

 Open access • Journal Article • DOI:10.2139/SSRN.2557457

Incorporating Order-Flow into Optimal Execution — [Source link](#)

Álvaro Cartea, Sebastian Jaimungal

Institutions: University of Oxford, University of Toronto

Published on: 12 Jun 2015 - Social Science Research Network

Topics: Order (exchange), Trading strategy, Algorithmic trading, Pairs trade and Volume-weighted average price

Related papers:

- [A Closed-Form Execution Strategy to Target Volume Weighted Average Price](#)
- [Trading strategies for stock pairs regarding to the cross-impact cost](#)
- [Optimal VWAP Tracking](#)
- [Order Execution and Limit Order Book](#)
- [Enhancing Trading Strategies with Order Book Signals](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/incorporating-order-flow-into-optimal-execution-2lolwkwd2b>

Incorporating Order-Flow into Optimal Execution[☆]

Forthcoming: Mathematics and Financial Economics

Álvaro Cartea^a, Sebastian Jaimungal^b

^a*Department of Mathematics, University of Oxford, Oxford, UK
Oxford-Man Institute of Quantitative Finance, Oxford, UK*

^b*Department of Statistical Sciences, University of Toronto, Toronto, Canada*

Abstract

We provide an explicit closed-form strategy for an investor who executes a large order when market order-flow from all agents, including the investor's own trades, has a permanent price impact. The strategy is found in closed-form when the permanent and temporary price impacts are linear in the market's and investor's rates of trading. We do this under very general assumptions about the stochastic process followed by the order-flow of the market. The optimal strategy consists of an Almgren-Chriss execution strategy adjusted by a weighted-average of the future expected net order-flow (given by the difference of the market's rate of buy and sell market orders) over the execution trading horizon and proportional to the ratio of permanent to temporary linear impacts. We use historical data to calibrate the model to Nasdaq traded stocks and use simulations to show how the strategy performs.

Keywords: Order-Flow, Algorithmic Trading, High Frequency Trading, Acquisition, Liquidation, Price Impact.

1. Introduction

When institutional investors come to market to execute large orders they bear direct and indirect costs. Direct costs consist of fees and other transaction costs which are paid to brokers and exchanges, are easy to compute and known upfront. Indirect costs are price impact costs that result from the traders' actions in the market and are very difficult to quantify even after

[☆]SJ would like to thank NSERC and GRI for partially funding this work. AC acknowledges the research support of the Oxford-Man Institute for Quantitative Finance and the hospitality of the Finance Group at Saïd Business School.

Email addresses: `alvaro.cartea@maths.ox.ac.uk` (Álvaro Cartea),
`sebastian.jaimungal@utoronto.ca` (Sebastian Jaimungal)

executing the order. In most cases price impact has an adverse effect on execution price: lower revenues if liquidating shares or higher costs if acquiring shares. In this paper we determine the optimal execution strategy when the price impact of market order-flow, which is given by the trading activity of all market participants, including the investor's own trades, affects the midprice of the asset and, in addition, how the investor's trades affect the execution prices she receives.

Price impact can be temporary or permanent. Market orders (MOs) that walk the limit order book (LOB) have price impact because they obtain average execution prices worse than the best quote and are considered temporary because it is assumed that: the LOB will replenish very quickly and, after restocking, there is no effect on the midprice of the asset. On the other hand, investors' trading activity has a permanent effect on midprices because traders' order-flow, i.e. the number and the size of MOs, conveys information that is impounded on the midprice of the asset. In general, when there is positive net order-flow (more buy than sell volume) the midprice tends to drift up, and when there is negative net order-flow (more sell than buy volume) the midprice tends to drift down. Thus, when an investor is executing a large order this adds one-sided pressure to the market's order-flow and this adds directional pressure on midprices: up if the investor is buying or down if selling.

In this paper we assume a general stochastic process for order-flow and provide an explicit closed-form expression for the optimal execution strategy when permanent and temporary impact are linear in the overall market's rate of trading, and the investor's execution rate respectively. The optimal strategy consists of an Almgren-Chriss execution strategy adjusted by a weighted-average of the future expected net order-flow over the execution trading horizon and proportional to the ratio of permanent to temporary linear impacts.

The rest of this paper is organized as follows. In Section 2 we use historical data to estimate the permanent impact effect of net order-flow on midprices and use LOB data to estimate temporary price impact. In Section 3 we set the model for the execution strategy, present the general stochastic process followed by order-flow, and derive the optimal execution strategy. Finally, in Section 4 we use simulations to show the performance of the strategy and Section 5 concludes.

2. Order-Flow, Trading Activity, and Price Impact

In this section we present statistics and parameter estimates for permanent and temporary price impact for stocks traded in Nasdaq for the year 2014. To focus the discussion, and for ease of presentation, in this section we present the results for four stocks – in the Appendix we show results for 17 stocks, see Tables A.7 and A.8. In the first row of Table 1 we show the parameter estimate \hat{b} for permanent price impact, and in the second row (in parenthesis)

	INTC	PNC	CNQR	BXS
\hat{b}	1.09×10^{-6} (0.54×10^{-6})	1.99×10^{-5} (0.88×10^{-5})	0.89×10^{-4} (0.58×10^{-4})	2.72×10^{-5} (1.57×10^{-5})
\hat{k}	0.44×10^{-6} (2.37×10^{-7})	1×10^{-5} (0.49×10^{-5})	0.76×10^{-4} (0.61×10^{-4})	1.72×10^{-5} (1.35×10^{-5})
$\widehat{\frac{b}{k}}$	2.5 (0.76)	2.11 (1.05)	1.1 (0.56)	1.73 (0.92)
ADV	5,489,503 (2,542,732)	382,020 (139,898)	164,904 (101,910)	92,525 (54,475)
Midprice	30.26 (4.37)	84.67 (3.22)	105.96 (16.42)	22.88 (1.55)
σ (open-to-close)	16.34%	15.25%	29.07%	26.46%
σ (5 min)	17.87%	17.95%	38.23%	36.47%
λ^+	438.86 (248.85)	206.79 (69.29)	96.47 (58.39)	47.64 (22.19)
$\mathbb{E}[\eta^+]$	1048.99 (351.13)	136.35 (16.16)	128.79 (34.09)	143.36 (33.42)
λ^-	453.91 (264.63)	209.22 (71.7)	99.96 (59.38)	48.19 (24.09)
$\mathbb{E}[\eta^-]$	1013.83 (306.58)	143.73 (16.95)	136.23 (49.1)	143.23 (34.24)

Table 1: Permanent and temporary price impact parameters for Nasdaq stocks, average volume of MOs, average midprice, σ volatility (annualized) of price returns using open-to-close prices and 5-minute windows, hourly mean arrival of MOs λ^\pm , and average volume of MOs $\mathbb{E}[\eta^\pm]$. The standard deviation of the estimate is shown in parentheses. Data are from Nasdaq 2014.

we show its standard deviation. We assume a linear relationship between net order-flow and changes in the midprice, thus for every trading day we perform the regression

$$\Delta S_n = b \mu_n + \varepsilon_n, \quad (1)$$

where $\Delta S_n = S_{n\tau} - S_{(n-1)\tau}$ is the change in the midprice, μ_n is net order-flow defined as the difference between the volume of buy and sell MOs during the time interval $[(n-1)\tau, n\tau]$, and ε_n is the error term (assumed normal). In the empirical analysis we choose $\tau = 5 \text{ min}$.¹ The table shows the mean and standard deviation of the daily estimate for b by first removing the upper and lower 0.5% tails and then carrying out a robust linear regression on the model (1).

In the third row of the table we show the parameter estimate for temporary impact. To estimate this parameter, which we denote by k , we assume that temporary price impact is linear in the rate of trading so the difference between the execution price that the investor

¹See Almgren et al. (2005) for a discussion on linear market impact using proprietary execution data.




Figure 1: Order-Flow and effect on the drift of midprice of INTC. 1st panel: daily estimates for b for every trading day in 2014. 2nd panel: histogram of net order-flow. 3rd panel: expected net order-flow (with error bars) conditioned on a given price change being observed. 4th panel: histogram of amalgamated residuals from the daily robust regression of order-flow onto price changes (together with a gaussian fit shown in red).

receives and the midprice is $k \Delta Q$, where ΔQ is the volume of trades. In the continuous time setting $\Delta Q = \nu \Delta t$, when the agent trades at a speed of ν over a small time interval Δt . To estimate the impact parameter k , we take a snapshot of the buy side of the LOB each second, determine the price per share for various volumes (by walking through the LOB), compute the difference between the price per share and the best quote at that time, and perform a linear regression. The slope of the linear regression is an estimate of the temporary price impact per share at that time. We do this for every second of every trading day and in the table we report the mean and standard deviation of these daily estimates when we exclude the first and last half-hour of the trading day, and remove the upper and lower 0.5% tails. Although here we only use the buy side of the LOB, because later we discuss in detail a liquidation strategy, including both sides does not affect the results, the order of magnitude of the estimated k is statistically the same for both sides of the book. Moreover, the fifth row shows the mean and standard deviation of the daily ratio b/k . We observe that CNQR shows the smallest ratio of 1.1 and INTC shows the largest at 2.5 – at the end of this section we discuss this ratio in more detail.

The table also reports the average daily volume (ADV) of shares for each stock, the average midprice, and the (annualized) volatility (σ) of price returns calculated using open-to-close prices and employing five-minute windows (to remove any excess spurious volatility due to microstructure noise). In addition, we report the average (hourly) number of buy and sell MOs, λ^+ and λ^- respectively, and the mean volume of MOs, $\mathbb{E}[\eta^+]$ and $\mathbb{E}[\eta^-]$ respectively. For example, in Nasdaq, INTC receives on an hourly basis 439 market buy orders, with an average of 1,049 shares per order. The parameter estimates for MO arrival and mean volume are statistically the same for both sides of the LOB. This is expected over long time scales. There could be days or periods of the day where there is more activity on either the buy or sell side, but over longer time scales buy and sell MOs are symmetric.

Moreover, to showcase the variability of the permanent price impact parameter, the first panel of Figure 1 depicts the estimate of b for each day of 2014 –the dashed line shows the average \hat{b} . The p -values corresponding to \hat{b} for all stocks are significant to less than 0.01%




Figure 2: An illustration of how the temporary impact may be estimated from snapshots of the LOB using INTC on Nov 3, 2014. The first panel is at 11:00am, the second from 11:00am to 11:01am and the third contains the entire day.

when all 5-minute windows for the full year are amalgamated. On any given day, however, the parameter may be less significant (i.e. the p -value may be larger). The second panel shows a histogram of the five-minute net order-flow using all the data in 2014. The third panel shows the expected net order-flow (with error bars) conditioned on a given price change being observed.² Finally, the fourth panel shows a histogram of amalgamated residuals from the daily robust regression of order-flow onto price changes and a gaussian fit shown in red solid line.

As already shown by the regression results there is a positive relationship between net order-flow and price changes. The third panel of Figure 1 provides further support of the finding that when net order-flow is positive (negative), that is more (less) buy than sell MOs, the midprice tends to increase (decrease). We see that assuming a linear relationship between price changes and net order-flow is plausible for a wide range of midprice changes. Only in the two extremes, very high or very low price changes, does the relationship fail to be linear. Nevertheless, there are fewer observations in the tails as shown by the histogram in the top-right panel and this is also shown by the confidence intervals around the estimates.

In Figure 2 we show a snapshot of the LOB (top left panel) for INTC on Nov 3, 2014 at 11am together with the price impact per share (measured from the best quote) that an MO of various volumes would face as it walks through the buy side of the LOB for every second from 11:00 to 11:01 as well as a linear regression (the dash line) with intercept set to zero. Naturally, the impact function fluctuates within the minute, and with it the impact that trades of different size have. The linear regression provides an approximation of the temporary impact during that one minute. The third picture shows how the slope of this linear impact model fluctuates throughout the entire day. We see that the largest impact tends to occur in the morning, then this impact flattens and stays flat throughout the day, and towards the end of the day it lessens. Such a pattern is seen in a number of assets across

²For the year 2014, 99% of the 5 minute price changes for INTC were within the range $[-0.1, 0.1]$.

all days.

The analysis above looks at temporary and permanent effects separately, but their joint dynamics is a relevant quantity in execution algorithms. Liquidation and acquisition strategies take into account the trade-off between costs that stem from walking the book and the permanent impact. In particular, when both types of impact are linear in the rates of trading, this trade-off is in part captured by the ratio b/k . On the left-hand side of Figure 3 we show a scatter plot of the daily pair (b, k) for INTC, where it is clear that there is a positive relationship between temporary and permanent impact. Days of high (low) permanent impact are usually those in which MOs of the same volume must deplete more (less) levels of the book. The right-hand side of the figure depicts a histogram of b/k which shows that this ratio ranges between 0.5 and 5 and is symmetric around 2.5 approximately.




Figure 3: Price Impact INTC using daily observations for 2014.

Furthermore, everything else equal, stocks with a high b/k ratio should be liquidated in a different fashion from those with a low ratio. When the ratio is high, the optimal strategy reacts more abruptly to current and expected net-order flow. For example, in a liquidation programme, if expected net order-flow is negative (negative trend in the midprice), it is optimal to execute shares quicker to avoid the adverse effect of the stock price declining. If the expected net order-flow is positive, it is better to slow down and liquidate fewer shares to take advantage of the positive trend in the midprice. This result is clearly borne out by some of the strategies we develop below – see for instance (14) and the discussion that follows.

Finally, Table 2 shows the correlation between b and k , and their skewness. We observe that temporary and permanent impacts are positively correlated and are positively skewed. Intuitively, one expects these two impacts to be positively correlated because execution algorithms, like the one proposed here, must trade off permanent and temporary impact.

	INTC	PNC	CNQR	BXS
$corr(b, k)$	0.72	0.00	0.61	0.41
$skew(b)$	1.13	0.57	0.33	1.19
$skew(k)$	1.59	6.79	1.13	2.97

Table 2: Skewness of b and k and their correlation.

3. The Model

The investor must choose the speed at which she sends MOs to liquidate \mathfrak{N} shares over a trading horizon T . Here we focus on the liquidation problem – the setup for the acquisition problem is similar. We denote the liquidation speed, which is what the investor controls, by $\nu = (\nu_t)_{\{0 \leq t \leq T\}}$ and inventory is $Q^\nu = (Q_t^\nu)_{\{0 \leq t \leq T\}}$, which is affected by how fast she trades, and satisfies

$$dQ_t^\nu = -\nu_t dt, \quad Q_0^\nu = \mathfrak{N}. \quad (2)$$

The investor’s speed of trading affects prices in two ways. One is a permanent impact that affects the midprice and the other a temporary impact that only affects the prices she obtains when she sells shares.

Permanent impact. We assume that when the investor sends MOs to liquidate shares, her actions induce a downward drift in the midprice. Similarly, the order-flow from the rest of the market also affects the midprice dynamics, see also Cartea and Jaimungal (2015b). In particular, the speed at which other traders send buy and sell MOs, denoted here by $\mu^+ = \{\mu_t^+\}_{0 \leq t \leq T}$ and $\mu^- = \{\mu_t^-\}_{0 \leq t \leq T}$, jointly assumed Markov, cadlag and bounded \mathbb{P} –a.s., also exert upward and downward pressure on the midprice respectively.³ In this manner, the action of the investor’s trades and other traders’ actions are treated symmetrically. Thus, the midprice process $S^\nu = (S_t^\nu)_{\{0 \leq t \leq T\}}$, satisfies the stochastic differential equation (SDE)

$$dS_t^\nu = (g(\mu_t^+) - g(\mu_t^- + \nu_t)) dt + \sigma dW_t, \quad S_0^\nu = S, \quad (3)$$

where $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denotes the permanent impact on the midprice as a result of order-flow, $W = (W_t)_{\{0 \leq t \leq T\}}$ is a standard Brownian motion, and μ^\pm are independent⁴ of W .

Here, the Brownian component W represents the changes in the midprice due to the reshuffling of limit orders in the absence of MOs. This reshuffling occurs at different levels of

³It is possible to allow order-flow to be Markov in not only μ^\pm by introducing other exogenous processes that also drive order-flow in a Markov manner, but we opt to leave this generalization out of the current analysis.

⁴It is possible to incorporate dependence, however, the resulting strategy does not change, only the filtration on which strategies are adapted and certain expectations will be computed using this extended filtration. To keep the analysis compact we opt to leave this dependence out.

the LOB including the best quotes to reflect the market's view on the fundamental value of the asset. Moreover, one can also view the effect of net order-flow as inducing a trend (short-lived) in the midprice of the asset due to excess buy or sell pressure. The effect of trends in the midprice have been studied in the algorithmic trading literature, see for instance Lorenz and Schied (2013), Cartea et al. (2014), Cartea et al. (2013), Cartea and Jaimungal (2013).

We remark that although in (3) we allow for generality in the choice of the price impact function g , the set of choices should only contain those that do not allow for price manipulation. Below we assume that $g(x)$ is linear in x which precludes price manipulation or arbitrage. For a general discussion on the shape and assumptions of permanent price impact see Almgren (2003), Gatheral (2010), Alfonsi et al. (2010) Gatheral et al. (2012).

Temporary impact. At any one time, the number of shares displayed and available in the market at the best bid ($S_t^\nu - \frac{1}{2}\Delta$) is limited – here $\Delta \geq 0$ is the bid-ask spread which is assumed to be a constant. A large MO will walk the book, so that the average price per share obtained will be worse than the current bid price. This is captured in our model, as an order of size νdt obtains an execution price per share of

$$\hat{S}_t^\nu = S_t^\nu - \left(\frac{1}{2}\Delta + f(\nu_t)\right), \quad (4)$$

where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denotes the temporary impact that the investor's trading action has on the price she can execute the trade at. In this setup the impact of the MO is temporary and only affects the price received by the trader who sent the MO, and, moreover, the LOB recovers infinitely fast to its state previous to the arrival of the MO, see Almgren (2003), Alfonsi et al. (2010), Kharroubi and Pham (2010), Gatheral et al. (2012), Schied (2013), Guéant and Lehalle (2013) for further discussions and generalizations to relax the resilience assumption. See also Graewe et al. (2013) for a discussion of the asymptotic behaviour of non-linear impact near the end of the trading horizon, as well as smoothness of the solutions. Moreover, as Frei and Westray (2013) point out, given the extremely low predictive accuracy of market impact models (typically $< 5\%$ R^2), the cost of increased complexity arising from moving away from a linear model would outweigh any gains from better describing market impact.

Finally, the investor's cash process, X_t^ν satisfies the SDE

$$dX_t^\nu = \hat{S}_t^\nu \nu_t dt, \quad X_0^\nu = x. \quad (5)$$

Performance criteria and value function. The investor's performance criteria is

$$H^\nu(t, x, S, \boldsymbol{\mu}, q) = \mathbb{E}_{t,x,S,\boldsymbol{\mu},q} \left[X_T + Q_T^\nu \left(S_T^\nu - \frac{1}{2}\Delta - \alpha Q_T^\nu \right) - \phi \int_t^T (Q_u^\nu)^2 du \right], \quad (6)$$

where $\boldsymbol{\mu} = \{\mu^+, \mu^-\}$, and the expectation operator $\mathbb{E}_{t,x,S,\boldsymbol{\mu},q}[\cdot]$ represents expectation conditioned on (with a slight abuse of notation) $X_t = x$, $S_{t-} = S$, $\mu_{t-}^+ = \mu^+$, $\mu_{t-}^- = \mu^-$ and $Q_t = q$,

and her value function is

$$H(t, x, S, \boldsymbol{\mu}, q) = \sup_{\nu \in \mathcal{A}} H^\nu(t, x, S, \boldsymbol{\mu}, q), \quad (7)$$

where \mathcal{A} is the set of admissible strategies consisting of \mathcal{F} -predictable processes such that $\int_0^T |\nu_u| du < +\infty$, \mathbb{P} -a.s..

The right-hand side of the performance criteria (6) contains three terms. The first is the investor's terminal cash from liquidating the shares throughout the trading horizon. The second is the proceeds that the investor receives when liquidating any remaining inventory Q_T^ν at the end of the strategy. This left over inventory is liquidated at the midprice S_T^ν and picks up the costs associated to crossing the spread, liquidity taking fees, and market impact, which is captured by the liquidation penalty parameter $\alpha \geq 0$.

Finally, the third term is the running penalty $\phi \int_t^T (Q_u^\nu)^2 du$ where $\phi \geq 0$ is the inventory penalty parameter. This penalty does not affect the investor's revenues but affects the optimal liquidation rate. When the value of the inventory penalty parameter ϕ is high, the liquidation speed will be high at the beginning of the trading horizon, so this parameter could be interpreted as an urgency parameter. Including this running inventory penalty is also justified in a setting where the agent considers model uncertainty – i.e. she is ambiguity averse. Cartea et al. (2013) show that including the running penalty is equivalent to the agent considering alternate models with stochastic drifts, but penalizes those models using relative entropy. In that setting, the higher the value of ϕ , the less confident is the agent about the trend of the midprice.

3.1. The Resulting Dynamic Programming Equation

We simplify the setup of the problem by assuming that the price impact functions described above are linear in the rate of trading. Thus in the rest of the paper we assume that $g(x) = bx$ $f(x) = kx$ where b, k are non-negative constants. For notational simplicity, we define the net order-flow $\mu_t = \mu_t^+ - \mu_t^-$, which excludes the investor's own trading rate, and write the midprice process as

$$dS_t^\nu = b(\mu_t - \nu_t) dt + \sigma dW_t. \quad (8)$$

The dynamic programming principle for the value function suggests that $H(t, x, S, \boldsymbol{\mu}, q)$ satisfies the dynamic programming equation (DPE):

$$0 = \left(\partial_t + \frac{1}{2} \sigma^2 \partial_{SS} \right) H + \mathcal{L}^\mu H - \phi q^2 + \sup_\nu \left\{ \left(\nu \left(S - \frac{1}{2} \Delta - k \nu \right) \partial_x + b(\mu - \nu) \partial_S - \nu \partial_q \right) H \right\}, \quad (9)$$

for $q > 0$, subject to the terminal and boundary conditions

$$H(T, x, S, \boldsymbol{\mu}, q) = x + q \left(S - \frac{1}{2} \Delta \right) - \alpha q^2,$$

and \mathcal{L}^μ is the generator of the process $\boldsymbol{\mu}$.

Proposition 1. Solving the DPE. Assuming the order-flow processes, $\boldsymbol{\mu} = \{\mu^+, \mu^-\}$, are Markov, cadlag, and bounded \mathbb{P} -a.s., the DPE (9) admits the solution

$$H(t, x, S, \boldsymbol{\mu}, q) = x + q \left(S - \frac{1}{2} \Delta \right) + h_0(t, \boldsymbol{\mu}) + q h_1(t, \boldsymbol{\mu}) + q^2 h_2(t), \quad (10)$$

where

$$h_2(t) = -\sqrt{k\phi} \frac{\zeta e^{\gamma(T-t)} - e^{-\gamma(T-t)}}{\zeta e^{\gamma(T-t)} - e^{-\gamma(T-t)}} - \frac{1}{2} b, \quad (11a)$$

$$h_1(t, \boldsymbol{\mu}) = b \int_t^T \left(\frac{\zeta e^{\gamma(T-u)} - e^{-\gamma(T-u)}}{\zeta e^{\gamma(T-t)} - e^{-\gamma(T-t)}} \right) \mathbb{E}_{t, \boldsymbol{\mu}} [\mu_u] du, \quad (11b)$$

$$h_0(t, \boldsymbol{\mu}) = \frac{1}{4k} \int_t^T \mathbb{E}_{t, \boldsymbol{\mu}} [h_1^2(t, \boldsymbol{\mu}_u)] du, \quad (11c)$$

with the constants

$$\gamma = \sqrt{\frac{\phi}{k}}, \quad \text{and} \quad \zeta = \frac{\alpha - \frac{1}{2}b + \sqrt{k\phi}}{\alpha - \frac{1}{2}b - \sqrt{k\phi}}. \quad (11d)$$

PROOF. See Appendix A.

The above proposition provides us with a candidate solution to our original control problem (7) and we interpret it as follows. The term $x + q \left(S - \frac{1}{2} \Delta \right)$ is the book-value (marked-to-market at the midprice minus the half-spread) of the investor's remaining inventory, plus the cash from the sales so far. The function $h_0(t, \boldsymbol{\mu}) + q h_1(t, \boldsymbol{\mu}) + q^2 h_2(t)$ represents the excess value of trading optimally from this point onwards and is independent of the midprice. Indeed, this additional term can be interpreted as the investor's marginal reservation price, i.e. the price at which she is willing to sell the outstanding inventory so that her value function remains unchanged.

The following theorem shows that this candidate solution is indeed the one we seek and also provides us with the explicit form of the optimal control.

Theorem 2. Verification. The candidate value function provided in Proposition 1 is indeed the solution to the optimal control problem (7). Moreover, the trading speed is given by

$$\nu_t^* = \gamma \frac{\zeta e^{\gamma(T-t)} + e^{-\gamma(T-t)}}{\zeta e^{\gamma(T-t)} - e^{-\gamma(T-t)}} Q_t^{\nu^*} - \frac{b}{2k} \int_t^T \left(\frac{\zeta e^{\gamma(T-u)} - e^{-\gamma(T-u)}}{\zeta e^{\gamma(T-t)} - e^{-\gamma(T-t)}} \right) \mathbb{E} [\mu_u | \mathcal{F}_t^\mu] du, \quad (12)$$

is the admissible optimal control we seek, and \mathcal{F}_t^μ denotes the natural filtration generated by $\boldsymbol{\mu}$.

PROOF. Here we provide a sketch of the proof. The right hand side of (10) is a classical solution and standard results imply that it suffices to check that the feedback control is indeed an admissible strategy. From the feedback form of the optimal control in (A.1) we have $\nu^* = -\frac{1}{2k} (b q + h_1 + 2 q h_2)$, and the explicit form of $h_2(t)$ and $h_1(t, \boldsymbol{\mu})$, thus we find the expression above for ν^* . Since μ^\pm are assumed to be bounded from above a.s., this implies that ν^* is bounded from above and below a.s. and therefore admissible.

□

The first term of the optimal speed of trading (12) corresponds to the classical Almgren-Chriss solution (at the end of this section we present the Almgren-Chriss strategy), while the second term adjusts the speed based on the weighted average of the future expected net order-flow over the remainder of the trading window $[t, T]$. The strategy gives more weight to expected net order-flow near the current time t and the contribution of future expected net order-flow decreases as it approaches the end of the trading horizon. If this weighted average future order-flow is positive (which would occur, e.g., if the current order-flow is positive and hence biased towards buying), then the investor slows down to take advantage of the upward trend in prices that the excess positive net order-flow will have. The opposite holds if order-flow is negative. This dependence on order-flow becomes less important as maturity approaches ($t \rightarrow T$), and the investor instead focuses on completing her execution using the Almgren-Chriss strategy.

An interesting limiting case of the optimal strategy is one where the investor is adamant that all shares are liquidated at the end of the trading horizon. This is obtained if the liquidation penalty becomes infinite, $\alpha \rightarrow \infty$, and the optimal trading speed simplifies to

$$\lim_{\alpha \rightarrow \infty} \nu_t^* = \gamma \frac{\cosh(\gamma(T-t))}{\sinh(\gamma(T-t))} Q_t^{\nu^*} - \frac{b}{2k} \int_t^T \frac{\sinh(\gamma(T-u))}{\sinh(\gamma(T-t))} \mathbb{E}[\mu_u | \mathcal{F}_t^\mu] du. \quad (13)$$

Another interesting limiting case is if the above limit is followed by letting $\phi \rightarrow 0$:

$$\lim_{\phi \rightarrow 0} \lim_{\alpha \rightarrow \infty} \nu_t^* = \frac{1}{(T-t)} Q_t^{\nu^*} - \frac{b}{2k} \int_t^T \frac{(T-u)}{(T-t)} \mathbb{E}[\mu_u | \mathcal{F}_t^\mu] du. \quad (14)$$

In this double-limit, where the agent liquidates all shares by the end of the trading horizon and does not penalize running inventory, the strategy consists of time-weighted-average-price (TWAP) adjusted by the weighted average of net order-flow over the remaining life of the strategy.

The two limiting cases provide further insight into how the strategy depends on the pair b, k . In strategy (14) it is clear that the higher is the ratio b/k , the more sensitive is the strategy to net order-flow over the remaining trading horizon. Moreover, strategy (13) can be interpreted similarly: if we first fix the ratio ϕ/k (recall that $\gamma = \sqrt{\phi/k}$), then the sensitivity of the strategy with respect to expected net order-flow depends on the ratio b/k .

Our model nests the Almgren-Chriss model in continuous-time. Throughout this paper we label the Almgren-Chriss strategy as the optimal strategy when order-flow does not affect the midprice dynamics, but the agent's own trades do have a permanent effect on the price as described by

$$dS_t^\nu = -b \nu_t dt + \sigma dW_t, \quad (15)$$

and the execution price receives a linear impact, so that $\hat{S}_t^\nu = S_t^\nu - k\nu_t$. In this case the optimal liquidation speed is

$$\nu_t^{AC} = \gamma \frac{\zeta e^{\gamma(T-t)} + e^{-\gamma(T-t)}}{\zeta e^{\gamma(T-t)} - e^{-\gamma(T-t)}} Q_t^{\nu*}, \quad (16)$$

with the constants γ and ζ as in (11d), see Almgren and Chriss (2001) and Cartea et al. (2015).

There are other studies that incorporate order-flow in algorithmic trading. In the context of optimal execution the recent work of Bechler and Ludkovski (2014) and Cartea et al. (2015) employ order-flow information to develop optimal execution models that take into account market impact and informational costs, see also Gatheral and Schied (2011) and Cartea and Jaimungal (2013). And in market making see the early work of Avellaneda and Stoikov (2008), and more recently, Guéant et al. (2012), Cartea and Jaimungal (2015b), Guilbaud and Pham (2013), Cartea et al. (2014), Cartea and Jaimungal (2015a), and Fodra and Labadie (2012).

4. Simulations of the Strategy with Order-Flow

In this section we perform simulations to show the behavior of the optimal liquidation strategy provided in (13) which guarantees full execution by the end of the trading horizon. We assume that the buy and sell rate of order-flow μ_t^\pm satisfy the SDEs

$$d\mu_t^\pm = -\kappa \mu_t^\pm dt + \eta_{1+L_t^\pm}^\pm dL_t^\pm, \quad (17)$$

where L_t^\pm are independent Poisson processes (assumed independent of all other processes as well) with equal intensity λ , $\{\eta_1^\pm, \eta_2^\pm, \dots\}$ are non-negative i.i.d. random variables, with finite first moment, independent from all processes. This assumption implies that the buy and sell order-flow arrive independently at Poisson times with rate λ , and induce an increase in the order-flow rate by η^\pm , and the effect of the jumps in order-flow revert to zero at the exponential speed κ .

The solutions to (17), for $s > t$, are

$$\mu_s^\pm = e^{-\kappa^\pm(s-t)} \mu_t^\pm + \int_t^s e^{-\kappa^\pm(s-u)} \eta_{1+L_{u-}^\pm}^\pm dL_u^\pm,$$

so that

$$\mathbb{E}[\mu_s^\pm | \mathcal{F}_t^\mu] = e^{-\kappa^\pm(s-t)} (\mu_t^\pm - \psi^\pm) + \psi^\pm,$$

where the constants

$$\psi^\pm = \frac{1}{\kappa^\pm} \lambda^\pm \mathbb{E}[\eta^\pm].$$

Inserting this model assumption and computing the expectation and integrals appearing in the full liquidation strategy from (13), results in the optimal trading strategy

$$\lim_{\alpha \rightarrow \infty} \nu_t^* = \gamma \frac{\cosh(\gamma(T-t))}{\sinh(\gamma(T-t))} Q_t^{\nu^*} - \frac{b}{2k} \left[\ell_1^+(t) (\mu_t^+ - \psi^+) - \ell_1^-(t) (\mu_t^- - \psi^-) + \ell_0(t) (\psi^+ - \psi^-) \right],$$

where

$$\ell_0(t) = \frac{1}{\gamma} \frac{\cosh(\gamma(T-t)) - 1}{\sinh(\gamma(T-t))},$$

and

$$\ell_1^\pm(t) = \frac{1}{2} \left(\frac{e^{\gamma(T-t)} - e^{-\kappa^\pm(T-t)}}{\kappa^\pm + \gamma} - \frac{e^{-\gamma(T-t)} - e^{-\kappa^\pm(T-t)}}{\kappa^\pm - \gamma} \right) \Bigg/ \sinh(\gamma(T-t)).$$

In the simulations we use the parameters for the seventeen stocks in Tables A.7 and A.8. We assume that the trading horizon is $T = 1$ hour, the running inventory penalty parameter is $\phi = 10 \times k$, the liquidation target \mathfrak{N} shares is set to 1% of the expected traded volume over the trading window (including the investor's trades), and run 10,000 simulations.

For each stock we present the performance of the strategy for different assumptions about (i) how often the agent executes trades, and (ii) the rate of decay of shocks to order-flow survive which is dictated by the parameter κ in (17). In all, we investigate four trading environments:

- I.** The agent executes orders at the same frequency as the market, so that she trades at a rate of λ^- . Moreover, order-flow decays at rate of total MO activity $\kappa = \lambda^+ + \lambda^-$.
- II.** The agent once again executes orders at the same frequency as the market. In this environment, however, order-flow decays on the time scale of 1 minute, so that $\kappa = 60$.
- III.** The agent executes trades only once per minute, while order-flow decays at rate of total MO activity $\kappa = \lambda^+ + \lambda^-$.
- IV.** The agent executes trades only once per minute, and order-flow decays on the time scale of 1 minute, so that $\kappa = 60$.

Finally, we point out that in this model we assume that the LOB is infinitely resilient, so as soon as an MO is filled, the book replenishes back to its original form. This is a simplistic assumption that biases the results (in favour of any execution strategy, not only the one developed here), so one way of reducing this bias in the simulations is to pace the investor's trading activity in tandem with the rest of the market (as in trading environments I and II) to allow the LOB to replenish – which is specific to each stock.




Figure 4: Optimal trading in the presence of order-flow for INTC. The dashed lines show the Almgren-Chriss solution and the solid lines in the heat-maps are the 5% and 95% quantiles of the inventory/trading speed at each point in time.

4.1. Simulation of liquidation strategy

Before showing the general results for the 17 stocks we analyze, we show in more detail the liquidation strategy for INTC. In this case the parameters are

$$k = 4.42 \times 10^{-7}, \quad b = 1.09 \times 10^{-6}, \quad \mathfrak{N} = 4,042,$$

$$\lambda^{\pm} = 567, \quad \kappa = 567, \quad \eta \sim \text{Exp}(353 \times \kappa),$$

$$\sigma = 0.23, \quad S_0 = 53.41,$$

and $\eta \sim \text{Exp}(\eta_0)$ denotes the exponential distribution with mean size $\mathbb{E}[\eta] = \eta_0$. The simulation results reported are based on the simulated model using the calibrated parameters for the corresponding stock, and real performance may differ since there are other market effects not captured by our model.

The first panel in Figure 4 shows one path of the net order-flow, and the second shows the trading speed where the dashed line shows the Almgren-Chriss liquidation speed given by (16). Clearly, when the strategy incorporates order-flow, the liquidation speed becomes more volatile to take advantage of changes in the drift induced by the price impact of net order-flow. In extreme cases, when the rate of market buy orders is much larger than that of market sell orders, the optimal strategy is to purchase shares to take advantage of the expected increase in the midprice. For this set of parameters, the strategy spent about 2% of the time (across all simulations and time points) with $\nu_t^* < 0$.

The third and fourth panels in Figure 4 show heat-maps of the inventory path and the trading speed. In both pictures the dashed lines show the Almgren-Chriss solution and the solid lines the heat-maps are the 5% and 95% quantile paths. Observe that although on average the strategy is close to that of Almgren-Chriss, it deviates from it throughout the trading horizon. Moreover, although the strategy converges to Almgren-Chriss as we approach the end of the trading horizon, as seen from (12), these deviations persist very close to T because the investor's inventory holdings are different from those of an agent who has executed Almgren-Chriss throughout the entire trading horizon.




Figure 5: The savings per share (INTG) measured in bps, $(X_T^{\nu^*} - X_T^{AC})/X_T^{AC} \times 10^4$, from following the optimal strategy relative to Almgren-Chriss. Dashed red lines indicate the 5%, 25%, 50%, 75% and 95% quantiles moving left to right.

Figure 5 shows a histogram of the financial performance of the strategy relative to Almgren-Chriss. This performance is measured in basis points (bps) using

$$\frac{X_T^{\nu^*} - X_T^{AC}}{X_T^{AC}} \times 10^4, \quad (18)$$

where X_T^{AC} is the cash obtained from running the Almgren-Chriss strategy with the same level of ϕ . This is a measure of relative performance, and as such, when we change the volatility σ , this relative performance remains economically indistinguishable for the set of stocks we look at.

4.2. Performance of strategy: simulations

In Tables 3 and 4 we summarize the performance of the strategy for all stocks. In the first table we employ the strategy ν^* as in (13) and in the second table we employ the truncated strategy $\max(\nu^*, 0)$ so there are no repurchases. We show the mean (in bps) of the savings per share using (18). The simulation parameters are taken from Tables A.7 and A.8. The tables show the results for the trading environments I, II, III, and IV discussed above. The tickers have been organised in descending order by activity rate (last column).

In both tables, columns 2 to 5 show the results for trading environments I to IV. Table 3 shows that for strategy ν^* , the stock which exhibits the largest savings per share is ARCC with over 110 bps, and the one with the lowest performance is EBAY with savings of around 0.1 bps per share. Moreover, note that ARCC is, for the four types of trading environments we run, always the top performer and EBAY the bottom performer, however, the ordering of all other assets is not always preserved.

Similarly, Table 4 shows the performance of the truncated strategy $\max(\nu^*, 0)$ where the top performer is also ARCC and the bottom performer is EBAY. It is also interesting to see that for every stock the truncated strategy delivers higher savings than the strategy ν^* which allows for speculative repurchases in the liquidation programme. Recall that the optimal

Strategy ν^*							
Ticker	I. $\kappa = \hat{\lambda}^+ + \hat{\lambda}^-$	II. $\kappa = 60$	III. $\kappa = \hat{\lambda}^+ + \hat{\lambda}^-$	IV. $\kappa = 60$	\hat{b}		
	$N = \hat{\lambda}^-$	$N = \hat{\lambda}^-$	$N = 60$	$N = 60$	$(\times 10^{-6})$	$\frac{\hat{b}}{k}$	$\hat{\lambda}^-$
EBAY	0.1	0.4	1.4	0.2	4.5	1.8	585
INTC	0.4	1.4	5.0	1.9	1.1	2.5	454
DOW	0.5	1.5	5.4	2.2	6.9	2.3	401
CSCO	1.4	3.3	10.9	6.3	0.6	2.6	308
COST	1.3	2.9	9.9	5.7	25.3	2.5	294
AXP	1.0	2.1	6.8	3.8	17.2	2.0	273
AMAT	3.2	5.7	18.6	12.2	1.6	2.7	248
PNC	2.6	4.4	13.4	8.8	19.9	2.1	209
AA	11.6	14.4	36.0	28.0	1.5	2.5	153
FMER	36.3	37.9	67.8	60.1	8.7	2.8	106
CNQR	9.1	9.4	18.6	14.7	88.6	1.1	100
ROC	3.9	4.0	8.1	6.0	43.4	0.8	93
IMGN	25.6	24.6	40.6	34.2	15.9	1.8	87
DK	10.2	9.9	16.3	13.4	45.5	1.3	86
ARCC	110.7	106.2	99.3	96.2	2.8	2.2	55
BXS	106.9	101.2	76.4	75.8	27.2	1.7	48
AYI	7.5	6.8	4.0	3.7	202.7	0.5	41

Table 3: Relative performance of the strategy (in basis points). Statistics are calculated per share using (18) and stock specific parameters are as in Tables A.7 and A.8.

strategy derived above does not require the liquidation speed to be non-negative, so this modification is ad-hoc and not the optimal solution to a constrained problem.

These two tables are expanded in the Appendix (see Tables A.9, A.10, A.11, A.12), where in addition to showing the mean performance (in bps) of the savings per share, we also show the standard deviation and quantiles. Moreover, the tables also show the percentage of the time (across all simulations and time points) with $\nu_t^* < 0$, as well as the percentage of the runs where the execution strategy underperforms Almgren-Chriss.

Finally, we also examine the performance of the strategy for different liquidation targets. In Table 5 we show the results for trading environment I when \mathfrak{N} is 0.5%, 1.0%, and 2.0% of expected traded volume, and the investor employs the (no-repurchase) liquidation strategy $\max(\nu^*, 0)$. We observe that as the liquidation target increases, the performance, relative to Almgren-Chriss, decreases. Note that the optimal strategy (12), or the particular one used in the simulations (14), consists of two terms where only the first one depends on inventory. Thus, we see that the higher is the liquidation target, the ‘closer’ is ν^* to the liquidation strategy of Almgren-Chriss. The intuition behind this result is that, for a fixed trading horizon, the higher is the liquidation target, the less able is the strategy to take advantage of shocks to order-flow – the strategy has fewer opportunities to slow down or speed up when order-flow induces a trend in the midprice because more shares must be liquidated over the

Strategy $\max(\nu^*, 0)$							
Ticker	I. $\kappa = \hat{\lambda}^+ + \hat{\lambda}^-$	II. $\kappa = 60$	III. $\kappa = \hat{\lambda}^+ + \hat{\lambda}^-$	IV. $\kappa = 60$	\hat{b}		
	$N = \hat{\lambda}^-$	$N = \hat{\lambda}^-$	$N = 60$	$N = 60$	$(\times 10^{-6})$	$\frac{\hat{b}}{k}$	$\hat{\lambda}^-$
EBAY	0.1	0.4	2.0	0.2	4.5	1.8	585
INTC	0.5	1.6	7.1	2.2	1.1	2.5	454
DOW	0.7	1.8	7.8	2.5	6.9	2.3	401
CSCO	2.4	4.2	17.0	8.0	0.6	2.6	308
COST	2.2	3.9	15.4	7.2	25.3	2.5	294
AXP	1.5	2.6	10.6	4.7	17.2	2.0	273
AMAT	6.0	7.6	27.8	15.9	1.6	2.7	248
PNC	4.7	5.7	20.6	11.3	19.9	2.1	209
AA	20.7	20.1	54.9	39.0	1.5	2.5	153
FMER	58.0	52.5	101.6	85.8	8.7	2.8	106
CNQR	14.4	12.3	28.1	19.2	88.6	1.1	100
ROC	5.9	4.8	11.7	7.5	43.4	0.8	93
IMGN	39.0	32.4	58.7	45.4	15.9	1.8	87
DK	15.5	12.4	24.4	17.3	45.5	1.3	86
ARCC	155.2	142.9	141.4	130.3	2.8	2.2	55
BXS	149.3	136.6	112.5	103.8	27.2	1.7	48
AYI	10.1	8.6	5.0	4.6	202.7	0.5	41.0

Table 4: Relative performance of truncated strategy (in basis points). Statistics are calculated per share using (18) and stock specific parameters are as in Tables A.7 and A.8.

trading window.

4.3. Stock characteristics and strategy performance

Are there stock characteristics which are correlated with the performance of the strategy? The results discussed above, and the more detailed tables in the appendix, do not show a clear-cut link between execution strategies that incorporates order-flow information and the classical Almgren-Chriss. Here we run an Ordinary Least Squares regression to understand how performance is related to the ratio b/k and activity rate λ^- , that is we perform the regression

$$bp_i = c_0 \frac{\hat{b}}{\hat{k}_i} + c_1 \hat{\lambda}_i^- + \varepsilon_i, \quad (19)$$

where i is the ticker, bp_i is the strategy's performance relative to Almgren-Chriss as discussed above, and ε the error term. We run this regression for the results using ν^* and the truncated strategy $\max(\nu^*, 0)$ for the case $\kappa = \lambda^+ + \lambda^-$ and $N = \hat{\lambda}^-$. In other words we run a regression of column 2 in Tables 3 and 4 on the explanatory variables in columns 7 and 8. Table 6 below shows the coefficient results and the confidence intervals. We see that performance is positively related to the ratio b/k and negatively related to MO activity. We also ran other specifications which contained an intercept but we could not reject the null hypothesis that the intercept is zero.

	Liquidation Target		
	(% expected volume)		
	0.5%	1.0%	2.0%
EBAY	0.7	0.1	0.0
INTC	4.0	0.5	0.1
DOW	4.9	0.7	0.1
CSCO	11.9	2.4	0.3
COST	12.0	2.2	0.2
AXP	9.6	1.5	0.2
AMAT	21.2	6.0	0.8
PNC	20.5	4.7	0.6
AA	52.4	20.7	4.2
FMER	125.6	58.0	18.2
CNQR	52.5	14.4	2.2
ROC	29.8	5.9	0.7
IMGN	100.5	39.0	9.6
DK	56.8	15.5	2.7
ARCC	333.1	155.2	57.6
BXS	324.9	149.3	50.2
AYI	47.2	10.1	1.3

Table 5: Relative performance of the strategy (in basis points) for different liquidation targets as percentage of expected trading volume over trading horizon. Liquidation strategy $\max(\nu^*, 0)$ and trading environment I, i.e. $\kappa = (\lambda^+ + \lambda^-)$, $N = \lambda^-$.

That the strategy's performance is positively related to the ratio b/k is expected because, as already discussed after (14), the higher is the ratio b/k the more sensitive is the strategy to net order-flow. This result is consistent across the four trading environments we examine in the simulations.

	c_0	c_1
ν^*	23.34 [9.17, 37.50]	-0.13 [-0.25, -0.02]
$\max(\nu^*, 0)$	34.09 [14.59, 53.58]	-0.2 [-0.35, -0.04]

Table 6: Regression of strategy performance on the ratio of impacts $\hat{\frac{b}{k}}$ and MO activity $\hat{\lambda}^-$.

5. Conclusions

We show how to incorporate the permanent impact of market order-flow in an execution strategy. The strategy is found in closed-form when the permanent and temporary price impacts are linear in the market's and investor's rates of trading. We do this under very general assumptions about the stochastic process followed by the order-flow of the market.

The optimal strategy consists of an Almgren-Chriss execution strategy adjusted by a weighted-average of the expected net order-flow (given by the difference of the markets rate of buy and sell market orders) over the execution trading horizon and proportional to the ratio of permanent to temporary linear impacts. At the beginning of the trading window the strategy places more emphasis on expected order-flow over the remaining trading horizon and, as the end of the trading horizon approaches, the strategy is closer to Almgren-Chriss.

We use Nasdaq data to calibrate the model to five stocks and use simulations to show how the strategy performs. We benchmark the strategy against execution prices obtained from employing the Almgren-Chriss strategy. The mean revenue per share is always higher than that of Almgren-Chriss and, depending on the ticker in approximately 50% to 95% of the runs, the strategy outperforms Almgren-Chriss. Finally, we show that the performance of the strategy is positively related to the ratio of permanent to temporary impact (b/k), and inversely related to the MO activity of the stock.

Appendix A. Proof of Proposition 1

To solve (9) we make the ansatz

$$H(t, x, S, \boldsymbol{\mu}, q) = x + q \left(S - \frac{1}{2} \Delta \right) + h(t, \boldsymbol{\mu}, q),$$

and upon substitution of the ansatz in the DPE above we see that $h(t, \boldsymbol{\mu}, q)$ satisfies

$$\partial_t h + \mathcal{L}^\mu h + b \mu q - \phi q^2 + \sup_{\nu} \{ -k \nu^2 - (b q + \partial_q h) \nu \} = 0,$$

subject to the terminal condition $h(T, \boldsymbol{\mu}, q) = -\alpha q^2$, and the optimal liquidation speed in feedback form is

$$\nu^* = -\frac{1}{2k} (b q + \partial_q h). \quad (\text{A.1})$$

Upon substitution back into the DPE we find that h satisfies the non-linear partial-integral differential equation (PIDE)

$$(\partial_t + \mathcal{L}^\mu) h + b \mu q - \phi q^2 + \frac{1}{4k} (b q + \partial_q h)^2 = 0. \quad (\text{A.2})$$

Due to the existence of linear and quadratic terms in q in (A.2), and its terminal conditions, we expect $h(t, \boldsymbol{\mu}, q)$ to be a quadratic form in q , and we assume the ansatz

$$h(t, \boldsymbol{\mu}, q) = h_0(t, \boldsymbol{\mu}) + q h_1(t, \boldsymbol{\mu}) + q^2 h_2(t, \boldsymbol{\mu}).$$

Inserting this into (A.2) and collecting like terms in q leads to the following coupled system of PIDEs

$$(\partial_t + \mathcal{L}^\mu) h_0 + \frac{1}{4k} h_1^2 = 0, \quad (\text{A.3a})$$

$$(\partial_t + \mathcal{L}^\mu) h_1 + b\mu + \frac{1}{2k} h_1 (b + 2h_2) = 0, \quad (\text{A.3b})$$

$$(\partial_t + \mathcal{L}^\mu) h_2 - \phi + \frac{1}{4k} (b + 2h_2)^2 = 0, \quad (\text{A.3c})$$

subject to the terminal conditions

$$h_0(T, \boldsymbol{\mu}) = 0, \quad h_1(T, \boldsymbol{\mu}) = 0, \quad h_2(T, \boldsymbol{\mu}) = -\alpha.$$

To solve for h_2 we note that since Equation (A.3c) for h_2 contains no source terms in μ and its terminal condition is independent of μ , the solution must be independent of μ , i.e. h_2 is a function only of time. In this case, (A.3c) is an ODE of Riccati type and can be solved explicitly:

$$h_2(t, \boldsymbol{\mu}) = \chi(t) - \frac{1}{2}b, \quad \text{where} \quad \chi(t) = \sqrt{k\phi} \frac{1 + \zeta e^{2\gamma(T-t)}}{1 - \zeta e^{2\gamma(T-t)}},$$

with the constants γ and ζ :

$$\gamma = \sqrt{\frac{\phi}{k}}, \quad \text{and} \quad \zeta = \frac{\alpha - \frac{1}{2}b + \sqrt{k\phi}}{\alpha - \frac{1}{2}b - \sqrt{k\phi}}.$$

Now we turn to solving (A.3b) which is a linear PIDE for h_1 where $h_2 + \frac{1}{2}b$ acts as an effective discount rate and $b\mu$ is a source term. The general solution of such an equation can be represented using the Feynman-Kac theorem. Thus we write

$$h_1(t, \boldsymbol{\mu}) = b \mathbb{E}_{t, \boldsymbol{\mu}} \left[\int_t^T \exp \left\{ \frac{1}{k} \int_t^u (h_2(s) + \frac{1}{2}b) ds \right\} \mu_u du \right]$$

which can be simplified to

$$h_1(t, \boldsymbol{\mu}) = b \int_t^T \left(\frac{\zeta e^{\gamma(T-u)} - e^{-\gamma(T-u)}}{\zeta e^{\gamma(T-t)} - e^{-\gamma(T-t)}} \right) \mathbb{E}_{t, \boldsymbol{\mu}} [\mu_u] du. \quad (\text{A.4})$$

Finally, we can solve for $h_0(t, \boldsymbol{\mu})$ by again noticing it is a linear PDE with non-linear source term and a straight forward application of Feynman-Kac, and interchanging integration and expectation, we obtain (11c).

□

The remainder of the appendix contains tables with estimated parameters and simulation results.

	EBAY	INTC	DOW	CSCO	COST	AXP	AMAT	PNC
\hat{b}	0.45×10^{-5} (1.95×10^{-6})	1.09×10^{-6} (0.54×10^{-6})	0.69×10^{-5} (2.69×10^{-6})	0.63×10^{-6} (2.57×10^{-7})	2.53×10^{-5} (1.05×10^{-5})	1.72×10^{-5} (0.73×10^{-5})	1.6×10^{-6} (0.64×10^{-6})	1.99×10^{-5} (0.88×10^{-5})
\hat{k}	2.5×10^{-6} (1.03×10^{-6})	0.44×10^{-6} (2.37×10^{-7})	3.01×10^{-6} (1.22×10^{-6})	2.44×10^{-7} (0.92×10^{-7})	1.06×10^{-5} (0.53×10^{-5})	0.89×10^{-5} (0.43×10^{-5})	0.6×10^{-6} (2.89×10^{-7})	1×10^{-5} (0.49×10^{-5})
$\frac{\hat{b}}{\hat{k}}$	1.82 (0.65)	2.5 (0.76)	2.32 (0.87)	2.59 (0.81)	2.5 (1.11)	2.01 (0.88)	2.74 (0.84)	2.11 (1.05)
ADV	2,506,048 (1,275,581)	5,489,503 (2,542,732)	1,453,332 (809,608)	5,598,870 (2,429,781)	540,008 (190,770)	574,606 (250,467)	2,621,837 (1,097,946)	382,020 (139,898)
Midprice	53.41 (2.4)	30.26 (4.37)	49.48 (3.04)	24.21 (1.69)	121.37 (9.32)	89.78 (3.01)	20.85 (2.06)	84.67 (3.22)
σ^{oc}	17.26%	16.34%	23.43%	14.18%	12.17%	18.74%	23.19%	15.25%
σ^{5m}	19.52%	17.87%	24.31%	15.74%	18.23%	19.81%	24.32%	17.95%
λ^+	549.41 (305.23)	438.86 (248.85)	389.76 (223)	306.7 (115.94)	294.98 (99.8)	281.1 (117.12)	248.86 (115.58)	206.79 (69.29)
$\mathbb{E}[\eta^+]$	358.36 (78.62)	1048.99 (351.13)	282.28 (61.59)	1500.35 (440.2)	139.89 (16.61)	159.27 (21.05)	861.38 (248.14)	136.35 (16.16)
λ^-	585.13 (330.21)	453.91 (264.63)	400.67 (226.67)	307.99 (131.62)	293.71 (95.71)	273.04 (114.59)	247.66 (107.72)	209.22 (71.7)
$\mathbb{E}[\eta^-]$	347.18 (75.41)	1013.83 (306.58)	286.41 (62.53)	1376.49 (302.65)	141.38 (17.88)	158.29 (20.17)	833.01 (212.85)	143.73 (16.95)

Table A.7: Permanent and temporary price impact parameters for Nasdaq stocks, average volume of MOs, average midprice, σ volatility of returns (open-to-close and 5-minute windows), mean arrival per hour of MOs λ^\pm , and average volume of MOs $\mathbb{E}[\eta^\pm]$. In parenthesis we report standard deviation of estimated parameter. Data are from Nasdaq 2014.

	AA	FMER	CNQR	ROC	IMGN	DK	ARCC	BXS	AYI
\hat{b}	1.49×10^{-6} (0.61×10^{-6})	0.87×10^{-5} (0.34×10^{-5})	0.89×10^{-4} (0.58×10^{-4})	0.43×10^{-4} (3.15×10^{-5})	1.59×10^{-5} (0.83×10^{-5})	0.45×10^{-4} (2.21×10^{-5})	2.83×10^{-6} (1.01×10^{-6})	2.72×10^{-5} (1.57×10^{-5})	2.03×10^{-4} (1.03×10^{-4})
\hat{k}	0.59×10^{-6} (2.76×10^{-7})	3.14×10^{-6} (1.31×10^{-6})	0.76×10^{-4} (0.61×10^{-4})	0.56×10^{-4} (0.46×10^{-4})	1×10^{-5} (0.68×10^{-5})	0.48×10^{-4} (0.52×10^{-4})	1.32×10^{-6} (0.59×10^{-6})	1.72×10^{-5} (1.35×10^{-5})	0.66×10^{-3} (0.79×10^{-3})
$\hat{\frac{b}{k}}$	2.52 (0.83)	2.82 (1.07)	1.1 (0.56)	0.83 (0.6)	1.75 (0.76)	1.33 (0.95)	2.2 (0.8)	1.73 (0.92)	0.51 (0.41)
ADV	1,889,223 (1,357,437)	339,908 (150,626)	164,904 (101,910)	147,946 (101,952)	256,366 (146,433)	128,864 (63,486)	385,168 (208,690)	92,525 (54,475)	57,324 (31,295)
Midprice	14.44 (2.05)	19.17 (1.5)	105.96 (16.42)	76.1 (3.71)	12.19 (2.53)	30.83 (2.15)	16.96 (0.74)	22.88 (1.55)	128.21 (9.15)
σ^{oc}	27%	21.44%	29.07%	19.2%	50.41%	77.51%	11.77%	26.46%	31.42%
σ^{5m}	27.88%	68.61%	38.23%	20.89%	71.51%	77.99%	19.95%	36.47%	75.87%
λ^+	154.65 (95.24)	104.04 (40.52)	96.47 (58.39)	96 (59.88)	90.61 (40.09)	80.26 (38.21)	54.83 (30.39)	47.64 (22.19)	41.57 (22.34)
$\mathbb{E}[\eta^+]$	953.22 (327.01)	254.01 (66.27)	128.79 (34.09)	114.76 (15.02)	218.09 (52.44)	118.25 (15.13)	583.79 (195.54)	143.36 (33.42)	106.51 (16.53)
λ^-	152.89 (98.75)	105.58 (44.23)	99.96 (59.38)	93.21 (60.43)	87.21 (34.68)	86.34 (41.71)	55.46 (29.42)	48.19 (24.09)	40.66 (23.69)
$\mathbb{E}[\eta^-]$	990.34 (355.48)	247.77 (51.32)	136.23 (49.1)	117.9 (17.5)	221.55 (52.66)	117.12 (15.3)	545.1 (161.34)	143.23 (34.24)	107.41 (17.81)

Table A.8: Permanent and temporary price impact parameters for Nasdaq stocks, average volume of MOs, average midprice, σ volatility of returns (open-to-close and 5-minute windows), mean arrival per hour of MOs λ^\pm , and average volume of MOs $\mathbb{E}[\eta^\pm]$. In parenthesis we report standard deviation of estimated parameter. Data are from Nasdaq 2014.

I. Order-flow decay $\kappa = (\hat{\lambda}^+ + \hat{\lambda}^-)$, and strategy's trading frequency $N = \hat{\lambda}^-$																		
		EBAY	INTC	DOW	CSCO	COST	AXP	AMAT	PNC	AA	FMER	CNQR	ROC	IMGN	DK	ARCC	BXS	AYI
		Strategy ν^*																
quantile	mean	0.1	0.4	0.5	1.4	1.3	1.0	3.2	2.6	11.6	36.3	9.1	3.9	25.6	10.2	110.7	106.9	7.5
	stdev	1.1	2.4	2.9	5.7	5.8	5.1	9.2	9.6	24.5	60.6	26.4	18.1	48.7	30.0	171.6	170.4	38.3
	5%	-1.7	-3.4	-4.1	-7.3	-7.5	-6.9	-10.8	-11.9	-22.0	-44.5	-29.1	-24.0	-41.5	-33.1	-92.3	-102.3	-50.0
	25%	-0.3	-0.7	-0.8	-1.1	-1.2	-1.2	-1.0	-1.7	0.0	4.8	-2.8	-3.9	2.4	-3.2	15.5	13.3	-9.0
	50%	0.1	0.3	0.3	0.8	0.7	0.6	2.0	1.4	7.4	24.8	5.5	2.2	17.8	6.3	72.7	67.5	3.7
	75%	0.6	1.5	1.7	3.8	3.6	3.1	7.1	6.7	21.4	62.3	19.7	11.8	44.8	22.6	178.3	171.8	23.7
	95%	1.9	4.4	5.4	11.1	11.0	9.7	19.2	19.4	56.4	149.1	57.1	34.9	114.4	62.2	425.7	429.4	70.9
		% $t : \nu_t^* < 0$	0%	2%	2%	4%	4%	3%	6%	5%	9%	13%	6%	4%	10%	6%	14%	3%
		% $X_T < X_T^{AC}$	41%	39%	39%	36%	37%	39%	31%	36%	25%	18%	31%	38%	22%	33%	15%	39%
		Strategy $\max(\nu^*, 0)$																
quantile	mean	0.1	0.5	0.7	2.4	2.2	1.5	6.0	4.7	20.7	58.0	14.4	5.9	39.0	15.5	155.2	149.3	10.1
	stdev	1.1	2.3	2.8	5.7	5.6	4.8	9.4	9.5	24.9	59.5	25.8	17.1	47.4	29.0	167.0	165.7	35.8
	5%	-1.6	-3.0	-3.6	-5.2	-5.5	-5.4	-5.5	-7.7	-5.3	2.1	-17.8	-18.5	-12.5	-21.5	4.5	-3.2	-40.6
	25%	-0.3	-0.5	-0.6	-0.4	-0.6	-0.8	0.6	-0.4	4.6	15.8	0.3	-2.4	9.2	-0.5	40.2	35.8	-7.1
	50%	0.1	0.3	0.4	1.3	1.1	0.7	3.9	2.7	13.9	40.6	8.7	3.0	26.5	9.1	105.5	102.1	4.3
	75%	0.6	1.5	1.8	4.7	4.4	3.5	9.8	8.6	30.3	82.0	24.3	13.2	56.7	26.7	218.3	212.7	24.7
	95%	1.9	4.5	5.6	12.8	12.4	10.3	24.1	22.1	68.9	178.4	63.9	37.3	132.0	71.8	476.3	476.7	71.7
		% $t : \nu_t^* < 0$	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
		% $X_T < X_T^{AC}$	41%	38%	38%	30%	32%	36%	20%	28%	10%	3%	24%	35%	11%	26%	3%	38%

Table A.9: Relative performance of the strategy (in basis points). Statistics are calculated per share using (18) and stock specific parameters are as in Tables A.7 and A.8.

II. Order-flow decay $\kappa = 60$, and strategy's trading frequency $N = \hat{\lambda}^-$																		
		EBAY	INTC	DOW	CSCO	COST	AXP	AMAT	PNC	AA	FMR	CNQR	ROC	IMGN	DK	ARCC	BXS	AYI
		Strategy ν^*																
quantile	mean	0.4	1.4	1.5	3.3	2.9	2.1	5.7	4.4	14.4	37.9	9.4	4.0	24.6	9.9	106.2	101.2	6.8
	stdev	1.0	2.9	3.2	6.2	5.9	4.9	9.8	9.1	23.1	56.2	21.6	13.4	41.5	24.5	158.3	160.2	33.9
	5%	-0.7	-1.1	-1.6	-2.1	-2.2	-2.3	-3.5	-4.2	-9.5	-23.5	-15.5	-12.9	-23.7	-20.6	-72.9	-86.5	-43.5
	25%	-0.2	-0.1	-0.3	-0.2	-0.3	-0.5	-0.1	-0.7	0.5	3.9	-2.4	-3.0	0.7	-3.4	14.5	11.1	-8.6
	50%	0.2	0.6	0.7	1.2	1.0	0.8	2.6	1.7	7.8	21.6	4.6	2.1	14.5	5.6	67.2	60.6	3.6
	75%	0.7	1.8	2.2	4.2	3.8	2.7	8.3	6.0	21.6	59.3	16.0	8.4	39.1	18.5	165.4	159.4	20.6
	95%	2.1	7.3	8.0	16.4	15.4	12.2	25.9	22.9	60.2	147.6	51.5	27.3	105.8	55.3	403.3	406.3	62.1
	$\%t : \nu_t^* < 0$	2%	5%	5%	8%	7%	5%	9%	7%	11%	14%	7%	4%	10%	6%	14%	13%	3%
$\%X_T < X_T^{AC}$	35%	30%	32%	29%	31%	34%	26%	32%	24%	19%	32%	38%	24%	34%	16%	18%	40%	
		Strategy $\max(\nu^*, 0)$																
quantile	mean	0.4	1.6	1.8	4.2	3.9	2.6	7.6	5.7	20.1	52.5	12.3	4.8	32.4	12.4	142.9	136.6	8.6
	stdev	1.0	3.0	3.4	6.8	6.6	5.2	10.6	9.8	25.1	59.5	23.0	13.8	42.8	25.3	161.3	159.3	32.4
	5%	-0.7	-1.0	-1.5	-1.9	-2.0	-2.2	-3.0	-3.8	-6.3	-7.2	-13.3	-11.9	-15.8	-17.3	-0.3	-19.4	-36.7
	25%	-0.2	-0.1	-0.2	0.0	-0.1	-0.3	0.6	-0.2	3.6	13.0	-0.8	-2.5	5.6	-2.0	35.2	30.7	-7.4
	50%	0.2	0.7	0.8	1.8	1.5	1.0	4.2	2.5	12.7	35.8	6.5	2.5	21.1	7.2	95.8	90.4	3.9
	75%	0.7	2.1	2.6	6.2	5.6	3.5	11.7	8.5	29.8	74.1	19.9	9.2	48.8	21.0	201.9	196.7	21.6
	95%	2.1	7.9	8.7	18.4	17.3	13.1	29.2	25.7	70.0	171.4	57.3	29.9	116.8	59.7	460.1	443.8	63.4
	$\%t : \nu_t^* < 0$	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
$\%X_T < X_T^{AC}$	35%	28%	30%	24%	26%	32%	20%	26%	14%	8%	27%	36%	16%	30%	5%	8%	39%	

Table A.10: Relative performance of the strategy (in basis points). Statistics are calculated per share using (18) and stock specific parameters are as in Tables A.7 and A.8.

III. Order-flow decay $\kappa = (\lambda^+ + \lambda^-)$, and strategy's trading frequency $N = 60$																		
		EBAY	INTC	DOW	CSCO	COST	AXP	AMAT	PNC	AA	FMER	CNQR	ROC	IMGN	DK	ARCC	BXS	AYI
		Strategy ν^*																
quantile	mean	1.4	5.0	5.4	10.9	9.9	6.8	18.6	13.4	36.0	67.8	18.6	8.1	40.6	16.3	99.3	76.4	4.0
	stdev	17.3	29.4	30.5	43.3	41.8	35.4	54.2	47.1	76.8	116.8	51.0	33.8	77.8	46.1	154.0	125.4	20.7
	5%	-26.8	-42.6	-44.1	-58.3	-57.1	-50.0	-67.7	-62.2	-84.4	-102.3	-59.6	-45.4	-71.7	-54.9	-91.9	-79.1	-26.0
	25%	-7.0	-9.4	-9.5	-9.9	-10.4	-10.3	-6.8	-9.5	2.7	11.3	-4.6	-8.2	4.9	-5.3	12.9	8.4	-4.7
	50%	1.2	4.4	4.8	9.3	8.3	5.9	14.9	10.6	27.4	50.0	13.4	5.6	30.0	11.8	64.5	47.3	1.8
	75%	9.9	19.3	20.4	31.6	30.1	23.9	44.3	35.9	70.5	117.4	41.7	23.9	73.4	37.5	159.9	122.6	11.5
	95%	29.8	53.8	56.2	82.8	79.5	65.3	110.5	92.3	171.6	276.0	107.9	65.9	176.0	97.2	388.4	315.0	40.1
	$\%t: \nu_t^* < 0$	1%	2%	3%	5%	5%	3%	7%	6%	10%	14%	7%	4%	10%	6%	14%	13%	3%
$\%X_T < X_T^{AC}$	45%	40%	39%	35%	35%	38%	30%	33%	23%	17%	29%	36%	21%	31%	16%	18%	40%	
		Strategy $\max(\nu^*, 0)$																
quantile	mean	2.0	7.1	7.8	17.0	15.4	10.6	27.8	20.6	54.9	101.6	28.1	11.7	58.7	24.4	141.4	112.5	5.0
	stdev	16.3	26.5	27.4	38.3	36.9	31.6	46.9	41.7	66.7	106.3	47.0	31.2	70.2	43.3	150.8	125.8	20.2
	5%	-24.0	-32.8	-33.4	-38.5	-38.8	-36.7	-37.5	-38.1	-24.8	2.9	-34.2	-32.7	-23.2	-33.5	3.9	-3.7	-23.1
	25%	-6.6	-7.5	-7.5	-4.3	-5.1	-6.8	2.3	-2.0	13.5	28.6	1.8	-4.7	15.0	0.3	37.1	27.4	-4.2
	50%	1.2	4.7	5.2	11.4	10.3	6.9	19.4	13.2	39.3	71.0	18.5	7.0	41.7	16.3	95.5	76.7	1.9
	75%	9.9	19.6	20.6	34.8	32.2	25.1	48.7	39.2	81.9	143.0	48.0	25.7	86.8	42.7	198.7	160.1	12.1
	95%	29.7	54.5	56.8	86.3	81.9	67.1	115.8	98.3	184.1	315.6	118.2	69.1	195.0	105.9	439.5	354.0	41.0
	$\%t: \nu_t^* < 0$	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
$\%X_T < X_T^{AC}$	45%	39%	38%	30%	31%	35%	22%	27%	11%	4%	22%	33%	11%	25%	3%	6%	39%	

Table A.11: Relative performance of the strategy (in basis points). Statistics are calculated per share using (18) and stock specific parameters are as in Tables A.7 and A.8.

IV. Order-flow decay $\kappa = 60$, and strategy's trading frequency $N = 60$																		
		EBAY	INTC	DOW	CSCO	COST	AXP	AMAT	PNC	AA	FMER	CNQR	ROC	IMGN	DK	ARCC	BXS	AYI
		Strategy ν^*																
quantile	mean	0.2	1.9	2.2	6.3	5.7	3.8	12.2	8.8	28.0	60.1	14.7	6.0	34.2	13.4	96.2	75.8	3.7
	stdev	7.4	14.2	15.0	24.1	23.2	18.9	33.7	28.3	56.7	97.1	37.8	23.2	63.3	35.7	143.1	119.8	18.5
	5%	-11.8	-19.6	-20.5	-28.1	-27.5	-24.0	-34.1	-31.2	-45.7	-60.9	-35.6	-27.0	-45.7	-34.6	-71.7	-64.8	-22.2
	25%	-3.8	-5.8	-6.0	-6.8	-6.9	-6.5	-6.3	-6.7	-2.2	6.5	-5.6	-6.6	0.3	-6.0	12.1	8.0	-4.7
	50%	0.2	1.3	1.4	4.2	3.7	2.3	7.7	5.5	17.4	38.2	8.8	3.6	22.5	8.1	61.3	45.7	1.8
	75%	4.1	8.3	9.0	16.4	15.5	12.1	26.0	20.6	50.4	97.5	28.9	15.7	57.7	27.2	149.0	118.4	10.5
	95%	12.5	25.8	27.1	49.0	46.0	36.4	72.9	60.1	132.7	246.9	83.6	46.5	151.0	77.9	368.4	304.1	35.7
$\%t : \nu_t^* < 0$		0%	2%	2%	4%	4%	3%	6%	5%	9%	13%	6%	4%	10%	6%	14%	13%	3%
$\%X_T < X_T^{AC}$		49%	45%	44%	39%	40%	42%	34%	37%	27%	20%	32%	39%	25%	34%	16%	18%	41%
		Strategy $\max(\nu^*, 0)$																
quantile	mean	0.2	2.2	2.5	8.0	7.2	4.7	15.9	11.3	39.0	85.8	19.2	7.5	45.4	17.3	130.3	103.8	4.6
	stdev	7.3	14.0	14.8	24.1	23.2	19.0	34.5	28.7	57.4	100.3	38.3	23.3	63.7	36.2	146.4	120.2	18.5
	5%	-11.7	-18.5	-19.1	-25.1	-24.6	-22.1	-28.7	-27.5	-30.9	-19.3	-28.9	-23.8	-30.0	-28.2	0.3	-13.9	-20.5
	25%	-3.8	-5.5	-5.6	-5.5	-5.6	-5.8	-3.4	-4.8	5.1	20.2	-1.9	-5.4	7.3	-3.2	31.5	23.7	-4.2
	50%	0.2	1.3	1.5	4.9	4.3	2.7	9.5	6.7	25.4	59.0	11.3	4.1	29.6	10.5	86.5	70.2	1.9
	75%	4.1	8.4	9.1	17.5	16.7	12.7	29.4	22.4	61.3	124.5	33.6	16.5	71.1	30.8	183.4	149.9	11.2
	95%	12.5	26.1	27.8	52.5	49.7	38.0	80.9	65.2	151.1	276.9	91.6	49.6	166.3	84.7	419.5	332.4	37.3
$\%t : \nu_t^* < 0$		0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
$\%X_T < X_T^{AC}$		49%	44%	44%	37%	37%	41%	30%	33%	18%	8%	27%	37%	17%	30%	5%	8%	40%

Table A.12: Relative performance of the strategy (in basis points). Statistics are calculated per share using (18) and stock specific parameters are as in Tables A.7 and A.8.

References

- 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. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. Applied Mathematical Finance 10(1), 1–18.
- Almgren, R. and N. Chriss (2001). Optimal execution of portfolio transactions. Journal of Risk 3, 5–40.
- Almgren, R., C. Thum, E. Hauptmann, and H. Li (2005). Direct estimation of equity market impact. Risk 18, 5752.
- Avellaneda, M. and S. Stoikov (2008, November). High-frequency trading in a limit order book. Quantitative Finance 8, 217–224.
- Bechler, K. and M. Ludkovski (2014). Optimal execution with dynamic order flow imbalance. arXiv preprint arXiv:1409.2618.
- Cartea, Á., R. Donnelly, and S. Jaimungal (2013). Algorithmic trading with model uncertainty. SSRN: <http://ssrn.com/abstract=2310645>.
- Cartea, Á., R. F. Donnelly, and S. Jaimungal (2015). Enhancing trading strategies with order book signals. Available at SSRN 2668277.
- Cartea, Á. and S. Jaimungal (2013). Optimal execution with limit and market orders. Forthcoming Quantitative Finance.
- Cartea, Á. and S. Jaimungal (2015a, Winter). Order-flow and liquidity provision. ARGO 2(5), 7–11.
- Cartea, Á. and S. Jaimungal (2015b). Risk metrics and fine tuning of high frequency trading strategies. Mathematical Finance 25(3), 576–611.
- Cartea, Á., S. Jaimungal, and D. Kinzebulatov (2013, December). Algorithmic trading with learning. SSRN eLibrary, <http://ssrn.com/abstract=2373196>.
- Cartea, Á., S. Jaimungal, and J. Penalva (2015). Algorithmic and High-Frequency Trading (1st ed.). Cambridge: 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.
- Fodra, P. and M. Labadie (2012). High-frequency market-making with inventory constraints and directional bets. arXiv preprint arXiv:1206.4810.

- Frei, C. and N. Westray (2013). Optimal execution of a VWAP order: a stochastic control approach. Mathematical Finance, n/a–n/a.
- 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., A. Schied, and A. Slynko (2012). Transient linear price impact and Fredholm integral equations. Mathematical Finance 22(3), 445–474.
- Graewe, P., U. Horst, and E. Séré (2013). Smooth solutions to portfolio liquidation problems under price-sensitive market impact. arXiv preprint arXiv:1309.0474.
- Guéant, O. and C.-A. Lehalle (2013). General intensity shapes in optimal liquidation. Mathematical Finance, Forthcoming.
- Guéant, O., C.-A. Lehalle, and J. Fernandez-Tapia (2012). Dealing with the inventory risk: a solution to the market making problem. Mathematics and Financial Economics, 1–31.
- Guilbaud, F. and H. Pham (2013). Optimal high frequency trading with limit and market orders. Quantitative Finance 13, 79–94.
- Kharroubi, I. and H. Pham (2010). Optimal portfolio liquidation with execution cost and risk. SIAM Journal on Financial Mathematics 1(1), 897–931.
- Lorenz, C. and A. Schied (2013). Drift dependence of optimal trade execution strategies under transient price impact. Finance and Stochastics 17(4), 743–770.
- Schied, A. (2013). Robust strategies for optimal order execution in the Almgren-Chriss framework. Applied Mathematical Finance 20(3), 264–286.