



King's Research Portal

DOI:

[10.1080/1350486X.2022.2161588](https://doi.org/10.1080/1350486X.2022.2161588)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Donnelly, R. (2022). Optimal Execution: A Review. *Applied Mathematical Finance*, 29(3).
<https://doi.org/10.1080/1350486X.2022.2161588>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Optimal Execution: A Review[☆]

Ryan Donnelly^a

^a*Department of Mathematics, King's College London
Strand, London, WC2R 2LS, UK*

Abstract

This review article is intended to collect and summarize many of the results in the field of optimal execution over the last twenty years. In doing so, we describe the general workings of the limit order book so that the sources of costs and risks which need to be optimized are understood. The initial models considered propose simple dynamics for prices which allow easily computable strategies which maximize risk-adjusted profits. Subsequently, the review is divided into two major parts. The first explores several works which investigate how optimal liquidation strategies are modified to account for more complex dynamics, namely other stochastic or non-linear factors. The second presents optimal trading strategies when the agent utilizes benchmarks in addition to risk-adjusted wealth, or when she has objectives beyond optimal liquidation.

Keywords: algorithmic trading, price impact, stochastic optimization

1. Introduction

In the last few decades, the actions taken by participants in financial markets have become increasingly based on quantitative analysis and algorithms rather than any human decision making process. The motivation for this shift comes from attempting to minimize costs associated with trading, or to minimize risks involved in holding financial assets, or both. Establishing a mathematical framework that models the setting of trading in a financial market allows one to quantify the trade-off between minimizing costs and minimizing risks, thereby tailoring a trading strategy to one's preferences optimally. In this review article, we provide an overview of many developments and results related to optimal execution of trades.

[☆]The author would like to thank Fayçal Drissi, Matthew Lorig, and Álvaro Cartea for their comments and suggestions which have improved this manuscript.

Email address: ryan.f.donnelly@kcl.ac.uk (Ryan Donnelly)

We begin with a description of the limit order book (LOB), and then proceed with a summary of some preliminary models in discrete time, in particular those of Bertsimas and Lo (1998) and Almgren and Chriss (2001). These pieces will provide a groundwork for the type of problems considered in this review, but most of the focus will be on analogous models and results in continuous time. After this groundwork is established, we review developments which take more realistic or generalized underlying dynamics, different or more general performance criteria, and other aspects of financial trading and optimization.

The primary mathematical tools that are used to solve many of the problems contained herein are the Dynamic Programming Principle (DPP) and the Hamilton–Jacobi–Bellman (HJB) equation, although sometimes other approaches such as variational techniques are necessary. The full solutions to optimizations are omitted, but the original sources of the results or standard references on stochastic optimal control should provide any details which are skipped (see for example Yong and Zhou (1999), Fleming and Soner (2006)).

We focus on modeling and optimization results of a theoretical nature, but many of the primary sources that are cited contain empirical analysis to motivate modeling choices. Some of these motivational empirical results will be provided when warranted, but the discussion about these results and their interpretation will be limited.

2. The Limit Order Book

This section describes the general mechanisms of the LOB. An understanding of how transactions take place within LOBs motivates some of the modeling choices in the next section. Many of these modeling choices may still be appropriate in other market settings that do not employ a LOB, nevertheless a specific context may help with interpretation of results that we present later.

Most electronic exchanges allow participants to submit several different types of orders, but these can be categorized broadly as either passive limit orders (LOs) or aggressive market orders (MOs), so for clarity we describe a simplified form of LOB mechanisms containing only the most basic of these two types of orders.

LOs consist of a price, a volume, and an indicator to either buy or sell the asset. After an LO is submitted, it sits in the LOB until it is either canceled or filled against an MO at a later time. The cancellation or filling of an LO can be either full or partial, meaning either the entire volume of the order is removed from the LOB or some amount of the original volume remains in the LOB. If a subsequent LO is placed with the same price, then both of these




Figure 1: Blue bars represent buy LOs, red bars represent sell LOs. Left panel: the highest price of all buy LOs and the lowest price of all sell LOs are the best bid and ask, respectively. Middle panel: a new LO is placed at a price where orders already exist at the back of the queue. Right panel: the volume at this price is increased. Analogous changes would result from an order cancellation, and thus the state of the LOB may change through time, even if no trades are taking place.

orders will exist simultaneously in the LOB at this price, and the order of execution of the orders depends on the matching rules of the exchange. The matching rules we describe below correspond to a price-time priority of orders, meaning the LOs form a queue at each price and the LOs to be filled are determined according to their position in this queue, known as the First In First Out (FIFO) rule.¹ When several LOs of different types (buys and sells) are sitting in the LOB, there are two price levels which correspond to the highest price across all buy LOs and the lowest price across all sell LOs, referred to as the best bid and best ask, respectively. The average of the best bid and best ask is referred to as the midprice of the asset. Figure 1 shows a graphical representation of the state of the LOB at a single point in time, followed by the state of the LOB after a single LO is placed.

An MO consists of a volume and indicator to either buy or sell the asset. When an MO is submitted it is matched against resting LOs of the opposite buy/sell type, and the price of the transaction is determined by the price level of all filled LOs up to the volume specified by the MO. In Figure 2 we show a graphical representation of an MO filling several LOs to illustrate how the volume of the MO and the volumes and prices of the LOs determine the overall transaction price of this trade.

With this structure of the LOB in mind, it is clear that the midprice of the asset may change through time as LOs are placed and canceled near the best bid and ask. Additionally, the effect that MOs have on prices can be interpreted in multiple ways. We generally divide the impact of MOs into temporary and permanent, although they are not fully distinct and

¹Other matching rules such as Pro Rata are also used, but the intricate details of how they differ from FIFO will not be necessary for the types of models considered here.




Figure 2: Left panel: the LOB immediately before an MO is submitted. Middle panel: a sell MO is submitted with volume greater than that of the best bid, so it must ‘walk the book’ into the next best price. Right panel: the MO has filled LOs in the book, so there is now less volume available at the two relevant prices.

there may be other impact phenomena which don’t fall into either category. Temporary price impact refers to the relationship between the volume of the MO and the average execution price of the trade. For larger MOs, multiple LOs at different levels are executed, thus the average price of the trade is a function of volume. Permanent price impact refers to the relationship between the volume of the MO and the midprice process at future times after the execution. For larger MOs, more LOs are removed from the LOB resulting in a greater change of the midprice which will persist into the future, affecting the price of subsequent trades.

The LOB structure discussed above is an idealization of real markets where there are several types of orders that can be submitted. Usually, order types can be loosely grouped by whether they add or remove liquidity from the market, and we think of them as LOs or MOs, respectively, in our idealized framework. These other order types include stop-loss, stop-limit, immediate-or-cancel, fill-or-kill, and all-or-none.

For extensive analysis on LOB dynamics, including and beyond price-impact effects, see Biais et al. (1995), Rinaldo (2004), Almgren et al. (2005), Cont (2011), Cont et al. (2014), Cartea et al. (2015), Lehalle and Mounjid (2017), Bechler and Ludkovski (2017), Cartea et al. (2018a), and Sirignano (2019). For more in depth outline of the operation of LOBs in financial markets and of advanced order types, see Cartea et al. (2015) and Guéant (2016).




Figure 3: Left panel: the LOB as volume density. Middle panel: a sell MO is submitted with volume equal to the area of the yellow region. Right panel: the MO has filled LOs in the book, and both buy and sell LOs are shortly thereafter placed to establish a new impacted midprice.

3. Modeling Trades, Price Impact, and Objectives

3.1. Discrete-time Model

The discrete time model that we present here is adapted from the models first established in Bertsimas and Lo (1998) and Almgren and Chriss (2001), and which have consistently formed the basis for most models in discrete time optimal execution since then. While Bertsimas and Lo (1998) takes into account additional considerations such as informed price changes and more realistic dynamics, the initial model they propose to get intuition about optimal execution is equivalent to what we give below.

We make a conceptual change to the description of the LOB in the previous section to motivate many of the assumptions made in this model. The main change we make is that we consider the LOB to be based on a continuum of prices rather than a discrete set, and the LOs form a volume density, rather than a collection of individual queues at each price. Thus, our graphical representation including the arrival of an MO changes to that of Figure 3.

Note that in the right-hand panel of Figure 3, the new midprice is located between (not necessarily at the midpoint) the old midprice and the lowest executed price of the MO from the middle panel. We assume that immediately after the MO, there is a flurry of activity in the LOB which includes many new LO placements to establish this new midprice. This activity could potentially also include LO cancellations resulting in a new midprice beyond the lowest executed price. In addition, there is no gap left in the LOB and indeed the overall shape is merely a translation of the original shape. This will be reflected in the model below through the fact that our impact functions depend only on the size of the MO.

We work on a finite time horizon equal to N and index time by $n \in \{0, \dots, N\}$. Let $Q =$

$\{Q_n\}_{n=0}^N$ be the agent's inventory process at each point in time. Denote by V_n the size of the agent's submitted MO at time n , so we have the relation

$$Q_{n+1} = Q_n + V_n. \quad (1)$$

Let $S = \{S_n\}_{n=0}^N$ be the midprice observable in the LOB at each time. In specifying the dynamics of this process, we wish to include a component contributed by the MOs submitted by the agent, as well as a component which comprises the actions of all other agents participating in the market. To this end, let $\epsilon = \{\epsilon_n\}_{n=1}^N$ be a sequence of independent random variables with mean zero, and let f be an increasing function with $f(0) = 0$. Then we set

$$S_{n+1} = S_n + f(V_n) + \epsilon_{n+1}. \quad (2)$$

Thus, if the agent does not trade ($V_n = 0$), the midprice changes randomly by ϵ_{n+1} from the perspective of the agent, where this change is due to the actions of other agents. If the agent submits a buy order ($V_n > 0$) then the next midprice is still random, but will tend to be greater than if there were no trade, and larger buy orders will have a larger effect on increasing this future midprice. Similar effects hold as a result of the agent submitting a sell order, and so the function f captures the permanent price impact effects of the agent's trades.

Let $X = \{X_n\}_{n=0}^N$ be the amount of cash held by the agent through time. In specifying the dynamics of X we include a term which represents a temporary price impact effect as discussed in Figure 2. Let g be an increasing function with $g(0) = 0$ and set

$$X_{n+1} = X_n - \widehat{S}_n V_n, \quad (3)$$

$$\widehat{S}_n = S_n + g(V_n). \quad (4)$$

The quantity \widehat{S}_n represents the average price per share of the agent's executed MO, given that the volume of the MO is V_n . A larger MO volume will result in executions farther from the midprice with buy orders executing at higher prices and sell orders at lower prices, and so the function g captures temporary price impact as desired. With an average execution price of \widehat{S}_n and a total volume of V_n , the change in the agent's cash account is $-\widehat{S}_n V_n$, hence the relation in (3). It is straightforward to show that if the shape of the LOB is given by a curve as in Figure 3, then the function g is uniquely determined by this shape.

3.2. Continuous-time Model

Although many of the problems we summarize in this survey can be formulated and solved in discrete-time, most recent work is conducted in continuous-time. Here we replace equations (1) to (4) with analogous continuous time dynamics. Equivalently, this could be done by

considering the limiting behaviour of the discrete time model as the number of time steps $N \rightarrow \infty$ (the functions f and g , and the distributions of ϵ_n would also have to be scaled appropriately, see Almgren and Chriss (2001)).

We work on a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{0 \leq t \leq T}, \mathbb{P})$ over the finite time interval $[0, T]$. The inventory, midprice, and cash processes are now represented by their continuous-time counterparts $Q^\nu = \{Q_t^\nu\}_{0 \leq t \leq T}$, $S^\nu = \{S_t^\nu\}_{0 \leq t \leq T}$, and $X^\nu = \{X_t^\nu\}_{0 \leq t \leq T}$, and we replace their discrete time dynamics with

$$dQ_t^\nu = \nu_t dt, \quad (5)$$

$$dS_t^\nu = f(\nu_t) dt + dM_t, \quad (6)$$

$$dX_t^\nu = -\widehat{S}_t^\nu \nu_t dt, \quad (7)$$

$$\widehat{S}_t^\nu = S_t^\nu + g(\nu_t), \quad (8)$$

where $M = \{M_t\}_{0 \leq t \leq T}$ is a martingale with respect to the filtration $\{\mathcal{F}_t\}_{0 \leq t \leq T}$. With these dynamics the trading strategy $\nu = \{\nu_t\}_{0 \leq t \leq T}$ represents the speed of trading, not the volume of a sequence of trades as in Section 3.1.

The form of equations (5) to (8) does not constrain the midprice to be necessarily positive. While this is usually a desirable feature in a model, the time scales for which typical optimal executions are implemented are short enough that the probability of observing large price deviations into negative values is very small. If the time scale is long enough to warrant a positive price restriction, then geometric dynamics may be more appropriate and realistic. See Gatheral and Schied (2011), Becherer et al. (2018), and Colaneri et al. (2020) for analysis of optimal execution problems with geometric dynamics for the price process. In addition, neither the midprice nor inventory will be constrained to lie on a grid as they would be in reality. This assumption will have little consequence if the typical inventory and midprice that we consider are large. If the relevant inventory levels are typically close to zero for a particular execution, a properly discretized model would be more appropriate over the continuous rate of trading implied by (5).

3.3. Performance Criterion

Now that we have specified the dynamics of price changes and the agent's inventory and wealth, we must propose a goal for the agent so that she may seek an optimal trading strategy. One of the most common types of goals we have in mind is an agent that holds a block of shares and wishes to liquidate them before time T . In doing so, she will have to consider the effect that her trades will have on the immediate transaction price due to temporary price

impact, as well as future prices due to permanent price impact. Additionally, we may impose a penalty on the agent if she does not fulfill her goal of full liquidation by time T . Thus, we specify that the agent attempts to maximize her expected terminal wealth at time T subject to a liquidation penalty. In Bertsimas and Lo (1998), the authors frame the liquidation problem in terms of minimizing execution costs. While the objective function in (9) below is framed in terms of maximizing wealth, it is equivalent to minimizing costs.

Let the initial cash holdings, inventory, and midprice be x_0 , q_0 , and S_0 respectively. Given a trading strategy $\nu = \{\nu_t\}_{0 \leq t \leq T}$ which is sufficiently integrable and adapted to the filtration $\{\mathcal{F}_t\}_{0 \leq t \leq T}$, and an increasing function ℓ with $\ell(0) = 0$, we define the agent's performance criterion as

$$H^\nu(x_0, q_0, S_0) = \mathbb{E} \left[X_T^\nu + Q_T^\nu (S_T^\nu - \ell(Q_T^\nu)) \right]. \quad (9)$$

The terminal wealth in (9) consists of the agent's accumulated cash, X_T^ν , as well as the value of the remaining Q_T^ν shares of the asset which must be liquidated at the prevailing midprice S_T^ν with an associated penalty $\ell(Q_T^\nu)$. This liquidation penalty is akin to the temporary price impact effect which arises from walking the book as discussed previously.

Although we are taking the perspective of an agent liquidating her holdings, this framework can also be applied to an agent wishing to acquire a block of shares. This can be done by setting her initial inventory holdings equal to $Q_0 = -N$, where N is the number of shares she is trying to acquire by time T . This is based on the realization that liquidating a negative position requires the agent to purchase the shares so that her terminal position is zero.

With this performance criterion in mind, define the agent's value function

$$H(x_0, q_0, S_0) = \sup_{\nu \in \mathcal{A}} H^\nu(x_0, q_0, S_0), \quad (10)$$

where the supremum is taken over an admissible set \mathcal{A} of sufficiently integrable $\{\mathcal{F}_t\}$ -adapted processes. We would like to determine a specific control ν^* such that the agent's performance function is maximized over all possible controls. That is, we desire a trading speed ν^* such that

$$H(x_0, q_0, S_0) = H^{\nu^*}(x_0, q_0, S_0). \quad (11)$$

One common method to solve such an optimization problem is to introduce the dynamic value function (with a slight abuse of notation)

$$H(t, X_t, Q_t, S_t) = \sup_{\nu \in \mathcal{A}} \mathbb{E} \left[X_T^\nu + Q_T^\nu (S_T^\nu - \ell(Q_T^\nu)) \middle| \mathcal{F}_t \right], \quad (12)$$

and then resort to methods of dynamic programming to solve for both H and ν^* . For more details on these methods, see Pham (2009), Björk (2009), Touzi and Tourin (2013), and Cartea et al. (2015).

In order to make these methods more tractable and allow for easier analysis, we assume for now that all of the relevant price impact functions are linear in the speed of trading. That is

$$f(\nu) = b\nu, \quad g(\nu) = k\nu, \quad \ell(\nu) = \alpha\nu, \quad (13)$$

with $b > 0$, $k > 0$, and $\alpha > 0$. Regarding the visualization in Figure 3, this makes the implicit assumption that the shape of the LOB is a flat line.

Proposition 1. *The dynamic value function in (12) is given by*

$$H(t, x, q, S) = x + qS + h_2(t)q^2, \quad (14)$$

$$h_2(t) = \frac{-k(2\alpha - b)}{2k + (2\alpha - b)(T - t)} - \frac{b}{2}, \quad (15)$$

and the optimal trading strategy in feedback form is

$$\nu^*(t, q) = \frac{-(2\alpha - b)}{2k + (2\alpha - b)(T - t)}q. \quad (16)$$

We should assume $2\alpha > b$ for otherwise there will be values of T and t which make the expression in (16) undefined.² Henceforth, we always assume parameter values are chosen so that the resulting expressions are well defined.

If we substitute the liquidation strategy from Proposition 1 into equation (5), we arrive at an ordinary differential equation (ODE) governing the dynamics of the agent's inventory, which can be solved to give

$$Q_t^{\nu^*} = \frac{2k + (2\alpha - b)(T - t)}{2k + (2\alpha - b)T}q_0, \quad (17)$$

and we see that the resulting inventory process is deterministic, decreasing linearly through time towards a terminal value of

$$Q_T^{\nu^*} = \frac{2k}{2k + (2\alpha - b)T}q_0. \quad (18)$$

²the opposite inequality means the original optimization problem in (12) may lack concavity, and the agent could make infinite profits through price manipulations.

We can ensure that the agent fully liquidates the initial position of shares by placing an infinitely large penalty on remaining inventory represented by the limit $\alpha \rightarrow \infty$. In this case the optimal strategy in (16) and the resulting inventory path in (17) become

$$\lim_{\alpha \rightarrow \infty} \nu^*(t, q) = -q \quad (19)$$

$$\lim_{\alpha \rightarrow \infty} Q_t^{\nu^*} = \frac{T-t}{T} q. \quad (20)$$

The trading strategy and resulting inventory in (19) and (20) can be thought of as a benchmark called time-weighted average price (TWAP). If the trades of this strategy, which occur at a constant speed over the trading interval, are marked to the midprice, S , rather than the execution price, \widehat{S} , then the average price of all trades by the agent is the average of the midprice process given by $\frac{1}{T} \int_0^T S_t^\nu dt$. Many of the trading strategies which appear later can be thought of as incorporating a correction to the TWAP strategy.

While the strategy in (16) is the strategy which minimizes expected costs, there is no consideration of risk involved in the execution. In Almgren and Chriss (2001), the authors allow the agent to also account for the variance of her performance in addition to the expectation. This allows the agent to tailor her performance by managing the trade-off between maximizing expected wealth and minimizing variance of wealth. This mean-variance approach is incompatible with dynamic programming methods, but the authors only consider deterministic trading strategies allowing solutions to be found more easily.

We incorporate risk considerations with a different approach to the mean-variance analysis, but with the additional assumption that price innovations are Gaussian. Under this assumption the class of optimal trading strategies coincide with Almgren and Chriss (2001). To this end, we assume $M = \sigma W$ in equation (6) where W is a Brownian motion. We also suppose the agent adopts a constant absolute risk-aversion (CARA) utility function so that her dynamic value function is

$$H(t, X_t, Q_t, S_t) = \sup_{\nu \in \mathcal{A}} \mathbb{E} \left[-\exp \left\{ -\gamma \left(X_T^\nu + Q_T^\nu (S_T^\nu - \alpha Q_T^\nu) \right) \right\} \middle| \mathcal{F}_t \right], \quad (21)$$

for some value of risk-aversion parameter $\gamma > 0$.

Proposition 2. *The dynamic value function in (21) is given by*

$$H(t, x, q, S) = -\exp \left\{ -\gamma \left(x + qS + h_2(t) q^2 \right) \right\}, \quad (22)$$

$$h_2(t) = \frac{\beta}{2} \frac{\psi^- e^{-\omega(T-t)} - \psi^+ e^{\omega(T-t)}}{\psi^- e^{-\omega(T-t)} + \psi^+ e^{\omega(T-t)}} - \frac{b}{2}, \quad (23)$$

with

$$\beta = \sqrt{2k\gamma\sigma^2}, \quad \omega = \sqrt{\frac{\gamma\sigma^2}{2k}}, \quad \psi^\pm = \beta \pm 2\alpha \mp b,$$

and the optimal trading strategy in feedback form is

$$\nu^*(t, q) = \frac{(2h_2(t) + b)}{2k} q. \quad (24)$$

The optimal trading strategy in Proposition 2 is equivalent to one which maximizes an alternative performance criteria given by

$$\mathbb{E} \left[X_T^\nu + Q_T^\nu (S_T^\nu - \alpha Q_T^\nu) - \frac{\gamma}{2} \int_0^T \sigma^2 (Q_t^\nu)^2 dt \right]. \quad (25)$$

The running quadratic inventory penalty in (25) is a popular choice to incorporate in optimal execution problems rather than using a utility function as in (21) (see for example Guilbaud and Pham (2013), Barger and Lorig (2019), and Cartea et al. (2019)). However, these two approaches would not be equivalent if the random innovations in price were not given by a Brownian motion or if there are other stochastic factors in the model. When there are more factors, this running penalty term can be considered to contribute inventory risk, while the agent remains risk-neutral towards other dynamics (see Cartea et al. (2017)).

Substituting the optimal strategy from Proposition 2 into equation (5) yields again an ODE which can be solved, resulting in the agent's inventory process through time

$$Q_t^{\nu^*} = \frac{\psi^- e^{-\omega(T-t)} + \psi^+ e^{\omega(T-t)}}{\psi^- e^{-\omega T} + \psi^+ e^{\omega T}} q_0. \quad (26)$$

If we wish to consider an agent who must finish the full liquidation by time T , we can enforce this by taking the limit $\alpha \rightarrow \infty$ which results in

$$\lim_{\alpha \rightarrow \infty} \nu^*(t, q) = -\omega \coth(\omega(T-t)) q, \quad (27)$$

$$\lim_{\alpha \rightarrow \infty} Q_t^{\nu^*} = \frac{\sinh(\omega(T-t))}{\sinh(\omega T)} q_0. \quad (28)$$

In Figure 4 we show the optimal inventory process for several values of the risk-aversion parameter γ . We note that for larger values of γ , the agent uses a trading strategy which liquidates faster towards the beginning of the time interval. This is because the agent's risk is largest when her inventory is largest, and to lower that risk she desires to lower her inventory quickly. This comes at an increased cost due to the temporary price impact. When the agent is risk-neutral (and has a resulting inventory path from (17)), she trades at a constant rate




Figure 4: Optimal inventory process for various values of risk-aversion parameter γ . The curve corresponding to $\gamma = 0$ is from equation (17), the remaining curves are from equation (26). Left panel: $\alpha = 0.05$. Right panel: $\alpha = 1$. Other parameter values are $T = 1$, $\sigma = 1$, $k = 0.01$, and $b = 0.01$.

through time which is determined only by the trade-offs between temporary price impact, permanent price impact, and the liquidation penalty. The optimal strategies of Propositions 1 and 2 are the result of many simplifying assumptions regarding the dynamics of trading in a market as well as the agent’s goals, but they still offer a benchmark to which many other strategies may be compared. In the following sections, we explore work that expands upon the previous results by considering more realistic market dynamics, alternate goals of market participants, or the inclusion of other phenomena.

4. Alternative Dynamics

We begin our further exploration by considering works which make alternative assumptions about the dynamics proposed in equations (5) to (8) in order to model more general or more realistic market settings.

4.1. Stochastic Temporary and Permanent Price Impact

First we revisit one of our assumptions on the nature of price impact, both permanent and temporary. In equations (6) and (8), the relation between the agent’s trading and price dynamics are homogeneous in time – identical actions and different points in time will have identical effects on the asset price. In terms of the graphical representation of the (idealized continuous) LOB, this means that its general shape is always the same, albeit translated horizontally as the midprice changes. In reality, market liquidity varies through time meaning the shape of our graphical representation should also fluctuate. We consider in some detail

the work Barger and Lorig (2019) which allows price impact effects to be stochastic, replacing price and wealth dynamics with

$$dS_t^\nu = \hat{f}(b_t) \nu_t dt + \sigma dW_t, \quad db_t = \mu_b(b_t) dt + \eta_b(b_t) dB_t^{(1)}, \quad (29)$$

$$dX_t^\nu = -\widehat{S}_t^\nu \nu_t dt, \quad (30)$$

$$\widehat{S}_t^\nu = S_t^\nu + \hat{g}(k_t) \nu_t, \quad dk_t = \mu_k(k_t) dt + \eta_k(k_t) dB_t^{(2)}, \quad (31)$$

where $B^{(1)}$ and $B^{(2)}$ are Brownian motions independent of W , but have instantaneous correlation with each other equal to $\rho \in [-1, 1]$. Furthermore, the functions \hat{g} , μ_k , and η_k should be chosen so that $\hat{g}(k_t) > 0$ for all t (otherwise temporary price impact would work in the agent's favour and she could make unbounded profits). In this model, price impact depends linearly on trading speed (as was assumed in Propositions 1 and 2), but the magnitude of the effect is driven by the stochastic processes b and k . The graphical LOB representation is therefore given by a flat line which fluctuates through time in the vertical direction. When $\hat{f}(b_t)$ and $\hat{g}(k_t)$ decrease, indicating a low impact state, the flat line level of the LOB moves up and there is high liquidity in the market.

The agent's performance criterion is taken as stated in (25) with a running penalty on inventory, and so as discussed earlier this can be interpreted as the agent being sensitive to the risk associated with midprice changes, but when considering the stochastic nature of price impact she only accounts for the potential benefit and detriment in terms of expected wealth. These changes to the assumed dynamics complicate the HJB analysis which was used to arrive at the conclusions of our previous propositions, so the authors resort to asymptotic approximations of the dynamic value function and optimal trading speed. The authors show that a zeroth order approximation in this setting can be obtained by the same expression as in (24) by making the substitutions $b \mapsto \hat{f}(b)$ and $k \mapsto \hat{g}(k)$. This means the agent acts according to the previously examined optimal strategy, but replaces the price impact parameters with the corresponding value of what is observed at any particular time.

The first order correction to this strategy does not have a simple expression in general. But if the agent is forced to liquidate her position by time T ($\alpha \rightarrow \infty$) and if she does not impose the running inventory penalty ($\gamma = 0$), then a first order approximation is given by

$$\hat{\nu}_1^*(t, q, b, k) = \hat{\nu}_0^*(t, q, b, k) - \frac{1}{2} \frac{\mu_k(k) \hat{g}'(k)}{\hat{g}(k)} q - \frac{1}{6} \frac{\mu_b(b) \hat{f}'(b)}{\hat{g}(k)} (T - t) q, \quad (32)$$

where the term $\hat{\nu}_0^*$ is understood to be the strategy in (24) along with the substitutions $b \mapsto \hat{f}(b)$ and $k \mapsto \hat{g}(k)$. Note that at any fixed time t with given impact parameters $\hat{f}(b_t)$ and $\hat{g}(k_t)$, this first order strategy still liquidates the asset at a speed proportional to the remaining inventory. The two correction terms stem from the two different types of impact

and can be interpreted in terms of what the agent expects the impact to be in the future, thereby acting accordingly. If we suppose $\mu_k(k)$ and $\hat{g}'(k)$ are both positive, then in the short term, the temporary price impact is expected to increase, so the agent is incentivized to trade faster now to avoid higher transaction costs. Similarly, if we suppose $\mu_b(b)$ and $\hat{f}'(b)$ are both positive, then permanent price impact is expected to increase, and the agent should trade faster now so that her remaining pool of shares does not have their value negatively affected. This term depends on the remaining trading time $T - t$ because the effects of permanent price impact as modeled here accumulate over time, and so become insignificant at the end of the trading horizon. Similar arguments hold accordingly if there are changes in sign of $\mu_k(k)$ and $\hat{g}'(k)$. Both of these terms are scaled by the instantaneous temporary impact $1/\hat{g}(k)$ because any expected benefits from trading faster must be weighed against the instantaneous costs of acquiring those benefits.

Numerous other works have also investigated the effects of stochastic liquidity on optimal execution. A similar model is proposed in Almgren (2012) where there is no permanent price impact, but the midprice volatility is allowed to be stochastic rather than constant. This work also recognizes that substituting appropriate parameters into a strategy such as (24) with the observed stochastic counterparts at time t ($k \mapsto k_t$ and $\sigma \mapsto \sigma_t$) is a plausible strategy. In general, it is found that the agent trades quickly when price impact is low, volatility is high, or the end of the trading period approaches. Both Cheridito and Sepin (2014) and Ackermann et al. (2021) take a fully discrete-time approach to this problem, each considering different aspects of the dynamics to be stochastic and classifying the optimal trading strategies in terms of recursion equations. Other works with elements of stochastic price impact are Ishitani and Kato (2014), Becherer et al. (2018), Fruth et al. (2019), Fouque et al. (2022), and Cartea et al. (2022), some of which will be discussed in a subsequent section on resilience and transient price impact.

4.2. Stochastic Volatility of Midprice

In the models described in Section 3, the only exogenous source of randomness comes from the midprice innovations which were assumed to have constant volatility. In reality, the volatility will change through time due to a variety of factors in the LOB. There is generally increased activity at the beginning and end of the trading day, both in terms of the number of LOs being placed and canceled, and in the number of MOs submitted. Each of these events can change the midprice, so we end up with seasonality of volatility in the LOB over the course of the day. Additionally, many events in the LOB are self-exciting, in particular there is a large spike in the intensity of events for a short time after MOs, and so spikes in volatility may also

be observed around these times.³ The changing volatility results in different levels of risk for any given inventory level, and accounting for this can help a risk-averse agent improve her performance.

We have already made brief mention of Almgren (2012) and Cheridito and Sepin (2014) which consider models with stochastic volatility of the midprice. Here we provide some discussion of Chan and Sircar (2015) which gives approximations to the optimal trading strategy in fast or slow variation volatility regimes, which correspond to the seasonality and spike effects described previously. To this end, we replace the midprice dynamics in equation (6) with

$$dS_t = \sigma(Z_t^\epsilon) dW_t, \quad (33)$$

where σ is a function of the stochastic process Z^ϵ , which determines the instantaneous volatility of the midprice. Permanent price impact is not considered in Chan and Sircar (2015), but an additional drift term representing an observable signal is included which we have suppressed in our present discussion. We return to the modeling of observable signals in Section 5.1.

When analyzing a slow volatility factor, the dynamics of Z are taken as

$$dZ_t^\epsilon = \epsilon \mu_Z(Z_t^\epsilon) dt + \sqrt{\epsilon} \eta_Z(Z_t^\epsilon) dB_t, \quad (34)$$

where B is a standard Brownian motion independent of W , whereas the fast volatility factor uses dynamics

$$dZ_t^\epsilon = \frac{1}{\epsilon} \mu_Z(Z_t^\epsilon) dt + \frac{1}{\sqrt{\epsilon}} \eta_Z(Z_t^\epsilon) dB_t. \quad (35)$$

In each case ϵ will act as a parameter of asymptotic expansion which allows for approximations of the optimal strategy when the value of the parameter is small. Chan and Sircar (2015) consider multiscale stochastic volatility with slow and fast factors simultaneously, but we only discuss the resulting effects individually. Additionally, they have assumed an infinitely long trading horizon with a performance criterion of the form

$$\mathbb{E} \left[\int_0^\infty -\frac{\gamma}{2} \sigma^2(Z_t^\epsilon) (Q_t^\nu)^2 - k \nu_t^2 dt \right], \quad (36)$$

which results in a trading strategy that does not explicitly depend on time and brings more focus to the effects of stochastic volatility. Otherwise, the general behaviour of the optimal strategy is largely unchanged from what would arise in a finite horizon. As may be expected

³See Cartea et al. (2018a) for some empirical investigation of LOB event frequency seasonality and self-excitation.

after the discussion in Section 4.1, the approximations are largely based on the strategy which maximizes (36) when volatility is assumed constant. This can be shown to be

$$\nu_0^*(q; \sigma^2) = -\sqrt{\frac{\gamma \sigma^2}{2k}} q,$$

which is equivalent to (24) in the limit $T \rightarrow \infty$.

Under slow scale volatility, the approximation of the optimal strategy is found through perturbation analysis of the HJB equation. The resulting strategy in feedback form is

$$\hat{\nu}^*(q, z) = \nu_0^*(q; \sigma^2(z)) = -\sqrt{\frac{\gamma \sigma^2(z)}{2k}} q,$$

which we note involves a similar substitution to previously discussed approximations, where parameters are replaced with their currently observed values through time. We also see that this approximation does not include an explicit dependence on the parameter ϵ . The smallest order corrections considered in Chan and Sircar (2015) involves a stochastic drift correlated to the slow scale volatility. In this case the appropriate correction is proportional to $\sqrt{\epsilon}$ as well as the correlation between the drift and volatility.

The fast scale volatility is treated with singular perturbation analysis to arrive at an approximate optimal control, with resulting approximation given by

$$\hat{\nu}^*(q, z) = \nu_0^*(q; \langle \sigma^2 \rangle) - \frac{\gamma q}{2k} \int_t^\infty \mathbb{E} \left[\sigma^2(Z_u^\epsilon) - \langle \sigma^2 \rangle \middle| Z_t^\epsilon = z \right] du, \quad (37)$$

where $\langle \sigma^2 \rangle$ is the long-term expected variance. This expression has contributions from two terms which represent a zeroth order approximation and first order correction. The zeroth order approximation makes a familiar parameter substitution $\sigma^2 \mapsto \langle \sigma^2 \rangle$, but the first order correction involves future expected deviations of the variance from its long-term expected value. If future variance is expected to be larger than the long-term expected variance, then the agent desires faster liquidation of her shares to avoid the associated inventory risk.

Other work of interest is Colaneri et al. (2020) where the authors assume that innovations in asset price are driven by a regime switching pure jump process of the form

$$dS_t = S_{t-} dR_t, \quad (38)$$

where R has compensating random measure $\eta(t, Y_{t-}, \nu_{t-}; dz) dt$ with Y a finite state Markov chain. While this is not strictly a stochastic volatility model because it lacks a diffusive term, the price innovations dR_t being driven by a Markov chain means that the price process will share some qualitative features of a stochastic volatility model. The dependence of the compensating measure on the trading strategy ν also captures permanent price impact effects. In this case, the more complicated nature of the model leads to numerical methods for computing the optimal strategy.




Figure 5: Ask side of the LOB showing resilience and transient price impact after a trade. The left panel shows minimal spread (best ask price coincides with reference price). In the middle panel, an MO walks the book and exhibits temporary price impact. Immediately after the MO in the right panel, the orders are filled, a new reference price is established due to permanent impact, and the new best ask price will gradually decay back to the new reference price. A subsequent order submitted immediately will be subjected to higher prices than an order submitted after waiting some time.

4.3. Resilience and Transient Impact

Thus far we have only described two types of price impact effects (permanent and temporary) which are included in the models of many algorithmic trading works. But when discussing the nature of price changes after the occurrence of a transaction, there are other dynamical effects at play. In particular, when multiple LOs are filled in the LOB due to a transaction, the resulting gap in the LOB is not replenished instantaneously as Figure 3 suggests. Instead it takes time for LOs to be submitted in place of those that were filled. The rate at which these orders are replenished is called resilience, and thus Figure 3 suggests a market with infinite resilience.

In Figure 5 we show (from the ask side of the book) how resilience might be interpreted in an idealized continuous LOB. After a transaction takes place (filling sell orders in the book), if no further transactions were to occur then these sell orders may gradually be replenished until a new stable best ask price is established. However, if a subsequent transaction were performed after only a short time, then it would be subjected to higher prices than if it had occurred after a long time. This difference in transaction price which depends on the amount of time between trades is called transient price impact.

We focus on describing Chen et al. (2019) in more detail, as the framework is most similar to what has been presented here so far. A general formulation of transient price impact can be incorporated into the dynamics by introducing an impact function and decay kernel, denoted by a strictly increasing function h and a non-increasing function G , respectively, and defining

the inventory, midprice, and cash processes by

$$Q_t^\nu = Q_0 + \int_0^t \nu_u du, \quad (39)$$

$$S_t^\nu = S_0 + \int_0^t b \nu_u du + \int_0^t h(\nu_u) G(|t - u|) du + \sigma W_t, \quad (40)$$

$$X_t^\nu = X_0 - \int_0^t \widehat{S}_u^\nu \nu_u du, \quad (41)$$

$$\widehat{S}_t^\nu = S_t^\nu + k \nu_t. \quad (42)$$

Here we have taken the permanent and temporary price impacts to be linear with proportionality constants b and k . The function h determines the magnitude of transient price impact and the decay kernel G controls how quickly the impact of trading decays through time. In Chen et al. (2019) the transient impact is linear with $h(\nu) = \lambda \nu$ and the decay is exponential with $G(|t - u|) = \exp(-\rho |t - u|)$. Thus, the parameter λ determines the distance from the new reference midprice to the best ask in the right panel of Figure 5, and the distance between the best ask and the new reference midprice decays exponentially. See Gatheral (2010) for some detailed discussion of other impact functions and decay kernels, and consequences of different models with respect to statistical arbitrage opportunities.

The agent is risk-neutral, but with the additional constraint that the inventory process must satisfy $Q_T^\nu = 0$. Thus, the performance criterion reduces to

$$H^\nu(x_0, q_0, S_0) = \mathbb{E}[X_T^\nu], \quad (43)$$

and the agent attempts to maximize this over trading speeds $\{\nu_t\}_{t \in [0, T]}$ which fully liquidate the inventory by time T . Recall that this type of constraint can be applied formally to the strategy given in Proposition 1 by taking $\alpha \rightarrow \infty$, and it is to the trading strategies in this limit that we shall compare the optimal strategy with transient impact. Other works mentioned at the end of this section allow for strategies that do not fully liquidate the position by incorporating a liquidation penalty.

Proposition 3. *The trading strategy which maximizes (43) is given by*

$$\nu_t^* = - \frac{\phi \left((\phi^2 k - \lambda \rho) (\rho \cosh(\frac{\phi T}{2}) + \phi \sinh(\frac{\phi T}{2})) + \lambda \rho^2 \cosh(\phi(t - \frac{T}{2})) \right)}{\phi \rho T \cosh(\frac{\phi T}{2}) (\phi^2 k - \lambda \rho) + \sinh(\frac{\phi T}{2}) (\lambda \rho (2\rho - \phi^2 T) + \phi^4 k T)} q_0, \quad (44)$$

which results in the inventory process

$$Q_t^{\nu^*} = q_0 - q_0 \left(\frac{A + Bt + C \sinh(\phi(t - \frac{T}{2}))}{2A + BT} \right), \quad (45)$$




Figure 6: Optimal trading speed and inventory process with transient price impact for various values of transient impact parameter λ and decay coefficient ρ . Top row panels: other parameter values are $k = 0.01$, $b = 0.01$, $\rho = 1$. Bottom row panels: other parameter values are $k = 0.01$, $b = 0.01$, $\lambda = 0.3$.

where

$$\phi = \sqrt{\rho \left(\rho + \frac{\lambda}{k} \right)}, \quad (46)$$

$$A = \frac{\lambda}{k} \sinh\left(\frac{\phi T}{2}\right), \quad (47)$$

$$B = \phi \rho \cosh\left(\frac{\phi T}{2}\right) + \phi^2 \sinh\left(\frac{\phi T}{2}\right), \quad (48)$$

$$C = \frac{\lambda}{k}. \quad (49)$$

In Figure 6 we show the optimal trading speed and resulting inventory for various values of the transient impact parameter λ (top row) and decay rate ρ (bottom row). When $\lambda = 0$ the optimal rate of trading is constant at q_0/T which is the TWAP strategy in (19). This is expected because $\lambda = 0$ corresponds to the same dynamics given in equations (5) to (8). When $\rho = 0$ we also obtain the TWAP strategy, which is consistent with intuition because if there is no resilience, the transient price impact becomes a permanent price impact, so the dynamics are equivalent to equations (5) to (8) with the substitution $b \mapsto b + \lambda$. The last case in which the optimal trading speed is constant is in the limit $\rho \rightarrow \infty$ which again corresponds to the dynamics of (5) to (8).

For parameter values which do not give rise to the TWAP strategy, the typical optimal behaviour is to trade at faster speeds closer to the beginning and end of the trading period, and at slower speeds closer to the middle of the trading period (compared to the constant rate of q_0/T). The reason for this is because the agent wants to minimize the time when they are trading under significant effect of transient price impact. A single large trade will suffer from temporary price impact, but it is not until subsequent trades are submitted that the effects of transient impact are encountered. By trading quickly early on, there will be more time spent in a state where the exponential resilience has brought the price back towards its unaffected level. Then, when the transient price impact has decayed significantly, the agent can again place large trades to complete the liquidation.

For other descriptions and works which model resilience and transient price impact, see Gatheral (2010), Alfonsi et al. (2010), Predoiu et al. (2011), Obizhaeva and Wang (2013), Guo and Zervos (2015), Graewe and Horst (2017), Becherer et al. (2018), Bank and Voß (2019), Fruth et al. (2019), Ekren and Muhle-Karbe (2019), Horst and Xia (2019), Forde et al. (2021), and Neuman and Voß (2022).

4.4. *Non-linear Price Impact*

Most of the results in previous sections rely on assuming that price impact functions are linear with respect to the rate of trading, even though the original models proposed in Section 3.1 and 3.2 have general impact functions. In this section we explore an optimization in which impact functions may be non-linear with respect to the rate of trading, with a focus on Almgren (2003).

This work begins with a model in discrete time which is based on that of Section 3.1, where the nature of temporary price impact is modified so that equations (1) to (4) become

$$Q_{n+1} = Q_n + V_n, \tag{50}$$

$$S_{n+1} = S_n + f(V_n) + \epsilon_{n+1}, \tag{51}$$

$$X_{n+1} = X_n - \widehat{S}_n V_n, \tag{52}$$

$$\widehat{S}_n = S_n + g(V_n) + h(V_n) \xi_{n+1}, \tag{53}$$

where $\{\xi\}_{n=1}^N$ is a sequence of independent random variables with mean zero, unit variance, and which are also independent from $\{\epsilon_n\}_{n=1}^N$. Thus, the temporary price impact is composed of the usual term, $g(V_n)$, which depends only on the trade submitted at time n , but also has a stochastic component given by $h(V_n) \xi_{n+1}$. This type of stochastic dependence of temporary impact could be the result of several factors, for example the latency between the time that

a trade decision is made and the time at which it is matched against passive orders during which the depth of available orders may have changed. See Cartea and Sánchez-Betancourt (2021a) which considers this stochastic delay directly.

The agent attempts to maximize a mean-variance style criterion with respect to her terminal wealth subject to the constraints $Q_N = 0$ and that the sequence $\{V_n\}_{n=0}^{N-1}$ is deterministic. The agent's terminal wealth is given by

$$X_N = X_0 + S_0 Q_0 + \sum_{n=0}^{N-1} Q_n f(V_n) + \sum_{n=0}^{N-1} Q_n \epsilon_{n+1} - \sum_{n=0}^{N-1} V_n g(V_n) - \sum_{n=0}^{N-1} V_n h(V_n) \xi_{n+1}, \quad (54)$$

which, under deterministic trading strategies, has expectation and variance

$$\mathbb{E}[X_N] = X_0 + S_0 Q_0 + \sum_{n=0}^{N-1} Q_n f(V_n) - \sum_{n=0}^{N-1} V_n g(V_n), \quad (55)$$

$$\mathbb{V}[X_N] = \sum_{n=0}^{N-1} \sigma^2 Q_n^2 + \sum_{n=0}^{N-1} V_n^2 h^2(V_n), \quad (56)$$

where $\sigma^2 = \mathbb{E}[\epsilon_n^2]$. Before optimizing the trade-off between these two terms, they are both replaced by their continuous-time analogues

$$\mathbb{E}[X_T] = X_0 + S_0 Q_0 + \int_0^T \left(Q_t f(\nu_t) - \nu_t g(\nu_t) \right) dt, \quad (57)$$

$$\mathbb{V}[X_T] = \int_0^T \left(\sigma^2 Q_t^2 + \nu_t^2 h^2(\nu_t) \right) dt. \quad (58)$$

The objective of the agent is to maximize the mean-variance criterion

$$H_0^\nu = \mathbb{E}[X_T] - \frac{\gamma}{2} \mathbb{V}[X_T]. \quad (59)$$

Recognizing that $\nu_t = Q_t'$ frames this computation into a standard problem of the calculus of variations. After classifying optimality in Almgren (2003), closed-form expressions for the strategy are found by allowing the horizon T to approach infinity. First, we investigate the effect of non-linear temporary impact with zero stochastic impact such that

$$f(\nu) = b\nu, \quad g(\nu) = k\nu^m, \quad h(\nu) = 0, \quad (60)$$

with $m > 0$. In this case the optimal inventory path is given by

$$Q_t^{\nu^*} = \begin{cases} q_0 \left(1 + \frac{1-m}{1+m} \frac{t}{T^*} \right)^{-\frac{1+m}{1-m}}, & \text{if } 0 < m < 1, \\ q_0 \exp\left\{ -\frac{t}{T^*} \right\}, & \text{if } m = 1, \\ q_0 \left(1 - \frac{m-1}{m+1} \frac{t}{T^*} \right)^{\frac{m+1}{m-1}}, & \text{if } m > 1 \text{ and } t \leq \frac{m+1}{m-1} T^*, \end{cases} \quad (61)$$




Figure 7: Optimal inventory paths from (61) for various values of parameter m in temporary impact function $g(\nu) = k \nu^m$. Other parameter values are $\sigma = 1$, $k = 0.01$, $b = 0.01$, $\gamma = 1$.

where

$$T^* = \left(\frac{2 m k q_0^{m-1}}{\gamma \sigma^2} \right)^{\frac{1}{m+1}}. \quad (62)$$

A comparison with (26) in the limit $T \rightarrow \infty$ shows that it coincides with the optimal strategy above when $m = 1$. This is expected because the price impact functions are the same and the optimal strategy which results from (21) is deterministic. The effects of the non-linear impact function g on the optimal strategy are shown in Figure 7. The other case which yields a closed-form expression for the inventory path is

$$f(\nu) = b \nu, \quad g(\nu) = k \nu, \quad h(\nu) = a, \quad (63)$$

with $a > 0$. The corresponding optimal inventory path is given by

$$q_0 \exp \left\{ -\frac{t}{T^*} \right\}, \quad (64)$$

where

$$T^* = \sqrt{\frac{2 k + \gamma a^2}{\gamma \sigma^2}}. \quad (65)$$

The effect of the stochastic component of price impact can be seen in Figure 8, which is to slow the rate of trading compared to the benchmark of (26). This is because trading quickly magnifies the risks associated with uncertainties in temporary price impact caused by the term $a V_n \xi_{n+1}$ in (53). The agent is willing to accept more risk associated with the uncertainty in future price in order to lower the certain costs and riskiness associated with temporary price impact by trading slower.




Figure 8: Optimal inventory paths from (64) for various values of parameter a in stochastic temporary impact function $h(\nu) = a$. Other parameter values are $\sigma = 1$, $k = 0.01$, $b = 0.01$, $\gamma = 1$.

Additional work related to optimal trading with non-linear impact functions can be found in Bouchard et al. (2011) and Colaneri et al. (2020). Various discussions are given in Gatheral and Schied (2013) related to the form of impact function and midprice dynamics.

5. Alternative Objectives

In each of the previous sections, the perspective of the optimization problem was that of an agent wishing to liquidate a position subject to various effects related to market dynamics. Within this section, we consider agents who are not necessarily attempting to liquidate some initial holdings, but instead behave according to an alternative objective. Inspection of the optimal trading strategies in previous sections shows that if the agent's inventory begins at zero, then she will do nothing over the course of the trading interval (as her position is already liquidated, any further trading will only be subjected to costs without benefit). Due to the alternative objectives we now consider, this result no longer holds.

5.1. Speculation, Trade Signals, and Statistical Arbitrage

From a modeling perspective, a straightforward method of introducing statistical arbitrage and the ability of an agent to speculate is to assume that the unaffected midprice of the asset is no longer a martingale. In each of the models considered in previous sections, if the agent doesn't trade then the midprice process is a martingale. If instead the unaffected price has non-zero expected change in value over some time interval, then any non-zero inventory

position will also be expected to change in value. The essence of a profitable strategy in this situation is to hold inventory positions which take advantage of these non-zero expected price changes, but to do so in an efficient manner so as to not suffer from price impact effects or a future change in the asset's expected drift.

We focus on Cartea and Jaimungal (2016b) which introduces a stochastic drift into the mid-price process. The interpretation of the drift in this work is that it arises from the permanent price impact of total order flow of all market participants, but it is statistically equivalent to having a short term price predictor. Let $\mu = \{\mu_t\}_{0 \leq t \leq T}$ be a stochastic process which satisfies the Markov property, and assume it is adapted to the filtration $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ we are working with. We specify the dynamics of the relevant processes to be

$$dQ_t^\nu = \nu_t dt, \quad (66)$$

$$dS_t^\nu = (\mu_t + b \nu_t) dt + \sigma dW_t, \quad (67)$$

$$dX_t^\nu = -\widehat{S}_t^\nu \nu_t dt, \quad (68)$$

$$\widehat{S}_t^\nu = S_t^\nu + k \nu_t, \quad (69)$$

where $b > 0$ and $k > 0$ are constants. Other than the inclusion of the drift μ_t in (67), these equations are equivalent to (5) to (8) (with f and g replaced by linear functions). The value function of our agent will correspond to the performance criterion given in (25), that is

$$H(t, X_t, Q_t, S_t, \mu_t) = \sup_{\nu \in \mathcal{A}} \mathbb{E} \left[X_T^\nu + Q_T^\nu (S_T^\nu - \alpha Q_T^\nu) - \frac{\gamma}{2} \int_t^T \sigma^2 (Q_u^\nu)^2 du \middle| \mathcal{F}_t \right], \quad (70)$$

with $\alpha > 0$ and $\gamma > 0$ constant. As discussed shortly after (25), this performance criterion can be interpreted as the agent expressing risk-aversion with respect to the random midprice changes, but risk-neutrality with respect to randomness associated with the stochastic drift process μ .

At this point, no specific dynamics for the process μ have been specified. As long as this stochastic drift term satisfies some technical integrability conditions, the value function and optimal trading strategy can be expressed in terms of expected future drift, as shown in the following Proposition.

Proposition 4. *The value function in (70) is given by*

$$H(t, x, q, S, \mu) = x + q S + h_0(t, \mu) + h_1(t, \mu) q + h_2(t) q^2, \quad (71)$$

$$h_2(t) = \frac{\beta}{2} \frac{\psi^- e^{-\omega(T-t)} - \psi^+ e^{\omega(T-t)}}{\psi^- e^{-\omega(T-t)} + \psi^+ e^{\omega(T-t)}} - \frac{b}{2}, \quad (72)$$

$$h_1(t, \mu) = \int_t^T \frac{\psi^- e^{-\omega(T-u)} + \psi^+ e^{\omega(T-u)}}{\psi^- e^{-\omega(T-t)} + \psi^+ e^{\omega(T-t)}} \mathbb{E}[\mu_u | \mathcal{F}_t] du, \quad (73)$$

$$h_0(t, \mu) = \frac{1}{4k} \int_t^T \mathbb{E}[h_1^2(t, \mu_u) | \mathcal{F}_t] du \quad (74)$$

with

$$\beta = \sqrt{2k\gamma\sigma^2}, \quad \omega = \sqrt{\frac{\gamma\sigma^2}{2k}}, \quad \psi^\pm = \beta \pm 2\alpha \mp b,$$

and the optimal strategy in feedback form is

$$\nu^*(t, q, \mu) = \frac{h_1(t, \mu)}{2k} + \frac{(2h_2(t) + b)}{2k} q. \quad (75)$$

A comparison of the optimal strategy (75) to the liquidation strategy without signals given in (24) shows that the inventory component of the feedback form is identical, but the strategy with signals has an additional component which comes from the agent's forecast of the expected future asset price built into the term $h_1(t, \mu)$. This term could naturally be compared with the more straightforward quantity $\int_t^T \mathbb{E}[\mu_u | \mathcal{F}_t] du$, which represents the expected change in the unaffected midprice from time t to T . While this itself yields a forecast of how much each unit of the agent's inventory is expected to change in value, there are other considerations that the agent must take into account (namely price impact effects, liquidation penalties, and risk) when taking advantage of this signal. This gives rise to the complicated discount factor within the integral which balances the trade-off between taking advantage of the signal while acknowledging that this position will be partially liquidated close to time T and that any inventory holdings have inherent risk of price changes. If we assume $2\alpha > b$,⁴ then inspection of this discount factor shows that it is decreasing from 1 as u ranges from t to T and remains positive. This indicates that the expected price changes in the immediate future as assigned more weight towards the trading strategy compared to expected price changes in the distant future. This arises because the agent would like to enter a position as early as possible to take advantage price changes over as long a time horizon as possible, but also because there is less risk associated with acting on immediate forecasts compared to distant ones.

⁴Refer to the discussion after Proposition 1.

There are numerous other works in which some form of trade signal or price predictor are incorporated. See for example Chan and Sircar (2015), Cartea et al. (2018a), Casgrain and Jaimungal (2019), Lehalle and Neuman (2019), Donnelly and Lorig (2020), Forde et al. (2021), Bergault et al. (2022), and Neuman and Voß (2022).

5.2. Hedging

The main focus of the previous section was not about optimally liquidating an asset position, but rather attempting to maximize profits that could be extracted from knowledge about expected price changes. Nevertheless, both optimal liquidation and speculative trading have the common theme that they try to minimize costs, or equivalently maximize profits. Here we consider an agent who has an additional objective which is not directly related to cost minimization, although that will still be a consideration which will influence her trading strategy.

In Bank et al. (2017), the authors consider an agent who is tasked with optimally tracking a hedge ratio which arises from attempting to minimize the risk associated with holding another security. Mathematically, they assume that there is a stochastic process $\xi = \{\xi_t\}_{0 \leq t \leq T}$ which indicates an inventory level which the agent intends to track efficiently. The dynamics of specified processes and the performance criterion used by the agent are given by

$$dQ_t^\nu = \nu_t dt, \quad (76)$$

$$dX_t^\nu = -(S_t + k \nu_t) \nu_t dt, \quad (77)$$

$$H^\nu(x_0, q_0, S_0) = \mathbb{E} \left[X_T^\nu + Q_T^\nu S_T - \frac{1}{2} \int_0^T (Q_t^\nu - \xi_t)^2 dt \right], \quad (78)$$

where the midprice process S is a martingale which is not subject to the effects of permanent price impact. Within the performance criterion, the first two terms reward the agent for trading optimally subject to temporary price impact. The second term gives incentive for the agent to closely track the indicated inventory level ξ .

With an optimization criterion in hand, the authors characterize the optimal trading strategy and value function. For this it is convenient to introduce the notation

$$\tau(t) = \frac{T-t}{\sqrt{2k}}, \quad (79)$$

$$\widehat{\xi}_t = \frac{1}{\sqrt{2k}} \mathbb{E} \left[\int_t^T \frac{e^{\tau(u)} + e^{-\tau(u)}}{e^{\tau(t)} - e^{-\tau(t)}} \xi_u du \middle| \mathcal{F}_t \right]. \quad (80)$$

Note the similarity between equations (80) and (73). In equation (73), the relevant quantity at any point in time is not the trade signal μ_t itself, but instead the optimal strategy depends on

the expected value of the signal at all times in the future, modified according to an integration kernel. A similar theme occurs in (80) where the relevant quantity to the optimal efficient hedging strategy is not the intended hedging target ξ_t , but a weighted average of the expected future value of the target. The precise way that the target hedge ratio enters the optimal trading strategy is summarized by the following proposition.

Proposition 5. *The value function corresponding to the performance criterion in (78) is given by*

$$\begin{aligned}
H(x_0, q_0, S_0) = & x_0 + q_0 S_0 - \sqrt{\frac{k}{2}} \tanh(\tau(0)) (q - \widehat{\xi}_0)^2 \\
& - \frac{1}{2} \mathbb{E} \left[\int_0^T (\xi - \widehat{\xi}_t)^2 dt \right] - \sqrt{\frac{k}{2}} \mathbb{E} \left[\int_0^T \tanh(\tau(t)) d\langle \widehat{\xi} \rangle_t \right], \tag{81}
\end{aligned}$$

and the trading strategy which achieves this performance yields an inventory process that satisfies

$$dQ_t = \frac{\tanh(\tau(t))}{\sqrt{2k}} (\widehat{\xi}_t - Q_t) dt, \quad Q_0 = q_0. \tag{82}$$

In equation (82) we see that at any time, the agent will trade such that her inventory moves towards the weighted target $\widehat{\xi}_t$ with speed being proportional to the current hedging error, with the coefficient being a function depending on the remaining time $T-t$ and the temporary price impact parameter k . This function increases with respect to the remaining time, indicating that the agent puts more effort into tracking the hedge ratio when there is a long amount of time remaining in the trading interval. Indeed, when close to the end of the trading horizon, the performance criterion has relatively small penalty associated with the remaining tracking error, so the agent deems it more efficient to minimize the losses due to trading rather than holding an accurate hedge.

In Bank et al. (2017) the authors also solve this hedge tracking problem when a terminal inventory constraint is imposed of the form $Q_T' = \Xi_T$, for a given $\Xi_T \in L^2(\Omega, \mathcal{F}_T, \mathbb{P})$. Other works related to efficient hedging include Almgren and Li (2016), Guéant and Pu (2017), and Ekren and Nadtochiy (2022) in which an agent holds a European option on the traded asset and attempts to maximize the expected utility or mean-variance performance of her terminal wealth. The work Cartea et al. (2020) is similar except the agent holds a contingent claim on a non-traded risky quantity which is hedged by a correlated asset, and provide closed form trading strategies which approximate the optimal hedge.

5.3. Targeting VWAP

Each of the previous sections considers an agent that measures her performance using a criterion which is internal. That is, there is no sense of measuring performance relative to other market participants or market state. In this section, we consider an agent who wishes to employ a trading strategy with performance measured by the average execution price of her trades relative to the volume weighted average price (VWAP). The time weighted average price (TWAP) can be captured by any trading strategy which trades at a constant speed over the time interval of interest, and such a strategy has already been seen in equations (19) and (20). Unlike TWAP, VWAP depends on the actions of all market participants and measures the average transaction price of all trades over a time period, and the uncertainty in the actions of all other market participants makes it a more difficult quantity to target.

Suppose that all trades over our time period of interest occur at a sequence of times denoted $\tau = \{\tau_n\}_{n=1}^N$, and that the trade at time τ_n has volume V_n which is executed when the midprice is S_{τ_n} .⁵ Then the VWAP for this time period is

$$\frac{\sum_{n=1}^N V_n S_{\tau_n}}{\sum_{n=1}^N V_n}. \quad (83)$$

Working in continuous time, if we denote total order-flow in the market by $\mu = \{\mu_t\}_{0 \leq t \leq T}$, then the appropriate expression for VWAP is

$$\frac{\int_0^T \mu_t S_t dt}{\int_0^T \mu_t dt}. \quad (84)$$

VWAP is a frequently used benchmark for a trading strategy because it represents the trade performance across the entire market, and additionally is easy to compute. When an agent is trying to optimally liquidate a block of shares with performance measured relative to VWAP, she must time her trades throughout the trading interval so that her average trade price is close to this target. The main difficulty is that VWAP is uncertain because there is randomness in future market order-flow. This difficulty is present in addition to trying to maintaining trade efficiency regarding price impact considerations. One method of attempting to target VWAP is to instead target the percentage of volume (POV) at each point in time. This is the approach taken in Cartea and Jaimungal (2016a) which we now focus on.

⁵We take all of these volumes to be positive whether they are buys or sells. One could alternatively consider the VWAP of only buy or sell trades depending on whether the agent is buying or selling, but typically all trades are considered in practice. Additionally, VWAP is typically measured with respect to the midprice without taking transaction or price impact effects into account.

Let $\mu^\pm = \{\mu_t^\pm\}_{0 \leq t \leq T}$ denote the order-flow of buy and sell orders in the market, both assumed to be positive. Then the total order-flow at time t including the agent's trading is given by $\mu^+ + \mu^- + |\nu_t|$. The midprice is assumed to be linearly impacted by all order-flow, and so its dynamics are

$$dS_t^\nu = b(\mu_t^+ - \mu_t^- + \nu_t) dt + dM_t, \quad (85)$$

where M is a martingale independent of μ^\pm . Similar to Section 5.1 we assume that μ^\pm have the Markov property with respect to our filtration. Instead of directly trying to target VWAP, the agent will attempt to track POV through time. Let $\rho \in (0, 1)$ be fixed and define $\chi_t^\nu = \rho(\mu^+ + \mu^- + |\nu_t|)$, and specify the agent's dynamic value function to be

$$H(t, X_t, Q_t, S_t, \mu_t^\pm) = \sup_{\nu \in \mathcal{A}} \mathbb{E} \left[X_T^\nu + Q_T^\nu (S_T^\nu - \alpha Q_T^\nu) - \varphi \int_t^T (|\nu_u| - \chi_u^\nu)^2 dt \middle| \mathcal{F}_t \right], \quad (86)$$

where $\varphi > 0$ is a penalty parameter for deviating from POV. To see why the inclusion of the POV target assists in targeting VWAP, suppose that the agent is only selling, meaning $\nu_t \leq 0$, and that her trading speed exactly matches the percentage of total order-flow, giving $\nu_t = -\rho(\mu_t^+ + \mu_t^- + |\nu_t|)$. Then the agent's trades, marked to the midprice as in (84), are executed at an average value of

$$\frac{\int_0^T \nu_t S_t^\nu dt}{\int_0^T \nu_t dt} = \frac{\int_0^T (\mu_t^+ + \mu_t^- + |\nu_t|) S_t^\nu dt}{\int_0^T (\mu_t^+ + \mu_t^- + |\nu_t|) dt}, \quad (87)$$

and thus she has exactly targeted VWAP. Unfortunately, by following this strategy she will likely fail to liquidate her initial inventory position, and so inclusion of the penalty term in (86) initiates a trade-off between targeting VWAP and efficiently liquidating her inventory as represented by the other terms in the performance criterion.

For the remainder of this analysis we assume that the agent's optimal trading strategy will consist only of selling, meaning we may replace occurrences of $|\nu_t|$ with $-\nu_t$, resulting in simplified computations. This assumption will likely be the case when the agent's initial inventory is large and positive and when the need for liquidation outweighs the potential benefits from buying shares of the asset due to price forecasts similar to the treatment in Section 5.1. Under this assumption, the value function and optimal strategy are given by the following Proposition.

Proposition 6. *The value function in (86) is given by*

$$H(t, x, Q, S, \mu^\pm) = h_0(t, \mu^\pm) + h_1(t, \mu^\pm) q + h_2(t) q^2, \quad (88)$$

$$h_2(t) = -(\hat{\varphi} + k) \left((T - t) + \zeta \right)^{-1} - \frac{b}{2} \quad (89)$$

$$h_1(t, \mu^\pm) = b \int_t^T \frac{(T - s) + \zeta}{(T - t) + \zeta} \mathbb{E}[\mu_s^+ - \mu_s^- | \mathcal{F}_t] ds + \frac{2 \hat{\varphi} \hat{\rho}}{(T - t) + \zeta} \int_t^T \mathbb{E}[\mu_s^+ + \mu_s^- | \mathcal{F}_t] ds \quad (90)$$

$$h_0(t, \mu^\pm) = \int_t^T \mathbb{E} \left[\frac{(h_1(s, \mu_s^\pm) - 2 \hat{\varphi} \hat{\rho} (\mu_s^+ + \mu_s^-))^2}{4 (\hat{\varphi} + k)} - \hat{\varphi} \hat{\rho}^2 (\mu_s^+ + \mu_s^-)^2 \middle| \mathcal{F}_t \right] \quad (91)$$

with

$$\zeta = \frac{2(\hat{\varphi} + k)}{2\alpha - b}, \quad \hat{\varphi} = \varphi(1 - \rho)^2, \quad \hat{\rho} = \rho/(1 - \rho),$$

and the optimal trading strategy in feedback form is

$$\nu^*(t, q, \mu^\pm) = - \left((T - t) + \zeta \right)^{-1} q + \frac{h_1(t, \mu^\pm) - 2 \hat{\varphi} \hat{\rho} (\mu^+ + \mu^-)}{2 (\hat{\varphi} + k)}. \quad (92)$$

In (92) the first term which depends on inventory resembles a strategy that targets TWAP, and the remaining term is an adjustment that attempts to profit from expected future price forecasts (as seen from the term which includes $\mathbb{E}[\mu_s^+ - \mu_s^- | \mathcal{F}_t]$) and the desire to closely target POV (as seen in all terms including $\mu^+ + \mu^-$). Further intuition about the optimal strategy can be gathered when considering first the limit $\alpha \rightarrow \infty$, then the limit $\varphi \rightarrow \infty$. The first of these limits ensures that the agent fully liquidates her position by time T , and the second limit drives her to match POV as closely as possible under the full liquidation constraint. Evaluating these limits gives

$$\lim_{\varphi \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \nu^*(t, q, \mu^\pm) = -\frac{1}{T - t} q - \hat{\rho} \left((\mu^+ + \mu^-) - \frac{1}{T - t} \int_t^T \mathbb{E}[\mu_s^+ + \mu_s^- | \mathcal{F}_t] ds \right). \quad (93)$$

The dependence on expected future price has disappeared because of the infinite weight associated with full liquidation and volume tracking. The form of makes it clear that the agent's trading strategy has a component which directly depends on the current total order flow $\mu^+ + \mu^-$, but the influence is adjusted downward due to the influence of future trades.

In Cartea and Jaimungal (2016a) it is also shown how an agent may target VWAP by tracking percentage of cumulative volume (POCV) rather than POV. Other works which consider an agent attempting to match her trade executions to market VWAP are Konishi (2002)

and Frei and Westray (2015). Both derive expressions for strategies which minimize the VWAP targeting error, but rely on future knowledge of the total market volume. A study of guaranteed VWAP contracts is conducted in Guéant and Royer (2014), in which an agent guarantees to provide an execution service at VWAP prices, but then mitigates the risk of having entered this contract. See Białkowski et al. (2008) and Mitchell et al. (2020) for empirical studies of strategies which target VWAP.

5.4. Other Objectives

There are numerous other criteria that an agent could consider for the purposes of evaluating the performance of her trading strategy. Here we briefly mention some lines of research that do not explicitly fall into any of the categories of the previous few sections.

In Neuman et al. (2020) the agent is considered to be a central bank, and she is tasked with maintaining a currency peg or currency target zone, the goal of which is to make sure that a particular foreign exchange rate remains within a specified domain. This is done by trading in the underlying currencies and utilizing the permanent price impact from the trades in such a way as to move the exchange rate to its desired value. This task has significant inventory risk associated with it which the central bank wishes to manage. In Neuman and Schied (2022), a similar currency peg context is considered, but in addition to a central bank maintaining the peg there is a strategic investor attempting to disrupt the bank's objective. This is modeled as a stochastic differential game and a Stackelberg equilibrium is established.

The operation of modern markets is strongly dependent on technology, and the frequency at which events occur mean that latency can have an affect on the performance of an execution strategy. From the time that an agent submits her intention to trade until that signal arrives at the matching engine of the exchange, the state of the limit order book may have changed in such a way that the agent's order cannot be filled. In Cartea and Sánchez-Betancourt (2021b) the agent acknowledges that latency can affect whether her submitted orders are executed as expected, and adjusts her performance criterion to account for the fraction of her trades that are filled. See also Moallemi and Sağlam (2013) for some empirical analysis of the effects of latency on optimal execution strategies.

6. Other Phenomena

Below is a small number of other styles of model which do not explicitly fall into any of the previous categories. Many of the ideas explored in the following sections could be adapted

in conjunction with the topics explored already. The referenced material should be consulted for the mathematical details used to approach these phenomena.

6.1. Multiple Agents

Each setting considered so far has incorporated the strategic actions of only one agent. Some of the previous models have incorporated total order flow of the market within the dynamics, but any such model does not arise as a result of an optimization procedure carried out by other agents. In reality, no single agent acts in isolation, and their actions will have an effect on the state of the market, thereby altering the actions of other agents. When modeling several agents simultaneously, we move from a stochastic optimization into a stochastic game setting.

In Schied and Zhang (2019) a competitive game is considered between a finite number of agents where they are restricted to trading at discrete times. Continuous-time counterparts are studied in Strehle (2017) and Neuman and Voß (2021), and in Voß (2022) a two agent setting is considered in which they must both efficiently track an inventory target.

Finite player stochastic differential games generally result in a set of coupled differential equations which can prove to be intractable. Modeling the problem as a mean-field game, appropriate for a large number of agents, often reduces the complexity of the computations. The optimal liquidation problem in a mean field setting is detailed in Huang et al. (2019) and Fu et al. (2021). In Casgrain and Jaimungal (2020) and Donnelly and Lorig (2020) a mean-field game is used to consider agents who have different methods of speculation about the future price of the asset.

6.2. Robustness

Each of the optimal trading strategies from previous sections depend on a specification of a dynamic model for the price process and other relevant quantities. But the performance of any trading strategy is subject to risks arising from using a misspecified model, thus it is desirable to account for these risks and attempt to formulate a trading strategy which is robust to model error.

In Schied (2013) the optimality of a trading strategy with respect to the underlying model is studied. In Nyström et al. (2014), Cartea et al. (2018b), and Horst et al. (2022), the agent's performance criterion accounts for the possibility that the dynamics evolve according to a different model, and she thereby adjusts her trading strategy to compensate.

6.3. Market Making and Liquidation with Limit Orders

The results which have been outlined in previous sections frame a trading strategy in terms of the speed or volume of trades. Keeping in mind the discussion about the LOB structure, an agent can only guarantee that a particular volume of trades be executed if she employs MOs. Here we give references which study the strategic placement of LOs instead of or in conjunction with MOs.

An agent who simultaneously places buy and sell LOs and attempts to maximize profits is considered in Avellaneda and Stoikov (2008). Further work on optimal market making is considered in Guilbaud and Pham (2013), Guéant et al. (2013), Cartea et al. (2014), Cartea and Jaimungal (2015b), and Cartea et al. (2018a).

The incorporation of LOs on only one side of the LOB for use in optimal liquidation is considered in Guéant et al. (2012), Bayraktar and Ludkovski (2014), Guéant and Lehalle (2015), and Cartea and Jaimungal (2015a).

7. References

- Ackermann, J., T. Kruse, and M. Urusov (2021). Optimal trade execution in an order book model with stochastic liquidity parameters. *SIAM Journal on Financial Mathematics* 12(2), 788–822.
- Alfonsi, A., A. Fruth, and A. Schied (2010). Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance* 10(2), 143–157.
- Almgren, R. (2012). Optimal trading with stochastic liquidity and volatility. *SIAM Journal on Financial Mathematics* 3(1), 163–181.
- Almgren, R. and N. Chriss (2001). Optimal execution of portfolio transactions. *Journal of Risk* 3, 5–40.
- Almgren, R. and T. M. Li (2016). Option hedging with smooth market impact. *Market Microstructure and Liquidity* 2(01), 1650002.
- Almgren, R., C. Thum, E. Hauptmann, and H. Li (2005). Direct estimation of equity market impact. *Risk* 18(7), 58–62.
- Almgren, R. F. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance* 10(1), 1–18.

- Avellaneda, M. and S. Stoikov (2008). High-frequency trading in a limit order book. *Quantitative Finance* 8(3), 217–224.
- Bank, P., H. M. Soner, and M. Voß (2017). Hedging with temporary price impact. *Mathematics and Financial Economics* 11(2), 215–239.
- Bank, P. and M. Voß (2019). Optimal investment with transient price impact. *SIAM Journal on Financial Mathematics* 10(3), 723–768.
- Barger, W. and M. Lorig (2019). Optimal liquidation under stochastic price impact. *International Journal of Theoretical and Applied Finance* 22(02), 1850059.
- Bayraktar, E. and M. Ludkovski (2014). Liquidation in limit order books with controlled intensity. *Mathematical Finance* 24(4), 627–650.
- Becherer, D., T. Bilarev, and P. Frentrup (2018). Optimal liquidation under stochastic liquidity. *Finance and Stochastics* 22(1), 39–68.
- Bechler, K. and M. Ludkovski (2017). Order flows and limit order book resiliency on the meso-scale. *Market Microstructure and Liquidity* 3(03n04), 1850006.
- Bergault, P., F. Drissi, and O. Guéant (2022). Multi-asset optimal execution and statistical arbitrage strategies under ornstein–uhlenbeck dynamics. *SIAM Journal on Financial Mathematics* 13(1), 353–390.
- Bertsimas, D. and A. W. Lo (1998). Optimal control of execution costs. *Journal of Financial Markets* 1(1), 1–50.
- Biais, B., P. Hillion, and C. Spatt (1995). An empirical analysis of the limit order book and the order flow in the paris bourse. *The Journal of Finance* 50(5), 1655–1689.
- Białkowski, J., S. Darolles, and G. Le Fol (2008). Improving vwap strategies: A dynamic volume approach. *Journal of Banking & Finance* 32(9), 1709–1722.
- Björk, T. (2009). *Arbitrage theory in continuous time*. Oxford University Press.
- Bouchard, B., N.-M. Dang, and C.-A. Lehalle (2011). Optimal control of trading algorithms: a general impulse control approach. *SIAM Journal on Financial Mathematics* 2(1), 404–438.
- Cartea, Á., R. Donnelly, and S. Jaimungal (2017). Algorithmic trading with model uncertainty. *SIAM Journal on Financial Mathematics* 8(1), 635–671.
- Cartea, Á., R. Donnelly, and S. Jaimungal (2018a). Enhancing trading strategies with order book signals. *Applied Mathematical Finance* 25(1), 1–35.

- Cartea, Á., R. Donnelly, and S. Jaimungal (2018b). Portfolio liquidation and ambiguity aversion. In *High-Performance Computing in Finance*, pp. 77–114. Chapman and Hall/CRC.
- Cartea, Á., R. Donnelly, and S. Jaimungal (2020). Hedging nontradable risks with transaction costs and price impact. *Mathematical Finance* 30(3), 833–868.
- Cartea, Á., F. Drissi, and M. Monga (2022). Decentralised finance and automated market making: Predictable loss and optimal liquidity provision. *Available at SSRN 4273989*.
- Cartea, Á., L. Gan, and S. Jaimungal (2019). Trading co-integrated assets with price impact. *Mathematical Finance* 29(2), 542–567.
- Cartea, Á. and S. Jaimungal (2015a). Optimal execution with limit and market orders. *Quantitative Finance* 15(8), 1279–1291.
- Cartea, Á. and S. Jaimungal (2015b). Risk metrics and fine tuning of high-frequency trading strategies. *Mathematical Finance* 25(3), 576–611.
- Cartea, Á. and S. Jaimungal (2016a). A closed-form execution strategy to target volume weighted average price. *SIAM Journal on Financial Mathematics* 7(1), 760–785.
- Cartea, Á. and S. Jaimungal (2016b). Incorporating order-flow into optimal execution. *Mathematics and Financial Economics* 10(3), 339–364.
- Cartea, Á., S. Jaimungal, and J. Penalva (2015). *Algorithmic and High-frequency Trading*. Cambridge University Press.
- Cartea, Á., S. Jaimungal, and J. Ricci (2014). Buy low, sell high: A high frequency trading perspective. *SIAM Journal on Financial Mathematics* 5(1), 415–444.
- Cartea, Á. and L. Sánchez-Betancourt (2021a). Optimal execution with stochastic delay. *Available at SSRN 3812324*.
- Cartea, Á. and L. Sánchez-Betancourt (2021b). The shadow price of latency: Improving intraday fill ratios in foreign exchange markets. *SIAM Journal on Financial Mathematics* 12(1), 254–294.
- Casgrain, P. and S. Jaimungal (2019). Trading algorithms with learning in latent alpha models. *Mathematical Finance* 29(3), 735–772.
- Casgrain, P. and S. Jaimungal (2020). Mean-field games with differing beliefs for algorithmic trading. *Mathematical Finance* 30(3), 995–1034.
- Chan, P. and R. Sircar (2015). Optimal trading with predictable return and stochastic volatility. *Available at SSRN 2623747*.

- Chen, Y., U. Horst, and H. H. Tran (2019). Portfolio liquidation under transient price impact—theoretical solution and implementation with 100 nasdaq stocks. *Available at SSRN 3504133*.
- Cheridito, P. and T. Sepin (2014). Optimal trade execution under stochastic volatility and liquidity. *Applied Mathematical Finance* 21(4), 342–362.
- Colaneri, K., Z. Eksi, R. Frey, and M. Szölgényi (2020). Optimal liquidation under partial information with price impact. *Stochastic Processes and their Applications* 130(4), 1913–1946.
- Cont, R. (2011). Statistical modeling of high-frequency financial data. *IEEE Signal Processing Magazine* 28(5), 16–25.
- Cont, R., A. Kukanov, and S. Stoikov (2014). The price impact of order book events. *Journal of Financial Econometrics* 12(1), 47–88.
- Donnelly, R. and M. Lorig (2020). Optimal trading with differing trade signals. *Applied Mathematical Finance* 27(4), 317–344.
- Ekren, I. and J. Muhle-Karbe (2019). Portfolio choice with small temporary and transient price impact. *Mathematical Finance* 29(4), 1066–1115.
- Ekren, I. and S. Nadtochiy (2022). Utility-based pricing and hedging of contingent claims in almgren-chriss model with temporary price impact. *Mathematical Finance* 32(1), 172–225.
- Fleming, W. H. and H. M. Soner (2006). *Controlled Markov processes and viscosity solutions*, Volume 25. Springer Science & Business Media.
- Forde, M., L. Sánchez-Betancourt, and B. Smith (2021). Optimal trade execution for Gaussian signals with power-law resilience. *Quantitative Finance*, 1–12.
- Fouque, J.-P., S. Jaimungal, and Y. F. Saporito (2022). Optimal trading with signals and stochastic price impact. *SIAM Journal on Financial Mathematics* 13(3), 944–968.
- Frei, C. and N. Westray (2015). Optimal execution of a VWAP order: a stochastic control approach. *Mathematical Finance* 25(3), 612–639.
- Fruth, A., T. Schöneborn, and M. Urusov (2019). Optimal trade execution in order books with stochastic liquidity. *Mathematical Finance* 29(2), 507–541.
- Fu, G., P. Graewe, U. Horst, and A. Popier (2021). A mean field game of optimal portfolio liquidation. *Mathematics of Operations Research* 46(4), 1250–1281.

- Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative Finance* 10(7), 749–759.
- Gatheral, J. and A. Schied (2011). Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework. *International Journal of Theoretical and Applied Finance* 14(03), 353–368.
- Gatheral, J. and A. Schied (2013). Dynamical models of market impact and algorithms for order execution. *Handbook on Systemic Risk*, Jean-Pierre Fouque, Joseph A. Langsam, eds, 579–599.
- Graewe, P. and U. Horst (2017). Optimal trade execution with instantaneous price impact and stochastic resilience. *SIAM Journal on Control and Optimization* 55(6), 3707–3725.
- Guéant, O. (2016). *The Financial Mathematics of Market Liquidity: From optimal execution to market making*, Volume 33. CRC Press.
- Guéant, O. and C.-A. Lehalle (2015). General intensity shapes in optimal liquidation. *Mathematical Finance* 25(3), 457–495.
- Guéant, O., C.-A. Lehalle, and J. Fernandez-Tapia (2012). Optimal portfolio liquidation with limit orders. *SIAM Journal on Financial Mathematics* 3(1), 740–764.
- Guéant, O., C.-A. Lehalle, and J. Fernandez-Tapia (2013). Dealing with the inventory risk: a solution to the market making problem. *Mathematics and Financial Economics* 7(4), 477–507.
- Guéant, O. and J. Pu (2017). Option pricing and hedging with execution costs and market impact. *Mathematical Finance* 27(3), 803–831.
- Guéant, O. and G. Royer (2014). VWAP execution and guaranteed VWAP. *SIAM Journal on Financial Mathematics* 5(1), 445–471.
- Guilbaud, F. and H. Pham (2013). Optimal high-frequency trading with limit and market orders. *Quantitative Finance* 13(1), 79–94.
- Guo, X. and M. Zervos (2015). Optimal execution with multiplicative price impact. *SIAM Journal on Financial Mathematics* 6(1), 281–306.
- Horst, U. and X. Xia (2019). Multi-dimensional optimal trade execution under stochastic resilience. *Finance and Stochastics* 23(4), 889–923.
- Horst, U., X. Xia, and C. Zhou (2022). Portfolio liquidation under factor uncertainty. *The Annals of Applied Probability* 32(1), 80–123.

- Huang, X., S. Jaimungal, and M. Nourian (2019). Mean-field game strategies for optimal execution. *Applied Mathematical Finance* 26(2), 153–185.
- Ishitani, K. and T. Kato (2014). Optimal execution for uncertain market impact: derivation and characterization of a continuous-time value function. In *2012 Recent Advances in Financial Engineering: Proceedings of the International Workshop on Finance 2012*, pp. 93–116. World Scientific.
- Konishi, H. (2002). Optimal slice of a VWAP trade. *Journal of Financial Markets* 5(2), 197–221.
- Lehalle, C.-A. and O. Mounjid (2017). Limit order strategic placement with adverse selection risk and the role of latency. *Market Microstructure and Liquidity* 3(01), 1750009.
- Lehalle, C.-A. and E. Neuman (2019). Incorporating signals into optimal trading. *Finance and Stochastics* 23(2), 275–311.
- Mitchell, D., J. Białkowski, and S. Tompaidis (2020). Volume-weighted average price tracking: A theoretical and empirical study. *IISE Transactions* 52(8), 864–889.
- Moallemi, C. C. and M. Sağlam (2013). On the cost of latency in high-frequency trading. *Operations Research* 61(5), 1070–1086.
- Neuman, E. and A. Schied (2022). Protecting pegged currency markets from speculative investors. *Mathematical Finance* 32(1), 405–420.
- Neuman, E., A. Schied, C. Weng, and X. Xue (2020). A central bank strategy for defending a currency peg. *Systems & Control Letters* 144, 104761.
- Neuman, E. and M. Voß (2021). Trading with the crowd. *arXiv preprint arXiv:2106.09267*.
- Neuman, E. and M. Voß (2022). Optimal signal-adaptive trading with temporary and transient price impact. *SIAM Journal on Financial Mathematics* 13(2), 551–575.
- Nyström, K., S. M. Ould Aly, and C. Zhang (2014). Market making and portfolio liquidation under uncertainty. *International Journal of Theoretical and Applied Finance* 17(05), 1450034.
- Obizhaeva, A. A. and J. Wang (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets* 16(1), 1–32.
- Pham, H. (2009). *Continuous-time stochastic control and optimization with financial applications*, Volume 61. Springer Science & Business Media.

- Preddoiu, S., G. Shaikhet, and S. Shreve (2011). Optimal execution in a general one-sided limit-order book. *SIAM Journal on Financial Mathematics* 2(1), 183–212.
- Ranaldo, A. (2004). Order aggressiveness in limit order book markets. *Journal of Financial Markets* 7(1), 53–74.
- Schied, A. (2013). Robust strategies for optimal order execution in the almgren–chris framework. *Applied Mathematical Finance* 20(3), 264–286.
- Schied, A. and T. Zhang (2019). A market impact game under transient price impact. *Mathematics of Operations Research* 44(1), 102–121.
- Sirignano, J. A. (2019). Deep learning for limit order books. *Quantitative Finance* 19(4), 549–570.
- Strehle, E. (2017). Optimal execution in a multiplayer model of transient price impact. *Market Microstructure and Liquidity* 3(03n04), 1850007.
- Touzi, N. and A. Tourin (2013). *Optimal stochastic control, stochastic target problems, and backward SDE*, Volume 29. Springer.
- Voß, M. (2022). A two-player portfolio tracking game. *Mathematics and Financial Economics* 16(4), 779–809.
- Yong, J. and X. Y. Zhou (1999). *Stochastic controls: Hamiltonian systems and HJB equations*, Volume 43. Springer Science & Business Media.