

The Shadow Price of Latency: Improving Intraday Fill Ratios in Foreign Exchange Markets*

Álvaro Cartea[†] and Leandro Sánchez-Betancourt[†]

Abstract. Latency is the time delay between an exchange streaming market data to a trader, the trader processing information and deciding to trade, and the exchange receiving the order from the trader. Liquidity takers face a moving target problem as a consequence of their latency in the marketplace. They send market orders with a limit price that aim at a price and quantity they observed in the limit order book (LOB), and by the time their order is processed by the exchange, prices could have worsened, so the order may not be filled, or prices could have improved, so the order is filled at a better price. In this paper we provide a model to compute the price that liquidity takers would be willing to pay to reduce their latency in the marketplace. To this end, we derive a latency-optimal strategy that specifies the limit price of liquidity taking orders to increase the chances of filling orders if, due to latency, prices or quantities in the LOB have worsened. The latency-optimal strategy balances the trade-off between the costs of walking the LOB and targeting a desired percentage of filled orders over a period of time. We employ the cost of improving fills with the latency-optimal strategy to compute the shadow price of latency. Finally, we use a proprietary data set of foreign exchange (FX) to compute the maximum price that an FX trader would be willing to pay for colocation and hardware to reduce their latency in the marketplace.

Key words. latency, fill ratio, high-frequency trading, algorithmic trading

AMS subject classifications. 93, 90, 60

DOI. 10.1137/19M1258888

1. Introduction. With the advent of computerized trading, the speed at which agents operate in electronic markets is down to milliseconds or microseconds. Speed to process information and make decisions is essential to the success of trading strategies and being faster than one's peers provides a valuable competitive edge. Traders require time to execute strategies and their instructions take time to arrive and to be processed by the exchange. This time delay, known as latency, affects the efficacy of trading strategies because, in the meantime, the exchange processes other instructions that update the limit order book (LOB) with information not known by the trader at the time she executed the strategy.

In foreign exchange (FX) markets, the supply of spot currency pairs displayed in the LOB might undergo thousands of updates over very short periods of time; see [Gould, Porter, and Howison \(2016\)](#). When the best bid and best ask prices in the LOB are changing very rapidly, liquidity takers face a moving target problem as a consequence of their latency in the marketplace. Liquidity takers send market orders with a limit price (MOsLP) that aim at a price and quantity they have observed in the LOB, but by the time their order is processed by the

*Received by the editors April 30, 2019; accepted for publication (in revised form) October 26, 2020; published electronically March 1, 2021.

<https://doi.org/10.1137/19M1258888>

[†]Mathematical Institute, University of Oxford, Oxford, OX2 6GG England (Alvaro.Cartea@maths.ox.ac.uk, leandro.sanchezbetancourt@maths.ox.ac.uk).

exchange, prices or quantities may have worsened, so the order cannot be filled; alternatively, prices or quantities may have improved, so the order is filled at better prices.

In this paper, we show how to compute the price that a liquidity taker is willing to pay to reduce their latency in the marketplace. We summarize the steps we take. One, we develop a latency-optimal strategy to improve the chances that liquidity taking orders are filled. Two, we compute the costs borne by the strategy. These costs consist of the slippage received to ensure that the order is filled, which is negative if prices or quantities worsen, or is positive if prices or quantities improve. Three, we build a function that links latency, fill ratios, and costs borne by the latency-optimal strategy. With this function, we compute the shadow price of latency that a trader would be willing to pay for colocation and hardware to reduce their latency in the marketplace.

Specifically, in our approach we show how the FX trader chooses the price limit of the MOLP in an optimal way to improve fill ratios over a period of time (days, weeks, months), while keeping orders exposed to receiving a price improvement. Traders can buffer the adverse effects of missing a trade by specifying a price limit in their market orders to increase the probability of filling the order when it is processed by the exchange. This price limit consists of the best price seen by the trader in the LOB, plus a discretionary buffer that determines the number of ticks the order can walk the LOB. This buffer does not preclude the order from being filled at better prices if the LOB is updated with more favourable prices or quantities.

Increasing fill ratios is costly. Everything else being equal, the chances of filling a MOLP increase if the order can walk the LOB. Thus, there is a trade-off between ensuring high fill ratios and the execution costs borne by the trading strategy. In our model, the problem solved by the trader balances this trade-off by optimizing the price limit specified in the MOLP, while targeting a fill ratio over a trading horizon. The trader's optimal strategy specifies the discretion for each transaction depending on the proportion of orders that have been filled, how far the strategy is from the target fill ratio, the cost of walking the LOB, and the volatility of the exchange rate.

We employ a proprietary data set of FX trades to analyze the performance of the optimal strategy developed here. The data are provided by LMAX Exchange—an electronic multilateral trading facility in the FX market (www.lmax.com). We use transaction data for two FX traders, trader 1 (T1) and trader 2 (T2), to compare the fill ratios they have achieved to those attainable with the optimal strategy derived in this paper. The data span a set of dates from December 2016 to March 2017. During this period both traders filled between approximately 80% and 90% of their liquidity taking orders in the currency pair US dollar and Japanese yen, which we refer to as the USD/JPY pair. Latency causes T1 and T2 to miss trades because by the time the exchange processes their orders, the best bid and best ask prices in the LOB are updated. The effect of latency on trade fills is exacerbated during times of heightened volatility in the exchange rate. When volatility is arranged in quartiles, we find that 36.5% (resp., 40.0%) of T1's (resp., T2's) unfilled trades occur in the top quartile of volatility.

Clearly, due to latency, trades are not filled because the market moves away from the prices and quantities that T1 and T2 are attempting to achieve. If T1 and T2 were to return to the market to complete unfilled trades, the prices they would receive are highly likely worse than those at which they attempted in the original trade. We compare the costs of the latency-optimal strategy to the mark-to-market costs of orders that walk the LOB until filled.

We assume these orders are sent 20 ms or 100 ms after the exchange receives the instruction to execute the orders that could not be filled. We find that the average mark-to-market cost for T1's missed trades is 2.33 and 2.88 ticks and for T2 is 3.18 and 3.06 ticks, 20 ms and 100 ms later, respectively (in the currency pair USD/JPY one tick is 10^{-3} JPY).^{1,2}

We employ the optimal strategy developed here to show the trade-off between increasing fill ratios and the costs incurred by the strategy. For example, for a particular choice of model parameters, we show that T1 and T2 can increase the percentage of filled trades, during the period 5 December 2016 to 31 March 2017, to 99% for both traders. The increase in the fill ratios is due to the discretion included in the liquidity taking orders to walk the book, which come at a cost. In this example, the average cost incurred by T1 to fill missed trades is 1.76 ticks and for T2 is 1.24 ticks. On the other hand, the mark-to-market average cost of filling the missed trades (which were filled by the optimal strategy) with market orders 20 ms and 100 ms later is 2.01 and 2.58 ticks, respectively, for T1, and 2.82 and 2.75 ticks, respectively, for T2.

The performance of the optimal strategy is more remarkable during times of heightened volatility of the exchange rate. In the top quartile of volatility, the average cost of filling missed trades is 1.88 ticks when T1 employs the optimal strategy, while the mark-to-market average cost of filling these missed trades with market orders 100 ms later is 3.04 ticks. Similarly, for T2 the average cost of filling missed trades with the optimal strategy is 1.86 ticks, while the mark-to-market average cost of filling the missed trades with a market order 100 ms later is 3.31 ticks.

Finally, we use the proprietary data to map various levels of latency to the corresponding percentage of filled orders. We use this mapping to calculate the shadow price of latency that T1 would be willing to pay to reduce latency in the marketplace. We show that T1 would be better off employing the latency-optimal strategy developed here, instead of investing in hardware and colocation services to reduce latency. The latency-optimal strategy is superior because it not only achieves the same fill ratios as those obtained with better hardware and colocation, but it scoops price improvements that stem from orders arriving with latency at the exchange. Note that if latency is reduced with hardware and colocation, T1's liquidity taking strategy would be less affected by the moving target problem, so fill ratios are high, but would benefit little from changes in the LOB that, due to latency, provide price improvements.

A number of authors have recently addressed various aspects of latency in electronic markets. Moallemi and Sağlam (2013) look at the cost of latency in high-frequency trading. The authors employ market data to quantify the costs of latency in equity markets and show that the latency costs are of the same order of magnitude as other trading costs such as commissions and exchange fees. The data employed are publicly available from the NYSE, which does not contain trader identification to analyze latency for individual market participants. Cartea, Jaimungal, and Sánchez-Betancourt (2019a) propose a theoretical framework to find the optimal trade-off between missing trades and the costs of walking the book when there is

¹The average mark-to-market cost is calculated by taking the difference between the cost of filling all the volume missed in the order and the cost that the trader initially attempted to pay to fill the order, and dividing by the missed volume of the order. This quantity is multiplied by 10^3 to express it in ticks.

²From February 2018, the tick size for the pair USD/JPY is 10^{-4} JPY in LMAX Exchange Tokyo.

latency in the marketplace. Our study, on the other hand, focuses on liquidity taking orders in the FX market. We develop trading strategies that minimize the adverse effects of latency and use these strategies to compute the latency of liquidity takers.

The work of [Stoikov and Waeber \(2016\)](#) shows how to execute a large order in electronic markets by employing the volume imbalance of the LOB to predict price changes and study the effect of latency in the efficacy of the execution strategy. The authors backtest their strategy on publicly available data for US Treasury bonds and show that the advantage of observing the LOB to compute volume imbalance dissipates quickly as latency increases. [Lehalle and Mounjid \(2017\)](#) employ data from Nasdaq-Omx and also find that as latency increases, the informational content in the volumes of the LOB diminishes.

Other strands of the literature discuss the relationship of market quality and the speed of market participants. See for example [Bayraktar and Ludkovski \(2010\)](#) and [Guéant \(2016\)](#) for trading in illiquid markets. [Barger and Lorig \(2019\)](#) model the rapid updates of the best quotes in the LOB to propose a model of stochastic price impact.

Our paper is the first to employ proprietary data to analyze the effect of latency on the fill rates of liquidity taking strategies in FX markets and to compute the shadow price of latency for particular traders. Latency is specific to each market participant, and the effects of latency depend on the type of strategy employed by each trader. Access to proprietary data of FX transactions allows us to quantify the effects of latency on the efficacy of trading strategies and to develop and test latency-optimal strategies to compute the shadow price of latency for types of market participants.

The remainder of the paper is organized as follows. In section 2 we define latency and propose various ways to measure the fill ratios of MOsLP. In section 3 we describe the proprietary data set that we employ and we present empirical evidence of the relationship between fill ratios and the volatility of the exchange rate. In section 4 we present the model of fill ratios and solve the trader's dynamic optimization problem, where we also assume that the trader makes her model of fill ratios robust to misspecification. In section 5 we show how we estimate model parameters, and in section 6 we show the performance of the trading strategy. Section 7 discusses the trade-off between latency and fill ratios and computes the shadow price of latency for T1 and section 8 concludes.

2. Latency and fill ratios.

2.1. Latency. Latency is defined as the time “delay between sending a message to the market and it being received and processed by the exchange. Sometimes the time it takes for the exchange to acknowledge receipt is also accounted for”; see [Cartea, Jaimungal, and Penalva \(2015\)](#). This concept is useful if traders are only interested in measuring how long it takes a message to reach the exchange once it has been sent. However, this definition is “narrow” because it only takes into account hardware capacity and colocation to measure latency and does not account for other factors that affect how quickly the trader can process and react to market information before sending an instruction to the exchange.

In this paper we employ a broader definition of latency. To the time delay considered in the narrow definition of latency we add the time it takes the trader to execute the trading strategy. Thus, latency consists of three layers of time delay associated with these events: (i) The exchange streams LOB data to the trader. (ii) The trader receives the data from the

exchange and processes this and other relevant information to make a decision. (iii) Once a decision is made, the trader sends a message to the exchange. The decision could be to trade using limit orders or MOsLP, or to amend or cancel a limit order already resting in the LOB; see Moallemi and Sağlam (2013) and Stoikov and Waeber (2016) who employ a similar definition of latency.

MOsLP consist of two types of orders: immediate-or-cancel (IoC) or fill-or-kill (FoK). IoC and FoK are orders in which the trader specifies the maximum (minimum) price they are willing to pay (receive) when buying (selling) currency pairs. IoC is an order to buy or sell currency pairs which must be executed immediately, obeying the price limit, and any portion of the volume of the order that cannot be filled at the desired price limit is canceled. FoK is an order to buy or sell currency pairs which must be executed immediately in full or canceled.

We provide an example of the steps in the life of a FoK order to illustrate the effect of latency on the fill of the trade.

Assume the exchange streams a quote to a trader at time t_0 . The trader receives the quote at time $t_1 > t_0$ and at time $t_2 > t_1$ the trader sends a buy FoK order for 100 lots of the currency pair USD/JPY. The limit price of the FoK order is the best ask streamed by the exchange at time t_0 . Finally, the exchange receives and processes the order at time $t_3 > t_2$, thus our measure of latency is $t_3 - t_0$. The order will be filled in full if at time t_3 the best ask in the book is equal to or lower than (i.e., price improvement) the best ask at time t_0 and there are at least 100 lots posted at that price, otherwise the order is not filled. In either case, the exchange sends a message notifying the trader of the outcome of the trade.

The effect of latency on the outcome of strategies depends on a number of factors that affect the supply and demand of liquidity in the FX market, including magnitude of the latency, type of order sent by trader, volatility of the midexchange rate, and trend in the midexchange rate.

On the supply side of liquidity, the number of updates in the LOB will determine if the trade will be filled at the price, or a better price, than that specified in the FoK order, or if the order is not filled because prices have worsened. Changes in the best quotes of the LOB are difficult to predict and depend on many factors such as news announcements and other idiosyncratic needs of the liquidity providers who send, update, and amend their limit orders in the LOB.

On the demand side of liquidity, the effect of latency will depend on the type of strategy employed by the trader. For example, for momentum trading strategies, i.e., buy before rates increase or sell before rates decrease, latency will increase the chances of missing trades and decrease the chances of obtaining price improvements.

2.2. Fill ratios and posttrade analysis. Fill ratios of orders sent by a trader to an exchange can be computed in a number of ways. A trader who only sends FoK orders will receive fills of all the volume or nothing, as opposed to IoC orders, which can be partially filled.

A measure of fill ratio is important because it summarizes how effective an agent's strategy is in completing trades. Here, we define three measures of fill ratio and show how to extend these measures to include relevant posttrade information such as (i) orders that were filled at better prices (i.e., price improvement), and (ii) orders that were not filled, but could have been filled at worse prices. Below we provide more details about the posttrade information we

include in the measures of fill ratio. Next, we describe the variables and notation we employ.

We denote by n the number of MOsLP employed to compute the trader's fill ratio. Let $\tau_1, \tau_2, \dots, \tau_n$ be the times at which the trader sends liquidity taking orders to the exchange and denote by $\ell_1, \ell_2, \dots, \ell_n$ the times at which the orders are processed by the exchange, $\ell_i > \tau_i$.

MOsLP specify an average limit price per unit, denoted by L_i , which is the maximum (minimum) average price the trader is willing to pay (receive) per unit of bought (sold) of the currency pair.³ When the exchange processes the trader's message, the order is filled if there is enough liquidity posted in the LOB at prices and volumes that obey the average price limit L_i . A MOLP buy (sell) with limit price $L_i = \infty$ ($L_i = 0$) will walk all levels of the LOB until the order is filled in full, in which case we denote by B_i the average price per unit of currency pair obtained by the order that is filled in full.

The quantity $L_i - B_i$ (which can be computed only after the exchange receives and processes the order) is important because it gives posttrade information about different outcomes such as (i) by how much did prices move to cause the liquidity taking order to miss its target, which we refer to as *potential slippage*, or (ii) by how much did the execution price improve, which we refer to as *price improvement*. For example, when the trader sends a buy order with an average price limit of L_i per unit of currency pair, there is a price improvement if $L_i - B_i > 0$. On the other hand, if the average price required to fill the order is higher than the average limit price the trader is willing to pay per unit of currency pair, then the potential slippage is $L_i - B_i < 0$ —note that in this second case the trader does not incur the potential slippage because the order is not filled in full due to the limit price.

Moreover, the variable $I_i \in \{-1, 1\}$ is used to indicate the direction of the liquidity taking order, where 1 denotes a buy order and -1 denotes a sell order. Finally, V_i denotes the volume of currency pair lots the trader attempts to fill, and V_i^f denotes the filled volume, clearly $V_i^f \leq V_i$.

We define three ways to compute a fill measure of n consecutive trades. These measures are based on (1) number of trades, (2) average volume of trades, and (3) volume of trades.

1.

$$(1) \quad F_n = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{I_i L_i \geq I_i B_i\}},$$

2.

$$(2) \quad F_n^{AV} = \frac{1}{n} \sum_{i=1}^n \frac{V_i^f}{V_i},$$

3.

$$(3) \quad F_n^V = \frac{\sum_{i=1}^n V_i^f}{\sum_{i=1}^n V_i},$$

³In practice, exchanges and over-the-counter facilities have their own rules regarding price limits. Some employ marginal prices as price limit (that is, the price limit applies to the last unit bought or sold), and others employ average prices of the order as we assume here. Our model works in both cases.

Table 1

Proprietary data and measures of fill ratios for period 2 December 2016 to 31 March 2017.

	Number of attempts N	Successful filled trades (incl. partial)	Successful filled trades (only full)	Avg. vol. attempted $\times 10^5$ USD	Avg. filled vol. $\times 10^5$ USD	Order type	F_N	F_N^{AV}	F_N^V
T1	116,912	106,878	106,878	4.5	4.1	FoK	90.78%	90.78%	90.54%
T2	68,129	61,440	59,303	1.2	1.0	IoC	87.04%	88.62%	79.00%

where $\mathbf{1}$ is the indicator function and the superscript V in F_n^V stands for volume and the superscript AV in F_n^{AV} stands for average volume.

One can enhance the three measures above by incorporating more information of the LOB, so strategies that monitor the fill ratio can also trade off price improvement and slippage. Specifically, one can add to the measures above the component

$$(4) \quad F(n) = \zeta \frac{1}{n} \sum_{i=1}^n I_i (L_i - B_i),$$

where ζ is a nonnegative parameter. Note that $F(n) \in \mathbb{R}$. The units of ζ are such that the measure $F(n)$ is unitless. Specifically, the units of ζ are $1/c$, where c is the unit of currency B in the pair A/B . For example, for the USD/JPY pair, c denotes ¥. For simplicity, in our analysis below we employ the fill measures in (1) and (2), although we aim to describe the general case when one adds $F(n)$ to measures (1) and (2).

3. Empirical analysis.

3.1. Data. We employ a set of proprietary data of transactions for T1 and T2 during the period 2 December 2016 to 31 March 2017 in the currency pair USD/JPY. Both traders are liquidity takers, i.e., they do not send limit orders to provide liquidity in the LOB during the period we study. We have the following information: side of order (buy or sell), time-stamp of when the exchange processed the order, volume, currency pair, type of order (FoK, IoC), limit price L_i , and whether the order was filled and at what price(s). Finally, we have the LOB for the pair USD/JPY at millisecond intervals. A summary of the main variables is given in Table 1.

From 2 December 2016 to 31 March 2017 the best prices and quantities were updated over 63 million times, i.e., a mean of 746,871 of daily updates with standard deviation of 218,298. During this period, both traders send orders that aim at the best bid or best ask price that was streamed to them by the exchange. T1 sent FoK orders with zero discretion and T2 sent IoC orders with zero discretion. For a buy MOLP we define “discretion” as the difference between the price limit and the best offer seen by the trader when she devised the strategy. Similarly, for a sell MOLP we define discretion as the difference between the best bid and the price limit of the MOLP.

3.2. Fill ratios and volatility. The liquidity provided in the LOB is constantly being updated as a result of news, changes in the willingness to trade of liquidity providers, and the arrival of liquidity taking orders; see, for example, [Almgren \(2012\)](#) for a discussion of trading

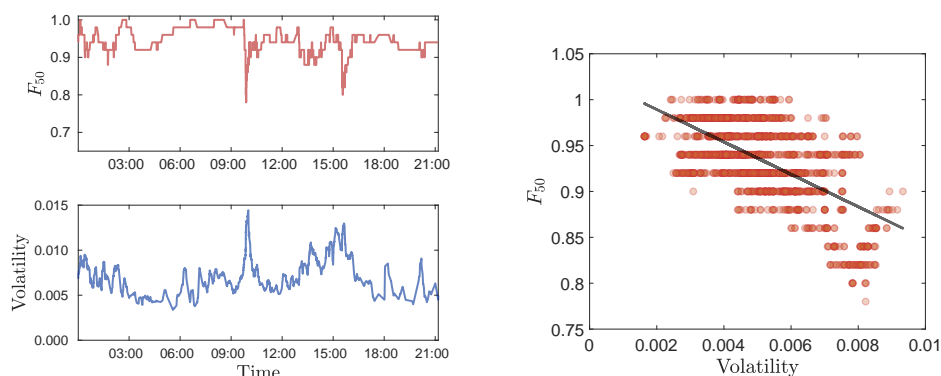


Figure 1. Left-hand top panel: fill ratio using measure (1) with $n = 50$. Left-hand bottom panel: volatility of microexchange rate. Right-hand panel: Scatter plot and ordinary least squares (OLS) regression of volatility and fill ratio. Data are from 26 January 2017.

with stochastic liquidity. Most of the updates of limit orders (cancellations, amendments, and arrival of new limit orders) take place in the best quotes and ticks closest to the prevailing midexchange rate. Traders with nonzero latency face a moving target problem when they attempt to hit bids and lift offers in the LOB. This problem is exacerbated during periods of high volatility of the exchange rate.

The top figure in the left-hand panel of Figure 1 shows the evolution of the fill ratio of T1's trade attempts, using measure (1) with $n = 50$, during 26 January 2017. The first observation in the figure is computed when the trader has attempted the 50th trade and the fill ratio is updated every time there is a new attempt.

The bottom figure in the left-hand panel shows the volatility of the midexchange rate over the same period. Throughout this paper volatility is computed as the standard deviation of the log-returns of the microexchange rate (sampled every 500 milliseconds) over a 10-minute rolling window.⁴ The microexchange rate at time t is denoted by m_t and given by

$$(5) \quad m_t = \frac{P_t^a Q_t^b + P_t^b Q_t^a}{Q_t^a + Q_t^b},$$

where Q_t^a and Q_t^b are the quantities available at the best ask price (denoted by P_t^a) and the best bid price (denoted by P_t^b), respectively.

From Figure 1, left-hand panel, we observe the negative correlation between volatility and the fill ratios achieved by T1's strategy, which always aims at the best bid or best offer price she observed before sending the MOLP. Clearly, volatility is high when the microexchange rate undergoes many changes and updates—this explains why the number of unfilled MOsLP (with zero discretion to walk the book) is higher when volatility is high.

⁴We chose the number of observations n to calculate the measure of fill ratio and the time-window to calculate the volatility of the midexchange rate as follows. For each day of the sample, we compute the correlation between the measure of fill ratio and the volatility of the midexchange rate for various values of n and various values of the time window. When $n = 50$ and the volatility window is ten minutes, the average of the daily correlations is smallest (i.e., is closest to -1).

The right-hand panel of Figure 1 shows a scatter plot of fill ratio and volatility. The solid line in the figure is obtained by performing an OLS regression of fill ratio on volatility.

4. Model. In this section we describe the model of fill ratios and the dynamic optimization problem solved by the trader. We develop the model of fills in three steps. First, we present a model for the dynamics of volatility of the microexchange rate. Second, we propose a model of fill ratios that includes the effect of volatility of the microexchange rate on fills. Third, we specify how MOsLP with discretion to walk the LOB affect the dynamics of fill ratios.

Volatility of microexchange rate. We denote the volatility of the microexchange rate, see (5), by $v = (v_t)_{\{0 \leq t \leq T\}}$ and assume it satisfies the stochastic differential equation (SDE)

$$(6) \quad dv_t = \kappa (\bar{v} - v_t) dt + \sigma_v v_t dW_t^v,$$

where $\kappa > 0$ is the exponential speed at which volatility reverts to its long-term level, denoted by $\bar{v} > 0$, $\sigma_v > 0$ is a dispersion parameter (i.e., volatility of volatility), and $W^v = (W_t^v)_{\{0 \leq t \leq T\}}$ is a standard Brownian motion.

In our analysis we consider other models for volatility of the microexchange rate. In Appendix B.3, we show results when volatility follows the Ornstein–Uhlenbeck process

$$(7) \quad dv_t^{\text{OU}} = \kappa^{\text{OU}} (\bar{v}^{\text{OU}} - v_t^{\text{OU}}) dt + \sigma_v^{\text{OU}} dW_t^v$$

and

$$(8) \quad dv_t^{\text{H}} = \kappa^{\text{H}} (\bar{v}^{\text{H}} - v_t^{\text{H}}) dt + \sigma_v^{\text{H}} \sqrt{v_t^{\text{H}}} dW_t^v;$$

see Heston (1993). We assume that the parameters κ^{H} , \bar{v}^{H} , σ_v^{H} satisfy the Feller condition, so that the process $(v_t^{\text{H}})_{\{0 \leq t \leq T\}}$ is strictly positive. However, the value of the volatility process in model (7) is not necessarily positive. In Appendix B we prove that the performance of the latency-optimal strategy is robust to the choice of the volatility model for the microexchange rate.

Fill ratio: Orders with no discretion to walk the LOB. The fill ratio of MOsLP (buy and sell) with no discretion to walk the LOB is denoted by $F = (F_t)_{\{0 \leq t \leq T\}}$ and satisfies the SDE

$$(9) \quad dF_t = (\lambda (\bar{F} - F_t) - b(v_t - \bar{v})) dt + \sigma_F dW_t^F.$$

Here, the parameter b determines the effect of the volatility of the microexchange rate on the fill ratio of the trader's strategy. The parameter $\lambda > 0$ is the exponential speed at which the fill rate reverts to the level $\bar{F} \in \mathbb{R}$, the dispersion parameter σ_F is a nonnegative constant,

and $W^F = (W_t^F)_{\{0 \leq t \leq T\}}$ is a standard Brownian motion uncorrelated to W^v .⁵ Recall that we propose three ways to compute fill ratios (see (1), (2), (3)), so here F refers to one of those choices.

The parameters λ , \bar{F} , σ^F , and the function $h(v)$ are specific to each trader in the FX market. Latency and type of liquidity taking strategy are the key characteristics that make the fill ratio dynamics specific to each trader. For example, the dynamics of the fill ratio for a trader with zero latency are as in (9) with parameters $\sigma_F = 0$, $\lambda > 0$, $\bar{F} = 1$, and $b = 0$.

The term $\sigma^F dW^F$ represents the combined effect of two sources of uncertainty in the dynamics of the fill ratio, both of which depend on latency. One, the volatility of the microexchange rate will affect fill ratios via the term $-b(v_t - \bar{v})$ in the drift of (9), and we observe noise around this effect. Two, between the time the trader is streamed LOB data and the time the exchange receives the trader's MOLP, the LOB will have processed instructions from other market participants, hence the fill of the order is uncertain.

Fill ratio: Orders with discretion to walk the LOB. Liquidity takers acknowledge that, however fast their computers are, or efficient their software is, or whether their hardware is colocated, the time to reach the market and execute orders is not zero. A trading strategy in which market orders aim at only the best bid or best ask prices in the LOB will hardly deliver fill ratios of 100%.

One alternative to compensate for not reaching the market in time is to specify a maximum (minimum) price the trader is willing to pay (receive) when sending a MOLP. As discussed above, the price limit consists of the best offer (resp., bid) price observed by the trader plus (resp., minus) the discretion δ . This discretion is the additional slippage she is willing to incur to fill the trade—the discretion to walk the LOB does not preclude the strategy from benefiting from price improvements. For example, if the trader sees the bid (ask) exchange rate at 1.000, then an immediate-execution limit sell (buy) order with discretion $\delta = 0.001$ will be filled by the exchange with limit buy (sell) orders resting in the LOB down (up) to a minimum (maximum) bid (offer) price of 0.999 (1.001).

When the trader includes positive discretion in her MOsLP she will exert upward pressure on the fill ratio of her strategy because the order can walk the LOB. Thus, we assume that the trader's fill ratio, denoted by the controlled process $F^\delta = (F_t^\delta)_{\{0 \leq t \leq T\}}$, follows the SDE

$$(10) \quad dF_t^\delta = \left(\lambda (\bar{F} - F_t^\delta) + a \delta_t - b(v_t - \bar{v}) \right) dt + \sigma_F dW_t^F,$$

with a and b nonnegative constants; when $a = 0$ we obtain (9).

⁵Here we assume that the Brownian motions W^F and W^v are independent because the impact of volatility v_t on the dynamics of the fill ratio is included in the drift of the fill ratio dynamics in (9). However, note that when the Brownian motions W^F and W^v are correlated, one can derive the latency-optimal strategy in closed form when the volatility of the midexchange rate is as in (6) or (7)—but not for the specification of volatility in (8). To show this, observe that when W^F and W^v are correlated with parameter ρ , the PDE in (18) would include an extra term, which for model (6) is $\sigma_F \sigma_v \rho v V_{fv}$, for model (7) is $\sigma_F \sigma_v \rho V_{fv}$, and for model (8) is $\sigma_F \sigma_v \rho \sqrt{v} V_{fv}$. The ansatz we proposed in (19) works for the volatility models (6) and (7). Specifically, for model (6), the term $\sigma_F \sigma_v \rho v h_5(t)$ collects in the ODE for h_3 which can be solved analytically. For model (7), the term $\sigma_F \sigma_v \rho h_5(t)$ collects in the ODE for h_0 which can be solved analytically. Finally, for model (8), the ansatz in (19) is not suitable and we cannot obtain a closed-form solution.

As usual, we work on a completed filtered probability space $(\Omega, \mathbb{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ with \mathcal{F}_t the natural filtration generated by the 2-dimensional Brownian motion $\mathbf{W} = (W^F, W^v)$, and we refer to \mathbb{P} as the trader's reference measure.

4.1. Model uncertainty. The trader acknowledges that the model of fill ratios (10) may be misspecified, so she considers alternative dynamics of the fill ratio to make the model robust to misspecification; see Cartea, Donnelly, and Jaimungal (2017). This ambiguity about model choice, or model uncertainty, affects the optimal strategy, i.e., affects the discretion of the orders she sends to the exchange. We incorporate the trader's ambiguity about model choice in two steps. First, we characterize alternative measures that describe the fill ratio dynamics considered by the trader, and then we determine how the trader chooses the reference measure \mathbb{P} or one of the alternative measures.

The trader considers candidate measures $\mathbb{Q}(x^F)$ that are equivalent to \mathbb{P} and characterized by the Radon–Nikodym derivative

$$(11) \quad \frac{d\mathbb{Q}(x^F)}{d\mathbb{P}} \Big|_t = \exp \left\{ -\frac{1}{2} \int_0^t (x_u^F)^2 du - \int_0^t x_u^F dW_u^F \right\},$$

where $(x_t^F)_{0 \leq t \leq T}$ is an \mathcal{F}_t -adapted process such that $(d\mathbb{Q}(x^F)/d\mathbb{P}|_t)_{0 \leq t \leq T}$ is a martingale. Thus, we denote by \mathfrak{Q} the class of alternative measures

$$\mathfrak{Q} = \left\{ \mathbb{Q}(x^F) \mid x^F \text{ is } \mathcal{F}\text{-adapted and } \left(\frac{d\mathbb{Q}(x^F)}{d\mathbb{P}} \Big|_t \right)_{0 \leq t \leq T} \text{ is a martingale} \right\}.$$

Next, the trader penalizes deviations from the reference measure using the relative entropy from t to T , which is given by

$$(12) \quad \mathcal{H}_{t,T}(\mathbb{Q} \mid \mathbb{P}) = \frac{1}{\varphi} \log \left(\frac{d\mathbb{Q}/d\mathbb{P}|_T}{d\mathbb{Q}/d\mathbb{P}|_t} \right).$$

Here the parameter φ is a nonnegative constant that represents the trader's degree of ambiguity aversion. If the trader is confident about the reference measure \mathbb{P} , then the ambiguity aversion parameter φ is small and any deviation from the reference model of fill dynamics is very costly. In the extreme $\varphi \rightarrow 0$, the trader is very confident about the reference measure, so she chooses \mathbb{P} because the penalty that results from rejecting the reference measure is too high. On the other hand, if the trader is very ambiguous about the reference model, considering alternative models results in a very small penalty. In the extreme $\varphi \rightarrow \infty$, deviations from the reference model are costless, so the trader considers the worst case scenario.

4.2. Value function. The trader's performance criterion under measure \mathbb{Q} is given by

$$\mathbb{E}^{\mathbb{Q}} \left[-\phi_\delta \int_0^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_0^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{0,T}(\mathbb{Q} \mid \mathbb{P}) \right],$$

and the trader's value function is

$$(13) \quad V(t, f, v) = \sup_{\delta \in \mathcal{A}_{t,T}} \inf_{\mathbb{Q} \in \mathfrak{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} \mid \mathbb{P}) \right],$$

where \hat{F} denotes the target fill ratio and $\hat{\delta}$ is a constant that denotes a fixed discretion to

target. Here, the penalty parameters ϕ_δ and ϕ_F are nonnegative constants. As the value of the parameter ϕ_δ increases, the trader submits orders with discretion closer to the target $\hat{\delta}$.⁶ Similarly, as the value of the parameter ϕ_F increases, the discretion to walk the book increases to ensure that the fill ratio hits the target \hat{F} . The last term on the right-hand side of the value function is the penalty in (12). Straightforward choices for the values of the parameters \hat{F} and $\hat{\delta}$ are one and zero, respectively.

Recall that the control variable δ denotes the discretion of the order to walk the LOB, which is measured from the best quote observed by the trader when receiving the quotes from the exchange.

Finally, recall that \mathcal{Q} represents all the measures equivalent to the reference measure \mathbb{P} , and the set of admissible strategies is

$$(14) \quad \mathcal{A}_{0,T} = \left\{ \delta = (\delta_t)_{\{0 \leq t \leq T\}} \mid \delta \text{ is } \mathcal{F}\text{-adapted, } \mathbb{E} \left[\int_0^T (\delta_s)^2 ds \right] < \infty \text{ and} \right.$$

$$(15) \quad \left. \delta_t = \mu(t, F_t^\delta, v_t) \text{ for } \mu \in \mathcal{M}_{0,T} \right\},$$

where

$$(16) \quad \mathcal{M}_{0,T} = \left\{ \mu : [0, T] \times \mathbb{R} \times \mathbb{R}^+ \mapsto \mathbb{R} \mid \exists K \in \mathbb{R} \text{ s.t. } \forall \mathbf{y}, \mathbf{z} \in \mathbb{R} \times \mathbb{R}^+ \text{ and } t \in [0, T], \right. \\ \left. |\mu(t, \mathbf{y}) - \mu(t, \mathbf{z})| \leq K |\mathbf{y} - \mathbf{z}| \text{ and } |\mu(t, \mathbf{y})|^2 \leq K^2 (1 + |\mathbf{y}|^2) \right\}.$$

We require the set of admissible strategies to satisfy the condition appearing in (15), so that the change of measure specified by (11) is valid.

By standard results, the Hamilton–Jacobi–Bellman–Isaacs (HJBI) equation associated with problem (13) is given by

$$(17) \quad V_t + \sup_{\delta} \inf_x \left((\lambda (\bar{F} - f) + a \delta - b(v - \bar{v}) - \sigma_F x) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{1}{2\varphi} x^2 \right. \\ \left. + \kappa (\bar{v} - v) V_v + \frac{1}{2} v^2 \sigma_v^2 V_{vv} - \phi_\delta (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0$$

with terminal condition $V(T, f, v) = 0$. Here, subscripts in the function V represent partial derivatives, for example, $V_t = \partial V / \partial t$.

Proposition 1. *The supremum and infimum in HJBI (17) are achieved (in feedback form) at*

$$\delta^* = \hat{\delta} + \frac{a V_f}{2 \phi_\delta} \quad \text{and} \quad x^* = \varphi \sigma_F V_f,$$

respectively.

⁶In the LMAX Exchange there are many traders who employ a fixed discretion (for all orders) to walk the book. For example, there are traders who for long periods of time, send all orders with a discretion of five ticks to walk the book. This is a simple strategy that decreases, due to latency, the chances of missing trades. In our setup, this fixed discretion for all trades is given by the parameter $\hat{\delta}$ in the performance criterion.

Proof. Apply first order conditions to the supremum and infimum terms in (17), and check second order conditions to verify that these are the maximizer and minimizer, respectively. ■

Substituting δ^* and x^* in (17), the HJBI becomes

$$(18) \quad \begin{aligned} V_t + (\lambda(\bar{F} - f) - b(v - \bar{v})) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{a^2 V_f^2}{4 \phi_\delta} - \frac{1}{2} \varphi V_f^2 \sigma_F^2 + a \hat{\delta} V_f \\ + \kappa(\bar{v} - v) V_v + \frac{1}{2} v^2 \sigma_v^2 V_{vv} - \phi_F (f - \hat{F})^2 = 0, \\ V(T, f, v) = 0. \end{aligned}$$

By inspection of the PDE (18) and its terminal condition, we propose the ansatz

$$(19) \quad V(t, f, v) = h_0(t) + h_1(t) f + h_2(t) f^2 + h_3(t) v + h_4(t) v^2 + h_5(t) f v,$$

which we substitute into (18) and obtain a coupled system of ordinary differential equations (ODEs). The system of ODEs can be solved explicitly to obtain closed-form solutions for the deterministic functions $h_0, h_1, h_2, h_3, h_4, h_5$; see Appendix A.1.

Theorem 1 (verification). Define the time-dependent deterministic functions $g_1, g_2, g_3 : [0, T] \mapsto \mathbb{R}$ as

$$g_1(t) = h_1(t) + 2 \hat{F} h_2(t) + \bar{v} h_5(t), \quad g_2(t) = -2 h_2(t), \quad g_3(t) = h_5(t),$$

and let the ambiguity aversion parameter obey the bound

$$\varphi < a^2 / 2 \sigma_F^2 \phi_\delta,$$

which guarantees that we find a solution for h_2 . Then, the optimal discretion to walk the LOB is

$$(20) \quad \delta_t^* = \hat{\delta} + \frac{a}{2 \phi_\delta} \left(g_1(t) + g_2(t) (\hat{F} - F_t^{\delta^*}) + g_3(t) (v_t - \bar{v}) \right)$$

and

$$(21) \quad x_t^* = \varphi \sigma_F \left(h_1(t) + 2 h_2(t) F_t^\delta + h_5(t) v_t \right).$$

The controls in (20) and (21) are admissible and (19) is the trader's value function in (13).

Proof. For the proof, see Appendix A.2. ■

We explain the intuition of the terms appearing in the optimal discretion (20). The first term on the right-hand side of (20) is the target discretion $\hat{\delta}$. For example, if the trader imposes an arbitrarily large penalty on deviations from this discretion target, i.e., penalty parameter $\phi_\delta \rightarrow \infty$, then the optimal discretion $\delta_t^* \rightarrow \hat{\delta}$.

The second term, $g_1(t)$, is the time-dependent baseline level of discretion the trader posts when her fill ratio is close to the target \hat{F} and volatility is around its long-term level \bar{v} .

The third term, $g_2(t) (\hat{F} - F_t^\delta)$, adjusts the strategy if the fill ratio is not on target. The function $g_2(t)$ is positive for $t < T$, so when the fill ratio is below (above) its target \hat{F} the discretion to walk the LOB is increased (decreased).

The last term of the optimal strategy shows how the optimal discretion depends on the volatility of the microexchange rate. The function $g_3(t)$ is positive for $t < T$, so when volatility is higher (lower) than its long-term level, the strategy increases (decreases) the discretion of the MOsLP.

Throughout this paper we assume that the volatility of the microexchange rate follows (6). In Appendix B.2 we derive the optimal strategy when volatility follows (7) and (8). Furthermore, as a robustness check, in section 6 we analyze the performance of the latency-optimal strategy for these two alternative volatility models.

4.2.1. Infinite-horizon latency-optimal strategy. We also derive the latency-optimal strategy when $T \rightarrow \infty$. In Appendix B.1 we present the infinite-horizon value function of the trader, where the optimal control is

$$(22) \quad \delta_t^* = \hat{\delta} + \frac{a}{2\phi_\delta} \left(c_1 + c_2 (\hat{F} - F_t^{\delta^*}) + c_3 (v_t - \bar{v}) \right)$$

with $c_1 = k_1 + 2\hat{F}k_2$, $c_2 = -2k_2$, and $c_3 = k_5$, and the constants k_1 , k_2 , k_5 are

$$k_2 = \begin{cases} \frac{2\lambda + \beta - \sqrt{(2\lambda + \beta)^2 + 8\alpha_1\phi_F}}{4\alpha_1} & \text{if } \alpha_1 > 0, \\ -\frac{\phi_F}{\beta + 2\lambda} & \text{if } \alpha_1 = 0, \\ \frac{2\lambda + \beta + \sqrt{(2\lambda + \beta)^2 + 8\alpha_1\phi_F}}{4\alpha_1} & \text{if } \alpha_1 < 0, \end{cases}$$

$$k_5 = \frac{2bk_2}{2k_2\alpha_1 - \lambda - \kappa - \beta}, \quad k_1 = \frac{2\phi_F\hat{F} + (2b\bar{v} + 2\lambda\bar{F} + 2a\hat{\delta})k_2 + \kappa\bar{v}k_5}{\lambda + \beta - 2k_2\alpha_1},$$

where $\alpha_1 = \frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2$, $\beta > 0$, is a discount factor (see (48)), and

$$\varphi \leq \frac{1}{2\sigma_F^2} \left(\frac{a^2}{\phi_\delta} + \frac{(2\lambda + \beta)^2}{4\phi_F} \right).$$

It is straightforward to check that c_2 and c_3 are positive, and k_2 is continuous in α . At the end of section 6 we discuss the performance of the infinite-horizon strategy.

5. Estimation of model parameters. In this section we show how to estimate the parameters of the models of volatility and fill ratio.

Volatility parameters. We discretize the continuous-time volatility model (6) and write

$$(23) \quad \tilde{v}_i = \tilde{v}_{i-1} + \kappa (\bar{v} - \tilde{v}_{i-1}) \Delta + \sigma_v \tilde{v}_{i-1} \sqrt{\Delta} \epsilon_i,$$

where $i \in \{1, 2, \dots, N-1\}$, $0 = t_1 < t_2 < \dots < t_N = T$, $T = 1$, $t_i = iT/N$, $\epsilon_i \sim \mathcal{N}(0, \sigma_v^2 \Delta)$, and $\Delta = T/N$.

Assume $\{\tilde{v}_i\}_{i \in \{1, 2, \dots, N-1\}}$ are positive and divide (23) by \tilde{v}_{i-1} , so that

$$\frac{\tilde{v}_i}{\tilde{v}_{i-1}} - 1 = -\kappa \Delta + \kappa \bar{v} \Delta \frac{1}{\tilde{v}_{i-1}} + \sigma_v \sqrt{\Delta} \epsilon_i,$$

which we write as the OLS problem

$$(24) \quad \mathbf{y} = \beta_0 \mathbf{1} + \beta_1 \mathbf{s} + \boldsymbol{\epsilon}.$$

Here \mathbf{y} is an $(N-1)$ -dimensional vector containing the dependent variables with entries $y_i = \tilde{v}_i/\tilde{v}_{i-1} - 1$, $\mathbf{1}$ is an $(N-1)$ -dimensional vector of ones, β_0 and β_1 are scalar parameters, \mathbf{s} is an $(N-1)$ -dimensional vector containing the regressors $s_i = 1/\tilde{v}_{i-1}$, and $\boldsymbol{\epsilon}$ is the vector of residuals with independent and identically distributed (i.i.d.) entries $\epsilon_i \sim \mathcal{N}(0, \sigma_v^2 \Delta)$.

To obtain parameter estimates for β_0 and β_1 we employ the volatility of the microexchange rate at 500 millisecond intervals (see subsection 3.2) and use a standard OLS routine to compute the parameter estimates $\hat{\beta}_0$ and $\hat{\beta}_1$. Now, because $\beta_0 = -\kappa \Delta$ and $\beta_1 = \kappa \bar{v} \Delta$, we obtain the estimates

$$(25) \quad \hat{\kappa} = -\hat{\beta}_0/\Delta, \quad \hat{\bar{v}} = \hat{\beta}_1/\hat{\beta}_0, \quad \text{and} \quad \hat{\sigma}_v = \sigma_\epsilon/\sqrt{\Delta},$$

where σ_ϵ is the standard deviation of the residuals in (24).

Fill ratio parameters. We discretize the continuous-time fill ratio model (10) and write

$$(26) \quad \tilde{F}_i - \tilde{F}_{i-1} = \lambda \bar{F} \Delta + a \delta_{i-1} \Delta - \lambda \Delta \tilde{F}_{i-1} - b (\tilde{v}_{i-1} - \bar{v}) \Delta + \sigma_F \sqrt{\Delta} \eta_i,$$

where, as above, $i \in \{1, 2, \dots, N-1\}$, $0 = t_1 < t_2 < \dots < t_N = T$, $T = 1$, $t_i = iT/N$, $\eta_i \sim \mathcal{N}(0, \sigma_\eta^2 \Delta)$, and $\Delta = T/N$.

We write (26) as the OLS problem

$$(27) \quad \mathbf{z} = \gamma_0 \mathbf{1} + \gamma_1 \mathbf{u} + \gamma_2 \mathbf{w} + \gamma_3 \mathbf{r} + \boldsymbol{\eta},$$

where the vector \mathbf{z} contains the dependent variables, $\gamma_0, \gamma_1, \gamma_2$, and γ_3 are scalar parameters, the vectors $\mathbf{u}, \mathbf{w}, \mathbf{r}$ are regressors, and $\boldsymbol{\eta}$ is the vector of residuals with i.i.d. entries $\eta_i \sim \mathcal{N}(0, \sigma_F^2 \Delta)$.

To estimate the impact parameter a in (26) we compute the effect of the discretion δ on fill ratios for the set of values $\delta_t = 0, 1, 2, 3, 4, 5$, as follows. Throughout the day we assume that all FoK trades sent by T1 included a fixed discretion of $\delta = 0$ to walk the LOB, and then repeat the analysis assuming $\delta = 1$, and $\delta = 2$, and so on. Thus, we have six scenarios with fixed discretion for the whole trading day. We employ LOB data to determine if the orders would have been filled and create six $(N-1)$ -dimensional vectors $\mathbf{z}_0, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4, \mathbf{z}_5$, where each vector corresponds to each scenario with fixed discretion. Next, we pool all the data so that the dependent variable on the left-hand side of the regression and the regressors on the right-hand side are $6 \times (N-1)$ -dimensional vectors. We run a standard OLS routine to compute the parameter estimates $\hat{\gamma}_0, \hat{\gamma}_1, \hat{\gamma}_2$, and $\hat{\gamma}_3$. Then, the parameters of the fill ratio dynamics are

$$(28) \quad \hat{\lambda} = -\hat{\gamma}_1/\Delta, \quad \hat{\bar{F}} = -\hat{\gamma}_0/\hat{\gamma}_1, \quad \hat{b} = -\hat{\gamma}_2/\Delta, \quad \hat{a} = \hat{\gamma}_3/\Delta, \quad \hat{\sigma}_F = \sigma_\eta/\sqrt{\Delta},$$

Table 2

Parameter estimates for T1's model, 4 January 2017. All parameters are statistically significant.

Parameter	\hat{v}	$\hat{\kappa}$	$\hat{\sigma}_v$	$\hat{\lambda}$	\hat{F}	$\hat{\sigma}_F$	\hat{a}	\hat{b}
Estimate	0.0073	12.3887	1.3894	58.7206	0.9331	0.1944	1.0027	527.44

where σ_η is the standard deviation of the errors in (27). To estimate model parameters for T2 we proceed in the same way and use the definition of fill ratio (2) because T2 sends IoC orders to the exchange.

Table 2 shows parameter estimates for T1 for the trading day starting at 22:05 of 3 January 2017 and ending at 22:00 on 4 January 2017, which for simplicity we refer to as the trading day of 4 January 2017.

The daily estimation of parameters $(\hat{\kappa}, \hat{\sigma}_v, \hat{v})$, $(\hat{F}, \hat{b}, \hat{\sigma}_F, \hat{\lambda})$, and (\hat{a}, \hat{b}) takes on average 1.5 seconds, 0.5 seconds, and 0.3 seconds, respectively. Hence, it takes under ten seconds to obtain the parameter estimates at the beginning of every trading day.

6. Performance of strategy. In this section we show the performance of the optimal discretion strategy developed above. In subsection 6.1 we employ two days of trading data for T1 to discuss features of the strategy. In subsection 6.2 we employ data during the period December 2016 to March 2017 to study the performance of the strategy for T1 and T2 and in subsection 6.3 we discuss the results of the infinite-horizon case and present robustness checks. Finally, in section 7 we use trade data to infer the latency of T1 and calculate the shadow price of latency that T1 would be willing to pay to reduce her latency in the marketplace.

6.1. Stylized features of latency-optimal strategy. We showcase features of the strategy for T1. We employ parameter estimates for 4 January 2017 to compute the optimal discretion for the trading day 5 January 2017. During 5 January 2017, T1 sent 3,043 FoK orders (1,601 buys and 1,442 sells) to the exchange. All orders were sent with a discretion of $\delta = 0$, and due to latency, T1 missed 279 trades, thus the fill ratio for the whole day was $F_{3043} = 90.83\%$.

Had T1 employed the strategy developed here, the optimal discretion in the MOsLP would have been as depicted in the top panel of Figure 2 and would have attained a fill ratio of $F_{3043}^{\delta^*} = 99.17\%$, i.e., miss 25 out of the 3,043 trade attempts.⁷ Compared to the fill ratio of the suboptimal strategy $\delta = 0$, which we refer to as the naive strategy, the latency-optimal strategy would have filled 254 more trades at an average cost of 1.74 ticks—recall that one tick in the USD/JPY pair is 10^{-3} JPY and that the naive and latency-optimal strategies receive the same price improvements. The average cost is calculated as the cost of filling the 254 trades, in JPY, divided by the total volume of currency pairs of the 254 trades, and then multiplied by 10^3 to express it in ticks. The bottom panel of the same figure shows the volatility of the microexchange rate. In the next subsection we analyze in more detail the costs incurred by the optimal strategy to fill missed trades.

The left-hand panel of Figure 3 shows a heatmap of the fill ratio $F_{3043}^{\delta^*}$ (i.e., 5 January 2017) for a range of parameters ϕ_F and ϕ_δ . We observe that when the penalty parameter ϕ_δ is low,

⁷Here, the fill measure $F_n^{\delta^*}$ denotes the measure F_n controlled by the optimal discretion δ^* .

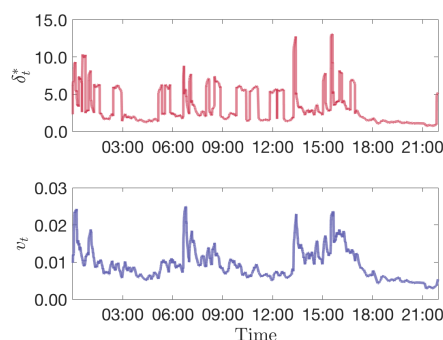


Figure 2. Parameters are $\hat{F} = 100\%$, $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $T = 5$ days, and other model parameters are in Table 2.

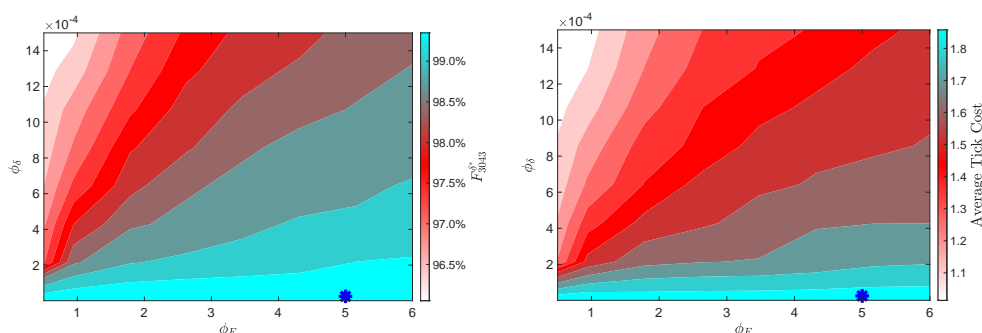


Figure 3. Fill ratio $F_{3043}^{\delta^*}$ for 5 January 2017. Parameters are $\varphi = 10^{-3}$, $\delta = 0$, $\hat{F} = 100\%$, $T = 5$ days, and other model parameters are in Table 2.

the strategy obtains the highest fill ratios because the trader is more willing to bear the costs of walking the LOB, i.e., incur more slippage. In the heatmap we also show, with a star, the pair $(\phi_\delta, \phi_F) = (2.5 \times 10^{-5}, 5)$, which are the parameters we use in the analysis of subsection 6.2.

The right-hand panel of the figure shows a heatmap of the average cost of the fill ratios attainable by the optimal strategy. The cost reported is the slippage of the order relative to T1's naive strategy.

Finally, Figure 4 shows the fill ratio $F_{3043}^{\delta^*}$ for the range $\varphi \in [0, a^2/\phi_\delta \sigma_F^2)$ of the ambiguity aversion parameter when the trader employs the optimal discretion strategy. Recall that the higher φ is, the less confident the trader is about the dynamics of F^δ . When $\varphi \rightarrow a^2/\phi_\delta \sigma_F^2$ the trader has very little confidence in her model of fill ratios, so the optimal strategy is to send orders with higher discretion to the exchange—the ambiguity averse trader is more willing to bear slippage costs to ensure high fill ratios. In Appendix A.1 we discuss the optimal control when $\varphi = a^2/2\phi_\delta \sigma_F^2$ and $a^2/2\phi_\delta \sigma_F^2 < \varphi$.

In the interest of space we only report results for T1 because those for T2 are similar. The computation time to implement the strategy is under 0.03 milliseconds per trade, which includes computing the volatility of the midexchange rate and the fill ratio.

6.2. Performance of strategy for traders T1 and T2. We employ trade data for T1 and T2 between 5 December 2016 and 31 March 2017 to analyze the performance of the latency-

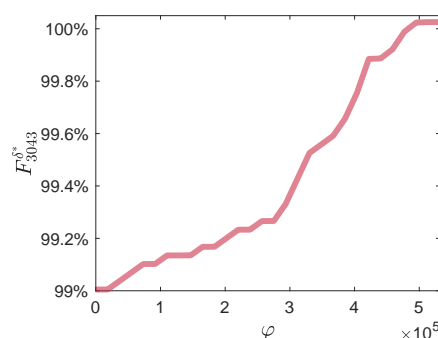


Figure 4. Fill ratio for optimal strategy. Parameters are $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\hat{F} = 100\%$, $T = 5$ days, and other model parameters are in Table 2.

optimal strategy. For each trading day we use the parameter estimates from the previous day and volatility of the exchange rate to compute the optimal discretion (20) in Proposition 1; recall that the volatility is calculated over a ten-minute window. We use LOB information to determine if the orders sent with the optimal discretion δ^* would have been filled and, if so, at what price(s).

For each day we employ the first 10 minutes of the trading day (22:05 to 22:15) to compute the first observation of volatility of the microexchange rate. The first fill ratio observation (F_{50} for T1 and F_{50}^{AV} for T2) is computed the first time there are 50 trades after 22:15. Throughout the period we study we employ 4,115 trades for T1 and 4,089 for T2 to compute the first observations of volatility and fill ratio for each trading day, and to estimate the parameters of the model for the first day we employ 2,339 trades for T1 and 772 for T2. Thus, the number of trades we employ to analyze the performance of the strategy is reduced to 110,458 for T1 and 63,268 for T2.

6.2.1. Performance of latency-optimal strategy for T1. Tables 3 and 4 provide information of the costs that T1 would have incurred had she employed the latency-optimal strategy and shows the mark-to-market costs of going back to the exchange to complete missed trades using market orders that walk the LOB until the order is filled in full.

During the period 5 December to 31 March 2017, T1 attempted 110,458 trades, all with discretion $\delta = 0$, and, due to latency, missed 9,483 trades, and received 3,938,559 JPY in price improvements. Had T1 employed market orders with infinite discretion, the cost of filling the 9,483 trades would have been 8,914,046 JPY; see row (a) in Table 3.⁸ The last two columns of the same row show the mark-to-market cost of market orders that walk the LOB 20 ms and 100 ms later to fill the 9,483 trades missed by the naive strategy—time is measured from the time-stamp recorded by the exchange when T1's order was received. For example, if T1 were to return to the market (20 ms and 100 ms later) to retry the missed trade using a market order that can walk the LOB until filled in full, then the costs to fill the 9,483 missed trades would have been 9,841,755 JPY and 12,158,945 JPY, respectively.

⁸A market order with infinite discretion walks the LOB until it is filled in full, that is, there is no price limit.

Table 3

Performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY. Computation times are under 10 seconds for each day.

		Missed by naive, filled by alternative strategies				
		Extra fills	Extra volume $\times 10^6$	Cost in JPY	Cost 20 ms later in JPY	Cost 100 ms later in JPY
(a)	$\delta_t = \infty \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
	δ^*				$\delta_t = \infty \forall t$	$\delta_t = \infty \forall t$
	Target					
(b)	$\hat{\delta} = 0$	8,318	3,718	6,528,965	7,459,198	9,601,222
(c)	$\hat{\delta} = 1$	8,603	3,839	6,807,416	7,770,984	9,959,191
(d)	$\hat{\delta} = 2$	8,758	3,906	6,999,804	7,950,543	10,185,240
(e)	$\hat{\delta} = 3$	8,938	3,982	7,232,545	8,181,285	10,447,450

Table 4

Performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (resp., 7) are the proportional increase of the values in columns 5 and 6 (resp., 5 and 7) in Table 3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY. Computation times are under 10 seconds for each day.

		Missed by naive, filled by alternative strategies					
		Extra fills	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks	Cost increase 20 ms later %	Cost increase 100 ms later %
(a)	$\delta_t = \infty \forall t$	9,483	2.11	2.33	2.88	10%	36%
	δ^*					$\delta_t = \infty \forall t$	$\delta_t = \infty \forall t$
	Target						
(b)	$\hat{\delta} = 0$	8,318	1.76	2.01	2.58	14%	47%
(c)	$\hat{\delta} = 1$	8,603	1.77	2.02	2.59	14%	46%
(d)	$\hat{\delta} = 2$	8,758	1.79	2.04	2.61	14%	46%
(e)	$\hat{\delta} = 3$	8,938	1.82	2.05	2.62	13%	44%

In rows (b) to (e) we show the performance of the latency-optimal strategy for discretion targets $\hat{\delta} \in \{0, 1, 2, 3\}$, respectively. For example, row (b) shows that a strategy with target discretion $\hat{\delta} = 0$ would have filled 8,318 more trades than the naive strategy. The cost of filling the 8,318 trades is 6,528,965 JPY and the mark-to-market costs to fill these trades, with market orders 20 ms and 100 ms later, are 7,459,198 JPY and 9,601,222 JPY, respectively.

Table 4 presents the average costs of filling missed trades in ticks (10^{-3} JPY). Column 3 in row (a) shows that the average mark-to-market cost of market orders with no price limit to fill

Table 5

Performance of strategy in volatility quartiles for T1. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Ticks are 10^{-3} JPY. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

Volatility quartiles %	Trades sent to exchange	Fills $\delta_t = 0 \forall t$		Fills $\delta_t^* \hat{\delta} = 0$		Missed by naive, filled by δ^* with $\hat{\delta} = 0$		
			%		%	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks
0–25	27,615	26,151	94.7%	27,274	98.8%	1.61	1.77	2.02
25–50	27,614	25,588	92.7%	27,303	98.9%	1.73	1.92	2.20
50–75	27,615	25,086	90.8%	27,339	99.0%	1.67	1.91	2.49
75–100	27,614	24,150	87.5%	27,377	99.1%	1.88	2.19	3.04

the 9,483 trades missed by the naive strategy is 2.11 ticks (i.e., $2.11 = 8,914,046/4,222 \times 10^3$). Row (b) shows that the average cost of filling 8,318 trades missed by the naive strategy, which were filled by the optimal strategy with discretion target $\hat{\delta} = 0$, would have been 1.76 ticks, and the mark-to-market cost to fill the 8,318 trades with market orders that walk the book until filled, submitted 20 ms or 100 ms later, is 2.01 or 2.58 ticks.

The last two columns in Table 4 show the “excess cost” (in percentage terms relative to optimal strategy’s attainable cost) of returning to the market, 20 ms and 100 ms later, to complete unfilled trades with market orders that walk the book until filled in full. For example, the costs of filling missed trades with market orders 20 ms (100 ms) later are between 13% and 14% (44% and 47%) more than the cost of filling them with the optimal strategy.

The costs of returning to the market to fill missed trades employing market orders with infinite discretion is, on average, greater than the costs of completing the trades, on the first attempt, employing the latency-optimal strategy. Thus, not only does the optimal strategy attain a very high fill ratio, it does it at a cost which is lower than the mark-to-market value of completing the trades on a second attempt. These savings justify the extra cost of optimally walking the LOB in the first trade attempt to achieve higher fill ratios than the fill ratios achieved by the naive strategy. Note that we are not imputing other costs that the trader may incur if a trade is not filled. For example, some trades could be one leg of a trade, which if not completed, will affect other trades or strategies; see [Cartea, Jaimungal, and Walton \(2019b\)](#).

Table 5 shows the results of the optimal strategy arranged in volatility quartiles when the discretion target is $\hat{\delta} = 0$. As discussed above, when the microexchange rate is more volatile, the chances of missing trades are greater than when volatility is low because the LOB is undergoing many more updates. The optimal strategy developed here counterbalances the adverse effects of high volatility on fill ratios, so that in periods of high volatility, everything else being equal, the trader sends orders with more discretion to walk the book. The table shows that as volatility increases, the strategy maintains a stable level of fills, while T1’s naive strategy misses more trades as volatility increases.

The last three columns in Table 5 show the cost of filling trades missed by T1’s naive strategy that are filled by the optimal strategy with $\hat{\delta} = 0$. The first of the last three columns

Table 6

Performance of optimal strategy for T2. The total number of trade attempts is 63,268. Cost-related quantities in ticks (10^{-3} JPY). Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies					
		Extra fills (incl. partial)	Extra fills (only full)	Extra volume $\times 10^6$ in JPY	Cost in JPY	Cost 20 ms later in JPY	Cost 100 ms later in JPY
(a)	$\delta_t = \infty \ \forall t$	7,038	7,038	1,455	2,294,999	4,628,993	4,455,516
δ^* Target						$\delta_t = \infty \ \forall t$	$\delta_t = \infty \ \forall t$
(b)	$\hat{\delta} = 0$	6,283	6,238	1,320	1,634,888	3,727,171	3,631,491
(c)	$\hat{\delta} = 1$	6,446	6,411	1,340	1,665,638	3,742,734	3,649,011
(d)	$\hat{\delta} = 2$	6,548	6,517	1,354	1,667,995	3,801,813	3,690,360
(e)	$\hat{\delta} = 3$	6,648	6,627	1,373	1,733,090	3,903,934	3,795,341

shows the cost incurred by the optimal strategy to execute missed trades for each volatility quartile. The penultimate and last columns of the table show the cost of returning to the market to fill the missed trades with market orders that walk the LOB until they are filled. As discussed earlier, the cost incurred by the latency-optimal strategy to fill trades missed by the naive strategy are lower than the costs of returning to the market 20 ms and 100 ms later. The gap between the costs of returning to the market 20 ms and 100 ms later and the costs incurred by the optimal strategy widens as volatility increases.

6.2.2. Performance of optimal strategy: T2. We repeat the analysis of the performance of the strategy using T2's trades. Tables 6 and 7 are presented in the same format as Tables 3 and 4, respectively.

During the period 5 December 2016 to 31 March 2017, T2 attempted 63,268 IoC trades, all with discretion $\delta = 0$ and, due to latency, missed 7,038 trades, including partial misses. The cost of achieving better fill ratios using the optimal strategy (for various discretion targets $\hat{\delta}$) is justified by the cost of returning to the exchange to fill missed trades 20 ms or 100 ms later with market orders with infinite discretion to walk the book. For example, had T2 employed the optimal strategy with discretion target $\hat{\delta} = 0$, only 755 trades would have been missed, as opposed to the 7,038 misses by the naive strategy. The average cost of filling the 6,283 trades missed by the naive strategy is 1.24 ticks, as opposed to 2.82 and 2.75 ticks, which are the mark-to-market costs of returning to the market 20 ms and 100 ms later.

Table 8 shows the results of the optimal strategy arranged in volatility quartiles when the discretion target is $\hat{\delta} = 0$. The interpretation is the same as that of Table 5 for T1. The cost incurred by the latency-optimal strategy to fill the trades missed by the naive strategy is lower than the costs of returning to the market 20 ms and 100 ms later. The gap between the costs of returning to the market 20 ms and 100 ms later and the costs incurred by the optimal strategy widen as volatility increases.

Table 7

Performance of optimal strategy for $T2$. The total number of trade attempts is 63,268. Cost-related quantities are weighted by volume and expressed in ticks (10^{-3} JPY). The metrics in column 6 (resp., 7) are the growth of the values in columns 5 and 6 (resp., 5 and 7) in Table 3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies				
		Extra fills (incl. partial)	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks	Growth 20 ms later %
(a)	$\delta_t = \infty \forall t$	7,038	1.58	3.18	3.06	102%
	δ^*					
	Target					
(b)	$\hat{\delta} = 0$	6,283	1.24	2.82	2.75	128%
(c)	$\hat{\delta} = 1$	6,446	1.24	2.79	2.72	125%
(d)	$\hat{\delta} = 2$	6,548	1.23	2.81	2.73	128%
(e)	$\hat{\delta} = 3$	6,648	1.26	2.84	2.76	125%

Table 8

Performance of strategy in volatility quartiles for $T2$. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

Volatility quartiles %	Trades sent to exchange	Fills % $\delta_t = 0 \forall t$ (incl. partial)		Fills % $\delta_t^*, \hat{\delta} = 0$ (incl. partial)		Missed by naive, filled by δ^* with $\hat{\delta} = 0$		
						Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks
0%–25%	15,817	14,950	94.5%	15,658	99.0%	0.94	1.90	1.87
25%–50%	15,817	14,545	92.0%	15,631	98.8%	1.12	2.72	2.68
50%–75%	15,817	14,226	89.9%	15,629	98.8%	1.25	2.75	2.92
75%–100%	15,817	13,329	84.3%	15,610	98.7%	1.86	3.51	3.31

6.3. Infinite-horizon strategy and robustness checks. Infinite-horizon: Tables 11 and 12, in Appendix B.1 which are the analogues of Tables 3 and 4, show the results of the latency-optimal strategy for the infinite-horizon case discussed in subsection 4.2.1. The results are broadly the same as those obtained for T1's strategy discussed above, where we assume $T = 5$ days; in the interest of space we do not report the results for T2.

Robustness to volatility parameter estimates. In our setup the investor makes her model of fill ratios robust to model misspecification. We also check the robustness of the performance of the latency-optimal to (i) parameter uncertainty of the volatility model (6), and (ii) different models of the volatility of the microexchange rate.

To check robustness to parameter uncertainty in (6), we look at eight scenarios that result from shocking the parameter estimates $\hat{\kappa}$, \hat{v} , $\hat{\sigma}_v$. The shocked parameters are: $\hat{\kappa}^u = 1.2 \hat{\kappa}$, $\hat{\kappa}^d = 0.8 \hat{\kappa}$, $\hat{v}^u = 1.2 \hat{v}$, $\hat{v}^d = 1.2 \hat{v}$, $\hat{\sigma}_v^u = 1.2 \hat{\sigma}_v$, and $\hat{\sigma}_v^d = 0.8 \hat{\sigma}_v$, where $\hat{\kappa}$, \hat{v} , $\hat{\sigma}_v$ are the

estimates we obtain with the methodology outlined in section 5. Next, we compute the optimal strategy δ^* with the eight triplets of parameters resulting from all combinations of $(\hat{\kappa}, \hat{v}, \hat{\sigma}_v) \in \{\hat{\kappa}^d, \hat{\kappa}^u\} \times \{\hat{v}^d, \hat{v}^u\} \times \{\hat{\sigma}_v^d, \hat{\sigma}_v^u\}$. In Appendix B.4.1 we report the results for $\hat{\delta} = 0$. We find that the fill ratios lie in the range 98.5% to 99% and the average tick cost is approximately 1.76 ticks in all cases.

Finally, we assume that the volatility of the microexchange rate follows (7) and (8) and compute the performance of the latency-optimal strategy; see Appendix B.3. We find that the results are broadly the same as those discussed above when volatility follows (6).

7. The cost of latency. One approach to mitigate the effects of latency on the efficacy of trading strategies is to invest on hardware and colocation services to reduce the time delays associated with the different components of the life cycle of a trade; see definition of latency in section 2. As latency shortens, traders that aim at the best prices they observe in the LOB will be on target more often. This improvement in the marksmanship of the trader increases the fill ratio of the strategy, but also reduces the chances of receiving price improvements.

Another approach is to devise latency-optimal strategies that lessen the adverse effects of time delays, e.g., failing to complete trades. Although these strategies do not reduce time delays in the life cycle of trade, they are a substitute for cutting down latency when the objective is to increase the percentage of MOsLP that are filled.

Neither approach is free. The first approach requires the trader to pay fixed costs to operate in the marketplace with lower latency, which increases the efficacy of the MOsLP but reduces opportunities of obtaining price improvements. In the second approach, costs accrue to the strategy when MOsLP walk the LOB, while benefits accrue because the latency-optimal approach scoops price improvements when the market, due to latency, moves in the trader's interest.

In this section we compute the shadow price that T1 would be willing to pay to reduce her latency in the marketplace. We proceed as follows. (i) Consider hypothetical traders who employ the same naive strategy (i.e., zero discretion to walk LOB) as that of T1, but each trader acts in the market with a different fixed latency—for simplicity we assume latency is not stochastic. (ii) Compute the fill ratio achieved by each hypothetical trader's naive strategy. (iii) Map levels of latency to fill ratios. (iv) Find the hypothetical trader whose fill ratio is the same as that achieved by T1's latency-optimal strategy developed above, and compute the price improvements of the hypothetical trader's strategy. (v) Compute the difference of the latency-optimal costs of walking the LOB and its price improvements. Thus, the difference between the costs in (iv) and (v) is the shadow price of latency that T1 would be willing to pay to employ a naive strategy (no discretion to walk LOB), which obtains the same fill ratio as that of the latency-optimal strategy.

7.1. Latency and fill ratios. Recall that the attributes we have for each order include type of trade, volume of the trade, price limit, and direction (buy or sell) of the trade. We also have the time-stamp of when the order is processed by the exchange (this time is denoted by t_3 in our discussion of latency in section 2 above). However, we do not know the time it took the trader to process information and make a decision before instructing the exchange, or the time when the trader sent the orders to the exchange—all of which affect the latency of the trader in the marketplace.

To map latency to fill ratios we devise fifty hypothetical traders who are clones of T1. The only difference between T1 and the hypothetical traders is that each clone has a different latency in the market. We denote each hypothetical trader by T_j with $j = 0, 1, 2, \dots, 50$, and assume that the latency of T_j is j ms, i.e., $t_3 - t_0 = j$ ms. Note that although it is impossible for a trader to operate in the marketplace with a latency of 0 ms, we include it in our list of hypothetical traders for comparison purposes.

In our approach we assume that the clones employ the same strategy as T1 regardless of their latency. However, this may not be the case because if T1's latency was lower than her actual latency, her trading strategy could be different from the one she employed during the period we analyze. For example, T1 may not be attempting potentially profitable trades because she acknowledges that her latency is too long to successfully complete trades that require more speed.

We use T_{10} to illustrate how we map a latency of 10 ms to a measure of fill ratio—the fill ratio of the other hypothetical traders is calculated in the same way. The fill ratio is computed as the percentage of filled orders in one trading day—this coincides with definition F_n , provided in (1), where n is the number of trades in the day. We assume that T_{10} replaces T1 in the market. Thus, T_{10} sends orders that have the same direction, volume, type (i.e., FoK), and discretion to walk the LOB, as those of T1, and assume the orders from T_{10} reach the LOB an instant before T1's trade was filled (for which we have a time-stamp), so the hypothetical trader and T1 face the same market conditions. Therefore, orders sent by T_{10} and T1 only differ in that T_{10} attempts to fill the trade at the best quote streamed by the exchange 10 ms before the exchange receives the order.

For example, to compute the fill ratio that T_{10} would have obtained on 5 January 2017, we employ the time-stamps (when the exchange processes the order) of the 3,043 orders sent by T1 on that day and compute the fill ratio that T_{10} would have achieved. We repeat this approach for each hypothetical trader on the same day and depict the results in the left-hand panel of Figure 5. As expected, there is a (monotonic) decreasing relationship between latency and fill ratio. In one extreme, when latency is zero, T_0 would have obtained a fill ratio of 100%. At the other extreme T_{50} 's fill ratio is 89%.

In the figure, the square marker shows the point where a latency of 27 ms is mapped to 90%, which is the fill ratio obtained by T1's naive strategy. The diamond marker in the upper-left corner of the figure shows the latency of a hypothetical trader who achieved a fill ratio of 99%. Note that this fill ratio corresponds (approximately) to the fill ratio T1 would have obtained ($F_{3043}^{\delta^*} = 98.73\%$) had she employed the latency-optimal strategy with targets $\hat{\delta} = 0$ and $\hat{F} = 100\%$. Thus, we say that, on 5 January 2017, T1's implied latency is 27 ms and the implied latency of the optimal strategy is 3 ms.

We repeat the analysis for each day between 5 December 2016 and 31 March 2017. The right panel of Figure 5 shows the implied latencies of T1's naive strategy (solid line) and that of the latency-optimal strategy with target $\hat{\delta} = 0$ (dotted line). The variations in “implied” latency are due to a number of factors, including the volatility of the best quotes which, as discussed above, has a considerable effect on fill ratios. For ease of presentation the y -axis is capped at 50 ms—there were two days when T1's implied latency was over 50 ms.

Table 9 summarizes the results. The top panel of the table shows the mean, standard deviation, and median of the latency implied by the naive strategy and by the latency-optimal

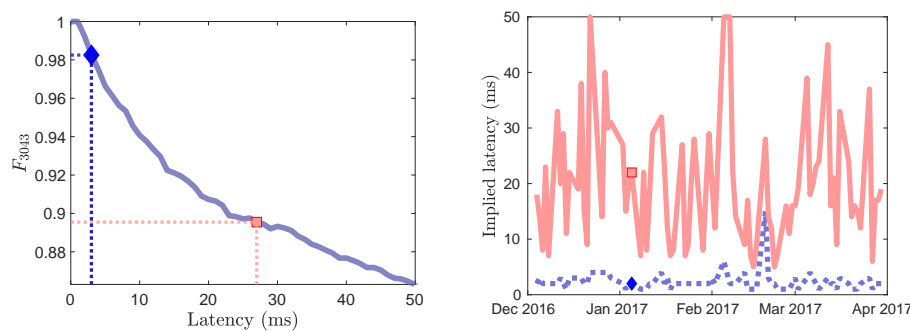


Figure 5. Left: Map of latency to fill ratio employing T1's trade time-stamps and LOB information on 5 January 2017. Square marker is T1's implied latency. Diamond marker is the implied latency of the optimal strategy. Right: implied latency for T1 and implied latency for latency-optimal strategy employing T1's trade time-stamps and LOB information during the period 5 December 2016 to 31 March 2017.

Table 9

Implied latency for T1. Period is 5 December 2016 to 31 March 2017, currency pair USD/JPY.

	Naive strategy	Optimal strategy δ^* target				$\delta_t = \infty$ $\forall t$
		$\hat{\delta} = 0$	$\hat{\delta} = 1$	$\hat{\delta} = 2$	$\hat{\delta} = 3$	
Latency						
mean	20.02	2.34	2.01	1.70	1.53	0
std	11.31	1.72	1.57	0.89	0.80	0
median	18.00	2.00	2.00	1.00	1.00	0
F_N						
mean	91.9%	98.9%	99.2%	99.3%	99.5%	100%
std	2.9%	0.4%	0.4%	0.4%	0.3%	0%
median	92.2%	98.9%	99.2%	99.4%	99.5%	100%

strategies with various discretion targets $\hat{\delta}$. The bottom panel provides summary statistics of the daily fill ratio obtained by the strategies. Here, the fill ratio F_N refers to the percentage of orders that were filled each day, and the table shows the mean, standard deviation, and median of the daily percentage of fills. For the latency-optimal strategy the mean fill is between 98.9% and 99.5% and for T1's naive strategy the mean fill is 91.0%. We observe that the standard deviation of the fill ratio of the optimal strategy is also considerably lower than that of the naive strategy. Therefore, the optimal strategy achieves much higher fill ratios than the naive strategy, and produces more stable results because the optimal strategy depends on the volatility of the microexchange rate.

7.2. Shadow price of latency. The shadow price of latency is the price that T1 would be willing to pay to reduce her latency, so a naive and the latency-optimal strategies obtain the same fill ratio. To compute this price we need the following: (i) Latency-optimal cost of walking the LOB and price improvements, and (ii) price improvements received by the hypothetical trader with latency equal to that implied by the latency-optimal strategy.

Table 10

Shadow price to reduce latency reported in JPY. Period is 5 December 2016 to 31 March 2017, currency pair USD/JPY.

	Naive strategy	Optimal strategy δ^* target				$\delta_t = \infty$ $\forall t$
		$\hat{\delta} = 0$	$\hat{\delta} = 1$	$\hat{\delta} = 2$	$\hat{\delta} = 3$	
Implied latency (ms)	20.02	2.34	2.01	1.70	1.53	0
Walk LOB cost price improvement	0 3,938,559	6,528,965 3,938,559	6,807,416 3,938,559	6,999,804 3,938,559	7,232,545 3,938,559	8,914,046 3,938,559
Price improvement with reduced latency	3,938,559	384,035	291,070	179,875	135,744	0
Shadow price to reduce latency	0	2,974,441	3,159,927	3,241,120	3,429,730	4,975,487

For example, on 5 January 2017, the latency-optimal strategy incurs 92,771 JPY in costs from walking the LOB and receives 69,305 JPY in price improvements. On the other hand, the naive strategy of the hypothetical trader with latency 3 ms, receives 7,592 JPY in price improvements. Thus, on 5 January 2017, T1 would be willing to pay 31,058 JPY to reduce her latency from 22 ms to 3 ms. We repeat the analysis for all trading days between 5 December 2016 and 31 March 2017 and compute T1's shadow price of latency for each day. Table 10 contains summary statistics of the results.

The table shows that T1 would be willing to pay approximately $3 \cdot 10^6$ JPY (around 30,000 USD) to reduce her latency from 20.02 ms to 2.34 ms. The upper bound of approximately $5 \cdot 10^6$ JPY (around 50,000 USD) for the shadow price of latency is given by the last column and last row of the table, where we show the costs of always employing market orders. These shadow prices are, according to market participants, lower than colocation, hardware, and other costs that traders would have to incur to reduce latency in the life cycle of a trade.

8. Conclusions. In this paper, we showed how to compute the price that FX traders are willing to pay to reduce their latency in the marketplace. To achieve this we developed a latency-optimal strategy to map the costs of the strategy to a reduction in latency.

In particular, we employed a set of proprietary data of FX transactions to show how latency affects the percentage of liquidity taking orders filled by a trader's strategy. We analyzed the trades of two liquidity takers over a four-month period and showed that approximately 10% of their orders were not filled due to latency.

To calculate the shadow price of latency of traders we derived a strategy that specifies the limit price of the market orders sent to the exchange. The price limit for each liquidity taking order balances the trade-off between the costs from walking the LOB and how far is the fill ratio from the trader's desired target. The trader's optimal strategy specifies the price limit for each transaction depending on the proportion of orders that have been filled, how far is the strategy from the target fill ratio, the cost of walking the LOB, and the volatility of the exchange rate. For example, everything else being equal, the optimal price limit to walk the LOB is higher (resp., lower) when the volatility of the exchange rate is above (resp., below) its long-term level.

We showed the performance of the latency-optimal strategy using trade data for two FX traders. We computed the costs incurred by the latency-optimal strategy to increase fill ratios and showed that the costs of the strategy are lower than the costs of returning to the market 20 ms or 100 ms later to fill a missed trade. Moreover, we showed that the latency-optimal strategy obtains similar fill ratios for different levels of volatility. This is in contrast to the performance of the naive strategy where the majority of the missed trades occur when the volatility of the exchange rate is high.

Finally, we employed the proprietary data set to build a function that maps latency of traders to fill ratios. We used this mapping to compute the shadow price of latency that traders are willing to pay to reduce latency in the marketplace. We computed the shadow price a particular trader is willing to pay to reduce her latency by investing in hardware, colocation, and other services. Our results showed that the trader would be better off employing the latency-optimal strategy developed here, instead of investing in hardware and colocation services to reduce latency.

Appendix A. Proofs and further results.

A.1. System of ODEs satisfied by ansatz (19). Here, we substitute the ansatz (19) into the PDE (18) to obtain a system of ODEs that we solve explicitly—we use these solutions in Theorem 1. Insert ansatz (19) into HJBI (18) and collect terms to obtain the following coupled system of six ODEs:

$$h'_0(t) + (b\bar{v} + a\hat{\delta} + \lambda\bar{F})h_1(t) + \sigma_F^2 h_2(t) + \kappa\bar{v}h_3(t) + \frac{1}{2} \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) (h_1(t))^2 - \phi_F \hat{F}^2 = 0, \quad (29)$$

$$h'_1(t) + (2b\bar{v} + 2\lambda\bar{F} + 2a\hat{\delta})h_2(t) - \lambda h_1(t) + \kappa\bar{v}h_5(t) \quad (30)$$

$$+ 2 \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) h_1(t)h_2(t) + 2\phi_F \hat{F} = 0, \quad (31)$$

$$h'_2(t) - 2\lambda h_2(t) + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) (h_2(t))^2 - \phi_F = 0, \quad (32)$$

$$h'_3(t) - b h_1(t) + (b\bar{v} + a\hat{\delta} + \lambda\bar{F})h_5(t) + 2\kappa\bar{v}h_4(t) - \kappa h_3(t) + \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) h_1(t)h_5(t) = 0, \quad (33)$$

$$h'_4(t) - b h_5(t) + (\sigma_v^2 - 2\kappa)h_4(t) + \frac{1}{2} \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) (h_5(t))^2 = 0, \quad (34)$$

$$h'_5(t) - 2b h_2(t) - (\lambda + \kappa)h_5(t) + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi\sigma_F^2 \right) h_2(t)h_5(t) = 0$$

with terminal conditions

$$h_0(T) = 0, h_1(T) = 0, h_2(T) = 0, h_3(T) = 0, h_4(T) = 0, h_5(T) = 0.$$

Here the notation $h'(t)$ is shorthand for $dh(t)/dt$. Next we provide solutions for h_2 , h_5 , and h_1 , which are the functions required to write the optimal control in closed form.

Explicit solution for h_2 . The function h_2 satisfies the Riccati ODE in (31) with terminal condition $h_2(T) = 0$ and its solution depends on the sign of the coefficient $\varphi \sigma_F^2 - a^2/2 \phi_\delta$. Therefore, we have the following:

Case 1: $\varphi < a^2/2 \sigma_F^2 \phi_\delta$, then

$$h_2(t) = c_1 \tanh(c_2 t + c_3) + c_4$$

with

$$\alpha_1 = \frac{a^2}{2 \phi_\delta} - \varphi \sigma_F^2, \quad c_1 = \frac{c_2}{2 \alpha_1}, \quad c_2 = -\sqrt{2 \alpha_1 \phi_F + \lambda^2},$$

$$c_3 = \operatorname{atanh}\left(-\frac{\lambda}{c_2}\right) + T \sqrt{2 \alpha_1 \phi_F + \lambda^2}, \quad c_4 = \frac{\lambda}{2 \alpha_1}.$$

The inequality $2 \alpha_1 \phi_F + \lambda^2 > 0$ follows from $\varphi < a^2/2 \sigma_F^2 \phi_\delta$, thus, c_2 is well defined.

Case 2: $\varphi = a^2/2 \sigma_F^2 \phi_\delta$, then

$$h_2(t) = \frac{\phi_F}{2 \lambda} \left(e^{-2 \lambda (T-t)} - 1 \right).$$

Case 3: $a^2/2 \sigma_F^2 \phi_\delta < \varphi$, then

$$h_2(t) = \tilde{c}_1 \tan(\tilde{c}_2 t + \tilde{c}_3) + \tilde{c}_4$$

with

$$\tilde{\alpha}_1 = \varphi \sigma_F^2 - \frac{a^2}{2 \phi_\delta}, \quad \tilde{c}_1 = \frac{\tilde{c}_2}{2 \tilde{\alpha}_1}, \quad \tilde{c}_2 = -\sqrt{2 \tilde{\alpha}_1 \phi_F - \lambda^2},$$

$$\tilde{c}_3 = -\operatorname{atan}\left(-\frac{\lambda}{\tilde{c}_2}\right) + T \sqrt{2 \tilde{\alpha}_1 \phi_F - \lambda^2}, \quad \tilde{c}_4 = \frac{\lambda}{2 \tilde{\alpha}_1},$$

when $(\lambda^2/\phi_F + a^2/\phi_\delta)/2 \phi_F^2 > \varphi$. We do not solve for h_2 when $(\lambda^2/\phi_F + a^2/\phi_\delta)/2 \phi_F^2 \geq \varphi > a^2/2 \sigma_F^2 \phi_\delta$ —throughout this paper we employ the solution of the ODE discussed in Case 1.

We use the integrating-factor technique to write

$$h_5(t) = -e^{\int_t^T (2 \alpha_1 h_2(u) - \kappa - \lambda) du} \int_t^T 2 b h_2(s) e^{-\int_s^T (2 \alpha_1 h_2(u) - \kappa - \lambda) du} ds$$

and, similarly, we write

$$h_1(t) = e^{\int_t^T (2 \alpha_1 h_2(u) - \lambda) du} \int_t^T (\kappa \bar{v} h_5(s) + e_1 h_2(s) + e_2) e^{-\int_s^T (2 \alpha_1 h_2(u) - \lambda) du} ds,$$

where $e_1 = 2 b \bar{v} + 2 \lambda \bar{F} + 2 a \hat{\delta}$ and $e_2 = 2 \phi_F \hat{F}$.

From now on we focus on Case 1 to obtain closed-form formulas for h_5 and h_1 ; Cases 2 and 3 are similar.

Explicit solution for h_5 . Function h_5 can be explicitly written as

$$h_5(t) = -\frac{e^{-\kappa(T-t)}}{\cosh(t c_2 + c_3)} \times \left[\left(\frac{e^{d_2(T-t)+d_3} - e^{d_3}}{d_2} \right) b(c_1 + c_4) + \left(\frac{e^{d_1(T-t)-d_3} - e^{-d_3}}{d_1} \right) b(c_4 - c_1) \right],$$

where the constants $d_1 = \kappa + c_2$, $d_2 = \kappa - c_2$, and $d_3 = c_3 + T c_2$.

Explicit solution for h_1 . Function h_1 can be solved explicitly to obtain

$$\begin{aligned} h_1(t) = & \frac{-\kappa \bar{v}}{\cosh(c_2 t + c_3)} \\ & \times \left[\frac{b(c_1 + c_4) e^{d_3}}{d_2} \left(\frac{e^{-\kappa(T-t)} - 1}{\kappa} - \frac{e^{-c_2(T-t)} - 1}{c_2} \right) \right. \\ & \quad \left. + \frac{b(c_4 - c_1) e^{-d_3}}{d_1} \left(\frac{e^{c_2(T-t)} - 1}{c_2} + \frac{e^{-\kappa(T-t)} - 1}{\kappa} \right) \right] \\ & + \frac{e_1 c_1}{c_2 \cosh(c_2 t + c_3)} (\cosh(c_2 T + c_3) - \cosh(c_2 t + c_3)) \\ & + \frac{e_1 c_4 + e_2}{c_2 \cosh(c_2 t + c_3)} (\sinh(c_2 T + c_3) - \sinh(c_2 t + c_3)). \end{aligned}$$

Finally, the functions h_0 and h_3 are obtained explicitly in a similar way using an integrating-factor technique. Note that to obtain the optimal controls explicitly we only require h_1 , h_2 , h_5 .

A.2. Proof of Theorem 1. Here we provide a class of controls that produces strong solutions for F^δ and v ; see Lemma 1. We employ Lemma 1 to prove the admissibility of the control δ^* (see Corollary 1). Then we prove that the control x^* is also admissible and conclude this section with the verification of the value function.

Lemma 1. Let $\delta_t = \mu(t, F_t, v_t)$ with $\mu \in \mathcal{M}_{0,T}$. Then, for all $t \in [0, T]$ there exist strong solutions to the SDEs of F^δ and v .

Proof. Let $\delta_t = \mu(t, F_t, v_t)$ for $\mu \in \mathcal{M}_{0,T}$. Thus, there exists $K \in \mathbb{R}$ such that for all $t \in [0, T]$, $\mathbf{y}, \mathbf{z} \in \mathbb{R}^2$,

$$(35) \quad |\mu(t, \mathbf{y}) - \mu(t, \mathbf{z})| \leq K |\mathbf{y} - \mathbf{z}| \quad \text{and} \quad |\mu(t, \mathbf{y})|^2 \leq K^2 (1 + |\mathbf{y}|^2).$$

Let $\mathbf{Z}_t^\delta = (F_t^\delta, v_t)^\top$, thus

$$d\mathbf{Z}_t^\delta = \left(A(t) + B(t) \mathbf{Z}_t^\delta + \boldsymbol{\mu}(t, \mathbf{Z}_t^\delta) \right) dt + C_t d\mathbf{W}_t$$

with $A(t)$, $B(t)$, C_t , and $\boldsymbol{\mu}$ given by

$$A(t) = \begin{bmatrix} \lambda \bar{F} + b \bar{v} \\ \kappa \bar{v} \end{bmatrix}, \quad B(t) = \begin{bmatrix} -\lambda & -b \\ 0 & -\kappa \end{bmatrix}, \quad C_t = \begin{bmatrix} \sigma_F & 0 \\ 0 & \sigma_v v_t \end{bmatrix}, \quad \boldsymbol{\mu}(t, \mathbf{Z}_t^\delta) = \begin{bmatrix} a \mu(t, \mathbf{Z}_t^\delta) \\ 0 \end{bmatrix}.$$

Define the functions $c : [0, T] \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $\sigma : [0, T] \times \mathbb{R}^2 \rightarrow \mathbb{R}^{2 \times 2}$, which are given by

$$c(t, \mathbf{Z}) = (A(t) + B(t) \mathbf{Z} + \mu(t, \mathbf{Z})) , \quad \sigma(t, \mathbf{Z}) = \begin{bmatrix} \sigma_F & 0 \\ 0 & \sigma_v v_t \end{bmatrix} .$$

Now, let $t \in [0, T]$, $\mathbf{Z}_1, \mathbf{Z}_2 \in \mathbb{R}^2$. We first verify the Lipschitz condition on the function c ,

$$\begin{aligned} |c(t, \mathbf{Z}_1) - c(t, \mathbf{Z}_2)| &= |B(t) (\mathbf{Z}_1 - \mathbf{Z}_2) + \mu(t, \mathbf{Z}_1) - \mu(t, \mathbf{Z}_2)| \\ &\leq |B(t)| |\mathbf{Z}_1 - \mathbf{Z}_2| + |\mu(t, \mathbf{Z}_1) - \mu(t, \mathbf{Z}_2)| \\ &\leq \bar{B} |\mathbf{Z}_1 - \mathbf{Z}_2| + K |\mathbf{Z}_1 - \mathbf{Z}_2| = \hat{K} |\mathbf{Z}_1 - \mathbf{Z}_2| \end{aligned}$$

for $\hat{K} = K + \bar{B}$, and $\bar{B} = \max_{t \in [0, T]} B(t)$, which is clearly attained. For the function σ , we have

$$|\sigma(t, \mathbf{Z}_1) - \sigma(t, \mathbf{Z}_2)| \leq \sigma_v |\mathbf{Z}_1 - \mathbf{Z}_2| .$$

We test the linear growth condition on the function c . We use triangle inequalities and simple bounds to obtain

$$\begin{aligned} |c(t, \mathbf{Z}_1)|^2 &= |A(t) + B(t) \mathbf{Z}_1 + \mu(t, \mathbf{Z}_1)|^2 \leq (|A(t)| + |B(t) \mathbf{Z}_1| + |\mu(t, \mathbf{Z}_1)|)^2 \\ &\leq 3 |A(t)|^2 + 3 |B(t)|^2 |\mathbf{Z}_1|^2 + 3 |\mu(t, \mathbf{Z}_1)|^2 \\ &\leq 3 \bar{A}^2 + 3 \bar{B}^2 |\mathbf{Z}_1|^2 + 3 K^2 (1 + |\mathbf{Z}_1|^2) \\ &\leq \bar{D}^2 + \bar{D}^2 |\mathbf{Z}_1|^2 + 3 K^2 (1 + |\mathbf{Z}_1|^2) \\ &= (\bar{D}^2 + 3 K^2) (1 + |\mathbf{Z}_1|^2) \end{aligned}$$

for $\bar{A} = \max_{t \in [0, T]} A(t)$, which is clearly attained. Finally, we take $\bar{D}^2 = \max \{\bar{A}^2, \bar{B}^2\}$.

On the other hand, for σ , we have

$$\begin{aligned} |\sigma(t, \mathbf{Z}_1)|^2 &\leq 2 \sigma_F^2 + 2 \sigma_v^2 |\mathbf{Z}_1|^2 \\ &\leq \bar{E}^2 (1 + |\mathbf{Z}_1|^2) \end{aligned}$$

for $\bar{E} = \max \{\sigma_F^2, \sigma_v^2\}$. Lipschitz and linear growth conditions follow by choosing

$$\bar{K} = \max \left\{ \bar{E}, \sqrt{\bar{D}^2 + 3 K^2}, \sigma_v, \hat{K} \right\} < \infty .$$

Thus the Lipschitz and linear growth conditions are satisfied; therefore, for all $t \in [0, T]$, there exists a strong solution to the SDE satisfied by \mathbf{Z}^δ , and subsequently F^δ and v_t . ■

Corollary 1. *Let $\delta \in \mathcal{A}_{t, T}$. Then, for all $t \in [0, T]$ there exist strong solutions to the SDEs of F^δ and v .*

This shows that as long as the control δ is admissible, F^δ and v have strong solutions.

Proof. Let $\delta \in \mathcal{A}_{t,T}$, thus there exists $\mu \in \mathcal{M}_{0,T}$ such that $\delta_t = \mu(t, F_t, v_t)$. By Lemma 1, for all $t \in [0, T]$ there exist strong solutions to the SDEs of F^δ and v . ■

Proof. The control δ^* is admissible.

Recall that δ^* is given by

$$\delta_t^* = \hat{\delta} + \frac{a}{2\phi_\delta} \left(h_1(t) + 2h_2(t) F_t^{\delta^*} + h_5(t) v_t \right).$$

Define $f : [0, T] \times \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ as

$$f(t, x, y) = \hat{\delta} + \frac{a}{2\phi_\delta} h_1(t) + \frac{a h_2(t)}{\phi_\delta} x + \frac{a h_5(t)}{2\phi_\delta} y,$$

so $\delta_t^* = f(t, F_t, v_t)$, and, since f is linear in x and y , and a continuous function of t , then $f \in \mathcal{M}_{0,T}$. The control δ^* is \mathcal{F} -adapted, because it is a continuous function of two \mathcal{F} -adapted processes. Then, we only need to prove that

$$\mathbb{E} \left[\int_0^T (\delta_s^*)^2 ds \right] < \infty.$$

Based on Young's inequality we find the following bound for the optimal control:

$$(\delta_t^*)^2 \leq 2 \left(\hat{\delta} + \frac{a}{2\phi_\delta} h_1(t) \right)^2 + \frac{2a^2}{\phi_\delta^2} (h_2(t))^2 (F_t^{\delta^*})^2 + \frac{a^2}{2\phi_\delta^2} (h_5(t))^2 (v_t)^2,$$

and because h_1 , h_2 , and h_5 are continuous functions in the closed interval $[0, T]$, they are bounded. Let U be an upper bound for $|h_1|$, $|h_2|$, and $|h_5|$ on $[0, T]$.

As $\delta_t^* = f(t, F_t, v_t)$ for $f \in \mathcal{M}_{0,T}$, by Lemma 1, we have that F^{δ^*} and v have strong solutions for all $t \in [0, T]$, thus

$$(36) \quad \mathbb{E} \left[\sup_{0 \leq t \leq T} (F_t^{\delta^*})^2 \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\sup_{0 \leq t \leq T} (v_t)^2 \right] < \infty.$$

Finally, using the upper bound U and the inequalities in (36), we conclude that

$$\begin{aligned} \mathbb{E} \left[\int_0^T (\delta_s^*)^2 ds \right] &\leq \mathbb{E} \left[\int_0^T 2 \left(\hat{\delta} + \frac{a}{2\phi_\delta} h_1(t) \right)^2 + \frac{2a^2}{\phi_\delta^2} (h_2(t))^2 (F_t^{\delta^*})^2 + \frac{a^2}{2\phi_\delta^2} (h_5(t))^2 (v_t)^2 \right] \\ &\leq \mathbb{E} \left[2T \left(\hat{\delta} + \frac{a}{2\phi_\delta} U \right)^2 + \frac{2a^2}{\phi_\delta^2} U^2 \sup_{0 \leq t \leq T} (F_t^{\delta^*})^2 + T \frac{a^2}{2\phi_\delta^2} U^2 \sup_{0 \leq t \leq T} (v_t)^2 \right] \\ &< \infty. \end{aligned} \quad \blacksquare$$

Proof. The control x^* is admissible.

This proof is similar to that of Theorem 1 in Cheridito, Filipović, and Kimmel (2007). Let $\delta \in \mathcal{A}_{t,T}$, then $\exists \mu \in \mathcal{M}_{t,T}$, such that $\delta_t = \mu(t, F_t, v_t)$. Recall that the optimal control x^* is given by

$$\begin{aligned} x_t^* &= \varphi \sigma_F \left(h_1(t) + 2 h_2(t) F_t^\delta + h_5(t) v_t \right) \\ &= \eta(t, F_t, v_t), \end{aligned}$$

which is adapted because it is a continuous function of adapted processes. Therefore, it remains to show that

$$Z_t = \exp \left\{ -\frac{1}{2} \int_0^t (x_u^*)^2 du - \int_0^t x_u^* dW_u^F \right\}$$

is a martingale. As the control x^* is a well-defined continuous process, Z is a well-defined local martingale with respect to \mathbb{P} and, subsequently, because Z is nonnegative, it follows that it is a supermartingale, and for it to be a martingale, we only need to prove that

$$\mathbb{E}^\mathbb{P} [Z_T] = 1.$$

Consider the function $\mu^\mathbb{Q}$ given by

$$\mu^\mathbb{Q}(t, F_t, v_t) = \mu(t, F_t, v_t) - \frac{\sigma_F}{a} \eta(t, F_t, v_t)$$

and, because the term $\eta(t, F_t, v_t)$ is continuous in time and linear in F and v , it belongs to $\mathcal{M}_{t,T}$, therefore, $\mu^\mathbb{Q} \in \mathcal{M}_{t,T}$.

Given that $\mu, \mu^\mathbb{Q} \in \mathcal{M}_{t,T}$, by Lemma 1, we know there is a strong solution for the SDEs given by

$$(37) \quad dF_t = (\lambda (\bar{F} - F_t) + a \mu(t, F_t, v_t) - b(v_t - \bar{v})) dt + \sigma_F dW_t^F,$$

$$(38) \quad d\tilde{F}_t = \left(\lambda (\bar{F} - \tilde{F}_t) + a \mu^\mathbb{Q}(t, \tilde{F}_t, v_t) - b(v_t - \bar{v}) \right) dt + \sigma_F dW_t^F.$$

Now define

$$(39) \quad \tilde{x}_t^* = \varphi \sigma_F \left(h_1(t) + 2 h_2(t) \tilde{F}_t^\delta + h_5(t) v_t \right)$$

$$(40) \quad = \eta(t, \tilde{F}_t, v_t).$$

For $n \in \mathbb{N}$, consider the stopping times τ_n and $\tilde{\tau}_n$:

$$(41) \quad \tau_n = \inf \{t > 0 \mid x_t^* \geq n\} \wedge T \quad \text{and} \quad \tilde{\tau}_n = \inf \{t > 0 \mid \tilde{x}_t^* \geq n\} \wedge T,$$

and because x^* and \tilde{x}^* are \mathbb{P} -a.s. finite, then

$$(42) \quad \lim_{n \rightarrow \infty} \mathbb{P}(\tau_n = T) = \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{\tau}_n = T) = 1.$$

Define for each $n \in \mathbb{N}$ the process

$$x_t^{(n)} = x_t^* \mathbf{1}_{\{t \leq \tau_n\}},$$

which satisfies the Novikov condition

$$\mathbb{E}^{\mathbb{P}} \left[\exp \left\{ \frac{1}{2} \int_0^T (x_u^n)^2 du \right\} \right] \leq \exp \left\{ \frac{n^2 T}{2} \right\} < \infty.$$

Then it follows that

$$\begin{aligned} \frac{d\mathbb{Q}(x^{(n)})}{d\mathbb{P}} \Big|_t &= \exp \left\{ -\frac{1}{2} \int_0^t (x_u^{(n)})^2 du - \int_0^t x_u^{(n)} dW_u^F \right\} \\ &= Z_t^{(n)} \end{aligned}$$

defines an equivalent measure to \mathbb{P} , and by Girsanov's theorem, $W_t^n = W_t^F + \int_0^t x_s^{(n)} ds$ is a Brownian motion under the measure $\mathbb{Q}(x^{(n)})$, abbreviated as \mathbb{Q}^n . Finally,

$$\begin{aligned} \mathbb{E}^{\mathbb{P}} [Z_T] &= \mathbb{E}^{\mathbb{P}} \left[\lim_{n \rightarrow \infty} Z_T \mathbf{1}_{\{\tau_n = T\}} \right] \\ &= \mathbb{E}^{\mathbb{P}} \left[\lim_{n \rightarrow \infty} Z_T^n \mathbf{1}_{\{\tau_n = T\}} \right] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}} [Z_T^n \mathbf{1}_{\{\tau_n = T\}}] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{Q}^n} [\mathbf{1}_{\{\tau_n = T\}}] \\ &= \lim_{n \rightarrow \infty} \mathbb{Q}^n (\tau_n = T), \end{aligned}$$

where the first two equalities hold because the control x^* is \mathbb{P} -a.s. finite, and by definition (41). The third equality follows from the monotone convergence theorem. The fourth equality is the change of measure imposed by Z_T^n .

Finally, starting at F_0 , the stopped process $F_{t \wedge \tau_n}$ under \mathbb{Q}^n is given by

$$\begin{aligned} F_{t \wedge \tau_n} &= F_0 + \int_0^{t \wedge \tau_n} \left(\lambda (\bar{F} - F_s) + a \mu(s, F_s, v_s) - \sigma_F x_s^{(n)} - b(v_s - \bar{v}) \right) ds + \int_0^{t \wedge \tau_n} \sigma_F dW_s^n \\ &= F_0 + \int_0^{t \wedge \tau_n} \left(\lambda (\bar{F} - F_s) + a \mu(s, F_s, v_s) - \sigma_F x_s^* - b(v_s - \bar{v}) \right) ds + \int_0^{t \wedge \tau_n} \sigma_F dW_s^n \\ &= F_0 + \int_0^{t \wedge \tau_n} \left(\lambda (\bar{F} - F_s) + a \mu^{\mathbb{Q}}(s, F_s, v_s) - b(v_s - \bar{v}) \right) ds + \int_0^{t \wedge \tau_n} \sigma_F dW_s^n \end{aligned}$$

and, similarly, the stopped process $\tilde{F}_{t \wedge \tilde{\tau}_n}$ under \mathbb{P} is given by

$$\tilde{F}_{t \wedge \tilde{\tau}_n} = F_0 + \int_0^{t \wedge \tilde{\tau}_n} \left(\lambda (\bar{F} - \tilde{F}_s) + a \mu^{\mathbb{Q}}(s, \tilde{F}_s, v_s) - b(v_s - \bar{v}) \right) ds + \int_0^{t \wedge \tilde{\tau}_n} \sigma_F dW_s^F.$$

Therefore $(\tilde{F}_{t \wedge \tilde{\tau}_n})_{0 \leq t \leq T}$, and $(F_{t \wedge \tau_n})_{0 \leq t \leq T}$ have the same distribution under \mathbb{Q}^n and \mathbb{P} , respectively, thus

$$\lim_{n \rightarrow \infty} \mathbb{Q}^n (\tau_n = T) = \lim_{n \rightarrow \infty} \mathbb{P} (\tilde{\tau}_n = T) = 1$$

and

$$\mathbb{E}^{\mathbb{P}}[Z_T] = \lim_{n \rightarrow \infty} \mathbb{Q}^n(\tau_n = T) = \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{\tau}_n = T) = 1,$$

which concludes the proof. ■

Finally, we present the proof of Theorem 1, i.e., we show that (19) is the trader's value function given in (13).

Proof. V is the value function of the trader's control problem in (13).

Let $\mathbb{Q}^* = \mathbb{Q}^*(\delta)$ be the measure inferred by $x_t^*(\delta) = \varphi \sigma_F (h_1(t) + 2h_2(t)F_t^\delta + h_5(t)v_t)$, in other words,

$$(43) \quad \frac{d\mathbb{Q}^*}{d\mathbb{P}} \Big|_T = \exp \left\{ -\frac{1}{2} \int_0^T \mathbf{x}_u^{*'} \mathbf{x}_u^* du - \int_0^T \mathbf{x}_u^* d\mathbf{W}_u \right\},$$

where $\mathbf{x}_t^* = (x_t^*, 0)$. Clearly $x_t^*(\delta)$ represents the infimum measure of (13) for an admissible δ .

The dynamics of v under \mathbb{Q}^* do not change, but the dynamics of F^{δ, \mathbb{Q}^*} do change. Specifically

$$\begin{aligned} dF_t^{\delta, \mathbb{Q}^*} = & \left(\lambda \bar{F} - \left(\lambda + 2\varphi \sigma_F^2 h_2(t) \right) F_t^{\delta, \mathbb{Q}^*} + a \delta_t + b \bar{v} \right. \\ & \left. - \left(b + \varphi \sigma_F^2 h_5(t) \right) v_t - \varphi \sigma_F^2 h_1(t) \right) dt + \sigma_F dW_t^{F, \mathbb{Q}^*}. \end{aligned}$$

As $V \in \mathcal{C}^{1,2,2}$, apply Itô's lemma under \mathbb{Q}^* , and consider $\delta \in \mathcal{A}_{t,T}$, hence

$$(44) \quad V(T, F_T^\delta, v_T) = V(t, f, v) + \int_t^T \partial_s V + \mathcal{L}^{\mathbb{Q}^*} V ds + \int_t^T \partial_v V \sigma_v v_s dW_s^{v, \mathbb{Q}^*} + \int_t^T \partial_f V \sigma_F dW_s^{F, \mathbb{Q}^*},$$

where W_s^{F*} and W_s^{v*} are standard \mathbb{Q}^* -Brownian motions. The function V is a solution of (17), and δ is an arbitrary control, thus

$$\partial_s V + \mathcal{L}^{\mathbb{Q}^*} V \leq \phi_\delta \left(\delta - \hat{\delta} \right)^2 + \phi_F \left(f - \hat{F} \right)^2 - \frac{1}{2\varphi} (x^*)^2.$$

Now, take expectations on both sides of (44) to obtain

$$\begin{aligned} 0 &= \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[V(T, F_T^\delta, v_T) \right], \\ &= V(t, f, v) + \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[\int_t^T \partial_s V + \mathcal{L}_s^{\mathbb{Q}^*} V ds \right], \\ &\leq V(t, f, v) + \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[\phi_\delta \int_t^T \left(\delta_s - \hat{\delta} \right)^2 ds + \phi_F \int_t^T \left(F_s^\delta - \hat{F} \right)^2 ds - \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds \right], \\ &= V(t, f, v) + \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[\phi_\delta \int_t^T \left(\delta_s - \hat{\delta} \right)^2 ds + \phi_F \int_t^T \left(F_s^\delta - \hat{F} \right)^2 ds \right. \\ &\quad \left. - \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds + \frac{1}{\varphi} \int_t^T x_s^* dW_s^{\mathbb{Q}^*} \right], \end{aligned}$$

therefore

$$\mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds - \frac{1}{\varphi} \int_t^T x_s^* dW_s^{\mathbb{Q}^*} \right] \leq V(t, f, v)$$

or, alternatively,

$$\begin{aligned} \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds - \frac{1}{2\varphi} \int_t^T (x_s^*)^2 ds - \frac{1}{\varphi} \int_t^T x_s^* dW_s \right] &\leq V(t, f, v), \\ \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] &\leq V(t, f, v), \end{aligned}$$

and because $V(t, f, v)$ is a bound for any $\delta \in \mathcal{A}_{t,T}$, we have that

$$\sup_{\delta \in \mathcal{A}_{t,T}} \mathbb{E}_{t,f,v}^{\mathbb{Q}^*} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q}^* | \mathbb{P}) \right] \leq V(t, f, v), \quad (45)$$

$$\sup_{\delta \in \mathcal{A}_{t,T}} \inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] \leq V(t, f, v),$$

which completes the first part of the proof.

Now, fix the strategy δ^* , then for any $\mathbb{Q} \in \mathcal{Q}$, because V is a solution of (17), and because the inf and sup can be interchanged in (17) (x and δ are decoupled in (17)), we have the inequality

$$\partial_s V + \mathcal{L}^{\mathbb{Q}} V \geq \phi_\delta (\delta^* - \hat{\delta})^2 + \phi_F (f - \hat{F})^2 - \frac{1}{2\varphi} (x)^2, \quad (46)$$

and, following the same steps as before together with (46), we have, under the measure \mathbb{Q} ,

$$\mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s^* - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^{\delta^*} - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] \geq V(t, f, v),$$

and because this is true for all $\mathbb{Q} \in \mathcal{Q}$, we write

$$\inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s^* - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^{\delta^*} - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] \geq V(t, f, v), \quad (47)$$

$$\sup_{\delta \in \mathcal{A}_{t,T}} \inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{t,f,v}^{\mathbb{Q}} \left[-\phi_\delta \int_t^T (\delta_s - \hat{\delta})^2 ds - \phi_F \int_t^T (F_s^\delta - \hat{F})^2 ds + \mathcal{H}_{t,T}(\mathbb{Q} | \mathbb{P}) \right] \geq V(t, f, v),$$

which completes the second part of the proof.

Combine (45) and (47) to obtain the desired result. ■

Appendix B.

B.1. Infinite-horizon value function. We solve the infinite-horizon problem for the trader, which is the analogue to (13) with $T \rightarrow \infty$ and here we include an exponential discount factor $\beta > 0$, so the problem is well defined. Thus, the trader's value function is

$$(48) \quad H(f, v) = \sup_{\delta \in \mathcal{A}_{0,\infty}} \inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{f,v}^{\mathbb{Q}} \left[- \int_0^\infty e^{-\beta s} \left(\phi_\delta (\delta_s - \hat{\delta})^2 ds + \phi_F (F_s^\delta - \hat{F})^2 \right) ds + \mathcal{H}^\beta(\mathbb{Q} | \mathbb{P}) \right],$$

where we now define the penalty as

$$\mathcal{H}^\beta(\mathbb{Q} | \mathbb{P}) = \frac{1}{\varphi} \log \left(\exp \left(-\frac{1}{2} \int_0^\infty e^{-\beta s} (x_s^F)^2 ds - \int_0^\infty e^{-\beta s} x_s^F dW_s^F \right) \right),$$

which is the infinite-horizon version of (11) and (12).

As above, the value function satisfies the HJBI

$$(49) \quad -\beta H + \sup_\delta \inf_x \left((\lambda(\bar{F} - f) + a\delta - b(v - \bar{v}) - \sigma_F x) H_f + \frac{1}{2} \sigma_F^2 H_{ff} + \frac{1}{2} \varphi x^2 + \kappa(\bar{v} - v) H_v + \frac{1}{2} v^2 \sigma_v^2 H_{vv} - \phi_\delta (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0,$$

which, after substituting the feedback controls, reduces to

$$-\beta H + (\lambda(\bar{F} - f) - b(v - \bar{v})) H_f + \frac{1}{2} \sigma_F^2 H_{ff} + \frac{a^2 V_f^2}{4 \phi_\delta} - \frac{1}{2} \varphi H_f^2 \sigma_F^2 + a \hat{\delta} H_f + \kappa(\bar{v} - v) H_v + \frac{1}{2} v^2 \sigma_v^2 H_{vv} - \phi_F (f - \hat{F})^2 = 0.$$

We use a similar ansatz to that in (19):

$$(50) \quad H(f, v) = k_0 + k_1 f + k_2 f^2 + k_3 v + k_4 v^2 + k_5 f v,$$

where $k_0, k_1, k_2, k_3, k_4, k_5$ are constants. Substitute the ansatz in the HJBI and collect terms to obtain the following system of equations:

$$\begin{aligned} -\beta k_0 + (b\bar{v} + a\hat{\delta} + \lambda\bar{F}) k_1 + \sigma_F^2 k_2 + \kappa\bar{v} k_3 + \frac{1}{2} \left(\frac{a^2}{2\phi_\delta} - \varphi \sigma_F^2 \right) (k_1)^2 - \phi_F \hat{F}^2 &= 0, \\ -\beta k_1 + (2b\bar{v} + 2\lambda\bar{F} + 2a\hat{\delta}) k_2 - \lambda k_1 + \kappa\bar{v} k_5 + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi \sigma_F^2 \right) k_1 k_2 + 2\phi_F \hat{F} &= 0, \\ -\beta k_2 - 2\lambda k_2 + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi \sigma_F^2 \right) (k_2)^2 - \phi_F &= 0, \\ -\beta k_3 - b k_1 + (b\bar{v} + a\hat{\delta} + \lambda\bar{F}) k_5 + 2\kappa\bar{v} k_4 - \kappa k_3 + \left(\frac{a^2}{2\phi_\delta} - \varphi \sigma_F^2 \right) k_1 k_5 &= 0, \\ -\beta k_4 - b k_5 + (\sigma_v^2 - 2\kappa) k_4 + \frac{1}{2} \left(\frac{a^2}{2\phi_\delta} - \varphi \sigma_F^2 \right) (k_5)^2 &= 0, \\ -\beta k_5 - 2b k_2 - (\lambda + \kappa) k_5 + 2 \left(\frac{a^2}{2\phi_\delta} - \varphi \sigma_F^2 \right) k_2 k_5 &= 0. \end{aligned}$$

Table 11

Performance of optimal strategy for T1 for finite- and infinite-horizon cases. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, $\beta = 0.01$, and $T = 5$ days in the finite-horizon case. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies				
		Extra fills	Extra volume $\times 10^6$	Cost in JPY	Cost 20 ms later in JPY	Cost 100 ms later in JPY
(a)	$\delta_t = \infty \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
$\delta^*, \hat{\delta} = 0$					$\delta_t = \infty \forall t$	$\delta_t = \infty \forall t$
Value function						
(b)	Finite horizon	8,318	3,718	6,528,965	7,459,198	9,601,222
(c)	Infinite horizon	8,317	3,718	6,527,965	7,458,198	9,600,222

Table 12

Performance of optimal strategy for T1 for finite- and infinite-horizon cases. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (resp., 7) are the proportional increase of the values in columns 5 and 6 (resp., 5 and 7) in Table 3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, $\beta = 0.01$, and $T = 5$ days in the finite-horizon case. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies					
		Extra fills	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks	Cost increase 20 ms later %	Cost increase 100 ms later %
(a)	$\delta_t = \infty \ \forall t$	9,483	2.11	2.33	2.88	10%	36%
$\delta^*, \hat{\delta} = 0$							
Value function							
(b)	Finite horizon	8,318	1.76	2.01	2.58	14%	47%
(c)	Infinite horizon	8,317	1.76	2.01	2.58	14%	47%

In Tables 11 and 12, which are the analogues of Tables 3 and 4, we compare the performance of the strategy in the finite- and infinite-horizon cases. The results are broadly the same—recall that in the finite-horizon case we assume that every day the latency-optimal strategy assumes $T = 5$.

B.2. Alternative models for volatility. If we model the volatility of the microexchange rate with the Ornstein–Uhlenbeck process (7), then the trader’s value function satisfies the HJBI

$$\begin{aligned}
 V_t + \sup_{\delta} \inf_x \left((\lambda(\bar{F} - f) + a\delta - b(v - \bar{v}^{\text{OU}}) - \sigma_F x) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{1}{2\varphi} x^2 \right. \\
 \left. + \kappa^{\text{OU}} (\bar{v}^{\text{OU}} - v) V_v + \frac{1}{2} v^2 (\sigma_v^{\text{OU}})^2 V_{vv} - \phi_\delta (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0.
 \end{aligned}
 \tag{51}$$

Table 13

Performance of optimal strategy for T1 for three volatility models of the microexchange rate. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies				
		Extra fills	Extra volume $\times 10^6$	Cost in JPY	Cost 20 ms later in JPY	Cost 100 ms later in JPY
(a)	$\delta_t = \infty \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
$\delta^*, \hat{\delta} = 0$						
Volatility model						
(b)	(6)	8,318	3,718	6,528,965	7,459,198	9,601,222
(c)	(7)	8,434	3,792	6,711,593	7,663,781	9,868,183
(d)	(8)	8,434	3,713	6,529,054	7,465,914	9,593,287

This HJBI is similar to (17), where the three parameters κ , \bar{v} , σ_v in (17) become κ^{OU} , \bar{v}^{OU} , σ_v^{OU} , respectively, and the term $\frac{1}{2} v^2 \sigma_v^2 V_{vv}$ becomes $\frac{1}{2} (\sigma_v^{\text{OU}})^2 V_{vv}$.

Clearly, the optimal controls that result from (51) have the same functional form as those we derived above in Proposition 1. To solve (51) we employ ansatz (19) and derive a system of ODEs similar to the one shown in Appendix A.1.

When we employ Heston's volatility model (8), the trader's value function satisfies the HJBI

$$(52) \quad V_t + \sup_{\delta} \inf_x \left((\lambda (\bar{F} - f) + a \delta - b(v - \bar{v}^H) - \sigma_F x) V_f + \frac{1}{2} \sigma_F^2 V_{ff} + \frac{1}{2} \varphi x^2 + \kappa^H (\bar{v}^H - v) V_v + \frac{1}{2} v^2 (\sigma_v^b)^2 V_{vv} - \phi_\delta (\delta - \hat{\delta})^2 - \phi_F (f - \hat{F})^2 \right) = 0.$$

As before, this HJBI is similar to (17), where the three parameters κ , \bar{v} , σ_v are now κ^H , \bar{v}^H , σ_v^H , and the term $\frac{1}{2} v^2 \sigma_v^2 V_{vv}$ is now $\frac{1}{2} (\sigma_v^H)^2 v V_{vv}$.

The optimal controls that result from (52) have the same functional form as those we derived above in Proposition 1. In Appendix B.3 we discuss the performance of the strategy for these two additional models of volatility.

B.3. Performance of the optimal strategy for alternative models of volatility. Table 13 presents the performance of the optimal strategy for three volatility models of the microexchange rate. We observe small differences with respect to the baseline model in row (b). Finally, Table 14 summarizes the results of Table 13 using average figures.

B.4. Parameter estimates and sensitivity analysis. Recall that we estimate daily parameters of the model for volatility during the period from December 2016 to March 2017, i.e., 83 trading days. In total, we employ around 170,000 observations of the microexchange rate because we sample prices every 500 milliseconds, and around 2,800 observations for the fill ratio because we sample fills and misses every 30 seconds. Here, we use measure F_n for T1,

Table 14

Performance of optimal strategy for T1 under different models for the volatility of the microexchange rate. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (resp., 7) are the proportional increase of the values in columns 5 and 6 (resp., 5 and 7) in Table 3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies					
		Extra fills	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks	Cost increase 20 ms later %	Cost increase 100 ms later %
(a)	$\delta_t = \infty \forall t$	9,483	2.11	2.33	2.88	10%	36%
$\delta^*, \hat{\delta} = 0$						$\delta_t = \infty \forall t$	$\delta_t = \infty \forall t$
Parameters							
(b)	(6)	8,318	1.76	2.01	2.58	14%	47%
(c)	(7)	8,395	1.77	2.02	2.60	14%	47%
(d)	(8)	8,307	1.76	2.01	2.58	14%	47%

Table 15

Number of days with at least one negative parameter estimate (out of 83 days).

Volatility model	(6): $\hat{\kappa}, \hat{v}, \hat{\sigma}_v$	(7): $\hat{\kappa}^{\text{OU}}, \hat{v}^{\text{OU}}, \hat{\sigma}_v^{\text{OU}}$	(8): $\hat{\kappa}^{\text{H}}, \hat{v}^{\text{H}}, \hat{\sigma}_v^{\text{H}}$
Days with at least one estimated negative parameter	1	7	7

Table 16

Number of days with at least one negative parameter estimate (out of 83 days).

	\hat{a}	\hat{b}	$\hat{\lambda}$	\hat{F}	$\hat{\sigma}_F$
Days with at least one estimated negative parameter	1	12	1	1	1

and measure F_n^{AV} for T2. In our analysis, the OLS regressions are not restricted to obtain nonnegative parameter estimates. In Tables 15 and 16 we report the number of occurrences when a parameter estimate is negative. In these few cases, to run the strategy, we replace the set of parameters with those of the previous day.

Table 16 shows the number of days in which the parameter estimates of the dynamics of fill ratio dynamics are negative.

B.4.1. Sensitivity to parameter estimates. In Tables 17 and 18 we show the performance of the optimal strategy δ^* when we stress the parameter estimates of the volatility model (6). For each day, we estimate the parameters $\hat{\kappa}$, \hat{v} , and $\hat{\sigma}_v$ as described in section 5 and analyze eight scenarios, where we shock each parameter by multiplying it by 1.2 and 0.80; see subsection 6.3. Next, we compute the optimal strategy δ^* with the eight triplets of parameters resulting from all the different combinations of $(\hat{\kappa}, \hat{v}, \hat{\sigma}_v) \in \{\hat{\kappa}^d, \hat{\kappa}^u\} \times \{\hat{v}^d, \hat{v}^u\} \times \{\hat{\sigma}_v^d, \hat{\sigma}_v^u\}$.

Tables 17 and 18 are the analogue of Tables 3 and 4.

Table 17

Sensitivity analysis (volatility parameters) of the performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are in JPY. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies				
		Extra fills	Extra volume $\times 10^6$	Cost in JPY	Cost 20 ms later in JPY	Cost 100 ms later in JPY
(a)	$\delta_t = \infty \forall t$	9,483	4,222	8,914,046	9,841,755	12,158,945
$\delta^*, \hat{\delta} = 0$					$\delta_t = \infty \forall t$	$\delta_t = \infty \forall t$
Parameters						
(b)	$\hat{\kappa}, \hat{v}, \hat{\sigma}_v$	8,318	3,718	6,528,965	7,459,198	9,601,222
(b.1)	$\hat{\kappa}^d, \hat{v}^d, \hat{\sigma}_v^d$	8,434	3,765	6,633,578	7,579,730	9,741,145
(b.2)	$\hat{\kappa}^d, \hat{v}^d, \hat{\sigma}_v^d$	8,434	3,765	6,633,578	7,579,730	9,741,145
(b.3)	$\hat{\kappa}^d, \hat{v}^u, \hat{\sigma}_v^d$	8,225	3,688	6,507,051	7,421,755	9,541,736
(b.4)	$\hat{\kappa}^d, \hat{v}^u, \hat{\sigma}_v^d$	8,225	3,688	6,507,051	7,421,755	9,541,736
(b.5)	$\hat{\kappa}^u, \hat{v}^d, \hat{\sigma}_v^d$	8,041	3,757	6,621,327	7,563,658	9,722,004
(b.6)	$\hat{\kappa}^u, \hat{v}^d, \hat{\sigma}_v^d$	8,041	3,757	6,621,327	7,563,658	9,722,004
(b.7)	$\hat{\kappa}^u, \hat{v}^u, \hat{\sigma}_v^d$	7,839	3,683	6,494,442	7,408,846	9,527,502
(b.8)	$\hat{\kappa}^u, \hat{v}^u, \hat{\sigma}_v^d$	7,839	3,683	6,494,442	7,408,846	9,527,502

Table 18

Sensitivity analysis (of volatility parameter) of the performance of optimal strategy for T1. The total number of trade attempts is 110,458. Cost-related quantities are weighted by volume and expressed in ticks, 10^{-3} JPY. The metrics in column 6 (resp., 7) are the proportional increase of the values in columns 5 and 6 (resp., 5 and 7) in Table 3. Parameters: $\phi_\delta = 2.5 \times 10^{-5}$, $\phi_F = 5$, $\varphi = 10^{-3}$, $\hat{F} = 100\%$, and $T = 5$ days. Period is from 5 December 2016 to 31 March 2017, currency pair USD/JPY.

		Missed by naive, filled by alternative strategies						
		Extra fills	Avg. cost in ticks	Avg. cost 20 ms later in ticks	Avg. cost 100 ms later in ticks	Cost increase 20 ms later %	Cost increase 100 ms later %	
(a)	$\delta_t = \infty \ \forall t$	9,483	2.11	2.33	2.88	10%	36%	
$\delta^*, \hat{\delta} = 0$							$\delta_t = \infty \ \forall t$	$\delta_t = \infty \ \forall t$
Parameters								
(b)	$\hat{\kappa}, \hat{v}, \sigma_v$	8,318	1.76	2.01	2.58	14%	47%	
(b.1)	$\hat{\kappa}^d, \hat{v}^d, \sigma_v^d$	8,434	1.76	2.01	2.59	14%	47%	
(b.2)	$\hat{\kappa}^d, \hat{v}^d, \sigma_v^d$	8,434	1.76	2.01	2.59	14%	47%	
(b.3)	$\hat{\kappa}^d, \hat{v}^u, \sigma_v^d$	8,225	1.76	2.01	2.59	14%	47%	
(b.4)	$\hat{\kappa}^d, \hat{v}^u, \sigma_v^d$	8,225	1.76	2.01	2.59	14%	47%	
(b.5)	$\hat{\kappa}^u, \hat{v}^d, \sigma_v^d$	8,041	1.76	2.01	2.59	14%	47%	
(b.6)	$\hat{\kappa}^u, \hat{v}^d, \sigma_v^d$	8,041	1.76	2.01	2.59	14%	47%	
(b.7)	$\hat{\kappa}^u, \hat{v}^u, \sigma_v^d$	7,839	1.76	2.01	2.59	14%	47%	
(b.8)	$\hat{\kappa}^u, \hat{v}^u, \sigma_v^d$	7,839	1.76	2.01	2.59	14%	47%	

REFERENCES

- R. ALMGREN (2012), *Optimal trading with stochastic liquidity and volatility*, SIAM J. Financial Math., 3, pp. 163–181.
- W. BARGER AND M. LORIG (2019), *Optimal liquidation under stochastic price impact*, Internat. J. Theoret. Appl. Finance, 22, pp. 1–28.
- E. BAYRAKTAR AND M. LUDKOVSKI (2010), *Optimal trade execution in illiquid markets*, Math. Finance, 21, pp. 681–701.
- Á. CARTEA, R. DONNELLY, AND S. JAIMUNGAL (2017), *Algorithmic trading with model uncertainty*, SIAM J. Financial Math., 8, pp. 635–671.
- Á. CARTEA, S. JAIMUNGAL, AND J. PENALVA (2015), *Algorithmic and High-Frequency Trading*, Cambridge University Press, Cambridge.
- Á. CARTEA, S. JAIMUNGAL, AND L. SÁNCHEZ-BETANCOURT (2019A), *Latency and Liquidity Risk*, preprint, arXiv:1908.03281.
- Á. CARTEA, S. JAIMUNGAL, AND J. WALTON (2019B), *Foreign exchange markets with last look*, Math. Finance, Econ., 13, pp. 1–30.
- P. CHERIDITO, D. FILIPOVIĆ, AND R. L. KIMMEL (2007), *Market price of risk specifications for affine models: Theory and evidence*, J. Finance, Econ., 83, pp. 123–170.
- M. D. GOULD, M. A. PORTER, AND S. D. HOWISON (2016), *The long memory of order flow in the foreign exchange spot market*, Market Microstruct. Liquid., 2, 1650001.
- O. GUÉANT (2016), *The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making*, Chapman & Hall/CRC Finance, Math. Ser. 33. CRC Press, Boca Raton, FL.
- S. L. HESTON (1993), *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, Rev. Finance, Stud. 6, pp. 327–343.
- C.-A. LEHALLE AND O. MOUNJID (2017), *Limit order strategic placement with adverse selection risk and the role of latency*, Market Microstruct. Liquid., 03, 1750009.
- C. C. MOALLEMI AND M. SAĞLAM (2013), *The cost of latency in high-frequency trading*, Oper. Res., 61, pp. 1070–1086.
- S. STOIKOV AND R. WAEBER (2016), *Reducing transaction costs with low-latency trading algorithms*, Quant. Finance, 16, pp. 1445–1451.