^2}$$

, where $\Sigma_{t-1} = Var(v|y_1, \dots, y_{t-1})$is the uncertainty about fundamental value v after period $t-1$. This means that as and when the informed trader trades, he risks leaking information to the market

[14] Although in the real world as and when an informed trader trades, he keeps revealing his information to the market participants and in such an endeavor, the market maker trades in a fashion to limit the informed trader's liquidation process thereby acting as a brake to it.

[15] Yes, there exist cases where the trader trades for reasons not pertaining to speculation for example portfolio rebalancing or liquidity provisions. In all such (other) scenarios, it may not be the most important consideration for the informed trader (which in our case in not truly informed now for his intentions are non-speculative (see Glosten and Milgrom , 1985) in nature to track the signaling process of the market maker.

maker and revealing his identity as well as intentions. Let us solve this problem backwards from the terminal condition all the way till time $t = 0$. Consider at time $t = T$, the informed trader liquidates his shares giving him the profit $\pi_T = x_T(v_T - p_{T-1})$. The market maker observes an order flow $y_T = x_T + u_T$. Let us define the information gap at time $t = T$ as $g_T = v - p_{T-1}$, because the noise u_T has mean zero, the insider's expected profit at time T is given as

$$E(\pi_T) = x_T(g_T - \lambda_T x_T)$$

The informed trader's problem is now finding a demand schedule x_T such that the expected value of his profit can be maximized as well as variance of the profit minimized. Hence the informed trader's objective now becomes

$$\max_{x_T}\left((x_T(g_{T-1} - \lambda_T x_T) - \gamma\lambda_T^2 x_T^2 \sigma_u^2\right)$$

In a more realistic situation, the market maker widens the spread to increase the price impact of the informed trader (see, for example Easley and O'Hara , 1997) . Taking a first order condition, x_T yields

$$x_t^* = \frac{g_T}{2(\lambda_T + \gamma\lambda_T^2\sigma_u^2)}$$

and

$$\beta_T = \frac{1}{2(\lambda_T + \gamma_T^2\sigma_u^2)}$$

Notice that if $\gamma = 0$ (risk-neutral), $\beta_T = \frac{1}{2\lambda_T}$ and if $\gamma > 0$, β_T has a smaller value meaning the trader limits his variance exposure and is less aggressive. For time $t = 1, \ldots, T-1$, we conjecture that the insider's value function is quadratic in the gap $g_t = v_t - p_{t-1}$ i.e. $V_t(g_t) = A_t g_t^2$ with A_t to be determined. The Bellman equation for the value function can be given as

$$V_t(g_t) = \max_{x_t}(\mathbb{E}[\pi_t|\epsilon_t] - \gamma Var(\pi_T|g_t) + \mathbb{E}[V_{t+1}(g_{t+1}|g_t, x_t])$$

6. Intentionality and Liquidation – The Informed Trader's Case

We now solve the Bellman equation recursively by working backward from the terminal condition. At period T, the value function $V_T(g_t) = x_t(g_t - \lambda_T x_T)$ which is just the profit from trading and is not dependent on the variance of the implementation shortfall[16]. Reworking the value function, we have

$$V_t(g_t) = \max_{x_t}\left(x_t(g_t - \lambda_t x_t) - \gamma\lambda_t^2 x_t^2 \sigma_u^2 + \frac{g_{t+1}^2}{4(\lambda_{t+1} + \gamma_{t+1}^2\sigma_u^2)}\right)$$

This can be recursively solved backward to obtain the optimal solution at each time period with the parameter λ_t given by

$$\lambda_t = \frac{\beta_t \Sigma_{t|t-1}}{\beta_t^2 \Sigma_{t|t-1} + \sigma_u^2}$$

[16] It may be just to criticize us for this for we consider a situation where there is no information event that may contribute to a change in fundamental value of the security but on another note, this case may be considered as a special case of Easley and O'Hara (1997) where there is an informed trader who trades the difference of the fundamental value and the market price. But in doing so, he chooses to camouflage his order flow and hence trades each of the shares at a different time. We assume that during the liquidation process, the fundamental value of the security remains unchanged for if it were to change, the informed trader was to have knowledge of this (by definition) and now change his strategy. Hence, there exists a liquidation strategy only if the informed trader is sure of the mispricing in the security. We consider such an assumption reasonable.

[Insert Fig. 4]

The result is a dynamic equilibrium in which informed traders begin with cautious, information-concealing trades and gradually accelerate as the trading horizon contracts and the marginal value of secrecy diminishes. The market maker, in turn, updates prices more aggressively as more information is revealed through order flow, amplifying the informed trader's price impact. It is evident that in equilibrium, the informed trader does not exploit his informational advantage in a single period. Instead, he strategically disseminates his trades over multiple periods to leverage his private information and reduce the likelihood of detection by the market maker. This order-splitting behavior arises from the trade-off between maximizing expected profits and minimizing adverse price impact. However, as the end of the trading horizon nears, the opportunity to profit from private information contracts necessitates more aggressive execution. In a single-period Kyle model, all private information about v (in that period) is partially revealed in that one shot. If σ_u^2 is finite and the insider is risk-neutral, some mispricing will remain because the insider does not want to push the price all the way to her private signal as she faces a price impact cost. The posterior variance update formula for the market maker may be given $\Sigma_{t|t} = \Sigma_{t|t-1} - \lambda_t$. Hence, the magnitude of the market maker's learning depends on β_t as well as σ_u^2. When σ_u^2 is large, it's harder for the market maker to separate informed trades from noise, so learning is slower. It thus is a very logical argument to conjecture that the informed trader chooses to trade in a scenario where there is a high probability of noise trading (see Weston, 2001).

Conversely, when β_t is large (the insider is more aggressive), more information is conveyed in the order flow, so posterior variance drops faster. Each time the market maker observes an aggregate order flow, she updates her beliefs about v_t. This reduces uncertainty about v_t. As Σ_t shrinks, the informational advantage held by the insider is eroded, effectively forcing him to realize a larger share of his remaining information before they are fully revealed into prices[17]. In the Kyle (1985) framework, the informed trader's central objective is to liquidate a position while minimizing the extent to which private information is revealed to the market maker, who infers the asset's fundamental value from observed order flow. The strategic challenge faced by the informed trader is thus to execute trades in a manner that does not immediately or excessively signal informational advantage, as this would prompt the market maker to adjust prices unfavorably. Early aggressive trading would directly signal private information about the asset's value, prompting the market maker to adjust their beliefs and set the price accordingly. By trading more cautiously at the beginning, the informed trader prevents large price movements and keeps the market maker's estimate of the asset value relatively stable. As the liquidation period progresses and the trader's position nears completion, the concern over signalling diminishes, especially if the time remaining for unwinding is short. In this context, the informed trader faces diminishing returns from delaying execution, as the window to manipulate the price without revealing information becomes narrower. Thus, the trader may accelerate trading towards the end of the liquidation period to minimize the time required to exit the position, reducing the risk of market impact from subsequent trades. Additionally, the trader's concern about preserving secrecy wanes, as any

[17] Another interesting feature here is that in equilibrium, since $Cov(v_t, y_t) \neq 0$, the posterior variance is strictly smaller than the prior.

residual information asymmetry becomes less significant relative to the remaining size of the position to be liquidated. Unlike the results discussed in the previous section—where price impact is typically treated as exogenous and time-invariant—this Kyle setting endogenizes market impact as a function of order flow informativeness. The market maker continuously updates their beliefs, dynamically adjusting the price impact parameter in response to perceived information leakage. If the informed trader observes a rising price impact over time, it may signal that the market maker is successfully inferring their private information. In such scenarios, the trader must adapt by tempering their trading intensity to reduce further leakage.

7. Conclusion

Our hybrid model thus highlights a fundamental difference which necessitates a reinterpretation of optimal execution as a problem not merely of cost minimization, but of belief management in an adversarial inference environment. In cases where the informed trader is certain their information is not being inferred (e.g., when trading for non-informational reasons such as portfolio rebalancing), a more aggressive strategy may remain viable. However, when the trader is acting on genuinely private information—such as knowing the asset is overvalued relative to fundamentals, the necessity of preserving informational advantage becomes paramount. Under such conditions, the trader may resort to obfuscation strategies, such as mimicking noise trader behavior, to disguise their intentions and mitigate the adverse effects of informational leakage. Thus, the informed trader's optimal strategy is conditional not only on market microstructure frictions but also on the purpose of trading and the degree of informational asymmetry. When trading is motivated by private information, the trader must adopt a cautious, adaptive execution schedule that balances liquidation efficiency with the risk of detection by an inference-driven market maker. This assumption in this model is rather very relaxed because this is how the results seem to point to such behavior. It may be possible to model a more complex dynamic for such existing scenarios but the intuition behind such an endeavor remains the same. How different parameter changes affect this model should be an interesting task for future research, but such an analysis may be reserved for another fruitful day.

References

Admati, A. R., & Pfleiderer, P. (1988). Selling and trading on information in financial markets. *The American Economic Review, 78*(2), 96–103.

Almgren, R., & Chriss, N. (2000). Optimal execution of portfolio transactions. *Journal of Risk, 3*(2), 5–39.

Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets, 5*(1), 31–56.

Amihud, Y., & Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics, 17*(2), 223–249.

Amihud, Y., Mendelson, H., & Pedersen, L. H. (2012). *Market liquidity: Asset pricing, risk, and crises*. Cambridge University Press.

Bertsimas, D., & Lo, A. W. (1998). Optimal control of execution costs. *Journal of Financial Markets, 1*(1), 1–50.

Bessembinder, H. (1994). Bid-ask spreads in the interbank foreign exchange markets. *Journal of Financial Economics, 35*(3), 317–348.

Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies, 22*(6), 2201–2238.

Chordia, T., Roll, R., & Subrahmanyam, A. (2001). Market liquidity and trading activity. *The Journal of Finance, 56*(2), 501–530.

Dubil, R. (2002). Optimal execution of portfolio transactions with trading-costs and value-at-risk constraints. *International Review of Economics & Finance, 11*(2), 147–166.

Foster, F. D., & Viswanathan, S. (1990). A theory of interday variations in volumes, variances, and trading costs in securities markets. *Review of Financial Studies, 3*(4), 593–624.

Foster, F. D., & Viswanathan, S. (1993). Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models. *The Journal of Finance, 48*(1), 187–211.

Glosten, L. R., & Harris, L. E. (1988). Estimating the components of the bid/ask spread. *Journal of Financial Economics, 21*(1), 123–142.

Guéant, O. (2016). *The financial mathematics of market liquidity: From optimal execution to market making*. CRC Press.

Guéant, O., Lehalle, C. A., & Tapia, M. (2012). Dealing with the inventory risk: A solution to the market making problem. *Mathematics and Financial Economics, 7*(4), 477–507.

Hasbrouck, J. (1991). Measuring the information content of stock trades. *The Journal of Finance, 46*(1), 179–207.

Hasbrouck, J., & Seppi, D. J. (2001). Common factors in prices, order flows, and liquidity. *Journal of Financial Economics, 59*(3), 383–411.

Hisata, Y., & Yamai, Y. (2000). Optimal execution of portfolio transactions with market impact. *Monetary and Economic Studies (Bank of Japan), 18*(3), 65–94.

Jain, P. C., & Joh, G. H. (1988). The dependence between hourly prices and trading volume. *Journal of Financial and Quantitative Analysis, 23*(3), 269–283.

Kirchner, C., & Schlag, C. (1998). The performance of informationless trading strategies in a limit order market. *European Financial Management, 4*(2), 225–242.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica, 53*(6), 1315–1335.

Kyle, A. S. (1989). Informed speculation with imperfect competition. *Review of Economic Studies, 56*(3), 317–356.

Lee, C., & Ready, M. J. (1991). Inferring trade direction from intraday data. *The Journal of Finance, 46*(2), 733–746.

Sadka, R. (2006). Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. *Journal of Financial Economics, 80*(2), 309–349.

Weston, J. P. (2001). Competition on the Nasdaq and the impact of recent market reforms. *The Journal of Finance, 56*(6), 2565–2598.

Figures and Tables

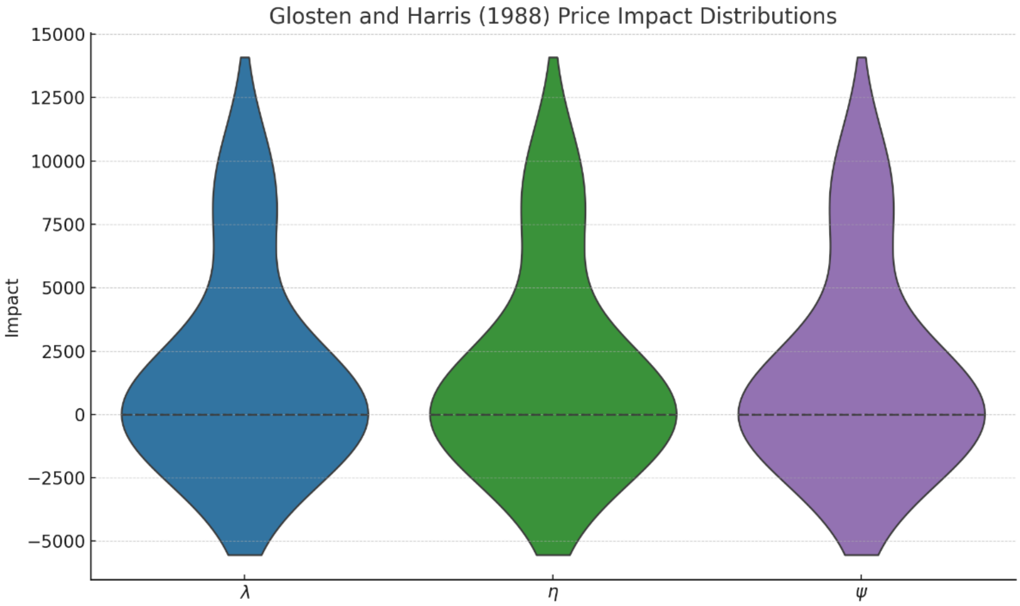

Fig. 1- Distributions of Estimated Price Impacts under the Glosten and Harris (1988) Model.
The violin plot displays the cross-sectional distributions of three key price impact components—permanent variable (λ), transitory variable (η), and permanent fixed (ψ)—as estimated from firm-level trade data on September 1, 2022. The impacts are derived from the extended Glosten and Harris (1988) model using TAQ data from NYSE and NASDAQ stocks. Each distribution reflects heterogeneity in firm-level execution costs, with fat tails and skewness indicating the presence of large idiosyncratic effects in trade-related price adjustments. All values are winsorized at the 1st and 99th percentiles to mitigate outlier distortion.

Fig. 2 - Pairwise Relationships among Glosten and Harris (1988) Price Impact Components

The above matrix of scatter plots with fitted linear trends illustrates the empirical correlations between the four estimated price impact coefficients—transitory fixed (ε), transitory variable (η), permanent variable (λ), and permanent fixed (ψ)—across firm-level data. The analysis is based on NYSE and NASDAQ stocks using TAQ data from September 1, 2022. The plots reveal interdependencies, including a negative correlation between ψ and λ, suggesting compensatory roles of fixed and variable permanent components in price discovery. Overall, the weak correlations involving η point to its higher idiosyncratic variation with structural price impact terms.

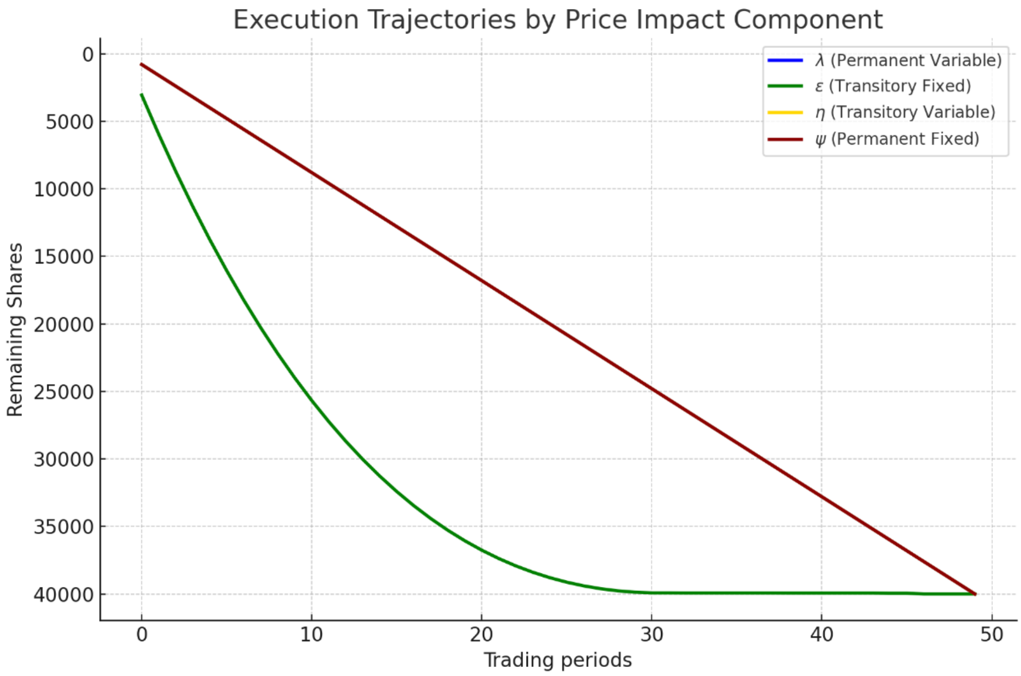

Fig 3. - Execution Trajectories by Impact Components

The above figure illustrates optimal execution paths under each estimated component of the Glosten and Harris (1988) price impact model. The trading trajectories are derived under isolated influence of the four distinct impact coefficients: permanent variable impact (λ), transitory fixed impact (ζ), transitory variable impact (ℓ), and permanent fixed impact (C). Each trajectory shows the optimal remaining inventory across 50 discrete trading periods for an initial position of 40,000 shares, minimizing execution cost based on the respective price impact parameter. The steeper, convex path under ζ reflects rapid early liquidation to minimize fixed transitory costs, while the linear trajectories under λ and C indicate uniform execution rates under constant marginal costs. All simulations assume absence of information asymmetry and zero risk aversion.

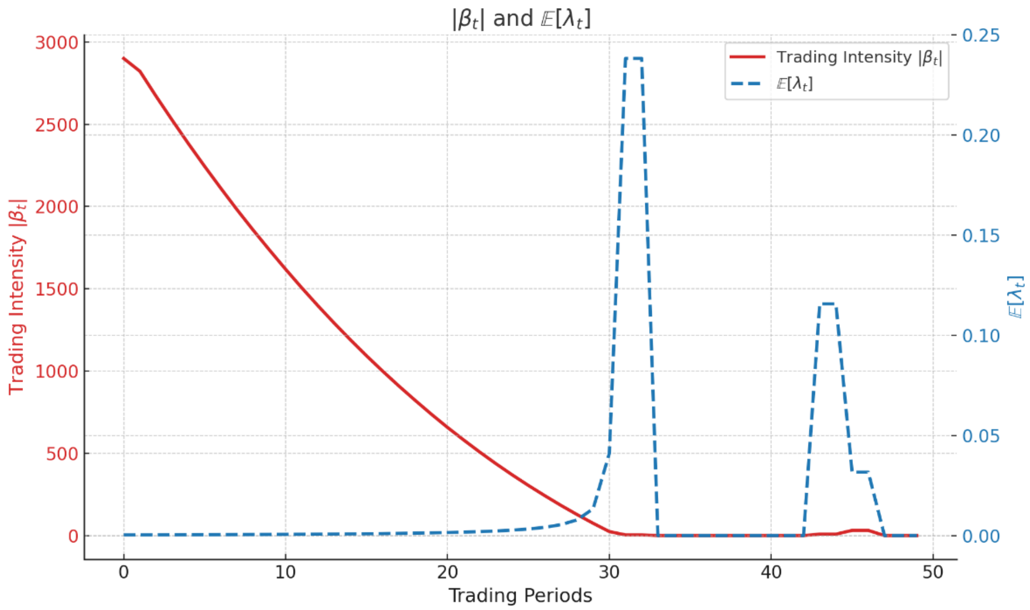

Fig. 4 - Time-Variation in Trading Intensity

The above figure compares the evolution of absolute trading intensity $|\beta t|$ and the conditional expectation of permanent variable price impact $\mathbb{E}[\lambda_t]$ simulated across 50 discrete trading periods. The left y-axis (red line) shows the optimal trading intensity derived from minimizing execution costs under dynamic market impact, while the right y-axis (blue dashed line) reflects the expected value of λ_t . Such an inverse relationship between $|\beta t|$ and $\mathbb{E}[\lambda_t]$ illustrates strategic liquidity timing, where traders decelerate execution when market depth deteriorates. Spikes in $\mathbb{E}[\lambda_t]$ signal adverse price impact conditions that discourage aggressive trading.

Table1 - Estimated Price Impacts

Price impacts are estimates of the Glosten and Harris (1988) model

$$\Delta p_t = C(D_t - \mathbb{E}_{t-1}[D_t]) + \lambda(D_t V_t - \mathbb{E}_{t-1}[D_t V_t]) + \zeta D_t + l\, D_t V_t + \eta_t$$

where Δp_t denotes the change in transaction price, $D_t \in \{-1,1\}$ indicates buyer- or seller-initiated trades,V_t is trade size and $\eta_t$captures residual microstructure noise and public information events. The coefficients C and λ represent fixed and variable permanent price impacts, while ζ and ℓ represent fixed and variable transitory components. The model is estimated using consolidated TAQ data for NYSE and NASDAQ stocks on September 1, 2022. Reported values are time-series means of stock-level cross-sectional estimates.

Variable	*Obs.*	*Mean*	*Std.*	*Min*	*Max*
C	8,549	0.00	0.01	-0.20	0.33
$\lambda(\times 10^6)$	8,549	1.11	71.30	-749.40	913.70
ζ	8,549	0.02	0.03	-0.17	0.31
$\ell\ (\times 10^6)$	8,549	-0.19	55.40	-450.00	427.80