

Unwinding Stochastic Order Flow: When to Warehouse Trades

Marcel Nutz*

Kevin Webster[†]

Long Zhao[‡]

October 21, 2023

Abstract

We study how to unwind stochastic order flow with minimal transaction costs. Stochastic order flow arises, e.g., in the central risk book (CRB), a centralized trading desk that aggregates order flows within a financial institution. The desk can warehouse in-flow orders, ideally netting them against subsequent opposite orders (internalization), or route them to the market (externalization) and incur costs related to price impact and bid-ask spread. We model and solve this problem for a general class of in-flow processes, enabling us to study in detail how in-flow characteristics affect optimal strategy and core trading metrics. Our model allows for an analytic solution in semi-closed form and is readily implementable numerically. Compared with a standard execution problem where the order size is known upfront, the unwind strategy exhibits an additive adjustment for projected future in-flows. Its sign depends on the autocorrelation of orders; only truth-telling (martingale) flow is unwound myopically. In addition to analytic results, we present extensive simulations for different use cases and regimes, and introduce new metrics of practical interest.

Keywords Internalization; Optimal Execution; Price Impact; Central Risk Book; Market Making

AMS 2010 Subject Classification 91G10

JEL Classification G24; C6

1 Introduction

This paper aims to model and solve the unwind problem for stochastic order flow; that is, an optimal execution-type problem where instead of a fixed position known at the initial time, the trading desk receives a dynamic flow of orders to be unwound.

Our work is motivated by the emergence of central risk books (CRBs) in most investment banks and several large trading firms over the last few years. The CRB is a trading desk aiming to minimize transaction costs by aggregating internal order flows. In an ideal situation, all other units within the organization (and possibly the asset class) interact exclusively with the CRB, instead of trading directly in the market. In [15, p. 2], CRBs are defined as “the places that (in theory) centralize the execution of trades for the whole bank.” As the total volume resulting from such centralization is very large,¹ transaction costs are of paramount importance. Orders across different business lines and strategies can have idiosyncratic directions and timing, leading to a stochastic order flow (*in-flow* orders from the desk’s perspective). For instance, an options desk’s delta hedging orders

*Depts. of Statistics and Mathematics, Columbia University, mnutz@columbia.edu. Research supported by NSF Grants DMS-1812661 and DMS-2106056.

[†]Dept. of Mathematics, Columbia University.

[‡]Dept. of Statistics, Columbia University, long.zhao@columbia.edu.

¹While public data seem to be unavailable, an equity CRB at a large institution would process a nontrivial percentage of the total trading volume in US equities, currently about \$500 billion daily [23].

are likely independent of prime brokerage clients' block trades. The CRB nets opposite in-flows against one another, a process named *internalization*. Remaining orders can either be routed to the market (*out-flow* orders), a process named *externalization*, or *warehoused*, meaning that the CRB temporarily keeps the position on its book. Out-flow orders incur transaction costs including transient and permanent price impact. The most favorable scenario for a warehoused order is to be netted against an opposite order later on, thus avoiding transaction costs. Even disregarding that possibility, warehousing can be beneficial to reduce costs; for instance, immediately externalizing a sizable block order would often lead to an unnecessarily large price impact.

Our model is also relevant to market makers. Similarly to banks, market makers face orders from clients ranging from retail traders to institutional funds, resulting in stochastic in-flow.² The distinction between in-flow and out-flow is most evident in foreign exchange (FX), where clients send orders directly to dealers, and dealers unwind their remaining inventory on a central venue (e.g., EBS), also called the lit market. The Bank of England's Fair and Effective Markets Review [7, p. 11] stressed the importance of internalization in this context: "Market participants have indicated that dealers with large enough market share now internalise up to 90% of their client orders."³ The market maker must decide which orders to warehouse and which to unwind, an aspect emphasized recently by [22]. Also recently, [10] suggested market makers leverage the same mechanisms as CRBs to unwind sizable inventory considering price impact. The present paper does not discuss how the market maker should set quotes, usually considered the core problem in the literature on market making (see [21] and the references therein). Our model is relevant to market makers specifically in the regime where the inventory gets too large to revert passively through limit orders. Then, market makers use dedicated algorithms to reduce the inventory aggressively. As the inventory simultaneously varies due to in-flow orders, this constitutes an unwind problem for a stochastic flow.

To the best of our knowledge, this is the first study to focus on the unwind problem for stochastic order flow. Compared with the extensive work on optimal execution, two essential differences arise. First, the final volume to be executed is unknown at the initial time. Instead, the in-flow of orders is a stochastic process and the unwind strategy must take into account projected future orders. Second, there are opportunities for orders to net, which introduces internalization benefits not directly modeled in the standard setting. Compared with the literature on market making, our in-flow is uncontrolled. Instead of studying quotes, we focus on minimizing the transaction costs of unwinding a potentially large inventory subject to stochastic shocks.

The study closest to ours, in terms of methodology and focus on transient price impact, is the concurrent work [41] on pre-hedging. In their model, a dealer learns about a potential trade through a client's request for quote. The probability of winning the trade depends on the dealer's quote, and there is uncertainty about the timing of the potential transaction. The authors determine the optimal pre-hedge strategy: the dealer starts to build inventory ahead of the potential trade, taking into account transaction costs and the expected size of the trade. Clearly, the potential trade plays a role similar to our stochastic in-flow. On the other hand, as there is only a single trade, internalization does not occur in their model. At the opposite extreme, [16] consider internalization in a queueing model. The dealer can skew their prices to encourage inventory-reducing customer

²Some market makers also use the name CRB. For instance, CEO D. Cifu of global market maker Virtu emphasized internalization opportunities at a Goldman Sachs Financial Services Conference [25]: "we're the only firm in the world that has a central risk book of retail order flow, prop flow, and now institutional agency customers. [...] So there's a real opportunity to be a real internalizer between the buy side and the retail side."

³The report further highlighted that interdealer trading as a share of overall FX market turnover declined by one third since the late 1990s, and attributed this to the growth of internalization [7, p. 59]. For reference, the Bank for International Settlements [11] estimated that interdealer trading in FX accounted for 1 trillion USD *daily*.

flow, but there is no option to externalize. While this model is quite different from ours, [16] gives insightful comments and data sources about internalization in the FX market. The more recent work [8, 9] combines internalization and externalization. Here a dealer sets quotes to attract flow and simultaneously hedges in a separate liquidity pool. The authors highlight the optimal behavior: pure internalization when inventory is small, aggressive externalization when inventory is large. The liquidity pool in this model bears instantaneous (and permanent impact) cost, but not the transient impact cost that is the main driving force in our study. A related equilibrium model is formulated in [6]. This model features clients (liquidity takers) and dealers who absorb their demands. Dealers trade competitively in the interdealer market and also in an open market subject to instantaneous transaction costs (but no impact cost). The study solves for the equilibrium price and highlights the endogenous price impact incurred by the clients. An early contribution with stochastic inventory shocks is [19] which considers optimal execution with market and limit orders. Limit orders are executed at random times, leading to stochastic shocks in the inventory. Another paper close to ours in terms of application is [22], where a foreign exchange market maker streams bespoke quotes to an informed and a noise trader, and also trades in the lit market. The market maker learns the signal of the informed trader from his trades and uses it to decide on quotes and externalization of her order flow, and on speculative trades. In particular, the model derives an unwind strategy for a specific stochastic order flow that is endogenous to the equilibrium.

By contrast, our approach is to model the stochastic order flow exogenously in a reduced but general form. That will allow us to analyze the trading metrics of the corresponding optimal strategy from an input–output perspective and tackle some of the key trading questions. We do not model incentives the desk can provide their clients to induce orders or the information the desk can glean from the clients’ orders. In particular, the desk we model does not engage in speculative trades considering client information.

This paper is organized as follows. In the remainder of the Introduction, Section 1.1 summarizes a selection of the high-level results, while Section 1.2 reviews background and modeling considerations around transaction costs and order flow. Section 2 details the mathematical model, then analytically derives the solution and some of its qualitative properties, with most proofs being reported in Appendices B and C. Section 3 defines a number of trading metrics and presents simulation studies on how they react to autocorrelation and volatility of the in-flow. Further numerical results are reported in Appendix A: relation to classical optimal execution; sensitivities to the spread cost parameter, the martingale driving the in-flow process, and the initial impact state; and finally an empirical study on autocorrelation of orders on the public trading tape. Appendix B gathers the proofs based on the dynamic programming approach while Appendix C proceeds through the stochastic maximum principle.

1.1 Synopsis

We propose and solve a flexible model for unwinding an exogenously given order flow, rich enough to serve as a practical order-scheduling algorithm. In particular, it allows for order flows with momentum or reversal, time-varying volatility and block trades. Moreover, it takes into account impact and spread costs with possibly time-varying liquidity parameters.

We model price impact by a generalized Obizhaeva–Wang model while instantaneous (spread) costs during limit order book (LOB) trading are modeled in reduced form by a quadratic penalty on trading speed. The desk thus trades continuously during the day, but also places block orders in the opening and closing auctions which only carry impact cost. The desk’s inventory must be flat after the close. See also Section 1.2 for further modeling background and terminology.

It turns out that our model can be transformed into a linear-quadratic stochastic control problem,

which is key to its tractability. The optimal intraday trading speed q_t takes the form

$$q_t = f_t X_t + g_t Y_t + h_t Z_t$$

where the time-varying coefficients f_t, g_t, h_t are explicitly determined through an ordinary differential equation (ODE), and even provided in closed form for the case of constant liquidity parameters. The coefficients describe the desk’s reaction to the current inventory X_t , the impact state Y_t , and the cumulative in-flow Z_t . We will show analytically that $h_t = 0$ when there is no in-flow after the initial position, meaning that $q_t = f_t X_t + g_t Y_t$ would be the solution in a deterministic optimal-execution framework. The additional term $h_t Z_t$ captures forecasted future in-flows. The sign of h_t depends on whether the in-flow exhibits momentum or reversal. For momentum, the adjustment leads to faster unwinding as more same-sided orders are expected; this result supports the “overtrading” of orders often seen in practice. The adjustment is zero if the in-flow is a martingale, thus justifying traders’ intuition that “truth-telling” flows can be executed myopically. The tractability of the model allows us to rigorously study various other properties by analytic means (see Section 2). This is useful not only to understand the general dependencies of the output, but also for a desk to recognize, e.g., if its algorithm incurs excessive model risk by chasing an arbitrage implicit in its liquidity estimates.

Transaction costs such as impact costs are difficult to study empirically because the unaffected price (that would have prevailed without the transaction) cannot be observed. Instead, key metrics are inherently model-dependent. Having a consistent, computable model for general in-flows opens the door to analyzing cost metrics side-by-side with model-independent quantities like the internalization rate. The most essential trading question is how unwind strategy and cost metrics depend on the characteristics of the in-flow; namely, autocorrelation and volatility. In our numerical simulations, we observe that all core trading metrics strongly depend on these characteristics, in line with the industry mantra, *know your client*. Some key insights are:

- Expected trading costs (per order notional) are minimized at a particular in-flow volatility. Below that threshold, additional volatility increases the internalization rate, whereas above the threshold, internalization plateaus and further volatility necessitates more aggressive trading, driving up costs. As a take-away, a desk may benefit from incentivizing additional order flow to increase internalization up to the threshold.
- Autocorrelation strongly affects all core trading metrics. Momentum requires more aggressive trading, increasing both impact costs and spread costs. The ratio between those costs is approximately stable. Reversion leads to more internalization and more warehousing until the close; the closing trade represents a larger fraction of the total trading volume.
- Our misspecification analysis shows that it is preferable to err on the *optimistic* side when estimating autocorrelation: overestimating momentum sharply increases costs by way of overly aggressive trading and missed netting opportunities, whereas trading too slowly only incurs a moderate additional cost by trading too much on the close. Misspecifying in-flow volatility bears no transaction cost since the unwind strategy’s feedback form is independent of volatility; only the calculation of expected costs is affected.

Section 3 and Appendix A present numerous further simulation studies, and also introduce new metrics of interest to practitioners.

Numerical experiments highlight that our model encompasses a wide spectrum of regimes, with scenarios resembling classical optimal execution and market making occurring as extreme cases. The regime depends on the realization of the random in-flow trajectory. To illustrate this, Figure 1 shows

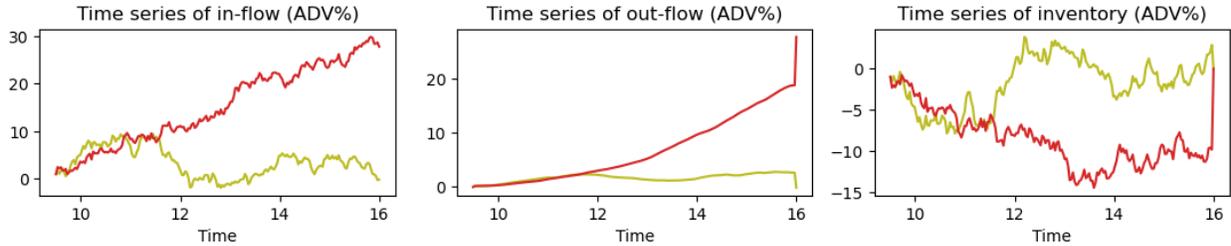


Figure 1: Two realizations of the same model illustrating extreme regimes.

two particular realizations from the same (martingale) in-flow process and the model’s corresponding optimal unwind strategy. One realization is closer to classical optimal execution: Most orders have the same sign and the strategy unwinds aggressively to reduce the inventory. Internalization occurs and leads to a smoothing between in and out-flow, but is not sufficient to control the inventory. Transaction costs are dominated by price impact. Such scenarios are frequent when the initial position is large compared to in-flow volatility, and when in-flow exhibits momentum. The other realization resembles a favorable market-making scenario: orders approximately cancel, the unwind strategy trades slowly most of the time, and the inventory never gets large. Spread costs are not negligible compared to impact costs. Such scenarios are frequent when in-flow volatility is large compared to the initial position, and when in-flow exhibits reversion. The relevance of models that can capture all regimes was previously emphasized by [16, p. 37]:

“An important message is that the polarisation between internalisers and externalisers is overly simplistic and that instead one needs to view liquidity providers across a continuum with passive internalisers at one end [...] and eventually externalisers at the other end.”

1.2 Background

Next, we discuss our model’s main ingredients. The trading metrics for its analysis are discussed in Section 3.1.

In-flow or *order flow*⁴ refers to the incoming orders from the desk’s clients. Practitioners propose stochastic processes for in-flow based on internal historical data. Depending on the organization, the clients, and the order options offered to them, the desk will observe different characteristics. The in-flow is generally stochastic, potentially with jumps if block orders are allowed, and there is a pronounced time-of-day effect: most orders arrive shortly after the open or before the close, as reflected in market volumes. For examples of intraday trading patterns, see [21, 13]. In-flow may exhibit momentum, for instance, if clients slice their orders to mask the total volume, or if they target momentum strategies. See [49] for an in-depth empirical analysis of market flow autocorrelation. Because splitting of meta orders is common in LOB trading, order flow on the public trading tape exhibits momentum, a finding that we confirm empirically for all S&P 500 stocks (see Section A.4). We emphasize that the situation might be quite different for a CRB with internal clients, say, which ideally (though maybe not in practice) deliver “truth-telling” (martingale) flow by being upfront about their volume. Less frequently, in-flow may exhibit mean reversion, for instance, if clients

⁴The term “order flow” is used differently in [20]. There, order flow represents the orders sent to the market by other market participants and the focus is on how those orders impact prices available during execution.

follow a mean-reversion strategy or if they tend to cancel orders mid-execution.⁵ Our model uses a stochastic process

$$dZ_t = -\theta_t Z_t dt + \sigma_t dW_t$$

with time-varying parameters to allow for a wide range of client scenarios. It exhibits momentum (reversion) if $\theta_t < 0$ (> 0). In practice, execution desks have ample historical data to determine suitable parameters and want to ensure their in-flow model reflects their client base. Indeed, different desks achieve different model parameters, with more sophisticated desks achieving milder order flow. For instance, a CRB with strong relationships may attract clients with less momentum. Another example is a low-latency market-maker: they could cancel limit orders and widen spreads when predicting more toxic (momentum) flow, effectively achieving a milder in-flow compared to the average market participant.

We use a Brownian motion driver $\sigma_t dW_t$ to keep the model readily understandable to a broader audience but show in Section 2.4 that the results extend to a general martingale driver dM_t with minor modifications, thus covering in-flows with jumps, self-excitement, etc. In the numerical experiments of Section 3, we use a driver of finite-variation so that trading metrics can be expressed “per order notional.”

Unwind trades or *out-flow* are the market-facing trades of the desk. We denote by Q_t the cumulative unwind trades up to time t ; this is the control variable in our model. An order scheduling algorithm typically uses block trades at the open and closing auctions and trades continuously on the limit order book during the day. The desk’s available actions may differ from the clients’ options for submitting orders. For example, clients might submit block orders intraday, leading to jumps in a CRB’s position that cannot be unwound immediately (at a reasonable cost). Another example of a client-triggered block trade is when a market maker’s resting order in a dark pool gets filled by a large mutual fund.

The desk’s *outstanding position* or *inventory* $X_t = Q_t - Z_t$ is the difference between the unwind trades and the in-flow. For simplicity, our model focuses on a single trading day and a desk that does not hold overnight positions. Hence, we impose the liquidation constraint $X_T = 0$ at the time horizon T corresponding to the end of the closing auction.

Transaction costs come in two forms: price impact and bid-ask spread. For CRBs servicing major financial institutions, trading costs are primarily driven by price impact, but spread costs are not negligible. For example, Nasdaq’s *The 2022 Intern’s Guide to Trading* [40] gives numbers for large-cap stocks of 30 basis points for price impact and 5 basis points for spread costs.

A price impact model generally consists of two elements. First, a push factor, also called Kyle’s lambda, describes a market fill’s immediate effect on the midprice. Second, a decay kernel describes how quickly this dislocation reverts in the absence of further fills. The interplay between push and reversion dictates the optimal trading speed. Empirically, traders observe a time-of-day effect for Kyle’s lambda: it is largest near the open and smallest near the close. See, e.g., [26] and [42]. See [48] for estimates of price impact in auctions.

The observed execution price P_t is thought of as the sum $P_t = S_t + Y_t$ of an “unaffected” price S_t (that would have prevailed without our trading) and a market impact process Y_t . Our model uses a generalized Obizhaeva–Wang (OW) model [45] with time-dependent coefficients,

$$dY_t = -\beta_t Y_t dt + \lambda_t dQ_t.$$

⁵The latter can occur if an execution desk allows for reductions of order size during execution but not for increases, as clients may then deliberately overstate the initial order and leverage the free option to reduce the order’s size.

Thus a trade of size δQ_t immediately dislocates the price by $\lambda_t \delta Q_t$, and the dislocation decays exponentially at rate β_t . This applies to block trades as well as to continuous trading. To the best of our knowledge, the generalized Obizhaeva–Wang model was first introduced in [27]. We use time-dependent but deterministic coefficients (in practice, these curves are typically fixed at the beginning of the day, based on data from previous trading sessions). A stochastic resilience β_t was considered in [33]. Both β_t and λ_t are stochastic in the formulation of [2] which in addition allows for a stochastic terminal condition. In principle, this allows to model stochastic in-flows, though the paper does not elaborate on that application. Other price impact models consider non-exponential kernels, as empirically studied in [14] and solved by [1], or non-linear price impact, as introduced by [14, 3] and studied by [30, 42, 35].

We do not model permanent impact (i.e., impact without decay). Because initial and terminal inventory are fixed, permanent impact only adds a constant to the cost and does not affect the optimal strategy (see [27, Proposition 3.3] for a detailed proof).

For costs related to the bid-ask spread, we use a straightforward model as in [5]: intraday trading incurs a quadratic instantaneous cost $\frac{1}{2} \varepsilon_t q_t^2$ on the trading speed $q_t = \dot{Q}_t$. While this forces intraday trading to be continuous, no such cost is charged for the auctions. We abuse the term “spread costs” in that this reduced-form cost stands in not just for the spread costs of market orders, but also for adverse selection incurred by limit orders and any other instantaneous costs that do not change the price in a persistent way.

If the desk nets orders of opposite sign instead of routing them to the market, it saves the associated spread costs entirely. We emphasize that the situation is different for impact costs, as this is often overlooked: Suppose a buy order pushes up the price. It incurs an impact cost, but a sell order immediately following it benefits from the increased price, so that the sum of the costs is approximately zero. In general, the cost saving due to netting is related to the impact decay that would have occurred between the trades.

To focus our study on the core internalization and externalization trade-off, and to provide results that are agnostic to the proprietary signals the desk has access to, we model the unaffected price S as a martingale. For considerations around toxic flows and proprietary trading, see for instance [22, 38].

2 Analytic Results: Optimal Strategy and Cost

This section contains our theoretical results. Section 2.1 details the mathematical formulation of the problem. Section 2.2 provides the optimal unwind strategy and its expected cost, for the case of in-flow driven by Brownian motion and possibly time-varying liquidity parameters. Section 2.3 presents the high-level steps of its derivation and further details; the technical proofs are reported in Appendix B. Section 2.4 generalizes the results to in-flows driven by a more general martingale (possibly with jumps, non-Markovian, etc.). On the other hand, Section 2.5 specializes to constant liquidity parameters and provides a closed-form solution for that case. The latter result is derived by a different approach than the rest (maximum principle instead of dynamic programming); the proof is reported in Appendix C.

2.1 Problem Formulation

We fix the time horizon $T > 0$ and bounded, measurable functions $\beta, \lambda, \varepsilon : [0, T] \rightarrow (0, \infty)$, $\theta : [0, T] \rightarrow \mathbb{R}$ and $\sigma : [0, T] \rightarrow [0, \infty)$, where ε, λ are also bounded away from zero. Moreover, λ is differentiable and its derivative $\dot{\lambda}$ is bounded. For convenience, we denote $\gamma_t := \log \lambda_t$. We also

introduce an additional value $\lambda_{0-} \in (0, \infty)$ representing the opening auction's liquidity parameter. We assume throughout that

$$\lambda_{0-} \leq \lambda_0 \quad \text{and} \quad 2\beta_t + \dot{\gamma}_t > 0, \quad (2.1)$$

which is trivially satisfied if the liquidity parameters are constant in time. We will see in Section 2.3 that these are no-arbitrage conditions, excluding profitable roundtrip trades and ensuring that our problem is convex. See [32] for estimates of price impact based on TAQ data and a comparison across continuous trading, closing auction, and opening auction.

Remark 2.1 (Price manipulation). Fast increases in liquidity can give rise to arbitrage, or *price manipulation*, as first noted by [36]: if one order pumps the price and a subsequent order of opposite sign has less price impact, a profitable roundtrip arises, unless the price has sufficiently reverted between the trades. We will see in Section 2.3 that (2.1) exclude such profitable (or free) roundtrips. For one of our comparative statics below (Proposition 2.3), we will require a strengthening of the second part of (2.1), namely that $\beta_t + \dot{\gamma}_t > 0$. This is related to additionally excluding “transaction-triggered price manipulation,” first studied by [4, 31]. See Remark 2.9 for further details, and Remark 2.10 for practical aspects. Note that the price impact parameter for the closing auction is $\lambda_T = \lim_{t \rightarrow T} \lambda_t$. A trader may be interested in specifying a smaller value for Kyle's lambda in the auction compared to LOB trading before the close. However, this would immediately give rise to price manipulation strategies. That may, of course, be an artifact of our simplified model for the auction—an area left for further investigation. (One could specify a value larger than λ_T , but this is not relevant in practice; see [32].)

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T]}, \mathbb{P})$ be a filtered probability space satisfying the usual conditions, endowed with a Brownian motion $W = (W_t)_{t \in [0, T]}$ and a càdlàg martingale $S = (S_t)_{t \in [0, T]}$. The process S is required to satisfy a minor integrability condition (see Footnote 11).

We define the in-flow process Z by $Z_0 = z \in \mathbb{R}$ and

$$dZ_t = -\theta_t Z_t dt + \sigma_t dW_t, \quad t \in [0, T]. \quad (2.2)$$

This process represents the aggregated order flow the desk faces. The sign convention follows the client's perspective; e.g., a positive increment corresponds to a buy order, causing a negative increment in the desk's outstanding position. The initial value z represents the orders present at the beginning of our problem.

Next, consider a càdlàg semimartingale $(Q_t)_{t \in [0, T]}$ with $Q_{0-} = 0$, representing the desk's cumulative unwind trades (out-flow). Here and below, certain processes are defined at time $t = 0-$ to allow for a jump at the initial time, in this case, a block trade $\Delta Q_0 = Q_0 - Q_{0-} = Q_0$ in the opening auction. Given Q , its price impact process Y is defined by $Y_{0-} = y$ and

$$dY_t = -\beta_t Y_t dt + \lambda_{t-} dQ_t, \quad t \in [0, T]. \quad (2.3)$$

Here and below, our convention is that integrals such as $\int_0^t \lambda_{s-} dQ_s$ are taken over the *closed* interval $[0, t]$, as for instance in [47]. In particular, a jump ΔQ_0 in the integrator leads to a jump of the integral at $t = 0$. In our model, Q will typically have block trades at the initial and terminal times, so that (2.3) gives rise to jumps in the impact process,

$$\Delta Y_0 = Y_0 - y = \lambda_{0-} Q_0 \quad \text{and} \quad \Delta Y_T = Y_T - Y_{T-} = \lambda_T (Q_T - Q_{T-}). \quad (2.4)$$

Thus, y represents the impact state when our problem begins (e.g., impact from the previous day's trading) and $\lambda_{0-} Q_0$ is the change in impact state due to the block trade in the opening auction.

In our model, the trading speed bears a quadratic cost on $(0, T)$. Hence unwind strategies (Q_t) trade in an absolutely continuous fashion on $(0, T)$, in addition to the block trades at $t = 0$ and $t = T$:

$$Q_t := J_0 + Q_t^c + J_T \mathbf{1}_{t=T}, \quad Q_t^c := \int_0^t q_s ds. \quad (2.5)$$

Specifically, we define an *admissible strategy* as a triplet (J_0, q, J_T) . Here $J_0 \in \mathbb{R}$ and $J_T \in L^2(\mathcal{F}_T)$. Moreover, $(q_t)_{t \in [0, T]}$ is a progressively measurable process such that $\mathbb{E}[\int_0^T q_t^2 dt] < \infty$ and such that the local martingale $\int_0^t Q_s^c dS_s$ is a true martingale. Finally,

$$J_0 + \int_0^T q_s ds + J_T = Z_T. \quad (2.6)$$

Thus J_0 and J_T are the sizes of the block trades ΔQ_0 and ΔQ_T , and q is the speed of absolutely continuous trading during regular hours. Finally, (2.6) is the liquidation constraint $Q_T = Z_T$. Occasionally it will be convenient to switch between seeing the strategy as a triplet (J_0, q, J_T) or a process Q . We then write $Q \equiv (J_0, q, J_T)$ to indicate that the relationship (2.5) holds.

Let \mathcal{A} denote the set of admissible strategies. Given $Q \equiv (J_0, q, J_T) \in \mathcal{A}$, the execution price is

$$P_t := \begin{cases} S_t + \frac{1}{2}(Y_{t-} + Y_t), & t \in \{0, T\}, \\ S_t + Y_t + \frac{1}{2}\varepsilon_t q_t, & t \in (0, T). \end{cases} \quad (2.7)$$

The returns from trading at the execution price P_t can be broken down into three parts.

(i) The frictionless wealth

$$S_0 J_0 + \int_0^T S_t q_t dt + S_T J_T$$

reflects returns from the unaffected price (not caused by the unwind strategy).

(ii) The *impact cost*

$$\frac{1}{2}(Y_{0-} + Y_0) J_0 + \int_0^T Y_t q_t dt + \frac{1}{2}(Y_{T-} + Y_T) J_T \quad (2.8)$$

reflects returns caused by the the price impact of the unwind strategy.

(iii) The *spread cost* (or more precisely, *instantaneous cost*)

$$\frac{1}{2} \int_0^T \varepsilon_t q_t^2 dt. \quad (2.9)$$

As described in Section 1.2, $t \in \{0, T\}$ corresponds to the opening and closing auctions. There is no spread cost, but there is price impact. The first part of (2.7) (and the jump terms in (2.8)) reflect that the Obizhaeva–Wang model prices a block trade at the average price over the volume of the block. For $t \in (0, T)$, we also have the instantaneous cost, and as trading is necessarily continuous, we can write Y_t instead of $\frac{1}{2}(Y_{t-} + Y_t)$ in the second part of (2.7) (and in the integral of (2.8)).

The expected execution cost of a given unwind strategy is

$$\mathcal{C}(J_0, q, J_T) := \mathbb{E} \left[\int_0^T P_t dQ_t \right] = \mathbb{E} \left[P_0 J_0 + \int_0^T P_t q_t dt + P_T J_T \right]. \quad (2.10)$$

The desk's aim is to minimize this cost,

$$\inf_{(J_0, q, J_T) \in \mathcal{A}} \mathcal{C}(J_0, q, J_T). \quad (2.11)$$

2.2 Optimal Strategy

The next result provides a semi-closed form solution for the optimal strategy and optimal cost of (2.11); i.e., explicit up to solving an ordinary differential equation (ODE). The ODE does not depend on the realization of the in-flow path, so that it can pre-computed in an implementation. For the case of constant liquidity parameters $\beta, \lambda, \varepsilon$, Section 2.5 will even provide a fully explicit solution; on the other hand, Section 2.4 will cover more general in-flow processes. We emphasize that the implementation of the below formulas is straightforward (even if some look lengthy).

Theorem 2.2. *The problem (2.11) has a unique solution. The optimal strategy $(J_0, q, J_T) \in \mathcal{A}$ is*

$$\begin{aligned} q_t &= f_t X_t + g_t Y_t^c + h_t Z_t, \quad t \in [0, T], \\ J_0 &= r^{-1}(g_{0-} + \eta_{0-})y + r^{-1}(-f_{0-} + h_{0-})z, \\ J_T &= Z_T - J_0 - \int_0^T q_t dt, \end{aligned} \tag{2.12}$$

where $\eta_{0-} = -\varepsilon_0^{-1}(1 - \lambda_{0-}/\lambda_0) \leq 0$ and $r = -f_{0-} - \lambda_{0-}(g_{0-} + \eta_{0-}) > 0$ are constants and

$$\begin{aligned} f_t &= -\varepsilon_t^{-1}(A_t + \lambda_t B_t), \\ g_t &= -\varepsilon_t^{-1}(B_t + \lambda_t C_t), \\ h_t &= -\varepsilon_t^{-1}(D_t + \lambda_t E_t) \end{aligned} \tag{2.13}$$

for $t \in \{0-\} \cup [0, T]$. Here, A_t, B_t, \dots are deterministic functions given by the Riccati ODE system (2.20) below, and $\varepsilon_{0-} := \varepsilon_0$, $A_{0-} := A_0$, $B_{0-} := B_0$, \dots . Finally, the state processes in (2.12) are

$$\begin{cases} dX_t = q_t dt - dZ_t, & X_0 = J_0 - z \\ dY_t^c = (-\beta_t Y_t^c + \lambda_t q_t) dt, & Y_0^c = y + \lambda_{0-} J_0 \\ dZ_t = -\theta_t Z_t dt + \sigma_t dW_t, & Z_0 = z. \end{cases} \tag{2.14}$$

The optimal expected cost (2.11) is

$$\mathcal{C}(J_0, q, J_T) = v(0, J_0 - z, y + \lambda_{0-} J_0, z) + \frac{1}{2} \left\{ (\lambda_{0-}^{-1} - \lambda_0^{-1}) (y + \lambda_{0-} J_0)^2 - \lambda_{0-}^{-1} y^2 + \mathbb{E}[S_T Z_T] \right\}$$

where v is the quadratic polynomial (2.18) with coefficients A_0, B_0, \dots .

The key formula (2.12) states that the optimal trading speed q_t is a linear combination of the auxiliary state processes (X_t, Y_t^c, Z_t) with time-varying coefficients (f_t, g_t, h_t) determined through the Riccati ODE. The coefficients are discussed in more detail in the next proposition. The processes (X_t, Y_t^c) in (2.14) are the inventory and impact state, but only up to the close. More precisely, they are equal to $(Q_t - Z_t, Y_t)$ on $[0, T)$ but are continuous at $t = T$, so that $(X_T, Y_T^c) = (Q_{T-} - Z_{T-}, Y_{T-})$ corresponds to the values immediately before the closing auction. Moreover, the jumps from the opening auction have been absorbed into the initial positions (X_0, Y_0^c) . As before, Z is the in-flow.

The optimal trading speed q_t in (2.12) is given by the feedback function $f_t x + g_t y + h_t z$; thus f_t reflects the reaction to outstanding position, g_t to impact state, and h_t to in-flow. The next proposition summarizes their key properties. (See Remark 2.9 for a discussion of the less important regime $\beta_t + \dot{\gamma}_t \leq 0$ related to transaction-triggered price manipulation.)

Proposition 2.3. *(i) Let $\beta_t + \dot{\gamma}_t > 0$. We have $f_t < 0$ for all $t < T$: a larger outstanding position causes faster selling.*

- (ii) We have $g_t < 0$ for all $t < T$: higher impact state causes slower selling.
- (iii) We have $h \equiv 0$ if $\theta \equiv 0$ (martingale in-flow). Let $\beta_t + \dot{\gamma}_t > 0$. Then h is monotone decreasing wrt. θ ; in particular, $h \geq 0$ if $\theta \leq 0$ (in-flow with momentum) and $h \leq 0$ if $\theta \geq 0$ (in-flow with reversion).
- (iv) The functions f and g depend on the liquidity parameters $(\varepsilon, \beta, \lambda)$, but not on the on the characteristics (θ, σ) of the in-flow. The function h additionally depends on θ . The volatility σ does not affect the strategy (in feedback form) but increases the execution cost $\mathcal{C}(J_0, q, J_T)$.

The feedback function quantifies some of our intuition on internalizing or externalizing orders. As seen in (i) and (ii), f_t promotes externalization when the inventory becomes sizable and g_t promotes internalization when the impact state is high (in absolute value). In view of (iii) and (iv), the contribution of order flow predictions becomes clear: in comparison with a standard optimal execution problem where the total order is given at $t = 0$, the unwind speed is adjusted to the random in-flow Z by adding the term $h_t Z_t$. If the in-flow has no drift ($\theta \equiv 0$), then no adjustment is necessary. Traders say in that case that the order flow is *truth-telling*; it does not predict future orders. Of course, the realization of q_t still depends on the realization of Z_t , but the optimal strategy is myopic in the sense that it does not anticipate *future* changes of the in-flow (because these changes have mean zero). By contrast, the strategy does adjust for projected future changes when there is momentum, as expected if clients are splitting meta orders, or when there is reversion (see also Section 1.2). The coefficient h_t promotes internalization if order flow is reverting, whereas in the more common case of momentum flow, h_t leads to additional externalization as one predicts more same-sided orders.

2.3 Reduction to Linear-Quadratic Control

To solve our problem, we first derive the optimal strategy and cost *after* an exogenously given block trade of arbitrary size J_0 in the opening auction. In particular, for $J_0 = 0$, we solve the problem in a market without opening auction (or, the problem started at a time after the opening auction). In a second step, we will then optimize the result of this problem over $J_0 \in \mathbb{R}$ (see Proposition 2.12) to solve the full problem (2.11).

To derive the optimal continuous strategy, we recast our problem: the liquidation constraint can be reformulated to say that the block trade J_T in the closing auction must match whatever is the outstanding position immediately before that auction. We can thus focus on the continuous trading on $(0, T)$ and let J_T be determined via the constraint (2.6). This effectively replaces the state constraint with a terminal cost: after the continuous trading, the desk is charged the cost of the block trade determined via the constraint. A priori, this terminal cost is intractable, but the next result shows that it can be transformed into a rather convenient form. More precisely, our problem is expressed as a tractable control problem of linear-quadratic (LQ) type.

Proposition 2.4. *Let $(J_0, q, J_T) \in \mathcal{A}$. The execution cost (2.10) can be stated as*

$$\begin{aligned} \mathcal{C}(J_0, q, J_T) = \frac{1}{2} \mathbb{E} \left\{ \int_0^T \left[\frac{2\beta_t + \dot{\gamma}_t}{\lambda_t} (Y_t^c)^2 + \varepsilon_t q_t^2 \right] dt + \frac{1}{\lambda_T} (\lambda_T X_T - Y_T^c)^2 \right. \\ \left. + \left(\frac{1}{\lambda_{0-}} - \frac{1}{\lambda_0} \right) (Y_0^c)^2 - \frac{y^2}{\lambda_{0-}} + S_T Z_T \right\}, \end{aligned} \quad (2.15)$$

where $(X_t, Y_t^c, Z_t)_{t \in [0, T]}$ are defined by (2.14).

The proof is stated in Appendix B. We note that the cost (2.15) no longer contains J_T . As J_0 is considered fixed at the moment, the last three terms in the cost (2.15) are independent of q and can thus be neglected for the purpose of finding the optimal continuous strategy q . (When $\lambda_0 = \lambda_{0-}$, those terms are independent of J_0 as well, but that is not crucial.) On a more technical note, we observe that the no-arbitrage conditions (2.1) are necessary to make (2.15) (strictly) convex. In particular, we have the following.

Corollary 2.5 (No free roundtrips). *Suppose that there is no in-flow ($Z \equiv 0$) and that the initial impact state is $y = 0$. If $(J_0, q, J_T) \in \mathcal{A}$ satisfies $\mathcal{C}(J_0, q, J_T) = 0$, then $J_0 = J_T = 0$ and $q = 0$ dt-a.e.*

Proof. Recall that $2\beta_t + \dot{\gamma}_t > 0$ by (2.1). If $J_0 \neq 0$, then $Y_0^c = y + \lambda_{0-}J_0 \neq 0$. As $t \mapsto Y_t^c$ is continuous, (2.15) shows that $\mathcal{C}(J_0, q, J_T) > 0$. Similarly, if $q \neq 0$, then $\mathcal{C}(J_0, q, J_T) > 0$ as $\varepsilon > 0$. Finally, if $J_0 = 0$ and $q = 0$, the liquidation constraint implies $J_T = 0$. \square

In summary, we have transformed our problem into a more standard LQ control problem with control q . The next result solves this problem using dynamic programming. As usual in the dynamic programming approach, we consider the problem from arbitrary initial states (x, y, z) and initial time t . Moreover we ignore the terms $(\lambda_{0-}^{-1} - \lambda_0^{-1})(Y_0^c)^2 - \lambda_{0-}^{-1}y^2 + S_T Z_T$ in the cost (2.15) as their expectation is the same for any admissible choice of q .

Proposition 2.6. *Fix $t \in [0, T)$ and let \mathcal{Q} be the set of progressively measurable processes q with $\mathbb{E}[\int_t^T q_s^2 ds]$. For $q \in \mathcal{Q}$ and $x, y, z \in \mathbb{R}$, define the state processes $(X_s, Y_s^c, Z_s)_{s \in [t, T]}$ by*

$$\begin{cases} dX_s = q_s ds - dZ_s, & X_t = x \\ dY_s^c = (-\beta_s Y_s^c + \lambda_s q_s) ds, & Y_t^c = y \\ dZ_s = -\theta_s Z_s ds + \sigma_s dW_s, & Z_t = z \end{cases} \quad (2.16)$$

and the value function $v(t, x, y, z)$ by

$$\inf_{q \in \mathcal{Q}} \frac{1}{2} \mathbb{E} \left\{ \int_t^T \left[\frac{2\beta_s + \dot{\gamma}_s}{\lambda_s} (Y_s^c)^2 + \varepsilon_s q_s^2 \right] ds + \frac{1}{\lambda_T} (\lambda_T X_T - Y_T^c)^2 \right\}. \quad (2.17)$$

Then

$$v(t, x, y, z) = \frac{1}{2} A_t x^2 + B_t xy + \frac{1}{2} C_t y^2 + D_t xz + E_t yz + \frac{1}{2} F_t z^2 + K_t, \quad (2.18)$$

where A_t, B_t, \dots, K_t are deterministic functions defined in Proposition 2.7 below. Moreover, the unique optimal control for (2.17) is given by the feedback function

$$q^*(s, x, y, z) = f_s x + g_s y + h_s z, \quad (2.19)$$

where

$$\begin{aligned} f_s &:= -\varepsilon_s^{-1} (A_s + \lambda_s B_s), \\ g_s &:= -\varepsilon_s^{-1} (B_s + \lambda_s C_s), \\ h_s &:= -\varepsilon_s^{-1} (D_s + \lambda_s E_s). \end{aligned}$$

The proof is stated in Appendix B. The following functions were used in (2.18).

Proposition 2.7. *The following Riccati ODE system has a unique solution on $[0, T]$:*

$$\begin{cases} \dot{A}_t = \varepsilon_t^{-1}(A_t + \lambda_t B_t)^2, & A_T = \lambda_T \\ \dot{B}_t = \varepsilon_t^{-1}(A_t + \lambda_t B_t)(B_t + \lambda_t C_t) + \beta_t B_t, & B_T = -1 \\ \dot{C}_t = \varepsilon_t^{-1}(B_t + \lambda_t C_t)^2 + 2\beta_t C_t - \lambda_t^{-1}(2\beta_t + \dot{\gamma}_t), & C_T = \lambda_T^{-1} \\ \dot{D}_t = \varepsilon_t^{-1}(A_t + \lambda_t B_t)(D_t + \lambda_t E_t) - \theta_t(A_t - D_t), & D_T = 0 \\ \dot{E}_t = \varepsilon_t^{-1}(B_t + \lambda_t C_t)(D_t + \lambda_t E_t) - \theta_t(B_t - E_t) + \beta_t E_t, & E_T = 0 \\ \dot{F}_t = \varepsilon_t^{-1}(D_t + \lambda_t E_t)^2 - 2\theta_t(D_t - F_t), & F_T = 0 \\ \dot{K}_t = -\sigma_t^2(A_t - 2D_t + F_t)/2, & K_T = 0. \end{cases} \quad (2.20)$$

Again, the proof is stated in Appendix B. Some observations regarding (2.20) are in order:

- (i) A, B, C form an autonomous system depending on the market impact parameters $\varepsilon_t, \beta_t, \lambda_t$ but not on the characteristics of the in-flow Z .
- (ii) D, E, F also depend on the mean reversion θ_t . If the in-flow is a martingale ($\theta \equiv 0$), then $D, E, F \equiv 0$.
- (iii) K also depends on the volatility of Z . If the in-flow is deterministic ($\sigma \equiv 0$), then $K \equiv 0$.

The following analytic result on A_t, B_t, \dots, K_t implies the properties of the optimal unwind speed reported in Proposition 2.3 by way of the formulas for f_t, g_t, h_t in (2.13).

Proposition 2.8. *Let A_t, B_t, \dots, K_t be the functions defined in (2.20).*

(i) *The matrix*

$$\begin{bmatrix} A_t & B_t \\ B_t & C_t \end{bmatrix}$$

is nonnegative definite, $0 \leq C_t \leq \lambda_t^{-1}$, and $B_t + \lambda_t C_t > 0$ for $t \in [0, T]$.

(ii) *Let $\beta_t + \dot{\gamma}_t > 0$. Then*

$$A_t + \lambda_t B_t > \max(D_t + \lambda_t E_t, 0) \quad \text{and} \quad B_t < \min(E_t, 0), \quad t \in [0, T].$$

(iii) *Let $\beta_t + \dot{\gamma}_t > 0$. Let $\tilde{\theta} : [0, T] \rightarrow \mathbb{R}$ be measurable and bounded, and define \tilde{D}, \tilde{E} like D, E but with $\tilde{\theta}$ instead of θ . If $\theta_t \geq \tilde{\theta}_t$ for $t \in [0, T]$, then*

$$D_t + \lambda_t E_t \geq \tilde{D}_t + \lambda_t \tilde{E}_t \quad \text{and} \quad E_t \leq \tilde{E}_t, \quad t \in [0, T].$$

In particular, $\theta_t \geq 0$, $t \in [0, T]$ implies $D_t + \lambda_t E_t \geq 0$ and $E_t \leq 0$.

The proof is stated in Appendix B. The work [33] on stochastic resilience (i.e., stochastic β_t) features a backward stochastic differential system with singular terminal condition that plays a similar role as our (much simpler) ODE system (2.20). We have benefited from several ideas in [33] for proving Proposition 2.8.

Remark 2.9 (Transaction-triggered price manipulation). Under the condition $\beta_t + \dot{\gamma}_t > 0$, Proposition 2.8 yields in particular that $f_t < 0$ for $t \in [0, T)$, which is the result reported in Proposition 2.3 (i). Then, a larger positive outstanding position leads to faster selling and a larger negative position leads to faster buying. If $\beta_t + \dot{\gamma}_t < 0$, the proof of Proposition 2.8 shows that $f_t > 0$ for $t \in [t_0, T)$, for some $t_0 < T$. Thus, the opposite behavior occurs: given a large positive outstanding position, the optimal strategy will initially add to the position, only to unwind later. This corresponds to the transaction-triggered price manipulation discussed in detail in [28, Theorem 8.4 and Corollary 8.5]. While the weaker condition $2\beta + \dot{\gamma} \geq 0$ from (2.1) rules out profitable round trips from zero initial position, it does not rule out profitable round trips from non-zero initial position.

If this seems unintuitive, considering the boundary case where $\beta_t + \dot{\gamma}_t = 0$ for $t \in [0, T]$ may be helpful. In this regime, the benefit of trading earlier to exploit the reversion β_t of the impact process exactly offsets the benefit of trading later to exploit the increase in liquidity (decay of λ_t) over time. Specifically, the proof of Proposition 2.8 shows that $f_t = h_t = 0$, whereas $g_t < 0$ for $t \in [0, T)$. Thus, the optimal trading speed (2.19) only depends on the current impact state, not on the current outstanding position. If the impact state at t is zero, no trading occurs over $[t, T)$; if the impact state at t is positive (negative), the optimal policy sells (buys) over $[t, T)$, to exploit the reversion of the impact process. In either case, the remaining outstanding position is unwound in the closing auction.

Remark 2.10 (Practical use of no-price manipulation conditions). Due to our framework’s inherent tractability, practitioners can directly apply no-price manipulation conditions on live data. Indeed, the condition only involves the liquidity curves β, λ , which the trading team typically builds from volume predictions or order book data. For instance, see [26] and [42], which fit parametric models of λ as functions of order book depth and market volumes, respectively.

Intuitively, the condition states that “liquidity cannot increase faster than impact reverts without triggering price manipulation opportunities.” If a given liquidity forecast breaches a no-price manipulation condition, traders can take two actions:

- (i) Switch to a simpler externalization strategy, e.g., VWAP or a constant parameter OW model. This avoids having the algorithm engage in odd trading patterns due to potential liquidity mirages.
- (ii) Investigate the liquidity surge that triggered the breach: Is there a data or model problem, or is the sudden liquidity caused by abnormal market behavior, e.g., order spoofing?

In that sense, no-price manipulation conditions are to trading algorithms as no-arbitrage conditions are to option pricing models. Models trade under the assumption they don’t exist, but traders monitor the market for unexplained behavior.

Next, we return to the derivation of the optimal strategy and cost. We can now deduce the optimal expected cost when the block trade in the opening auction is prescribed to be of size $j_0 \in \mathbb{R}$, by evaluating the auxiliary value function (2.18) at the corresponding initial conditions and adding back the constant terms of (2.15). This yields the following.

Corollary 2.11. *Given $j_0 \in \mathbb{R}$, the optimal execution cost is*

$$\begin{aligned} \inf_{\substack{(J_0, q, J_T) \in \mathcal{A} \\ J_0 = j_0}} \mathcal{C}(J_0, q, J_T) &= v(0, j_0 - z, y + \lambda_{0-} j_0, z) \\ &+ \frac{1}{2} \left\{ (\lambda_{0-}^{-1} - \lambda_0^{-1}) (y + \lambda_{0-} j_0)^2 - \lambda_{0-}^{-1} y^2 + \mathbb{E}[S_T Z_T] \right\}. \end{aligned} \quad (2.21)$$

The optimal trading speed in feedback form is given by (2.19), with the state processes started at time 0 with the initial condition $(x, y, z) = (j_0 - z, y + \lambda_{0-}j_0, z)$.

Finally, we determine the optimal size J_0 for the block trade in the opening auction.

Proposition 2.12. *The optimal initial block trade is*

$$J_0 = r^{-1}(g_{0-} + \eta_{0-})y + r^{-1}(-f_{0-} + h_{0-})z, \quad (2.22)$$

where $\eta_{0-} := -\varepsilon_0^{-1}(1 - \lambda_{0-}/\lambda_0) \leq 0$,

$$g_{0-} + \eta_{0-} < 0 \quad \text{and} \quad r := -f_{0-} - \lambda_{0-}(g_{0-} + \eta_{0-}) > 0.$$

If $\beta_t + \hat{\gamma}_t > 0$ for $t \in [0, T]$, then $-f_{0-} + h_{0-} > 0$.

The proof is stated in Appendix B. We see that the coefficients of y and z in (2.22) resemble the ones of the trading speed q and have similar comparative statics (cf. Proposition 2.3). That the coefficient of z contains f and not only h , is merely a consequence of our notation: the initial in-flow z also represents an outstanding position of $x := -z$ at the beginning of the auction, and now $-f_{0-}z = f_{0-}x$ is analogous to the formula of q .

2.4 In-Flows with Jumps

In this section we generalize our results to more general in-flow processes

$$dZ_s = -\theta_s Z_s ds + dM_s, \quad Z_0 = z, \quad (2.23)$$

where M_t is a general martingale instead of $\int_0^t \sigma_s dW_s$. Specifically, allowing a process with jumps is important as a CRB's in-flow may contain large block orders, and market makers may receive large block trades from dark pools. Self-exciting (hence non-Markovian) dynamics are also relevant. The main take-away is that the formula for the optimal strategy remains unchanged (similarly as, e.g., in [29]), whereas the expression of the optimal expected cost involves an integral of the quadratic variation of M . This generalizes our previous observations in Proposition 2.3, where the optimal strategy was independent of σ while the cost was an integral against $\sigma_s^2 ds$ (as seen in the formula for K in (2.20)).

Let M be a càdlàg, square-integrable martingale with $M_0 = 0$. We consider a variant of the control problem in Proposition 2.6, now only for the initial time $t = 0$ (since M was not assumed to be Markovian). To state its solution, we define the deterministic functions A_t, \dots, E_t as in (2.20), whereas K_t is generalized to the stochastic process

$$K_t := \mathbb{E} \left[\int_t^T \frac{1}{2} (A_s - 2D_s + F_s) d[M, M]_s \middle| \mathcal{F}_t \right]. \quad (2.24)$$

At a technical level, we also need to amend the definition of an admissible strategy and impose a mild integrability condition (see the proof).

Proposition 2.13. *Define the state processes X, Y^c as in (2.16) while Z is generalized to (2.23). Moreover, define A, \dots, E as in Proposition 2.7 whereas K is generalized to (2.24).*

(i) *The problem (2.17) admits a unique optimal strategy q^* , given by the same feedback function as in (2.19).*

(ii) *The value of (2.17) is given by*

$$V_0 := \frac{1}{2}A_0x^2 + B_0xy + \frac{1}{2}C_0y^2 + D_0xz + E_0yz + \frac{1}{2}F_0z^2 + K_0.$$

The proof is sketched in Appendix B.

2.5 Closed-form Solution for Constant Liquidity

In this section we assume that the liquidity parameters $\beta, \lambda, \varepsilon \in (0, \infty)$ are constant (and in particular $\lambda_{0-} = \lambda$). The coefficients θ_t, σ_t of the in-flow can still be time-dependent. In that setting, we can obtain the optimal strategy in closed form. This solution is based on a different approach, the stochastic maximum principle, which avoids the Riccati system (2.20). On an abstract level, the dynamic programming approach implies the conclusion of the stochastic maximum principle in our context. However, the stochastic maximum principle focuses on the optimal strategy along the optimal state trajectory, and sometimes leads to a closed-form solution even if dynamic programming does not.

Theorem 2.14. *Let $\beta, \lambda, \varepsilon \in (0, \infty)$ be constant and define also*

$$\tilde{\varepsilon} := \frac{\varepsilon\beta}{2\lambda}, \quad \kappa := \beta\sqrt{1 + \tilde{\varepsilon}^{-1}} > \beta.$$

Then the functions f_t, g_t, h_t of Theorem 2.2 are given by

$$f_t = \tilde{f}_t(e^{\kappa(T-t)} - 1)/dt, \tag{2.25}$$

$$g_t = \tilde{g}_t(e^{\kappa(T-t)} - 1)/dt, \tag{2.26}$$

$$h_t = f_t(1 - e^{-\int_t^T \theta_u du}) \tag{2.27}$$

for all $t \in [0, T]$, where

$$\begin{aligned} \tilde{f}_t &:= -[\beta^{-1} - (\kappa + \beta)^{-1}]e^{-\kappa(T-t)} - [\beta^{-1} + (\kappa - \beta)^{-1}], \\ \tilde{g}_t &:= \lambda^{-1}\tilde{f}_t - \lambda^{-1}(1 + e^{-\kappa(T-t)})(T-t) + 2\lambda^{-1}\kappa^{-1}(1 - e^{-\kappa(T-t)}), \\ d_t &:= e^{\kappa(T-t)} \{(\kappa - \beta)^{-1}[(\kappa - \beta)^{-1} + \beta^{-1} - \kappa^{-1}] + \beta^{-1}\kappa^{-1}\} \\ &\quad + e^{-\kappa(T-t)} \{(\kappa + \beta)^{-1}[-(\kappa + \beta)^{-1} + \beta^{-1} + \kappa^{-1}] + \beta^{-1}\kappa^{-1}\} \\ &\quad + (T-t)e^{\kappa(T-t)}(\kappa - \beta)^{-1} + (T-t)e^{-\kappa(T-t)}(\kappa + \beta)^{-1} + 4\tilde{\varepsilon}\beta^{-1}\kappa^{-1}. \end{aligned}$$

Moreover, $f_{0-} = f_0, g_{0-} = g_0, h_{0-} = h_0$.

The proof is reported in Appendix C. By specifying the functions f_t, g_t, h_t , Theorem 2.14 gives the optimal strategy q in feedback form (as well as the optimal block trade). This may seem surprising, as the maximum principle generally only yields the control in open-loop form, along the optimal trajectory. Indeed, to derive Theorem 2.14, the maximum principle is carried out from an arbitrary initial state (t, x, y, z) and then the feedback function is recovered by identifying the initial value of the open-loop control.

We note that this approach does not immediately yield a closed-form expression for the cost $\mathcal{C}(J_0, q, J_T)$ or the functions A_t, B_t, \dots in Theorem 2.2. That would require an integration that we cannot carry out in closed form.

3 Numerical Results

In this section, we study key features of our model and its optimal strategy by simulation. In addition, we introduce new metrics of interest to practitioners. Section 3.1 describes the methodology, model parameters, and trading metrics. The most essential trading question is how unwind strategy and cost metrics depend on the characteristics of the in-flow; namely, autocorrelation and volatility.

Sections 3.2 and 3.3 answer this question and show that different parameters give rise to different trading regimes ranging from optimal execution to market making. Our model covers the full range of scenarios. A crucial insight from the study is that practitioners must carefully model their in-flow to obtain useful results. Therefore, the mantra “know your client” is a recurring theme. Section 3.4 proposes a metric capturing how *monotone* a given in-flow path is. Monotonicity, in turn, correlates strongly with internalization and trading cost under the optimal trading strategy.

Several further experiments are reported in Appendix A. Section A.1 links our model to classical (deterministic) models of optimal execution and illustrates that the Obizhaeva–Wang model is recovered for small spread cost parameter $\varepsilon \rightarrow 0$. Section A.1 also studies the sensitivity with respect to ε , both in the deterministic and stochastic case. Section A.2 analyzes the sensitivity with respect to the martingale driver of the in-flow as the martingale changes from a finite-variation process with a small number of shocks to a continuous-time Brownian motion. Section A.3 studies the effect of a non-zero initial impact state on the trading strategy. Lastly, Section A.4 presents an empirical study on autocorrelation of orders on the public trading tape, confirming momentum for all S&P 500 stocks.

3.1 Numerical Setup

We focus on constant parameters throughout our experiments. While time-varying parameters are important in a practical implementation, constant parameters allow for more interpretable results and a straightforward way of varying parameter values in our experiments.

The theoretical results of the preceding section enable an efficient numerical implementation: the feedback form of the optimal strategy (i.e., the coefficients f_t, g_t, h_t and the block trade J_0) can be precomputed independently of the random inputs. As parameters are constant, we use the closed-form formulas outlined in Section 2.5. (The ODEs for the time-varying case are also straightforward to implement and fully precomputable). Given the optimal strategy, the simulation of the state processes (X, Y, Z) is straightforward by discretization of their SDEs. As a result, large-scale numerical experiments using many sample paths can be run without heavy-duty hardware. The code is provided on GitHub.⁶

One difference with the main part of Section 2 is that we focus on an in-flow process of finite variation. For practitioners, the most important metric in our context is the internalization rate, defined in (3.1) below. Its formula includes a normalization by the total variation of the in-flow process, or in other words, the cumulative size of all orders received. If the in-flow is modeled using a Brownian motion driver as in Theorem 2.2, the total variation is infinite and the internalization rate cannot be defined. Instead, our experiments use a discrete driver where the inventory is shocked by a block trade approximately every 20 minutes. We recall from Section 2.4 that our model is solvable for general martingale drivers. The shocks are taken to be i.i.d. Gaussian (any centered distribution could be used). While the main purpose is to have a well-defined internalization rate, this also adds realistic block orders to our simulation. Section A.2 discusses the robustness of our results with respect to this discretization.

3.1.1 Inputs and Methodology

We use standard units to normalize our variables across securities. The time unit is days; we set $T = 1$ so that the strategy trades daily over the interval $[0, 1]$.⁷ Denote by $[T]$ the time unit. All trading quantities are expressed in percent of Average Daily Volume (ADV%). For instance, a jump

⁶<https://github.com/KevinThomasWebster/UnwindingStochasticOrderFlow/blob/main/simulation.ipynb>

⁷In our plots, we nevertheless label time with the standard 9:30–16:00 hours.

in the in-flow of 2% means that the desk received an order of size 2% ADV. In-flow volatility of 10% indicates that the trading team expects 10% of the ADV to go through their desk over the course of the day. We denote by $[V]$ the volume unit.

Next, we describe the trading inputs and their default values in our simulations.

Impact and spread parameters.

- (i) Impact decay $\beta = 8$ has unit $[T]^{-1}$ and describes the speed at which impact reverts. The corresponding half-life is approximately 40 minutes.
- (ii) Kyle's lambda $\lambda = 0.2$ has unit $[V]^{-1}$ and describes an individual fill's impact.⁸
- (iii) The spread cost parameter $\varepsilon = 0.01$ has unit $[V]^{-2}[T]^{-1}$ and describes in reduced form the various instantaneous trading costs, such as the bid-ask spread and adverse selection.

These parameters were picked to match the overall statistics (spread and impact cost, volume on close) from Nasdaq's 2022 Intern's Guide to Trading [40] and align with the empirical estimates found in [24] and [42].

In-flow parameters. The in-flow process uses 20 equispaced i.i.d. Gaussian shocks (rather than a Brownian motion). The volatility parameter σ controls the variance of those shocks. We express σ in terms of daily quadratic variation to make it comparable to the Brownian motion case or a different number of intraday shocks.

- (i) The in-flow volatility $\sigma = 0.1$ has unit $[V][T]^{-1/2}$ and describes in-flow shocks.
- (ii) The in-flow's autocorrelation (Ornstein-Uhlenbeck) parameter θ has unit $[V][T]^{-1}$. We consider three cases: the default martingale case $\theta = 0$, the momentum case $\theta = -1$, and the reversal case $\theta = 1$. Therefore, in the momentum case, if the desk stopped trading the out-flow, their inventory would in expectation double over $\log(2) \approx 0.69$ days. Conversely, in the reversal case, if the desk stopped trading the out-flow, their inventory would naturally internalize by half over $\log(2) \approx 0.69$ days.

Initial conditions.

- (i) The initial inventory $x = -z = -0.1$ has unit $[V]$.
- (ii) The initial impact state $y = 0$ is unitless (but can be normalized in terms of stock volatility in an empirical study).

Non-zero initial conditions matter in practice. While the market-making literature often assumes zero initial position, the desk cannot assume that its portfolio is always trading in a steady state. The optimal execution literature obviously uses non-zero initial inventory but often assumes zero initial impact. In practice, impact from the previous day's trading cannot be neglected and is similar to an alpha signal: the strategy must take the reversal of prices into account (see also Section A.3). Our model and the provided code allow for non-zero initial conditions.

⁸For empirical studies, Kyle's lambda is often normalized for the stock's volatility. For our numerical study, we picked a stock with moderate volatility.

3.1.2 Trading Metrics

Following industry practice, we normalize trading metrics by the total variation of the in-flow (i.e., the cumulative order notional). This corresponds to the *client’s perspective* as the metric can be understood as “per order notional.”

In-flow, out-flow, inventory (ADV%). We plot the intraday time series of flows and inventories. As explained in the beginning of Section 3.1, we choose an in-flow with finite variation for our simulations.

Internalization (%). The internalization rate is defined [12, p. 29] as

$$\text{internalization} = 1 - \frac{\text{total variation of out-flow}}{\text{total variation of in-flow}}; \quad (3.1)$$

roughly speaking, it is the fraction of in-flow orders that were netted. From the client’s point of view, high internalization indicates that many orders were netted *instead* of trading in the market and thus that the desk was able to reduce trading costs. The internalization rate is upper bounded by one, but this bound is achieved only if there is no out-flow (or if the in-flow has infinite total variation). Zero internalization is achieved in particular if out-flows match in-flows, meaning that the desk routes all orders to the market. For a given in-flow, the highest internalization rate is obtained by a monotone out-flow (the desk only sells or only buys in the market), as that minimizes the total variation given the liquidation constraint. In principle, internalization can be negative, but it usually takes a value in $[0, 1]$ because the desk chooses an out-flow smoother than their in-flow.

Practitioners consider internalization rate as a proxy for saved trading costs (but see our discussion of Figure 5 below for a caveat). Compared to cost, internalization rate has the advantage of being model-free, which may explain why it is the most important metric for clients and regulators (cf. [12, Footnote 23]).

Internalization regret (%). We propose a novel metric called internalization regret which focuses on the *desk’s perspective* and measures how many trades were ex-post unnecessary. We define internalization regret as⁹

$$\text{internalization regret} = 1 - \frac{|\text{terminal in-flow}|}{\text{total variation of out-flow}}; \quad (3.2)$$

roughly speaking, it is one minus the ratio of minimally necessary over executed out-flow trades. Note that we use the absolute terminal in-flow $|Z_T|$ rather than the total variation of the in-flow process Z . Because the total variation of any out-flow satisfying the liquidation constraint is at least $|Z_T|$, internalization regret always takes values in $[0, 1]$. (We use the convention $0/0 := 1$.) It has value zero if and only if the out-flow is monotone (the desk only sells or only buys in the market), indicating that the desk captured all possible internalization. In particular, zero regret is always achieved by the naive strategy that warehouses all in-flow during the day and then externalizes the inventory on the close. On the other hand, a high regret indicates that the desk traded in the market where netting would (ex-post) have been possible.

In contrast to the internalization rate, the internalization regret is well-defined even if the in-flow has infinite variation (see also Section A.2).

⁹A similar definition would be $(\text{total variation of out-flow})/|\text{terminal in-flow}|$. That quantity is numerically ill-behaved as the total in-flow may be close to or equal to zero.

Impact, spread, total costs (bps). The absolute impact cost was defined in (2.8) in continuous time; it is discretized as

$$yJ_0 + \frac{\lambda}{2}J_0^2 + \sum_t Y_t q_t \Delta t + Y_{T-\Delta t} J_T + \frac{\lambda}{2}J_T^2.$$

As we measure trading costs from the clients’ perspective in this section, we report the absolute cost divided by the total variation of the in-flow. Similarly, the absolute spread (or instantaneous) cost defined in (2.9) is discretized as $\frac{1}{2} \sum_t \varepsilon q_t^2 \Delta t$ and divided by the total variation of the in-flow.

Closing trade (ADV%, %). The closing trade is the block trade placed in the closing auction. We express it in two units: either in “absolute” terms (ADV%) or as a proportion of the (total variation of the) out-flow (%). The former is more relevant to the client, as it directly compares to the in-flow. On the other hand, the latter captures how much the desk warehouses until the close compared to its overall trading activity.

3.2 Sensitivity to In-Flow Auto-Correlation

We first study how the optimal out-flow reacts to the autocorrelation parameter (θ). The figures show martingale, reversal, and momentum in-flow. Figure 2 shows a sample path of the trading strategy’s intraday evolution. As expected, the in-flow exhibits larger variations with momentum dynamics and lower variations with reversal dynamics. Therefore, the out-flow and impact state is most aggressive for momentum and least aggressive for reversal in-flow. Figure 3 shows the average path of the trading strategy. In expectation, the out-flow and impact time series are identical to the deterministic optimal execution case when accounting for the expected future in-flow (cf. Section A.1).

Parameter scans	θ	In-flow (ADV%)	Out-flow (ADV%)	Spread Cost (bps)	Impact cost (bps)	Closing trade (% total)	Internalization (%)
momentum	-1	61	31	4.9	42.6	17	51
martingale	0	46	15	1.7	14.5	21	68
reversal	1	52	9	0.5	4.8	27	84

Table 1: Sensitivity to flow autocorrelation (θ) for an initial inventory (z) and daily flow volatility (σ) of 10% ADV: average metrics.

Table 1 reports the average daily statistics. We observe that out-flows are significantly more sensitive to θ than in-flows. Consequently, in-flow autocorrelation strongly affects all core trading metrics:

- (i) Higher autocorrelation leads to higher costs. The impact-to-spread ratio remains stable. Therefore, higher out-flows drive the cost increase, which is proportionally shared by spread and impact.
- (ii) Lower autocorrelation leads to higher internalization. In particular, the trading strategy warehouses more inventory with lower autocorrelation, which leads to a larger closing trade as a proportion of out-flow.

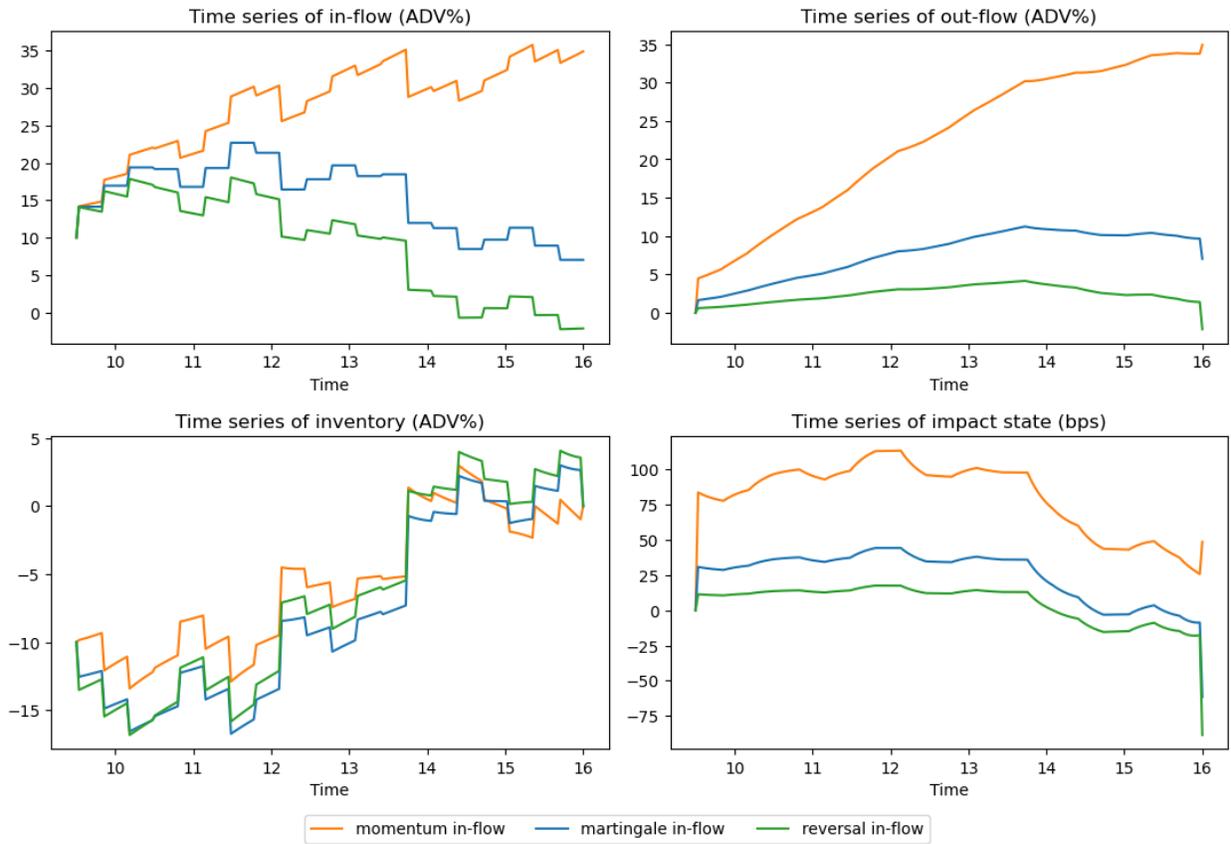


Figure 2: Sensitivity to flow autocorrelation (θ) for an initial inventory (z) and daily flow volatility (σ) of 10% ADV: sample path.

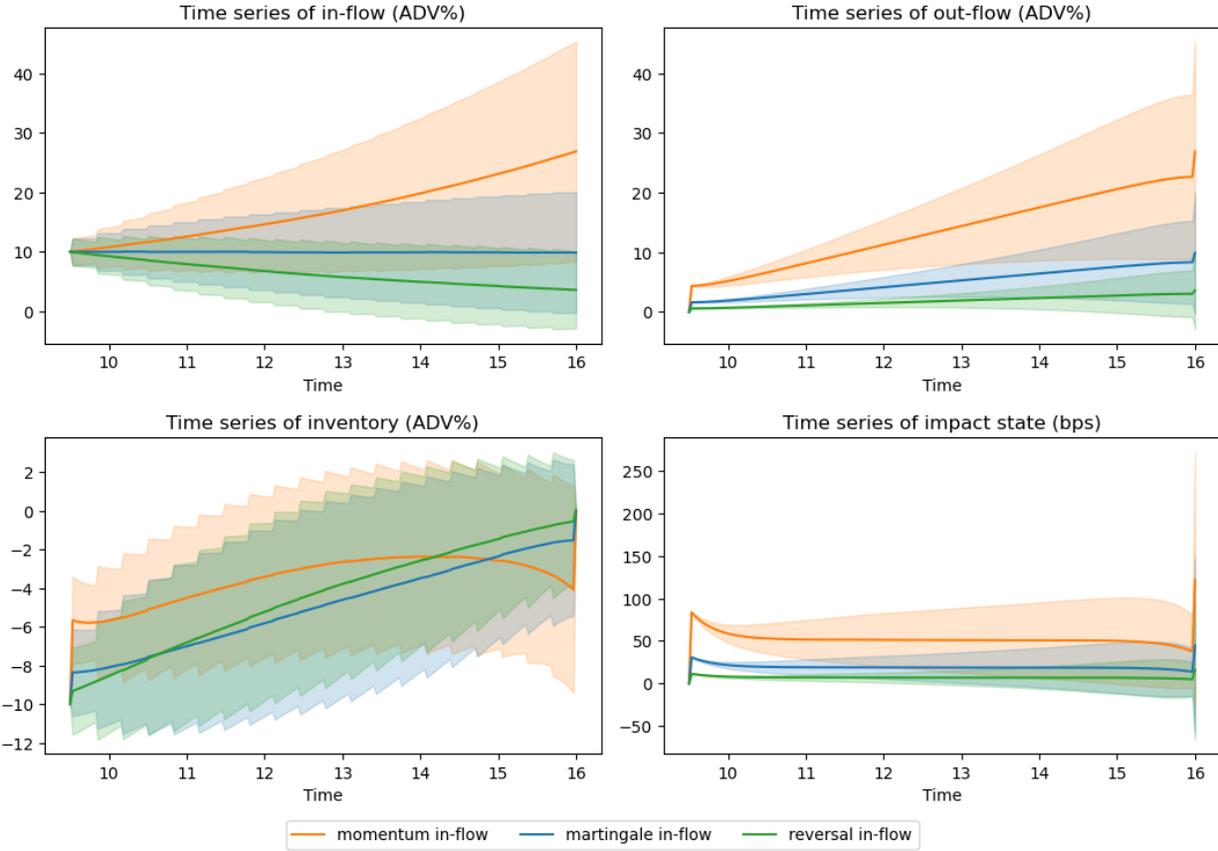


Figure 3: Sensitivity to flow autocorrelation (θ) for an initial inventory (z) and daily flow volatility (σ) of 10% ADV: average path. The error bars illustrate one standard deviation around the average path.

Figure 4 shows the distribution of daily trading metrics across samples of the in-flow path. As expected, distributions are wide, highlighting the need for stochastic control theory over deterministic optimization techniques: the deterministic optimal execution problem only describes the strategy's *average* behavior.

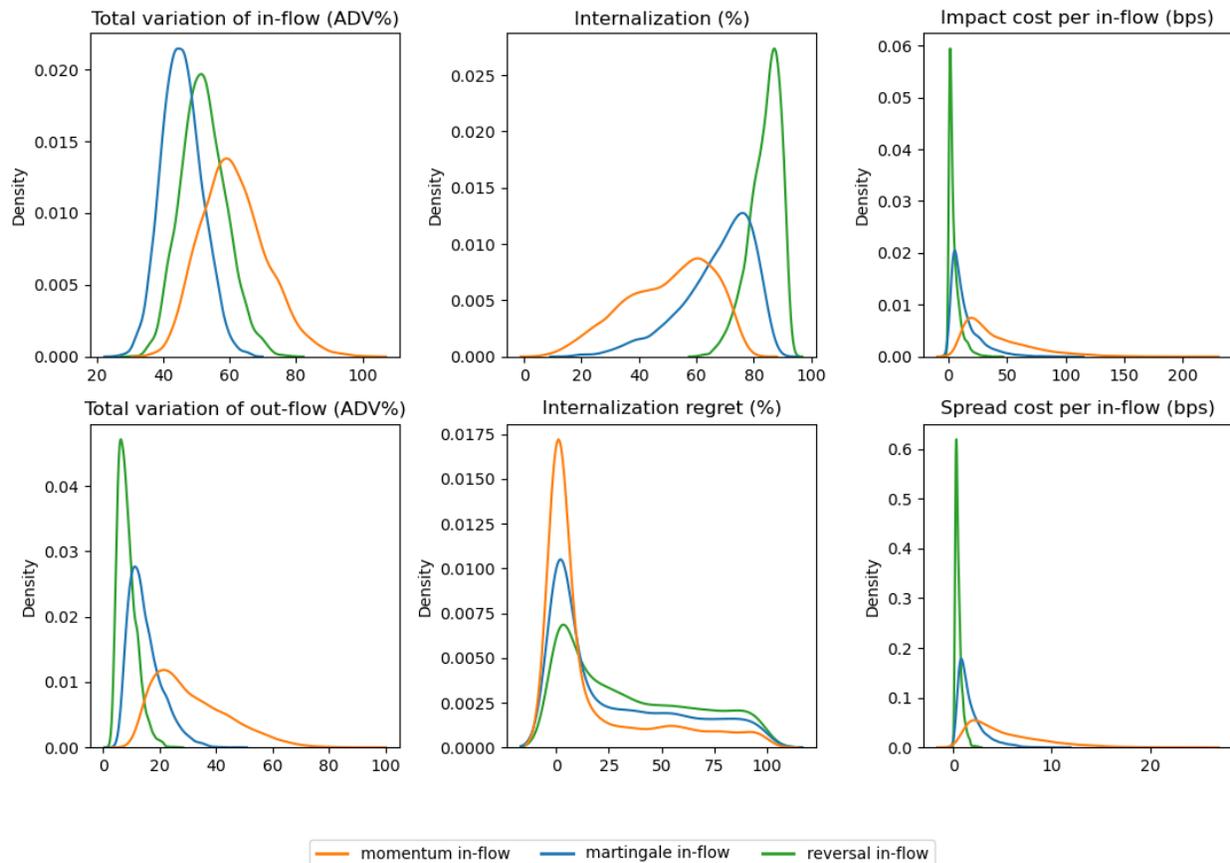


Figure 4: Sensitivity to flow autocorrelation (θ) for an initial inventory (z) and daily flow volatility (σ) of 10% ADV: metrics' distribution.

As the first panel of Figure 4 illustrates, martingale in-flow has the smallest total variation. Indeed, any drift must increase total variation, but momentum more so than reversal. This is relevant as the in-flow's total variation is used to normalize trading metrics. Nevertheless, internalization is highest in the reversal case and lowest in the momentum case. This is explained by the total variation of the out-flow: reversal naturally leads to netting while momentum requires more trading. We observe that the distribution of the internalization rate is significantly wider for momentum than for reversal. The impact cost shows largely the same characteristics as the internalization.

The distribution of the internalization regret has an atom at zero, decays away from zero, and becomes approximately uniform for large regret values. The mass at zero reflects monotone trading paths (the desk only sells or only buys). The concentration is strongest in the momentum case, where the probability of zero regret is about 23% and the probability of regret below 1% is about 46%, compared with (11%, 26%) in the martingale case and (4%, 13%) for reversal. (Density plotted outside $[0,100]$ is an artifact of kernel density estimation.) Generally, increasing autocorrelation

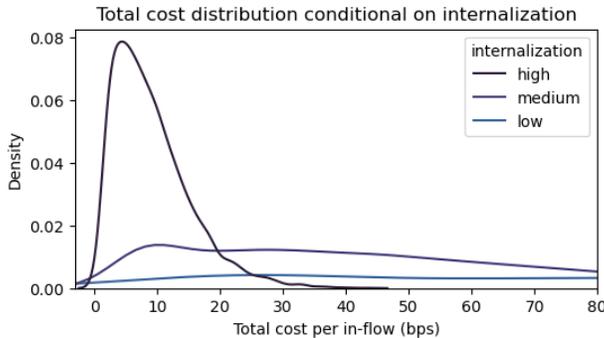


Figure 5: Total cost distribution conditional on internalization. Samples are taken across a wide range of in-flow parameters.

lowers internalization and increases costs.

Figure 5 highlights an important rule of thumb: high internalization is a strong proxy for low costs, regardless of the shape of the in-flow path. On the other hand, when internalization is low, costs can take a wide range of values depending on the in-flow path, meaning that internalization alone is not a suitable proxy for the cost.

3.3 Sensitivity to In-Flow Volatility

Figure 6 shows the sensitivity of key metrics to in-flow volatility (σ). We observe that the relation between initial inventory and in-flow volatility is key. In the regime where the initial inventory is significantly larger than the daily volatility of the in-flow, the trading strategy behaves like an optimal execution strategy. Conversely, if the in-flow volatility dominates the initial inventory, the trading strategy behaves like a steady-state market-making strategy. An important transition occurs between those regimes:

- (i) Before the critical value, increasing the in-flow volatility increases internalization and drives down (per in-flow) costs.
- (ii) After the critical value, internalization reaches a plateau. Therefore, in-flow volatility drives up total trading, and costs increase.

As an important consequence, a desk with low in-flow volatility may benefit from attracting additional in-flows, e.g., by onboarding new clients or posting more attractive prices, in order to increase internalization and reduce costs, whereas a desk with high volatility may benefit from lowering it. We observe that the transition happens when volatility and initial position are comparable in size, with the exact location depending on the in-flow autocorrelation. Table 2 quantifies the optimal in-flow vol to attract for a given in-flow autocorrelation. The in-flow vol scales linearly with the initial inventory, and, therefore, we quote the ratio.¹⁰

¹⁰The optimal in-flow volatility is insensitive to the martingale driver: there is no material difference between the diffusion and jump cases. The fact that the location of the cost minimum depends only on the ratio σ/z can be seen analytically by a scaling argument. If the initial impact state y is nonzero, y also needs to be scaled, so that the location of the minimum depends on the two ratios $\sigma : z : y$.

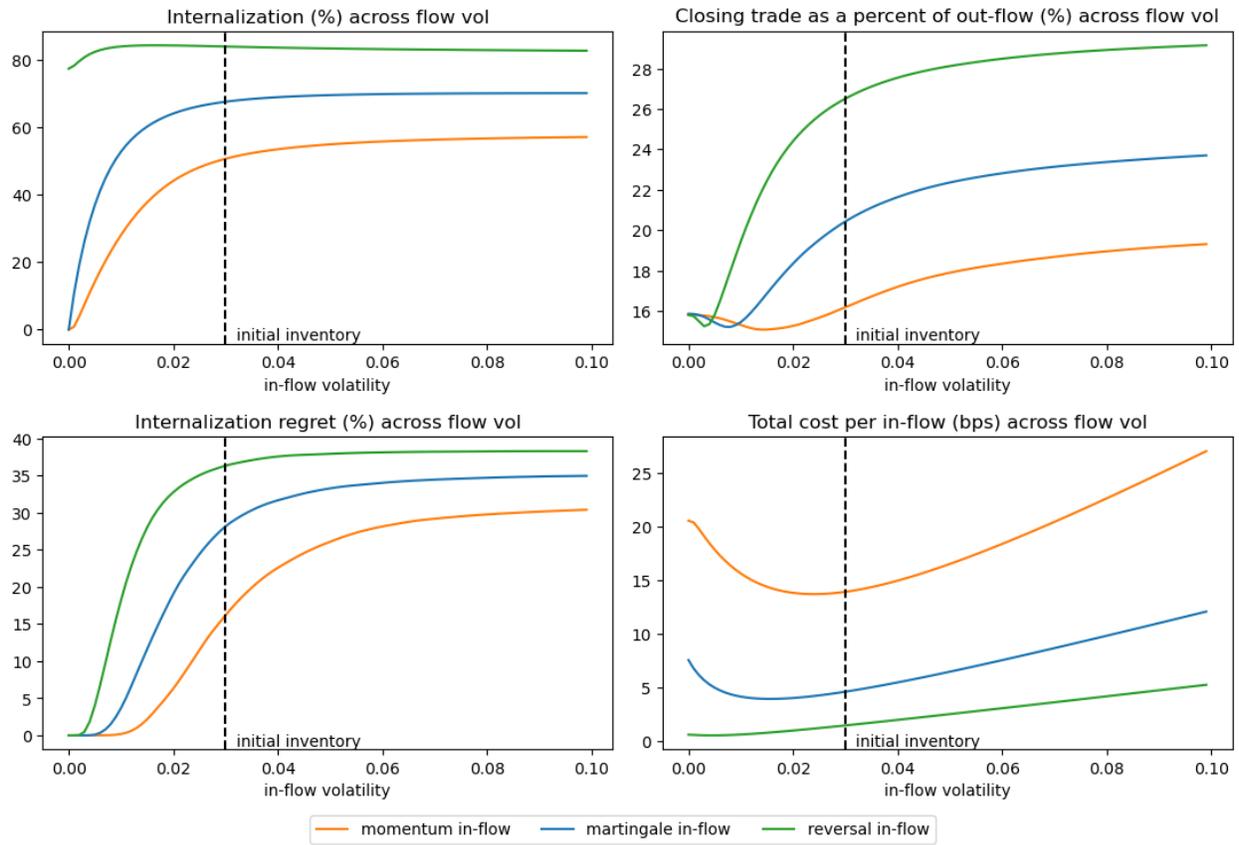


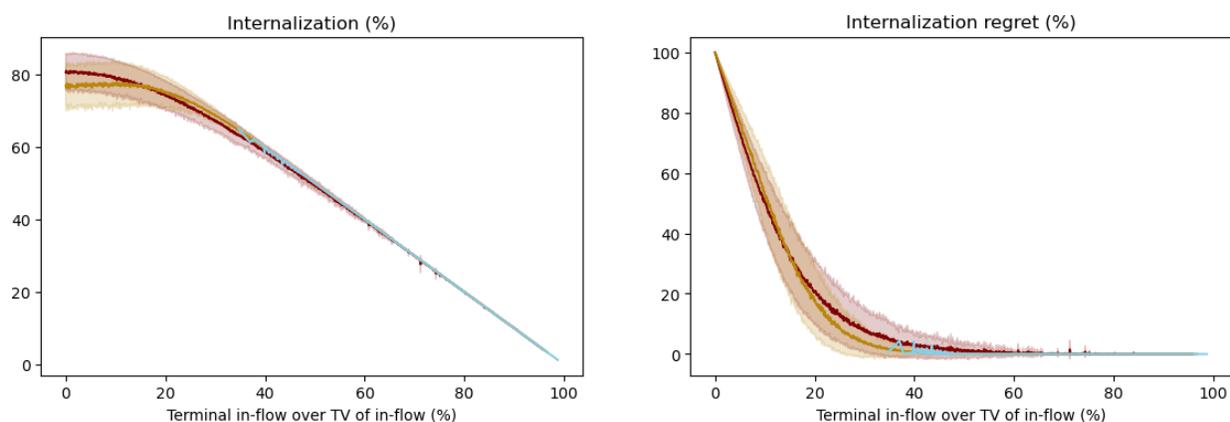
Figure 6: Average metrics' sensitivity to in-volatility (σ) for an initial inventory (z) of 3% ADV and different autocorrelation values (θ).

θ	-1	-0.5	0	0.5	1
optimal σ/z (%)	99	79	55	35	17
Internalization (%)	55	65	72	77	82

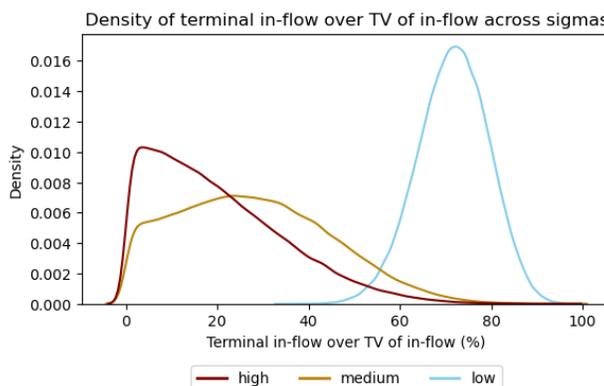
Table 2: Optimal in-flow volatility, as a percentage of inventory, across various θ .

3.4 Internalization and In-Flow Monotonicity

Opportunities for internalization depend on the in-flow path, which may trend in a single direction or net out naturally. To capture this, we propose a measure of monotonicity of in-flow paths, namely the terminal absolute value of in-flow normalized by the total variation of in-flow ($|Z_T|/\|Z\|_{TV}$). A value of 100% corresponds to an in-flow that is monotone; i.e., consist of only buy or only sell orders. Figure 7 shows that this metric is closely related to internalization under the optimal out-flow strategy. To produce Figure 7 (a), we sample 100'000 paths from our model and plot the internalization rate and regret of the out-flow versus the normalized terminal in-flow. The more “monotone” the in-flow, the lower the internalization rate and internalization regret. We emphasize that some parts of the x-axis correspond to more sample paths than others. Indeed, the distribution of $|Z_T|/\|Z\|_{TV}$ depends on the in-flow volatility (σ), as shown in Figure 7 (b).



(a) Normalized terminal in-flow ($|Z_T|/\|Z\|_{TV}$) is a strong predictor of internalization.



(b) In-flow volatility affects the distribution of the normalized terminal in-flow ($|Z_T|/\|Z\|_{TV}$).

Figure 7: (a) Monotonicity test: trading metrics as a function of the terminal in-flow state, as a fraction of in-flow TV. The error bars illustrate one standard deviation around the average. (b) The distribution of the proposed metric depends on in-flow volatility.

In practice, the shape of the in-flow path depends largely on the desk’s clients. The strong relationship in Figure 7 demonstrates that the nature of the clients—and not only the desk’s unwind strategy—is a key determinant for the internalization rate. Therefore, the trading team gains a

lot from understanding their clients. Furthermore, our proposed monotonicity metric is a strong, model-free client metric forecasting achievable internalization rates.

3.5 Misspecification Costs

This section quantifies the cost of *not* knowing your client, or more precisely, the cost of trading with the wrong in-flow parameter θ . Thanks to the analytic structure provided by Theorem 2.2, this misspecification cost can be computed in two steps: first, compute the strategy under a given (misspecified) estimate $\hat{\theta}$, then, simulate the state processes under a ground-truth value θ^* . The strategy becomes suboptimal unless $\hat{\theta} = \theta^*$.

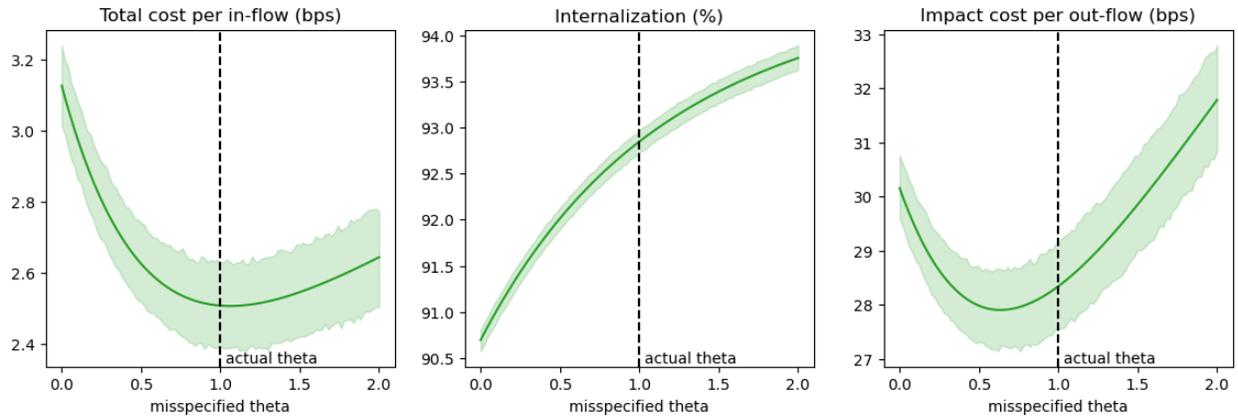
Figure 8 studies three misspecification examples with the following key conclusions:

- (i) Misspecification costs are highest for momentum and lowest for mean-reverting flow. Getting θ wrong by 100% increases costs by less than a third for mean-reverting flow. By contrast, misspecifying θ by 100% for momentum flow more than doubles costs.
- (ii) Misspecification costs are asymmetric: for a given θ^* , one would rather over- than underestimate θ . This is in line with the misspecification behavior for impact cost in [35]: generally, if a model parameter is uncertain, one prefers to err on the side which leads to more conservative trading.
- (iii) Both internalization and impact costs increase drastically when $\hat{\theta}$ moves away from θ^* . Therefore, both channels drive the misspecification costs.

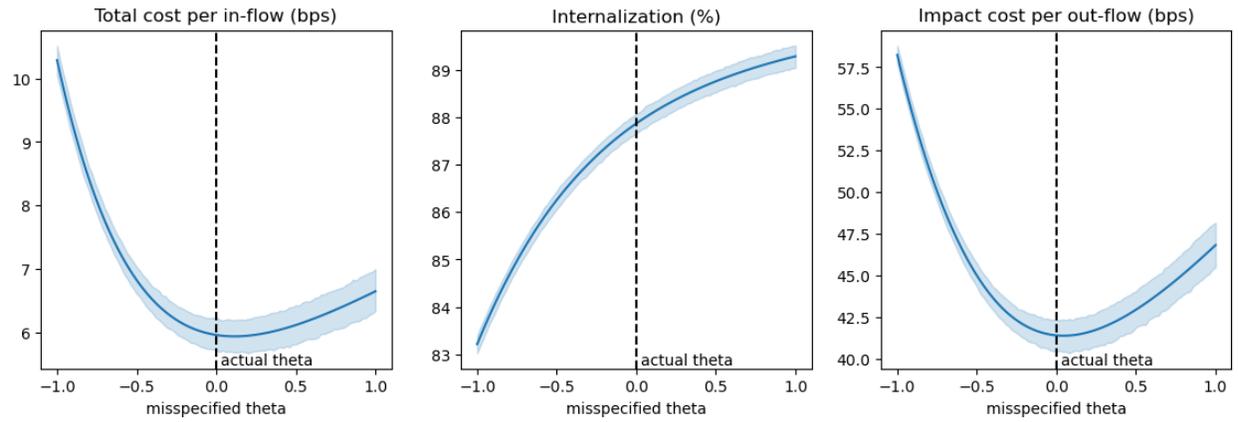
One concrete application of this misspecification study relates to estimates of in-flow autocorrelation. All else being equal, given a statistical confidence interval over θ (or any other measure of in-flow autocorrelation), it is best to deploy the most conservative trading strategy, that is, the strategy that underestimates in-flow momentum.

There is no misspecification cost for σ since the feedback form of the optimal strategy does not depend on σ (cf. Proposition 2.3 (iv)). Of course, the expected cost is nevertheless affected, hence the quote given to the client may bear a misspecification error.

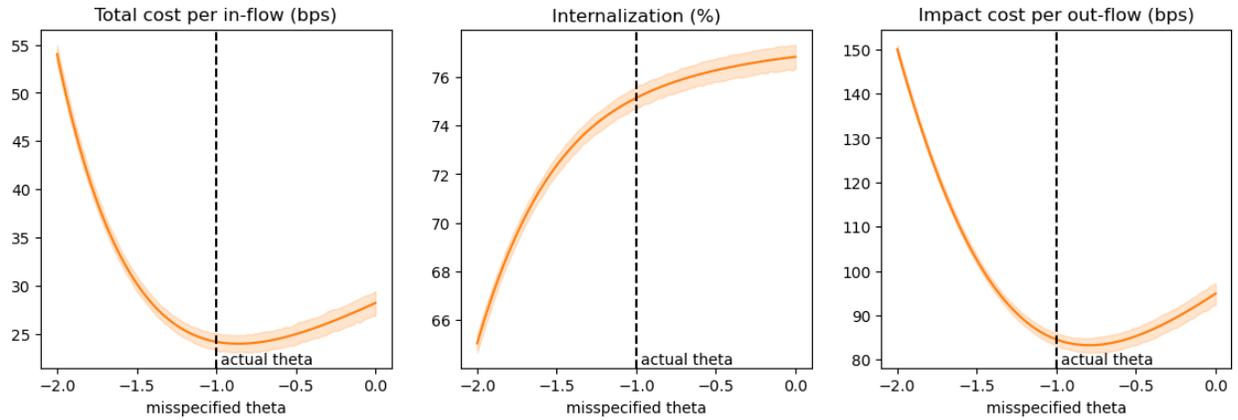
We do not conduct a misspecification study for the liquidity parameters; the existing results [44, 35, 43] apply to our setting.



(a) $\theta^* = 1$. The actual in-flow is mean-reverting.



(b) $\theta^* = 0$. The actual in-flow is a martingale.



(c) $\theta^* = -1$. The actual in-flow has momentum.

Figure 8: Simulated cost per in-flow, internalization, and impact per out-flow across various θ^* , $\hat{\theta}$.

A Further Numerical Results

A.1 Recovering Deterministic Results; Sensitivity to Spread Cost Parameter

This section serves two purposes. First, we show how our model links to the traditional optimal execution literature where in-flow is deterministic. Second, we analyze the sensitivity of our model to the spread cost parameter (ε). The setup is as described in Section 3.1.

The first experiment is best understood as a minimal extension of the traditional optimal execution literature. We consider the deterministic case with no in-flow apart from the initial order of $z = 10\%$ ADV (i.e., $\sigma = \theta = 0$). Figure 9 reproduces the optimal trading strategy of Obizhaeva and Wang in the limit as the spread cost parameter $\varepsilon \rightarrow 0$. As discussed in [50], the optimal impact state tracks a certain target value in the Obizhaeva–Wang model, with jumps at the start and end of the day to hit that target. In our model with instantaneous cost during LOB trading, the intraday impact state is smoothened, but the change in target is substantial only for high spread cost parameter.

In the out-flow and inventory, we observe that the size of the block trades in the auctions increases with the spread cost: as there is no spread cost during the auctions, trading a larger volume during the auctions is a way evade high spread cost in LOB trading, partly at the expense of higher impact costs. That also explains why the impact state initially overshoots the target value.

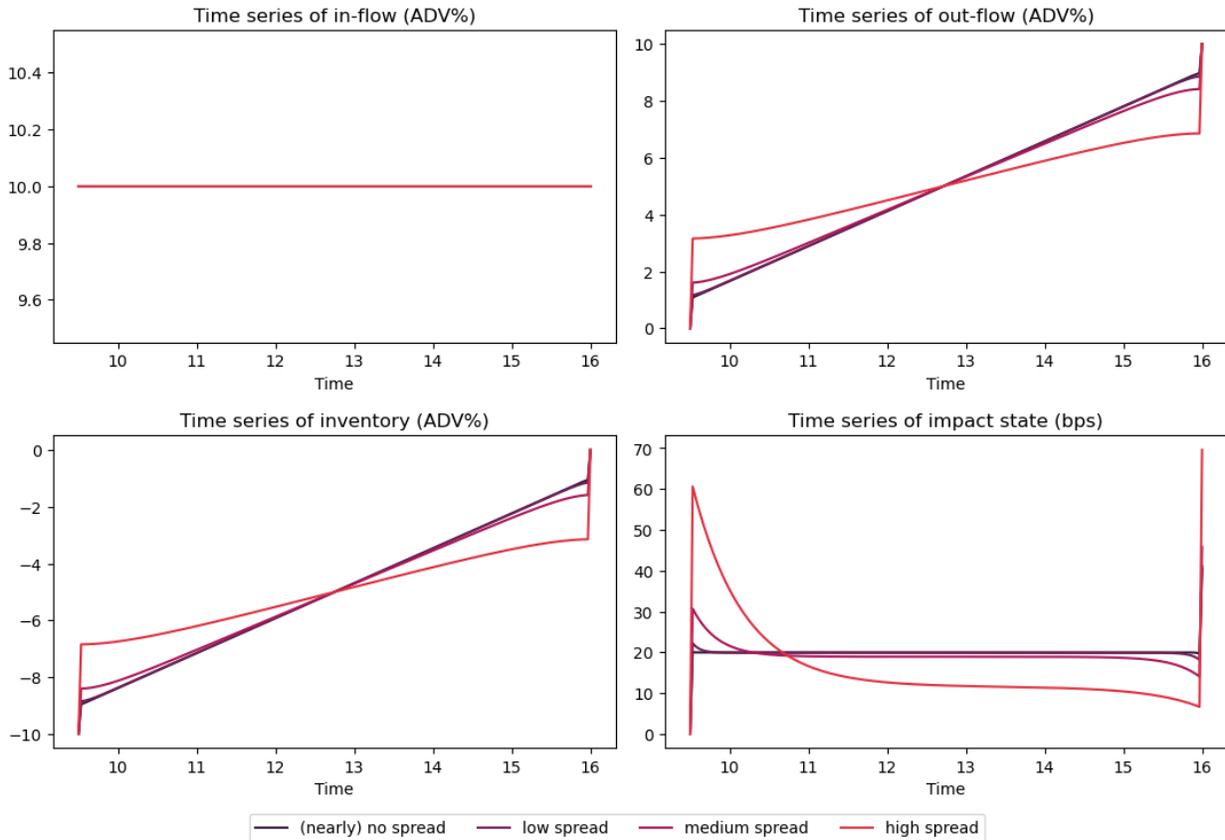


Figure 9: Sensitivity to ε of the optimal execution strategy for a 10% ADV order: path.

Table 3 quantifies some of these insights and highlights the importance of the impact-to-spread

cost ratio. When impact costs outweigh spread costs, ε 's effect is second order and does not meaningfully affect the impact state or trading speed. This is the typical regime for a desk dealing with large orders. However, once spread costs become comparable to impact costs, the trading regime changes. Overall costs increase dramatically and substantial part of the inventory is unwound by block trades on open and close.

Parameter scans	ε	In-flow (ADV%)	Spread Cost (bps)	Impact cost (bps)	Closing trade (ADV%)	Impact to spread ratio
(nearly) no spread	1e-4	10	0.0	21.1	1.0	668
low spread	1e-3	10	0.3	21.3	1.1	71
medium spread	1e-2	10	2.4	22.8	1.6	9
high spread	1e-1	10	7.5	36.9	3.2	5

Table 3: Sensitivity to ε of the optimal execution strategy for a 10% ADV order: metrics.

Next, we study the sensitivity of daily metrics to ε when the in-flow is random. Connecting to Section 3.2, Figure 10 also shows the sensitivity with autocorrelation (θ). As expected, a higher spread cost parameter ε leads to higher internalization, lower internalization regret, and larger closing trades as it slows down the continuous trading. This holds for all values of θ , with the ordering among reverting, martingale, and momentum flow being as expected. The main take-away, however, is that internalization and impact cost are relatively insensitive to ε as long as values are moderate. Internalization regret is only slightly more sensitive (note that ε is being varied over two orders of magnitude in the figure).

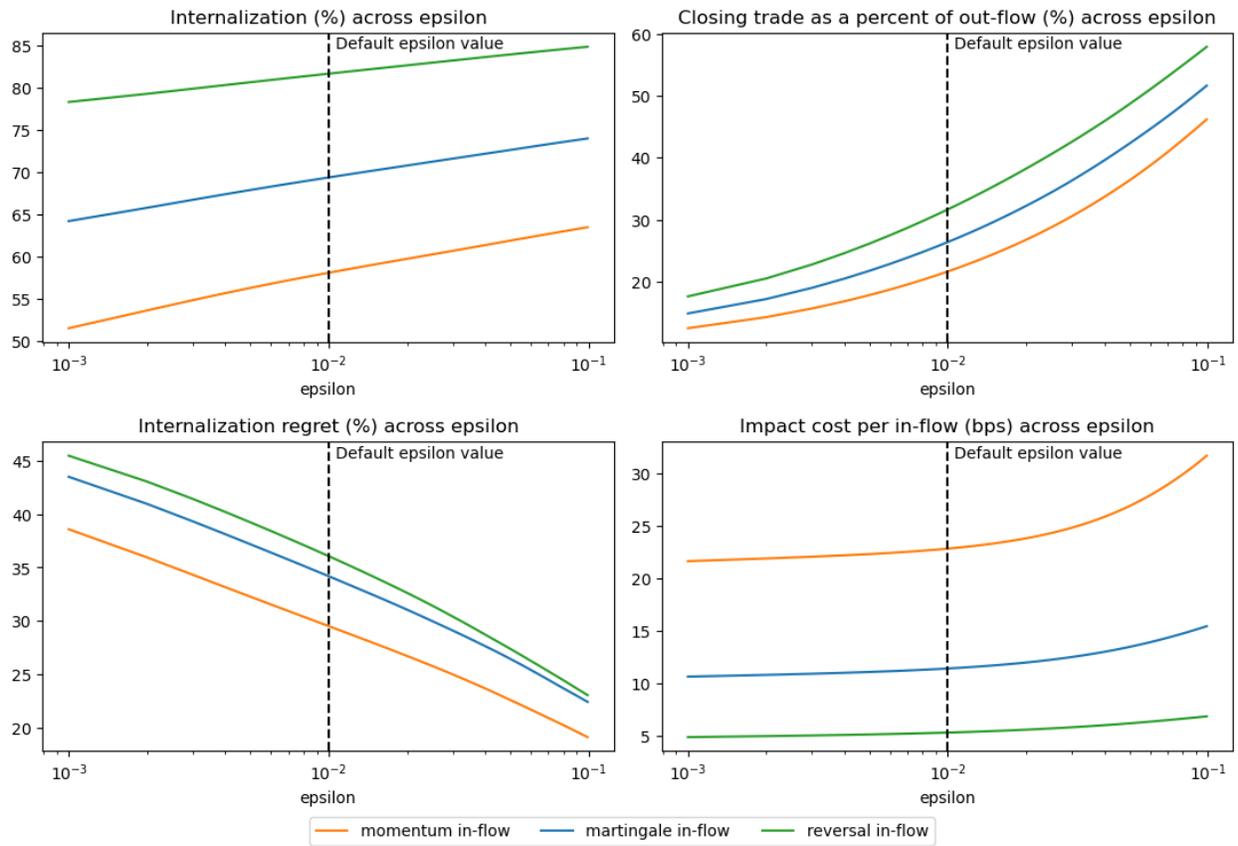


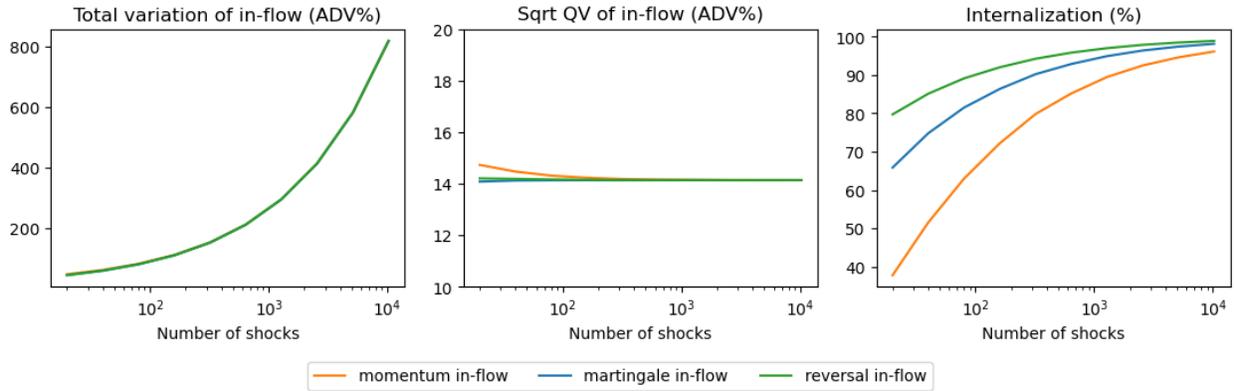
Figure 10: Average metrics' sensitivity to spread costs (ϵ) for different flow autocorrelations (θ), for an initial inventory (z) and daily in-flow volatility (σ) of 10% ADV.

A.2 Itô Diffusion Limit for In-Flows

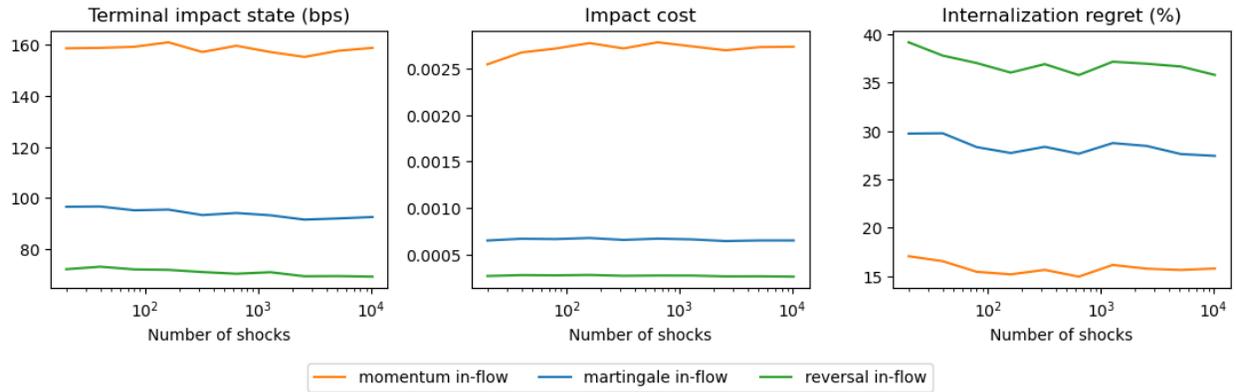
This section revisits the driving martingale of the in-flow process Z . As seen in Section 2.4, the model is solvable for finite-variation (jump) processes, diffusions, and combinations thereof. The feedback form of the optimal strategy does not depend on the driving martingale, but of course the actual trade path depends on the realization of the state processes. Our preceding numerical experiments use a finite-variation process so that trading metrics can be reported in a standard fashion. Here, we study how the metrics behave in the diffusion limit; i.e., as the number of jumps increases (while their size decreases such as to keep the expected quadratic variation constant).

Figure 11 shows the average metrics. The overarching conclusion is that the trading strategies as well as trading metrics remain similar in the diffusive limit, with the obvious exception of those metrics that explicitly depend on the total variation of in-flow. In particular, internalization regret is fairly stable, while the internalization rate tends to 100% due to its definition (3.1).

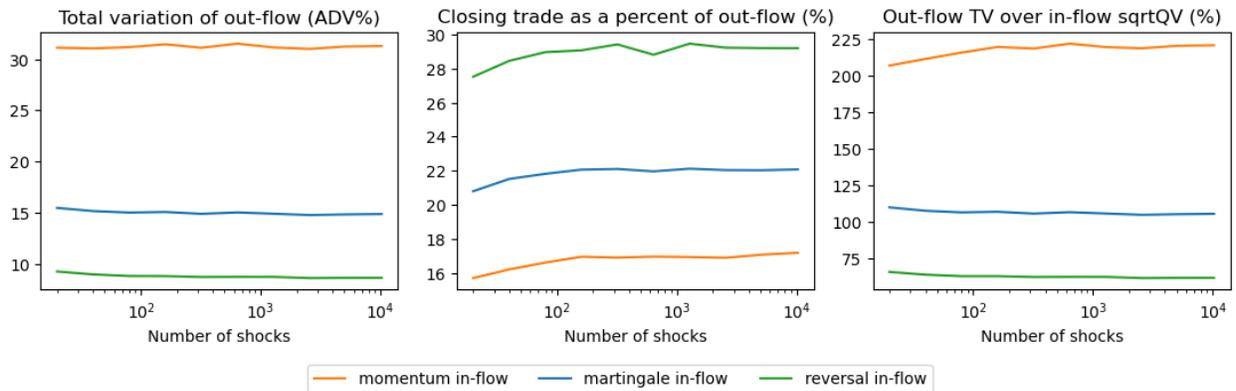
As an ad-hoc adaptation, we propose to use the square-root of the quadratic variation to replace the total variation for use in Brownian model, as it remains stable over the number of jumps.



(a) Total variation, quadratic variation, and internalization change in the passage to the diffusion limit.



(b) However, trading metrics that do not rely on total variation of in-flow remain stable.



(c) The square root of the quadratic variation of in-flow proxies well for the in-flow activity in the transition from jump process to Itô process.

Figure 11: Sensitivity of average trading metrics to the diffusive limit for in-flow.

A.3 Sensitivity to the Initial Impact State

The optimal execution literature often assumes zero initial impact, and then shows that trade paths are monotone. Clearly paths need not be monotone when the in-flow is stochastic, as the inventory can change sign. Beyond that, we emphasize in this section that the initial impact state matters (even in deterministic execution). This is straightforward but sometimes overlooked: The impact state forecasts reversal of the price, hence is equivalent to an alpha signal. Since the strategy maximizes P&L, it balances this signal with the other costs of unwinding, which can lead to trades that *add* to the inventory rather than reducing it. Analytically, the formulas for J_0 and q_t in Theorem 2.2 clearly show that trades can have either sign depending on the level of the impact state.

Trading based on past days' impact is far from an unwanted behavior in practice. Indeed, one challenge in trading systems is to correctly consider past days' trading in today's execution. For instance, [37, 34, 50] tackle the issue from a transaction cost analysis (TCA) perspective. Unfortunately, the common solution of ignoring past days' trades leads to biases and sub-optimal trading. [13] coins this problem "issuer bias" and highlights it as one of the four most common TCA and execution challenges among practitioners. While our solution is not unique in solving "issuer bias", it considers past days' trading with no additional effort within a consistent, tractable framework.

Figure 12 shows average trading paths from simulations with different initial impact states ($y = -30, 0, 30, 100$ bps) and martingale in-flow. A negative value is favorable in the sense that the price of a trade offsetting the initial in-flow $z > 0$ is reduced. Indeed, in the favorable simulation, the unwind strategy hedges a sizable part of z in the opening auction. For increasing values of y , the block trade is reduced; volume is deferred to exploit impact reversal. For high initial impact, reducing the impact state takes priority over reducing inventory and the opening trade adds to the initial inventory rather than hedging it. The strategy then trades aggressively during the day to unwind the acquired position.

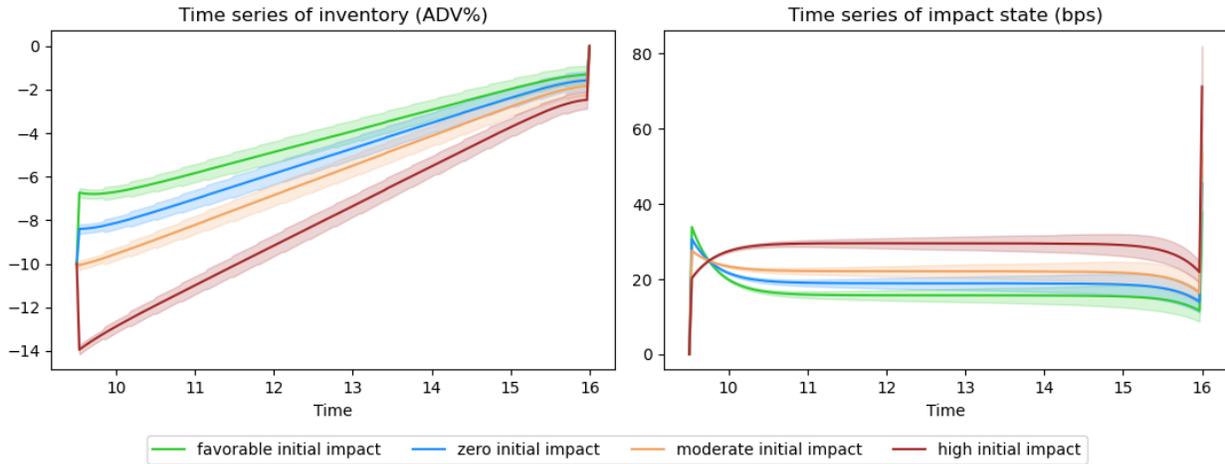


Figure 12: Sensitivity to initial impact state (y) for an initial inventory (z) and daily flow volatility (σ) of 10% ADV: average metrics.

A.4 A Representative Market-Maker Model for Stochastic Order flow

This section illustrates possible θ values from public market data. We preface this study with two important caveats to highlight that these estimates are purely presented for illustration purposes.

First, our estimates of market flow autocorrelation are simple, and there are many more sophisticated models and studies. Second, *know your client*. Trading teams should calibrate stochastic models on their own in-flows and compare those to the representative market maker; they should not assume their in-flows behave like the aggregate market. For instance, competitive market makers may capture milder order flows, while less competitive desks may only capture toxic order flow, leading to more inventory risk. Similarly, a CRB's orders may not include smaller retail trades present on the tape.

The representative market-maker approach is a standard model for stochastic order flow in financial markets. For instance, market-making applications in [21] and the references therein assume that a trading strategy can capture a certain proportion of the overall market order flow. Empirical papers leverage the representative market-maker approach to provide statistics on market flow for trading applications. For example, [18] quantify market-maker's inventories and the follow-up paper [17] tracks individual traders' flows. This section uses the representative market-maker approach to calibrate the proposed OU model for stochastic order flow on public market data. Therefore, the study discretizes the model for Z using a step size Δt . Furthermore, σ is normalized to one, as the representative market maker by definition captures 100% of the day's trading volume:

$$Z_{t+\Delta t} = -\theta Z_t \Delta t + W_{t+\Delta t} - W_t. \quad (\text{A.1})$$

Figure 13 plots the *doubling time* $\log(2)/\theta$ as a function of the discretization choice Δt . The doubling time is to momentum flow what the half-life is to reverting flow. For each stock in the S&P 500, the study estimates a single θ per stock over the year 2019, and the figure plots the corresponding distribution of doubling times across stocks. Our data source is LOBSTER [39], which provides all Nasdaq limit order book events in a standard table format and is aimed at academic research.

In line with [13] and the references therein, we find that the market order flow consistently exhibits momentum across the S&P 500 for time steps ranging from ten seconds to one hour. The corresponding doubling times are of the order of a couple of hours, to be compared to the roughly one-hour half-life of the OW model, as per [42]. Again, in line with [14], order flow autocorrelation is roughly of the same magnitude as impact's decay, leading to approximately martingale impact state.

Furthermore, our study finds that θ depends on the discretization choice, in line with [13, 49] and the references therein, where order flow is found to have long-range autocorrelation and, therefore, cannot be calibrated on multiple time scales with a single OU model. Intuitively, different market participants, e.g., mutual funds, hedge funds, HFTs, submit orders over different timescales, e.g., minutes to days, and the aggregate order flow autocorrelation reflects each timescale. Furthermore, [16] empirically link client timescales to liquidity providers' internalization strategies.

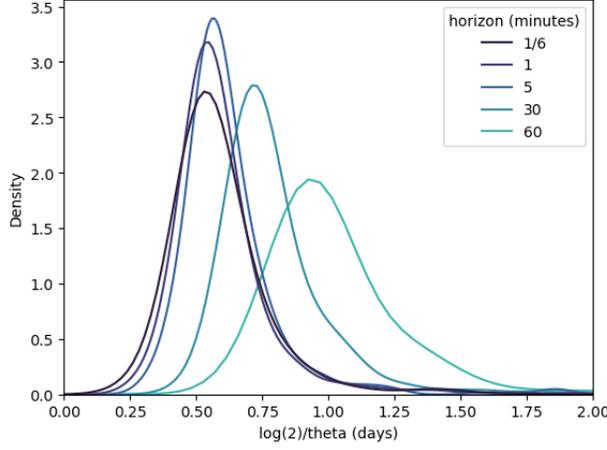


Figure 13: Doubling time $\log(2)/\theta$ as a function of the discretization choice Δt . The distribution covers the S&P 500 for the year 2019.

B Proofs for Sections 2.2–2.4: Dynamic Programming

Proof of Proposition 2.4. By its definition (2.10), $\mathcal{C}(J_0, q, J_T)$ equals the sum of two expectations. First, $\mathbb{E}[\int_0^T S_t dQ_t]$, which equals $\mathbb{E}[S_T Z_T]$ by Lemma B.1. And second,

$$\mathbb{E} \left[\int_0^T Y_t q_t dt + \frac{1}{2}(Y_{0-} + Y_0)J_0 + \frac{1}{2}(Y_{T-} + Y_T)J_T + \frac{1}{2} \int_0^T \varepsilon_t q_t^2 dt \right].$$

Here Lemma B.2 shows that

$$\begin{aligned} & \int_0^T Y_t q_t dt + \frac{1}{2}(Y_{0-} + Y_0)J_0 + \frac{1}{2}(Y_{T-} + Y_T)J_T \\ &= \int_0^T \frac{2\beta_t + \dot{\gamma}_t}{2\lambda_t} Y_t^2 dt + \frac{1}{2} \left(\frac{Y_T^2}{\lambda_T} - \frac{Y_0^2}{\lambda_0} + \frac{Y_0^2}{\lambda_{0-}} - \frac{y^2}{\lambda_{0-}} \right). \end{aligned} \quad (\text{B.1})$$

Moreover, using (2.4) and the liquidation constraint,

$$Y_T = Y_{T-} + \lambda_T J_T = Y_{T-} - \lambda_T X_T = Y_T^c - \lambda_T X_T.$$

Substituting this into (B.1) and noting that $Y_t = Y_t^c$ for $t \in [0, T)$, the claim follows. \square

Lemma B.1. *Let $Q \equiv (J_0, q, J_T) \in \mathcal{A}$. Then*

$$\mathbb{E} \left[\int_0^T Q_t dS_t \right] = \mathbb{E} \left[S_0 J_0 + \int_0^T S_t q_t dt + S_T J_T \right] = \mathbb{E}[S_T Z_T]. \quad (\text{B.2})$$

Proof. Note that $Q_t^c := \int_0^t q_s ds$ is a continuous process of finite variation. The integration-by-parts formula and $Q_0^c = 0$ then yield that

$$\int_0^T S_t q_t dt = \int_0^T S_t dQ_t^c = S_T Q_T^c - \int_0^T Q_t^c dS_t.$$

By the definition of admissibility, the last integral is a true martingale, hence has zero expectation. Admissibility also includes the liquidation constraint $J_0 + Q_T^c + J_T = Z_T$ and thus

$$S_0 J_0 + S_T Q_T^c + S_T J_T = S_0 J_0 + S_T (Z_T - J_0) = -(S_T - S_0) J_0 + S_T Z_T.$$

Here $\mathbb{E}[(S_T - S_0)J_0] = 0$ as S is a martingale and $J_0 \in \mathbb{R}$. The claim follows by combining the displays and taking expectation. \square

Lemma B.2. *Let $Q \equiv (J_0, q, J_T) \in \mathcal{A}$. Then*

$$\int_0^T Y_t q_t dt = \int_0^T \frac{2\beta_t + \dot{\gamma}_t}{2\lambda_t} Y_t^2 dt + \frac{1}{2} \left(\frac{Y_{T-}^2}{\lambda_T} - \frac{Y_0^2}{\lambda_0} \right). \quad (\text{B.3})$$

Moreover,

$$(Y_{0-} + Y_0)J_0 = \frac{Y_0^2}{\lambda_{0-}} - \frac{y^2}{\lambda_{0-}} \quad \text{and} \quad (Y_{T-} + Y_T)J_T = \frac{Y_T^2}{\lambda_T} - \frac{Y_{T-}^2}{\lambda_T}. \quad (\text{B.4})$$

Proof. We first focus on $t \in (0, T)$. Note that Y as defined in (2.3) is continuous and has finite variation on $(0, T)$ when Q is of the form (2.5). Using $\lambda_t = e^{\gamma t}$ and $q_t dt = e^{-\gamma t}(dY_t + \beta_t Y_t dt)$, we have

$$\int_0^T Y_t q_t dt = \int_{(0, T)} e^{-\gamma t} Y_t dY_t + \int_0^T e^{-\gamma t} \beta_t Y_t^2 dt.$$

Applying integration-by-parts to $e^{-\gamma t} Y_t^2$ over $t \in (0, T)$ yields that

$$2 \int_{(0, T)} e^{-\gamma t} Y_t dY_t = e^{-\gamma T} Y_{T-}^2 - e^{-\gamma 0} Y_{0+}^2 + \int_0^T \dot{\gamma}_t e^{-\gamma t} Y_t^2 dt.$$

Since Y, λ are right-continuous at $t = 0$ and λ is also left-continuous at $t = T$, the claim (B.3) follows by combining the displays. The second claim, (B.4), follows by specializing

$$(Y_t + Y_{t-})(Q_t - Q_{t-}) = \frac{1}{\lambda_{t-}} (Y_t + Y_{t-})(Y_t - Y_{t-}) = \frac{1}{\lambda_{t-}} (Y_t^2 - Y_{t-}^2)$$

to $t \in \{0, T\}$ and recalling that $\lambda_{T-} = \lambda_T$. \square

Proof of Proposition 2.6. Omitting arguments for brevity, the generator of the controlled process is

$$\begin{aligned} \mathcal{L}^q v &:= (q + \theta z) \partial_x v + (-\beta y + \lambda q) \partial_y v - \theta z \partial_z v + \frac{\sigma^2}{2} (\partial_{xx} v - 2\partial_{xz} v + \partial_{yy} v) \\ &= (\partial_x v + \lambda \partial_y v) q - \beta \partial_y v y + \theta (\partial_x v - \partial_z v) z + \frac{\sigma^2}{2} (\partial_{xx} v - 2\partial_{xz} v + \partial_{yy} v). \end{aligned}$$

The running cost is $c(t, x, y, z, q) = \frac{2\beta_t + \dot{\gamma}_t}{2\lambda_t} y^2 + \frac{\varepsilon_t}{2} q^2$, so that the HJB equation $-\partial_t v - \inf_q (\mathcal{L}^q v + c) = 0$ reads

$$\begin{aligned} -\partial_t v - \inf_{q \in \mathbb{R}} \left[(\partial_x v + \lambda_t \partial_y v) q + \frac{\varepsilon_t}{2} q^2 \right] + \beta_t \partial_y v y - \frac{2\beta_t + \dot{\gamma}_t}{2\lambda_t} y^2 \\ - \theta_t (\partial_x v - \partial_z v) z - \frac{\sigma_t^2}{2} (\partial_{xx} v - 2\partial_{xz} v + \partial_{yy} v) = 0. \end{aligned} \quad (\text{B.5})$$

The candidate optimal control in feedback form is given by the minimizer in (B.5),

$$q_t^* = -\varepsilon_t^{-1} (\partial_x v + \lambda_t \partial_y v),$$

which yields

$$-\inf_{q \in \mathbb{R}} \left[(\partial_x v + \lambda_t \partial_y v) q + \frac{\varepsilon_t}{2} q^2 \right] = \frac{1}{2\varepsilon_t} (\partial_x v + \lambda \partial_y v)^2.$$

If $v \in C^{1,2}([0, T] \times \mathbb{R}^3)$ is defined by (2.18) with A, \dots, F solving (2.20), one verifies by direct calculation that v solves the HJB equation (B.5) as well as the terminal condition

$$v(T, x, y, z) = \frac{1}{2} \left(\lambda_T x^2 - 2xy + \frac{y^2}{\lambda_T} \right). \quad (\text{B.6})$$

corresponding to the terminal cost. A standard verification argument then yields that v is the value function and q_t^* gives the optimal strategy.¹¹ Uniqueness of the optimal control holds by strict convexity of (C.1). \square

Proof of Proposition 2.7. We show that the Riccati system (2.20) has a unique solution on $[0, T]$ and that $A_t, C_t \geq 0$ for $t \in [0, T]$.

Consider first the (autonomous) ODE system for (A, B, C) . It can be stated as the matrix Riccati equation

$$\dot{\mathbf{P}}_t = \mathbf{P}_t \mathbf{B}_t \mathbf{N}_t^{-1} \mathbf{B}_t' \mathbf{P}_t - \mathbf{A}_t' \mathbf{P}_t - \mathbf{P}_t \mathbf{A}_t - \mathbf{C}_t, \quad t \in [0, T]$$

with the terminal condition

$$\mathbf{P}_T = \frac{1}{2} \begin{bmatrix} \lambda_T & -1 \\ -1 & 1/\lambda_T \end{bmatrix},$$

where

$$\mathbf{P}_t := \frac{1}{2} \begin{bmatrix} A_t & B_t \\ B_t & C_t \end{bmatrix}, \quad \mathbf{N}_t := \frac{1}{2} [\varepsilon_t]$$

$$\mathbf{A}_t := \begin{bmatrix} 0 & 0 \\ 0 & -\beta_t \end{bmatrix}, \quad \mathbf{B}_t := \begin{bmatrix} 1 \\ \lambda_t \end{bmatrix}, \quad \mathbf{C}_t := \begin{bmatrix} 0 & 0 \\ 0 & \frac{2\beta_t + \gamma_t}{2\lambda_t} \end{bmatrix}.$$

We note that \mathbf{P}_T is nonnegative definite. As the coefficients are bounded and ε_t was assumed to be bounded away from zero, [51, Theorem 2.1(i)(ii)] shows that there exists a unique solution \mathbf{P}_t and that \mathbf{P}_t is nonnegative definite for all for $t \in [0, T]$. In particular, this yields $A_t \geq 0$ and $C_t \geq 0$.

We can now consider A, B, C as given bounded functions. The ODE system for (D, E, F, K) in (2.20) is then linear with bounded coefficients, hence has a unique solution. \square

Proof of Proposition 2.8. (i) Nonnegative definiteness was established in the proof of Proposition 2.7. We show that $C_t \leq \lambda_t^{-1}$ and $B_t + \lambda_t C_t > 0$ for $t \in [0, T]$.

Recall that the solution to a linear, first-order ODE $\dot{y}_t = p_t y_t - \ell_t$ with terminal condition y_T is given by

$$y_t = e^{-\int_t^T p_s ds} y_T + \int_t^T e^{-\int_t^s p_u du} \ell_s ds. \quad (\text{B.7})$$

¹¹To be precise, this requires q^* to be admissible. While q^* is automatically square-integrable, we also need that $\int_0^t Q_s^c dS_s$ is a true martingale, where $Q_t^c := \int_0^t q_s^* ds$. This is an integrability condition on S , and is satisfied, e.g., if $dS_t = \nu_t dB_t$ where B is a Brownian motion and ν is bounded.

In particular, $y_t \geq 0$ for $t \in [0, T)$ if $y_T \geq 0$ and $\ell_t \geq 0$ for $t \in [0, T)$, and $y_t > 0$ for $t \in [0, T)$ if additionally either $y_T > 0$ or $\ell_t > 0$ for t in a non-degenerate interval of $[0, T)$.

We first prove that $C_t^\dagger := -C_t + \lambda_t^{-1} \geq 0$ for $t \in [0, T]$. Indeed,

$$\dot{C}_t^\dagger = -\dot{C}_t - \lambda_t^{-1} \dot{\gamma}_t = 2\beta_t C_t^\dagger - \varepsilon_t^{-1} (B_t + \lambda_t C_t)^2, \quad C_T^\dagger = 0.$$

Seeing this as a special case of (B.7) with $\ell_t := \varepsilon_t^{-1} (B_t + \lambda_t C_t)^2 \geq 0$, we deduce $C_t^\dagger \geq 0$ for $t \in [0, T]$.

Lastly, define $M_t := A_t + \lambda_t B_t$ and $N_t := B_t + \lambda_t C_t = B_t - \lambda_t C_t^\dagger + 1$. Then

$$\begin{aligned} \dot{N}_t &= \dot{B}_t - \lambda_t (\dot{C}_t^\dagger + \dot{\gamma}_t C_t^\dagger) \\ &= \varepsilon_t^{-1} M_t N_t + \beta_t B_t + \lambda_t \varepsilon_t^{-1} N_t^2 - 2\lambda_t \beta_t C_t^\dagger - \lambda_t \dot{\gamma}_t C_t^\dagger \\ &= (\varepsilon_t^{-1} M_t + \lambda_t \varepsilon_t^{-1} N_t + \beta_t) N_t - \lambda_t [(2\beta_t + \dot{\gamma}_t) C_t^\dagger + \beta_t C_t], \quad N_T = 0 \end{aligned}$$

where we used $B_t = N_t - \lambda_t C_t$ from the second to the third line. Recall that $C_t \geq 0, \beta_t > 0$ and $2\beta_t + \dot{\gamma}_t > 0$. As $C_T = \lambda_T^{-1} > 0$, for t close to T we have $C_t > 0$ and thus $\ell_t := (2\beta_t + \dot{\gamma}_t) C_t^\dagger + \beta_t C_t > 0$. Using again (B.7), we deduce $N_t > 0$ for $t \in [0, T)$.

(ii) Let $B_t^\dagger := -B_t$, $M_t := A_t + \lambda_t B_t$ and $N_t := B_t + \lambda_t C_t$. To show that $B_t < 0$ and $M_t > 0$ for $t \in [0, T)$, we first show that $(B_t^\dagger, M_t) \in [0, \infty)^2$. Indeed, we have the system

$$\begin{cases} \dot{B}_t^\dagger = \beta_t B_t^\dagger - \varepsilon_t^{-1} M_t N_t, & B_T^\dagger = 1 \\ \dot{M}_t = \varepsilon_t^{-1} (M_t + \lambda_t N_t) M_t - \lambda_t (\beta_t + \dot{\gamma}_t) B_t^\dagger, & M_T = 0, \end{cases} \quad (\text{B.8})$$

which is of the general form $(\dot{B}_t^\dagger, \dot{M}_t) = \Phi(t, B_t^\dagger, M_t)$. Here the vector field $-\Phi$ is inward-pointing on the boundary of the first quadrant $[0, \infty)^2$; that is, $-\Phi_1(t, b, m) \in [0, \infty)$ whenever $b = 0, m \geq 0$, and $-\Phi_2(t, b, m) \in [0, \infty)$ whenever $b \geq 0, m = 0$. In view of the terminal condition $(B_T^\dagger, M_T) = (1, 0) \in [0, \infty)^2$, it follows that $(B_t^\dagger, M_t) \in [0, \infty)^2$ for all $t \in [0, T]$. We can now apply (B.7) together with $\beta_t + \dot{\gamma}_t > 0$ and $B_T^\dagger = 1$ to deduce $(B_t^\dagger, M_t) \in (0, \infty)^2$ for all $t \in [0, T)$.

Next, we set $V_t^\dagger := M_t - (D_t + \lambda_t E_t)$, $W_t^\dagger := -B_t + E_t$ and consider

$$\begin{cases} \dot{V}_t^\dagger = (\varepsilon_t^{-1} M_t + \lambda_t \varepsilon_t^{-1} N_t + \theta_t) V_t^\dagger - \lambda_t (\beta_t + \dot{\gamma}_t) W_t^\dagger, & V_T^\dagger = 0 \\ \dot{W}_t^\dagger = (\theta_t + \beta_t) W_t^\dagger - \varepsilon_t^{-1} N_t V_t^\dagger, & W_T^\dagger = 1. \end{cases} \quad (\text{B.9})$$

Writing this as $(\dot{V}_t^\dagger, \dot{W}_t^\dagger) = \Phi(t, V_t^\dagger, W_t^\dagger)$, we see that $-\Phi$ is inward-pointing on the boundary of the first quadrant. Repeating the argument above, we deduce that $(V_t^\dagger, W_t^\dagger) \in (0, \infty)^2$ for all $t \in [0, T)$.

(iii) Lastly, we establish the comparison result regarding θ . Let $\tilde{\theta} : [0, T] \rightarrow \mathbb{R}$ is measurable and bounded, and define \tilde{D}, \tilde{E} like D, E but with $\tilde{\theta}$ instead of θ . Suppose that $\theta_t \geq \tilde{\theta}_t$ for $t \in [0, T]$, and define $(V_t^\dagger, W_t^\dagger), (\tilde{V}_t^\dagger, \tilde{W}_t^\dagger)$ and $\Phi, \tilde{\Phi}$ as in (B.9), where $(V_T^\dagger, W_T^\dagger) = (\tilde{V}_T^\dagger, \tilde{W}_T^\dagger) = (0, 1)$. Check that

$$\begin{aligned} -\Phi_1(t, v, w) &\leq -\tilde{\Phi}_1(t, \tilde{v}, \tilde{w}), \quad \text{for } 0 \leq v = \tilde{v}, \quad 0 \leq w \leq \tilde{w} \\ -\Phi_2(t, v, w) &\leq -\tilde{\Phi}_2(t, \tilde{v}, \tilde{w}), \quad \text{for } 0 \leq v \leq \tilde{v}, \quad 0 \leq w = \tilde{w}. \end{aligned}$$

It follows that $V_t^\dagger \leq \tilde{V}_t^\dagger$ and $W_t^\dagger \leq \tilde{W}_t^\dagger$ for $t \in [0, T]$, or equivalently,

$$D_t + \lambda_t E_t \geq \tilde{D}_t + \lambda_t \tilde{E}_t \quad \text{and} \quad E_t \leq \tilde{E}_t, \quad t \in [0, T],$$

as M_t and B_t do not depend on θ_t or $\tilde{\theta}_t$. In particular, we note that $\tilde{D}_t, \tilde{E}_t \equiv 0$ when $\tilde{\theta}_t \equiv 0$. \square

Proof of Proposition 2.12. We first note that the cost (2.21) is convex in j_0 . Indeed, it follows from Proposition 2.8 that $v(t, x, y, z)$ is convex in x and y , and thus $v(0, j_0 - z, y + \lambda_{0-}j_0, z)$ is convex in j_0 . The second term in (2.21) is also convex in j_0 as $0 < \lambda_{0-} \leq \lambda_0$.

Since (2.21) is quadratic in j_0 , we can compute $\partial_{j_0} V(j_0)$ where

$$V(j_0) := v(0, j_0 - z, y + \lambda_{0-}j_0, z) + \frac{1}{2} (\lambda_{0-}^{-1} - \lambda_0^{-1}) (y + \lambda_{0-}j_0)^2. \quad (\text{B.10})$$

Recall that

$$\partial_x v(0, x, y, z) = A_0 x + B_0 y + D_0 z, \quad \partial_y v(0, x, y, z) = B_0 x + C_0 y + E_0 z,$$

and

$$-\frac{f_{0-}}{\varepsilon_0} = A_0 + \lambda_{0-} B_0, \quad -\frac{g_{0-}}{\varepsilon_0} = B_0 + \lambda_{0-} C_0, \quad -\frac{h_{0-}}{\varepsilon_0} = D_0 + \lambda_{0-} E_0.$$

Define $\eta_{0-} := -\varepsilon_0^{-1}(1 - \lambda_{0-}/\lambda_0) \leq 0$. It is straightforward to find

$$\varepsilon_0^{-1} \partial_{j_0} v(0, j_0 - z, y + \lambda_{0-}j_0, z) = -(f_{0-} + \lambda_{0-}g_{0-})j_0 - g_{0-}y + (f_{0-} - h_{0-})z$$

and

$$\varepsilon_0^{-1} \partial_{j_0} V(j_0) = r j_0 - (g_{0-} + \eta_{0-})y + (f_{0-} - h_{0-})z \quad (\text{B.11})$$

where $r := -f_{0-} - \lambda_{0-}(g_{0-} + \eta_{0-})$. As $j_0 \mapsto v(0, j_0 - z, y + \lambda_{0-}j_0, z)$ is convex and quadratic, we must have $-f_{0-} - \lambda_{0-}g_{0-} \geq 0$. We show below that the inequality is strict, and thus $r > 0$ as $\eta_{0-} \leq 0$. It then follows that V is not affine, hence strictly convex, and the desired formula for the optimal j_0 follows by setting (B.11) to zero and solving for j_0 .

Suppose towards a contradiction that $-f_{0-} - \lambda_{0-}g_{0-} = 0$. Recall that f_{0-} and g_{0-} do not depend on y, z, θ , and consider the particular case where $y, z = 0$ and $\theta, \sigma \equiv 0$. The above calculation shows $\partial_{j_0} v(0, j_0, \lambda_{0-}j_0, 0) = 0$ and in particular $v(0, j_0, \lambda_{0-}j_0, 0) = v(0, 0, 0, 0) = 0$. In words, the round-trip strategy starting with an initial jump $j_0 \neq 0$ has zero cost, contradicting Corollary 2.5.

Next, we show $g_{0-} + \eta_{0-} < 0$. Indeed,

$$\begin{aligned} -\varepsilon_0(g_{0-} + \eta_{0-}) &= B_0 + \lambda_{0-}C_0 + (1 - \lambda_{0-}/\lambda_0) \\ &= B_0 + \lambda_0 C_0 + (1 - \lambda_{0-}/\lambda_0)(1 - \lambda_0 C_0) > 0, \end{aligned}$$

where we recall from Proposition 2.8 that $1 - \lambda_0 C_0 \geq 0$ and $B_0 + \lambda_0 C_0 > 0$.

Lastly, suppose that $\beta_t + \dot{\gamma}_t > 0$ for all $t \in [0, T]$. Then

$$\begin{aligned} -\varepsilon_0(f_{0-} - h_{0-}) &= (A_0 + \lambda_{0-}B_0) - (D_0 + \lambda_{0-}E_0) \\ &= (A_0 + \lambda_0 B_0) - (D_0 + \lambda_0 E_0) - (\lambda_0 - \lambda_{0-})(B_0 - E_0) > 0, \end{aligned}$$

where we again recall from Proposition 2.8 that $A_0 + \lambda_0 B_0 > D_0 + \lambda_0 E_0$ and $B_0 < E_0$. \square

Sketch of proof for Proposition 2.13. Given $q \in \mathcal{A}$, introduce a process V^q by $V_t^q = v(t, X, Y, Z_t)$, where v is defined in (2.18) and X, Y are the controlled states corresponding to q . Note that $V_0^q = v(0, x, y, z)$ is deterministic and independent of q ; it is the quantity V_0 in the statement. Consider also the process

$$\hat{V}_t^q := \frac{1}{2} \int_0^t \left[\frac{2\beta_s + \gamma'_s}{\lambda_s} (Y^q)_s^2 + \varepsilon_s q_s^2 \right] ds + V_t^q.$$

In view of (B.6), the expected cost of $q \in \mathcal{A}$ in (2.17) is given by $E[\hat{V}_T^q]$. A (lengthy) application of Ito's formula shows that the process \hat{V}^q is a local submartingale for any $q \in \mathcal{A}$, and a local martingale for $q = q^*$. This calculation uses the fact that (2.24) is equivalent to the generalized linear backward SDE

$$dK_t = -\frac{1}{2}(A_t - 2D_t + F_t) d[M, M]_t + dN_t, \quad K_T = 0,$$

where N is a martingale (namely, $N_t = \mathbb{E}[\int_0^T \frac{1}{2}(A_s - 2D_s + F_s) d[M, M]_s | \mathcal{F}_t]$). The definition of \mathcal{A} and an integrability condition on M are chosen such that $q^* \in \mathcal{A}$ and such that the local martingale terms are true martingales for any $q \in \mathcal{A}$. Then, the (sub)martingale properties yield

$$\mathbb{E}[\hat{V}_T^q] \geq \hat{V}_0^q = V_0 = \hat{V}_0^{q^*} = \mathbb{E}[\hat{V}_T^{q^*}].$$

This shows that q^* is optimal and that V_0 is the value. Uniqueness of the optimal control is clear from the strict convexity of the cost. \square

C Proofs for Section 2.5: Maximum Principle

Throughout this section, $\beta, \lambda, \varepsilon \in (0, \infty)$ are constant; in particular, $\lambda_{0-} = \lambda$. The proof of Theorem 2.14 is based on the stochastic maximum principle. Consider the control problem (2.17), i.e.,

$$\inf_{q \in \mathcal{Q}} \frac{1}{2} \mathbb{E} \left\{ \int_t^T \left[\frac{2\beta}{\lambda} (Y_s^c)^2 + \varepsilon q_s^2 \right] ds + \frac{1}{\lambda} (\lambda X_T - Y_T^c)^2 \right\} \quad (\text{C.1})$$

with the state processes $(X_s, Y_s^c, Z_s)_{s \in [t, T]}$ with initial time $t \in [0, T)$ and position (x, y, z) ,

$$\begin{cases} dX_s = q_s ds - dZ_s, & X_t = x \\ dY_s^c = (-\beta Y_s^c + \lambda q_s) ds, & Y_t^c = y \\ dZ_s = -\theta_s Z_s ds + \sigma_s dW_s, & Z_t = z \end{cases} \quad (\text{C.2})$$

(For brevity, we write Y instead of Y^c throughout this section.) We emphasize that (t, x, y, z) should be seen as fixed parameters in the subsequent arguments; they are introduced in order to recover the full form of the feedback strategy in Theorem 2.2.

The maximum principle (e.g., [46, Theorem 6.4.6]) yields a BSDE system on $[t, T]$,

$$\begin{cases} -dU_s^X = -V_s^X dW_s, & U_T^X = \lambda X_T - Y_T \\ -dU_s^Y = (-\beta U_s^Y + \frac{2\beta}{\lambda} Y_s) ds - V_s^Y dW_s, & U_T^Y = -X_T + \lambda^{-1} Y_T \\ -dU_s^Z = \theta_s (U_s^X - U_s^Z) ds - V_s^Z dW_s, & U_T^Z = 0, \end{cases} \quad (\text{C.3})$$

with optimality condition

$$U_s^X + \lambda U_s^Y + \varepsilon q_s^* = 0, \quad s \in [t, T] \quad (\text{C.4})$$

along the optimal state processes. Here $(q_s^*)_{s \in [t, T]}$ denotes the optimal strategy in open-loop form (i.e., as a function of $\omega \in \Omega$). In particular, (C.4) and the terminal conditions in (C.3) together yield that

$$q_T^* = 0. \quad (\text{C.5})$$

We note that the maximum principle applies rigorously in our context, as we have already established that the value function is smooth (e.g., [46, Theorem 6.4.7]).

Recall from Proposition 2.3 that the desired functions f_t, g_t, h_t do not depend on σ_t . Thus, we may specialize to $\sigma \equiv 0$ without loss of generality, rendering the in-flow process Z deterministic. The BSDE system (C.3) then reduces to the backward ODE system

$$\begin{cases} -\dot{U}_s^X = 0, & U_T^X = \lambda X_T - Y_T \\ -\dot{U}_s^Y = -\beta U_s^Y + \frac{2\beta}{\lambda} Y_s, & U_T^Y = -X_T + \lambda^{-1} Y_T \\ -\dot{U}_s^Z = \theta_s(U_s^X - U_s^Z), & U_T^Z = 0. \end{cases} \quad (\text{C.6})$$

(The above steps also apply in the case of time-varying coefficients, but we do not expect a closed-form solution in that situation.) We recall that

$$\tilde{\varepsilon} := \frac{\varepsilon\beta}{2\lambda}, \quad \kappa := \beta\sqrt{1 + \tilde{\varepsilon}^{-1}} > \beta, \quad (\text{C.7})$$

and note the relations

$$\frac{1}{\kappa \pm \beta} = \frac{1}{\beta(\sqrt{1 + \tilde{\varepsilon}^{-1}} \pm 1)} = \frac{\tilde{\varepsilon}}{\beta} \left(\sqrt{1 + \tilde{\varepsilon}^{-1}} \mp 1 \right) = \frac{\tilde{\varepsilon}}{\beta} \left(\frac{\kappa}{\beta} \mp 1 \right) \quad (\text{C.8})$$

which will be used repeatedly without further mention.

Proposition C.1. *If $\beta, \lambda, \varepsilon \in (0, \infty)$ are constant and $\sigma \equiv 0$, the optimal trading speed (in open-loop form) for (C.1) is*

$$q_s^* := C_t^1 e^{\kappa(T-s)} + C_t^2 e^{-\kappa(T-s)} - (C_t^1 + C_t^2), \quad s \in [t, T], \quad (\text{C.9})$$

where C_t^1 and C_t^2 are uniquely determined by the invertible linear system

$$\mathbf{M}_t \begin{bmatrix} C_t^1 \\ C_t^2 \end{bmatrix} = \begin{bmatrix} -\bar{x} \\ -\lambda^{-1}y \end{bmatrix}, \quad \mathbf{M}_t := \begin{bmatrix} a_t^{11} & a_t^{12} \\ a_t^{21} & a_t^{22} \end{bmatrix}, \quad (\text{C.10})$$

where

$$\begin{aligned} \bar{x} &= x + (1 - e^{-\int_t^T \theta_u du}) z \\ a_t^{11} &= \kappa^{-1} e^{\kappa(T-t)} - (T-t) + [(\kappa + \beta)^{-1} - \beta^{-1} - \kappa^{-1}] \\ a_t^{12} &= -\kappa^{-1} e^{-\kappa(T-t)} - (T-t) + [-(\kappa - \beta)^{-1} - \beta^{-1} + \kappa^{-1}] \\ a_t^{21} &= e^{\kappa(T-t)} (\kappa - \beta)^{-1} + \beta^{-1} \\ a_t^{22} &= -e^{-\kappa(T-t)} (\kappa + \beta)^{-1} + \beta^{-1}. \end{aligned}$$

This result will be obtained by solving the forward-backward system consisting of the (forward) state dynamics (C.2) and the backward system (C.6), which are coupled through the optimality condition (C.4).

Proof. Step 1: We first derive a second-order, linear, non-homogeneous ODE for q^* . From the first equation in (C.6), $U^X \equiv U_T^X$ is constant, and hence the optimality condition (C.4) yields

$$U_s^Y = -\frac{\varepsilon}{\lambda} q_s^* - \frac{1}{\lambda} U_T^X.$$

The second equation in (C.6) then reads

$$\frac{\varepsilon}{\lambda} \dot{q}_s^* = -\dot{U}_s^Y = \beta \left(\frac{\varepsilon}{\lambda} q_s^* + \frac{1}{\lambda} U_T^X \right) + \frac{2\beta}{\lambda} Y_s,$$

or, equivalently,

$$Y_s = \frac{1}{2} \left(\frac{\varepsilon}{\beta} \dot{q}_s^* - \varepsilon q_s^* - U_T^X \right). \quad (\text{C.11})$$

The second equation in (C.2) now implies that

$$\frac{1}{2} \left(\frac{\varepsilon}{\beta} \ddot{q}_s^* - \varepsilon \dot{q}_s^* \right) = -\frac{\beta}{2} \left(\frac{\varepsilon}{\beta} \dot{q}_s^* - \varepsilon q_s^* - U_T^X \right) + \lambda q_s^*$$

or

$$\ddot{q}_s^* - \beta^2 \left(1 + \frac{2\lambda}{\varepsilon\beta} \right) \dot{q}_s^* - \frac{\beta^2}{\varepsilon} U_T^X = 0. \quad (\text{C.12})$$

We also have the terminal condition $q_T^* = 0$ from (C.5). With $\tilde{\varepsilon}$ and κ as defined in (C.7), the solution to (C.12) can be written as

$$\begin{aligned} q_s^* &= C_t^1 e^{\kappa(T-s)} + C_t^2 e^{-\kappa(T-s)} - \frac{\beta}{2\lambda(1+\tilde{\varepsilon})} U_T^X \\ &= C_t^1 e^{\kappa(T-s)} + C_t^2 e^{-\kappa(T-s)} - (C_t^1 + C_t^2); \end{aligned} \quad (\text{C.13})$$

here

$$U_T^X = \frac{2\lambda}{\beta} (1 + \tilde{\varepsilon}) (C_t^1 + C_t^2) \quad (\text{C.14})$$

due to the terminal condition $q_T^* = 0$ and $C_t^1, C_t^2 \in \mathbb{R}$ will be determined in Step 3 below.

Step 2: We calculate the state processes X, Y, Z in (C.2) at $s \in [t, T]$. Note that

$$Z_s = z e^{-\int_t^s \theta_u du}. \quad (\text{C.15})$$

Then

$$\begin{aligned} X_s &= x + \int_t^s q_u du - (Z_s - z) \\ &= x + C_t^1 \left\{ \kappa^{-1} [e^{\kappa(T-t)} - e^{\kappa(T-s)}] - (s-t) \right\} \\ &\quad + C_t^2 \left\{ \kappa^{-1} [e^{-\kappa(T-s)} - e^{-\kappa(T-t)}] - (s-t) \right\} \\ &\quad + (1 - e^{-\int_t^s \theta_u du}) z \end{aligned} \quad (\text{C.16})$$

and

$$\begin{aligned} Y_s &= e^{-\beta(s-t)} y + \lambda e^{-\beta s} \int_t^s e^{\beta u} q_u du \\ &= e^{-\beta(s-t)} y \\ &\quad + \lambda C_t^1 \left\{ (\kappa - \beta)^{-1} [e^{\kappa(T-t) - \beta(s-t)} - e^{\kappa(T-s)}] - \beta^{-1} [1 - e^{-\beta(s-t)}] \right\} \\ &\quad + \lambda C_t^2 \left\{ (\kappa + \beta)^{-1} [e^{-\kappa(T-s)} - e^{-\kappa(T-t) - \beta(s-t)}] - \beta^{-1} [1 - e^{-\beta(s-t)}] \right\}. \end{aligned} \quad (\text{C.17})$$

In particular, the state variables at T are

$$\begin{aligned} X_T &= x + C_t^1 \left\{ \kappa^{-1} [e^{\kappa(T-t)} - 1] - (T-t) \right\} \\ &\quad + C_t^2 \left\{ \kappa^{-1} [1 - e^{-\kappa(T-t)}] - (T-t) \right\} + (1 - e^{-\int_t^T \theta_u du}) z \end{aligned} \quad (\text{C.18})$$

and

$$\begin{aligned} Y_T &= e^{-\beta(T-t)} y + \lambda C_t^1 \left\{ (\kappa - \beta)^{-1} [e^{(\kappa-\beta)(T-t)} - 1] - \beta^{-1} [1 - e^{-\beta(T-t)}] \right\} \\ &\quad + \lambda C_t^2 \left\{ (\kappa + \beta)^{-1} [1 - e^{-(\kappa+\beta)(T-t)}] - \beta^{-1} [1 - e^{-\beta(T-t)}] \right\}. \end{aligned} \quad (\text{C.19})$$

Step 3: We use the expressions from Step 2 to determine C_t^1 and C_t^2 . The derivation of (C.11) from the maximal principle (C.4) shows that the forward-backward system is solved only if Y satisfies

$$Y_s = \frac{\tilde{\varepsilon}\lambda}{\beta^2} \dot{q}_s^* - \frac{\tilde{\varepsilon}\lambda}{\beta} q_s^* - \frac{1}{2} U_T^X, \quad s \in [t, T]. \quad (\text{C.20})$$

Using the expressions of q_s^* in (C.13), U_T^X in (C.14) and Y_s in (C.17), we claim that (C.20) holds if and only if

$$C_t^1 \left[e^{\kappa(T-t)} (\kappa - \beta)^{-1} + \beta^{-1} \right] + C_t^2 \left[-e^{-\kappa(T-t)} (\kappa + \beta)^{-1} + \beta^{-1} \right] = -\lambda^{-1} y. \quad (\text{C.21})$$

To wit, we can group the terms in (C.20) by $e^{-\beta(s-t)}$, $e^{\kappa(T-s)}$, $e^{-\kappa(T-s)}$ and the remaining constants. The coefficient of $e^{-\beta(s-t)}$ being zero gives rise to (C.21); the other coefficients can be fully canceled thanks to (C.8), so they do not yield further conditions.

On the other hand, the first terminal condition in (C.6) reads

$$-U_T^X + \lambda X_T - Y_T = 0. \quad (\text{C.22})$$

Note that (C.20) at T reduces to

$$Y_T = \frac{\lambda \tilde{\varepsilon} \kappa}{\beta^2} (-C_t^1 + C_t^2) - \frac{1}{2} U_T^X, \quad (\text{C.23})$$

since $q_T^* = 0$. Plugging (C.14), (C.18), (C.23) into (C.22), we arrive at

$$\begin{aligned} &C_t^1 \left[(\kappa + \beta)^{-1} - \beta^{-1} + \kappa^{-1} [e^{\kappa(T-t)} - 1] - (T-t) \right] \\ &+ C_t^2 \left[-(\kappa - \beta)^{-1} - \beta^{-1} + \kappa^{-1} [1 - e^{-\kappa(T-t)}] - (T-t) \right] \\ &= -x - (1 - e^{-\int_t^T \theta_u du}) z \end{aligned} \quad (\text{C.24})$$

again with the help of (C.8).

The equations (C.21) and (C.24) forms the linear system (C.10). To calculate its determinant, note that

$$\begin{aligned} \det(\mathbf{M}_t) &= a_t^{11} a_t^{22} - a_t^{12} a_t^{21} \\ &= -a_t^{11} e^{-\kappa(T-t)} (\kappa + \beta)^{-1} - a_t^{12} e^{\kappa(T-t)} (\kappa - \beta)^{-1} + \beta^{-1} (a_t^{11} - a_t^{12}), \end{aligned}$$

where thanks to (C.8), we have

$$\begin{aligned} a_t^{11} - a_t^{12} &= \kappa^{-1}(e^{\kappa(T-t)} + e^{-\kappa(T-t)}) + [(\kappa + \beta)^{-1} + (\kappa - \beta)^{-1} - 2\kappa^{-1}] \\ &= \kappa^{-1}(e^{\kappa(T-t)} + e^{-\kappa(T-t)}) + 2(\tilde{\varepsilon}\kappa\beta^{-2} - \kappa^{-1}). \end{aligned}$$

We group the terms in $\det(\mathbf{M}_t)$ by $e^{\kappa(T-t)}$, $e^{-\kappa(T-t)}$, $(T-t)e^{\kappa(T-t)}$, $(T-t)e^{-\kappa(T-t)}$ and the constant terms. All coefficients are strictly greater than zero when $\tilde{\varepsilon}, \beta, \lambda > 0$, and hence $\det(\mathbf{M}_t) > 0$:

$$\begin{aligned} \det(\mathbf{M}_t) &= e^{\kappa(T-t)} \{(\kappa - \beta)^{-1}[(\kappa - \beta)^{-1} + \beta^{-1} - \kappa^{-1}] + \beta^{-1}\kappa^{-1}\} \\ &\quad + e^{-\kappa(T-t)} \{(\kappa + \beta)^{-1}[-(\kappa + \beta)^{-1} + \beta^{-1} + \kappa^{-1}] + \beta^{-1}\kappa^{-1}\} \\ &\quad + (T-t)e^{\kappa(T-t)}(\kappa - \beta)^{-1} + (T-t)e^{-\kappa(T-t)}(\kappa + \beta)^{-1} + 4\tilde{\varepsilon}\beta^{-1}\kappa^{-1}. \end{aligned}$$

In particular, the constant term is simplified from

$$\kappa^{-1}[-(\kappa + \beta)^{-1} + (\kappa - \beta)^{-1}] + 2\beta^{-1}(\tilde{\varepsilon}\kappa\beta^{-2} - \kappa^{-1}) = 4\tilde{\varepsilon}\beta^{-1}\kappa^{-1}$$

thanks to (C.8) and $\tilde{\varepsilon}\kappa^2\beta^{-2} = 1 + \tilde{\varepsilon}$. □

We can now derive Theorem 2.14.

Proof of Theorem 2.14. As mentioned above, by Proposition 2.3, we can assume without loss of generality that $\sigma \equiv 0$. We can then apply Proposition C.1. Rewriting the expression (C.13) at the initial time $s = t$ yields

$$\begin{aligned} q_t^* &= C_t^1(e^{\kappa(T-t)} - 1) - C_t^2(1 - e^{-\kappa(T-t)}) \\ &= (e^{\kappa(T-t)} - 1)(C_t^1 - e^{-\kappa(T-t)}C_t^2). \end{aligned} \tag{C.25}$$

From (C.10), we have

$$C_t^1 = -(a_t^{22}\bar{x} - \lambda^{-1}a_t^{12}y)/d_t, \tag{C.26}$$

$$C_t^2 = (a_t^{21}\bar{x} - \lambda^{-1}a_t^{11}y)/d_t, \tag{C.27}$$

where $d_t := \det(\mathbf{M}_t)$. Note that

$$C_t^1 - e^{-\kappa(T-t)}C_t^2 = (-\tilde{f}_t\bar{x} + \tilde{g}_ty)/d_t, \tag{C.28}$$

where

$$\begin{aligned} \tilde{f}_t &:= -(a_t^{22} + e^{-\kappa(T-t)}a_t^{21}) \\ &= -[\beta^{-1} - (\kappa + \beta)^{-1}]e^{-\kappa(T-t)} - [\beta^{-1} + (\kappa - \beta)^{-1}] < 0 \end{aligned}$$

and

$$\begin{aligned} \tilde{g}_t &:= \lambda^{-1}a_t^{12} + \lambda^{-1}e^{-\kappa(T-t)}a_t^{11} \\ &= -\lambda^{-1}(1 + e^{-\kappa(T-t)})(T-t) - \lambda^{-1}[\beta^{-1} + 2\kappa^{-1} - (\kappa + \beta)^{-1}]e^{-\kappa(T-t)} \\ &\quad - \lambda^{-1}[(\kappa - \beta)^{-1} + \beta^{-1} - 2\kappa^{-1}] < 0. \end{aligned}$$

In particular, it can be verified in the last line that the constant term is negative,

$$-\frac{1}{\beta} \left[\frac{1}{\sqrt{1 + \tilde{\varepsilon}^{-1}} - 1} + 1 - \frac{2}{\sqrt{1 + \tilde{\varepsilon}^{-1}}} \right] < 0.$$

We can further rewrite \tilde{g}_t as

$$\tilde{g}_t = \lambda^{-1} \tilde{f}_t - \lambda^{-1}(1 + e^{-\kappa(T-t)})(T - t) + 2\lambda^{-1}\kappa^{-1}(1 - e^{-\kappa(T-t)}).$$

Plugging (C.28) into (C.25) yields

$$q_t^* = (e^{\kappa(T-t)} - 1)(-\tilde{f}_t \bar{x} + \tilde{g}_t y) / d_t,$$

completing the proof. □

References

- [1] E. Abi Jaber and E. Neuman. Optimal liquidation with signals: the general propagator case. *Preprint*, 2022. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4264823.
- [2] J. Ackermann, T. Kruse, and M. Urusov. Reducing Obizhaeva-Wang type trade execution problems to LQ stochastic control problems. *Preprint arXiv:2206.03772v1*, 2022.
- [3] A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157, 2010.
- [4] A. Alfonsi, A. Schied, and A. Slynko. Order book resilience, price manipulation, and the positive portfolio problem. *SIAM J. Financial Math.*, 3(1):511–533, 2012.
- [5] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *J. Risk*, 3(2):5–39, 2001.
- [6] P. Bank, I. Ekren, and J. Muhle-Karbe. Liquidity in competitive dealer markets. *Math. Finance*, 31(3):827–856, 2021.
- [7] Bank of England. How fair and effective are the fixed income, foreign exchange and commodities markets? <https://www.bankofengland.co.uk/-/media/boe/files/paper/2014/how-fair-and-effective-are-the-fixed-income-foreign-exchange-and-commodities-markets.pdf>, 2014.
- [8] A. Barzykin, P. Bergault, and O. Guéant. Market-making by a foreign exchange dealer. *Risk*, 2022. August 2022. <https://www.risk.net/cutting-edge/banking/7952481/market-making-by-a-foreign-exchange-dealer>.
- [9] A. Barzykin, P. Bergault, and O. Guéant. Algorithmic market making in dealer markets with hedging and market impact. *Math. Finance*, 33(1):41–79, 2023.
- [10] P. Bergault, F. Drissi, and O. Guéant. Multi-asset optimal execution and statistical arbitrage strategies under Ornstein-Uhlenbeck dynamics. *SIAM J. Financial Math.*, 13(1):353–390, 2022.
- [11] BIS. Global foreign exchange market turnover in 2016. <https://www.bis.org/publ/rpfx16fxt.pdf>, 2016.
- [12] BIS. Triennial central bank survey of foreign exchange and OTC derivatives markets. Reporting guidelines for turnover. *Bank for International Settlement*, 2022. https://www.bis.org/statistics/triennialrep/2022survey_guidelinesturnover.pdf.
- [13] J.-P. Bouchaud, J. Bonart, J. Donier, and M. Gould. *Trades, Quotes and Prices*. Cambridge University Press, 2018.
- [14] J.-P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and response in financial markets: The subtle nature of random price changes. *Quantitative Finance*, 4(2):176–190, 2004.

- [15] S. Butcher. Is this the best trading job in a bank? or the worst? <https://www.efinancialcareers.co.uk/news/2018/02/central-risk-desks-banking>, 2018. Accessed 10 November 2021.
- [16] M. Butz and R. Oomen. Internalisation by electronic FX spot dealers. *Quant. Finance*, 19(1):35–56, 2019.
- [17] R. Carmona and L. Leal. Optimal execution with quadratic variation inventories. *SIAM J. Financial Math.*, 14(3):751–776, 2023.
- [18] R. Carmona and K. Webster. The self-financing equation in limit order book markets. *Finance and Stochastics*, 23(3):729–759, 2019.
- [19] A. Cartea and S. Jaimungal. Optimal execution with limit and market orders. *Quant. Finance*, 15(8):1279–1291, 2015.
- [20] A. Cartea and S. Jaimungal. Incorporating order-flow into optimal execution. *Math. Financ. Econ.*, 10(3):339–364, 2016.
- [21] A. Cartea, S. Jaimungal, and J. Penalva. *Algorithmic and High Frequency Trading*. Cambridge University Press, 2015.
- [22] A. Cartea and L. Sánchez-Betancourt. Brokers and informed traders: Dealing with toxic flow and extracting trading signals. *Preprint*, 2022. <https://ssrn.com/abstract=4265814>.
- [23] Cboe Global Markets. U.S. equities market volume summary. https://www.cboe.com/us/equities/market_share/, 2023. Accessed 25 August 2023.
- [24] Y. Chen, U. Horst, and H. Hai Tran. Portfolio liquidation under transient price impact – theoretical solution and implementation with 100 NASDAQ stocks. *Preprint*, 2019. <https://arxiv.org/abs/1912.06426>.
- [25] D. Cifu. Virtu Financial @ Goldman Sachs financial services conference, December 14, 2017. https://s21.q4cdn.com/422114427/files/doc_events/transcripts/Virtu-at-Goldman-Financial-Services-Conference-2017_vFinal.pdf, 2017. Accessed 25 August 2023.
- [26] R. Cont, A. Kukanov, and S. Stoikov. The price impact of order book events. *Journal of Financial Econometrics*, 12(1):47–88, 2013.
- [27] A. Fruth, T. Schöneborn, and M. Urusov. Optimal trade execution and price manipulation in order books with time-varying liquidity. *Mathematical Finance*, 24(4):651–695, 2013.
- [28] A. Fruth, T. Schöneborn, and M. Urusov. Optimal trade execution and price manipulation in order books with time-varying liquidity. *Mathematical Finance*, 24(4):651–695, 2013.
- [29] N. Garleanu and L. H. Pedersen. Dynamic portfolio choice with frictions. *J. Econ. Theory*, 165:487–516, 2016.
- [30] J. Gatheral, A. Schied, and A. Slynko. Exponential resilience and decay of market impact. *Econophysics of order-driven markets*, 2011.
- [31] J. Gatheral, A. Schied, and A. Slynko. Transient linear price impact and Fredholm integral equations. *Math. Finance*, 22(3):445–474, 2012.
- [32] A. Goyal, N. Jegadeesh, and Y. Wu. Price impact: Continuous trading, closing auctions, and opening auctions. *Preprint*, 2022. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4300417.
- [33] P. Graewe and U. Horst. Optimal trade execution with instantaneous price impact and stochastic resilience. *SIAM J. Control Optim.*, 55(6):3707–3725, 2017.

- [34] C. R. Harvey, A. Ledford, E. Sciulli, P. Ustinov, and S. Zohren. Quantifying long-term market impact. *The Journal of Portfolio Management*, 48(3):25–46, 2022.
- [35] N. Hey, J.-P. Bouchaud, I. Mastromatteo, J. Muhle-Karbe, and K. Webster. The cost of misspecifying price impact. *Preprint*, 2023. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4465282.
- [36] G. Huberman and W. Stanzl. Price manipulation and quasi-arbitrage. *Econometrica*, 72(4):1247–1275, 2004.
- [37] P. Kolm and N. Westray. What happened to the rest? A principled approach to clean-up costs in algorithmic trading. *Risk*, 2021. <https://www.risk.net/cutting-edge/investments/7865821/a-principled-approach-to-clean-up-costs-in-algo-trading>.
- [38] C.-A. Lehalle and E. Neuman. Incorporating signals into optimal trading. *Finance Stoch.*, 23(2):275–311, 2019.
- [39] Lobster. Lobster. <https://lobsterdata.com/index.php>. Accessed 10 December 2021.
- [40] P. Mackintosh. The 2022 Intern’s Guide to Trading. <https://www.nasdaq.com/articles/the-2022-interns-guide-to-trading>, 2022. Accessed 5 July 2022.
- [41] J. Muhle-Karbe and R. Oomen. Pre-hedging. *Preprint SSRN*, 2023. <https://ssrn.com/abstract=4499729>.
- [42] J. Muhle-Karbe, Z. Wang, and K. Webster. Stochastic liquidity as a proxy for nonlinear price impact. *Preprint*, 2022. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4286108.
- [43] E. Neuman, W. Stockinger, and Y. Zhang. An offline learning approach to propagator models. *Preprint*, 2023. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4564225.
- [44] E. Neuman and Y. Zhang. Statistical learning with sublinear regret of propagator models. *Preprint*, 2023. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4323112.
- [45] A. A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. *Journal of Financial markets*, 16(1):1–32, 2013.
- [46] H. Pham. *Continuous-Time Stochastic Control and Optimization with Financial Applications*, volume 61. Springer-Verlag, Berlin, 2009.
- [47] P. Protter. *Stochastic Integration and Differential Equations*. Springer, 2005.
- [48] M. Salek, D. Chellet, and I. M. Toke. Price impact in equity auctions: zero, then linear. *Preprint*, 2023. https://hal.science/hal-03938660/file/Auctions_BEDOFIH_average_impact_pre-print.pdf.
- [49] B. Toth, I. Palit, F. Lillo, and J. D. Farmer. Why is equity order flow so persistent? *Journal of Economic Dynamics and Control*, 51:218–239, 2015.
- [50] K. Webster. *Handbook of Price Impact Modeling*. CRC Press, Boca Raton, FL, 2023.
- [51] W. M. Wonham. On a matrix Riccati equation of stochastic control. *SIAM J. Control*, 6:681–697, 1968.