High Frequency Market Making*

Yacine Aït-Sahalia†
Department of Economics
Bendheim Center for Finance
Princeton University
and NBER

Mehmet Sağlam‡
Department of Finance
Carl H. Lindner College of Business
University of Cincinnati

This Version: October 27, 2016

Abstract

We propose an inventory-based model of market making where a strategic high frequency trader exploits his speed and informational advantages to place quotes that interact with low frequency traders. We characterize the optimal market making policy analytically, illustrate that it generates endogenous order cancellations, and compute the long-run equilibrium bid-ask spread and other liquidity measures. The model predicts that the high frequency trader provides more liquidity as he gets faster and shies away from it as volatility increases due to a higher risk of his stale quotes being picked by arbitrageurs. Competition with another liquidity provider increases the overall liquidity. Finally, we provide the first formal, model-based analysis of the impact of four widely discussed policies designed to regulate high frequency trading: imposing a transactions tax, setting minimum-time limits before quotes can be cancelled, taxing the cancellations of limit orders, and replacing time priority with a pro rata or random allocation. We find that these policies are largely unable to induce high frequency traders to provide robust liquidity.

Keywords: High Frequency Trading, Market Making, Liquidity, Order Cancellations, Competition for Order Flow, Tobin Tax, Order Resting Time, Order Cancellation Tax, Pro Rata Allocation.
JEL Classification: G10, G12, G14.

---

*We are grateful to Jonathan Brogaard, Arseniy Kukanov, Costis Maglaras, Albert Menkveld, Ciamac Moallemi, and Fabrice Rousseau for comments and suggestions, and to seminar and conference participants at UC San Diego, UC Riverside, University of Luxembourg, Harvard University, HEC Paris, MIT, University of Venice, University of Washington in Seattle, the Florence-Ritsumeikan Workshop on Stochastic Processes, the Humboldt-Copenhagen Conference, the AFFI Conference, the Banff High-Frequency Conference, the Okinawa Stochastic Processes Conference, MFA 2014, and the NYU Five Star Conference. An earlier version of the paper circulated under the title "High Frequency Traders: Taking Advantage of Speed."

†Email: yacine@princeton.edu
‡Email: mehmet.saglam@uc.edu

1. Introduction

High-frequency traders (HFTs) have become a potent force in many markets, representing between 40 and 70% of the trading volume in US futures and equity markets, and slightly less in European, Canadian and Australian markets (see Biais and Woolley (2011)). HFTs develop and invest in a trading infrastructure designed to analyze a variety of trading signals and send orders to the marketplace in a fraction of a second. The potential profit from any single transaction resulting from an execution may be very tiny, and be achieved ex ante with a probability only slightly above 50%, but HFTs rely on this process being repeated thousands, if not more, times a day. As the law of large numbers and the central limit theorem relentlessly take their hold, profits ensue and presumably justify the HFTs' large investment in trading technology.

This significant amount of market activity has been accompanied by theoretical research addressing some of issues raised by the rapid development of high frequency trading. In particular, Foucault et al. (2016) extend Kyle's model by incorporating heterogeneity in the speed of information processing. Foucault et al. (2013) develop a model in which HFTs choose the speed at which to react to news, based on a trade-off between the advantages of trading first compared to the attention costs of following the news. Biais et al. (2015) analyze the arms race and equilibria arising in a model where traders choose whether to invest in fast trading technologies. Jovanovic and Menkveld (2010) study the effect of high frequency trading activity on welfare and adverse selection costs. Cvitanić and Kirilenko (2010) study the distribution of prices in a market before and after the introduction of HFTs. Other relevant papers include Pagnotta and Philippon (2011), Cartea and Penalva (2012), Jarrow and Protter (2012) and Moallemi and Sağlam (2013).

This paper focuses on a specific type of high frequency trading, market making. HFTs have to a large extent become the de facto market makers, or providers of liquidity to the market. We model a fully dynamically optimizing high frequency market maker as in the classical inventory control problem of Amihud and Mendelson (1980) and Ho and Stoll (1981) for "traditional" market makers (see also Avellaneda and Stoikov (2008), Guilbaud and Pham (2013), Guéant et al. (2013), Cartea et al. (2014) and Hendershott and Menkveld (2014)). In the classical models, the advantages between providers and consumers of liquidity are split: while the market maker has access to the order book, his trading counterparties could potentially be better informed about the fundamental value of the asset, generating adverse selection risk for the market maker. We depart from the classical inventory literature by studying the optimization problem with features that are new to high frequency market making: instead of splitting the advantages between traders and market makers, we turn the tables by endowing HFTs with both speed and informational advantages. The fact that the HFT "holds all the cards" contrasts with the situation in the classical model of Glosten and Milgrom (1985), where the specialist is quoting to potentially better informed traders, and matches the current situation where HFTs are both faster and typically better informed than their trading counterparties. The informational advantage of the HFT in our model is microstructure-driven, rather than pertaining to the fundamental value of the asset.

We then characterize the optimal quoting policy of the HFT, initially in a monopolistic position, who trades against many uninformed low frequency traders (LFTs) with different propensity to trade, and different willingness to wait to get the best price. The bid-ask spread is endogenously determined as a result of the equilibrium between the HFT's optimal quoting strategy and the incoming LFT orders. The model rationalizes why HFTs may change or cancel their orders at a very rapid rate; determine how their endogenous inventory constraints help shape their order placement and cancellation strategies; how the HFT's provision of liquidity can be expected to change in different market environments, such as high volatility ones; and how competition among HFTs can be expected to affect the provision of liquidity and the resulting welfare of LFTs. Most of these questions inherently require a dynamic model to be addressed.

The paper provides several new results. First, we show that being faster translates into higher profits for the HFT, higher liquidity provision and higher cancellation rates in normal times. This is largely consistent with the view that has emerged out of both the academic literature on HFTs and many public policy and industry analyses, namely that HFTs improve market quality by providing liquidity, contributing to price discovery, improving market efficiency and easing market fragmentation. This prediction of the model is compatible with the empirical findings in Hendershott et al. (2011), Hasbrouck and Saar (2013), Chaboud et al. (2010) and Menkveld (2013).

Second, we analyze HFTs' optimal provision of liquidity and equilibrium bid-ask spread as a function of fundamental price volatility, over which the HFT holds no informational advantage. One of the main theoretical predictions of the paper is that the HFT's liquidity provision is U-shaped as a function of volatility, first increasing as volatility attracts more LFTs, but then decreasing when price volatility increases beyond a certain level. Since this is precisely when large unexpected orders are likely to hit, markets can become fragile in volatile times, with imbalances arising because of inventories that intermediaries used to, but are no longer willing to temporarily hold. This U-shaped prediction of the model is consistent with the empirical evidence. Brogaard et al. (2014) and Anand and Venkataraman (2016) find that HFT participation is higher on more volatile days. However, in extreme volatility events, such as flash crashes, Easley et al. (2011) and Kirilenko et al. (2010) show that some HFTs withdrew from their market-making roles.

Third, we show that the ability of HFTs to (imperfectly) predict the types of LFTs they face leads to a strategic widening of his quotes when an impatient trader is signaled to arrive, not unlike the ability of airlines to price-discriminate against business travelers. This results in a time-varying equilibrium bid-ask spread and is consistent with the common "unfairness" complaints of LFTs, whereby some orders are executed a penny away from the best prices despite what the trading screen suggested at the time the LFT submitted his order.

Fourth, we show that competition for order flow among HFTs results in splitting the rent extracted from LFTs, but that the overall liquidity provision increases and equilibrium bid-ask spreads decrease in the presence of competition, and LFTs tend to be better off.

Finally, we provide the first formal, model-based, analysis of the impact of four widely discussed, and in some cases already implemented, HFT policies or regulations: imposing a transaction tax on each trade; setting minimum rest times before limit orders can be cancelled; taxing the cancellations of limit orders; replacing price and time priority with a pro rata or random allocation. The built-in advantages of the HFTs in terms of speed and information prove hard to undo using these policies, and providing proper incentives to HFTs in terms of liquidity provision is also difficult. We find that both a transactions tax and a pro rata allocation scheme reduce the provision of liquidity. Imposing minimum rest times or cancellation taxes induces the HFT to quote more in low volatility environments, but then reduce his provision of liquidity when volatility is high. These policies lead to more liquidity when it is least needed, and less liquidity when it would be most needed. To summarize, the model predicts that these four policies will be ineffective as far as making the provision of liquidity by HFTs more robust, or countercyclical, across the volatility cycle.

The paper is organized as follows. Section 2 sets up the model. Section 3 starts by deriving the optimal quoting strategy of the HFT in a simple version of the model without price volatility, before solving the full model in Section 4. In Section 5, we study the situation where the HFT is competing with another market maker. Section 6 develops the implications of the model for market structure, while Section 7 analyzes in the context of the model the impact of possible HFT policies. Section 8 concludes. Proofs and technical results are in the Appendix.

2. The Model

Our modeling choices are informed by some of the main empirical stylized facts that are known about HFTs (see e.g., Brunetti et al. (2011)), particularly as they relate to market making. They include the fact that HFTs are recognizable by their high frequency of quoting updates, small size on each quote, use of low inventory as a primary risk control strategy and unwillingness to take directional bets. They also tend to place many limit orders, with only some actually leading to execution, and many cancelled; they appear to systematically beat the odds when trading against LFTs; and they seem to exploit order flow information and generate trading signals on a very short time scale rather than longer-run information about the fundamental value of the asset.

Technically, we use multiple, staggered, Poisson processes to represent the arrival of the various elements of the model: market orders by LFTs, signals to the HFT, and jumps in the asset's fundamental price. This modeling device keeps the analysis of the dynamic optimization problem facing the HFT tractable and flexible, and makes it possible to analyze different market environments by varying the parameters of the model. In the following sections, we describe each element of the model in detail.

2.1. HFT's Quotes and Endogenous Bid-Ask Spread

The HFT and a large number of uninformed LFTs are trading a single asset in an electronic limit order book. The HFT acts exclusively as a market maker, employing only limit orders to buy and sell. The quantity of each order, market or limit, is fixed at 1 share or lot. Small volume on each trade matches what is observed empirically in markets that are popular with HFTs, such as the S&P500 eMini futures. Generally speaking, the quantity exchanged in each transaction has been going down over time (see, e.g., Angel et al. (2015)).

We assume that the HFT can place limit orders at four discrete price levels around the fundamental value of the asset, X_t. The tick size in the market is $2C$ for some constant $C > 0$ (e.g., \$0.01). The minimum price at which the HFT is willing to sell the asset is $X_t + C$ with $C > 0$ while the minimum price at which he is willing to buy the asset is $X_t - C$. We will refer to these prices as the best ask and the best bid prices. The HFT can also choose to quote at price levels one tick away from the best ones, $X_t + 3C$ and $X_t - 3C$, and transactions will take place at these levels in the absence of quotes at the best bid and/or ask prices. We will refer to these prices as the second-best ask and the second-best bid prices.

These quoting decisions define the set of available actions to the HFT. Formally, $\ell_t^a = 1$ (resp. $\ell_t^b = 1$) will imply that the HFT has an active quote at the best ask (resp. bid) at time t, and similarly, $\ell_t^a = 2$ (resp. $\ell_t^b = 2$) will imply that the HFT has an active quote at the second-best ask (resp. bid) at time t. Finally, $\ell_t^a = 0$ (resp. $\ell_t^b = 0$) means that the HFT either chooses not to quote to sell (resp. buy) for inventory considerations or his most recent active order has been filled by a market order but the HFT has not yet submitted a new order due to the technological constraints. Figure 1 illustrates the possible quoting decisions by the HFT.[1] A solid bar in the figure indicates that the

[1]For simplicity, the HFT can only have at most one quote at a time on each side of the market. At the cost of increased complexity, we can of course extend the range of tick prices at which the HFT can provide quotes, as well as

HFT is quoting at that price level; there are nine possible quoting decisions.

The HFT is initially a monopolistic liquidity provider at tick levels $X_t \pm C$ and $X_t \pm 3C$. We will introduce below in Section 5 a duopoly model in which there is an additional nonstrategic liquidity provider quoting at the available bid and ask prices, competing with the HFT. We assume that the limit order book contains exogenous depth supplied by competing liquidity providers at tick levels $X_t + 5C$ and/or $X_t - 5C$. This assumption is needed only to define a valid spread measure should the HFT completely withdraw from providing liquidity ($\ell_t^b = 0$ or $\ell_t^a = 0$).

The bid-ask spread will be endogenously determined as a result of the equilibrium between the HFT's optimal quoting strategy and the incoming LFT orders, to be described in Section 2.3. The spread between the active quotes at the ask and the bid can be $2C$, $4C$, $6C$, $8C$ and $10C$, depending upon the quoting decisions by the HFT:

$$
\begin{array}{cc}
\text{HFT Quoting Decisions} & \text{Resulting Bid-Ask Spread} \\
\ell_t^a = 1 \text{ and } \ell_t^b = 1 & 2C \\
\ell_t^a = 1 \text{ and } \ell_t^b = 2 & 4C \\
\ell_t^a = 2 \text{ and } \ell_t^b = 1 & 4C \\
\ell_t^a = 1 \text{ and } \ell_t^b = 0 & 6C \\
\ell_t^a = 0 \text{ and } \ell_t^b = 1 & 6C \\
\ell_t^a = 2 \text{ and } \ell_t^b = 2 & 6C \\
\ell_t^a = 2 \text{ and } \ell_t^b = 0 & 8C \\
\ell_t^a = 0 \text{ and } \ell_t^b = 2 & 8C \\
\ell_t^a = 0 \text{ and } \ell_t^b = 0 & 10C
\end{array}
\tag{2.1}
$$

2.2. *Dynamics of the Underlying Asset Price and Volatility*

We assume that the asset's fundamental price is subject to exogenous variability in the form of pure compound Poisson jumps:

$$X_t = X_0 + \sum_{i=1}^{N_t^\sigma} Z_i, \tag{2.2}$$

where N_t^σ is a Poisson process with arrival rate σ counting the number of price jumps up to time t and Z_i is the stochastic magnitude of the ith jump with distribution

$$Z_i \sim \begin{cases} J & \text{with probability } 1/2, \\ -J & \text{with probability } 1/2. \end{cases} \tag{2.3}$$

the number of simultaneous quotes.

J takes values kC, where k is a positive integer. The dynamics of X_t are illustrated in Figure 2.

With this specification, the fundamental value of the asset is a martingale. The variance of the fundamental value changes is proportional to σ. Although it is in principle a misnomer to call σ "volatility" in light of the differences between discontinuous and continuous components of a semi-martingale, in the absence of a continuous price variation, jumps are the only source of price volatility in the model.

2.3. LFT Orders and Equilibrium

The HFT acts as an intermediary who executes incoming LFTs' orders through his own inventory. LFTs are exclusively liquidity takers. Their orders come with a price condition (like limit orders) combined with an immediacy condition (like market orders): they must either be immediately executed if the price condition is met or cancelled if it is not ("immediate-or-cancel" orders). Consider for instance, an LFT order to buy the asset at price $X + C$ or lower. If an HFT quote to sell is present at tick level $X + C$, the order will get immediately filled, and the HFT inventory will decrease by one share. If the HFT is only quoting at that instant to sell at $X + 3C$, or not at all on that side of the book, an LFT order to buy at $X + C$ or lower will not get executed and will be returned to the LFT as being cancelled; an LFT order to buy at $X + 3C$ or lower would get executed, however.

LFTs' orders to buy and sell arrive to the market as Poisson processes, in the form of a demand function $\lambda^B (X, Y, \sigma)$ to buy the asset and a supply function $\lambda^S (X, Y, \sigma)$ to sell the asset (representing the arrival rate or quantity expected, per unit of time) that depend upon both the fundamental value of the asset X, the transaction price $Y = X \pm kC$ (where k is an odd integer) quoted by the HFT and potentially the asset volatility as well. The micro foundations for this Poissonian assumption are similar to those in Garman (1976), with the addition of the fundamental value and asset price volatility: if a large number of agents demand and supply the asset, each acting independently in the timing of their orders, none being large, and none being able to generate an infinite number of orders per unit of time, then as the number of agents grow the aggregate demand and supply will converge to Poisson processes. It is natural to assume that $\partial\lambda^B/\partial X \geq 0$ and $\partial\lambda^B/\partial Y \leq 0$, and vice versa for λ^S. For simplicity, we assume that λ^B and λ^S depend on (X, Y) through $Y - X$ only, so the demand and supply functions become $\lambda^B (Y - X, \sigma)$ and $\lambda^S (Y - X, \sigma)$.

For the market to be in equilibrium, we require that $\lambda^B (0, \sigma) = \lambda^S (0, \sigma)$. Recognizing the discreteness of the possible prices, we also require that in equilibrium

$$\begin{cases} \lambda^B(-C,\sigma) = \lambda^S(C,\sigma) \equiv \lambda_A + \lambda_P + \lambda_I \\ \lambda^B(C,\sigma) = \lambda^S(-C,\sigma) \equiv \lambda_P + \lambda_I \\ \lambda^B(3C,\sigma) = \lambda^S(-3C,\sigma) \equiv \lambda_I \end{cases} \tag{2.4}$$

with λ_I, λ_P and λ_A parameters capturing the slope of the demand and supply curves near X. When Y is further away from X, the demand and supply curves satisfy for simplicity:

$$\begin{cases} \lambda^B(-kC,\sigma) = \lambda^S(kC,\sigma) = \lambda_A + \lambda_P + \lambda_I \text{ for } k \geq 3 \\ \lambda^B(kC,\sigma) = \lambda^S(-kC,\sigma) = 0 \text{ for } k \geq 5 \end{cases} \tag{2.5}$$

Note that λ_I, λ_P and λ_A as defined above are potentially functions of σ, but we leave this dependence implicit for now. We will return to this in Section 6.3.

LFTs arrive with equal probability on both sides of the market, to buy or sell the asset, as expressed by the equality between λ^B and λ^S at the relevant tick levels in (2.4-2.5). This equilibrium condition is necessary for the HFT not to be in a position to accumulate inventory over time, which would otherwise force him to stop quoting on one side of the market due to inventory aversion; this would happen if some imbalance existed between the arrival rates of the buyers and sellers. Note also that this notion of equilibrium is statistical; it holds on average over time but at any given instant a trader may or may not arrive, and do so on either side of the market, with the HFT's inventory absorbing the imbalance whenever the HFT is quoting and a transaction takes place against an incoming market order.

Equivalently, the arrival of the LFT orders in (2.4-2.5) can be interpreted as emanating from three types of LFTs, patient, impatient and arbitrageurs. Patient LFTs submit orders only at the best available bid or ask price i.e., at $X_t + C$ or $X_t - C$, and their order is only executed if there is an existing limit order (by the HFT) at these prices. By contrast, impatient LFTs are also willing to trade at the second best available quotes when that is all that is available, i.e., sell at $X_t - 3C$ and buy at $X_t + 3C$. Arbitrageurs are only willing to trade if they can buy the asset at a cheaper price than its fundamental value or sell the asset at a higher price than its fundamental value.

Patient, impatient and arbitrageurs arrive at random times according to Poisson processes with respective arrival rates λ_P, λ_I and λ_A; the sum of their arrivals produces the aggregate demand and supply functions λ^B and λ^S faced by the HFT and given in Figure 3. In practice, if λ_A is sufficiently high, then arbitrageurs, although formally classified as LFTs in the model, act more like high frequency traders themselves who are usually referred as high frequency bandits (Menkveld (2016)): with high probability, someone will quickly take liquidity from the HFT by buying at a price Y below X and selling at a price Y above X.

2.4. *Price Volatility, Stale Quotes and Adverse Selection*

We assume that the HFT can only alter his quotes at random times specified by a Poisson distribution with arrival rate μ. Although the HFT is in principle in charge of the timing of his quoting decisions, this technologically-induced randomness is consistent with the broader literature on modeling communication networks at ultra high frequencies. A higher value of μ allows the HFT to revise his quotes more frequently, and thus μ can be interpreted as the speed of the HFT. The HFT is by definition fast relative to any single LFT. However the arrival rates of the LFT orders $\lambda = (\lambda_P, \lambda_I, \lambda_A)$ result from the aggregation of many such LFTs, so we do not necessarily expect μ to be very large relative to λ. In any event, the model does not require assumptions about their relative values.

Consistent with the notion that the HFT trades on the basis of short-lived market microstructure information (see Section 2.5 below), as opposed to fundamental information regarding the true value of the asset, we assume that the HFT has no informational advantage regarding the exogenous jumps in the asset fundamental price, X_t, that were described in Section 2.2. So, when a price jump occurs, the HFT is as surprised as everyone else in the market and is stuck with his existing bid and ask quotes until the arrival of the next decision event (μ) at which point he can peg his limit order to the new mid-price: this is the time marked τ_2^μ in Figure 2.

If the HFT is much faster than the arbitrageurs and other LFTs (μ high relative to λ), he is unlikely to get caught with stale quotes by an incoming order, since a $\mu-$event is likely to have occurred before a $\lambda-$event, but should an arbitrageur's order nevertheless arrive in the interim, it will get executed against the HFT's prevailing pre-jump quotes, creating a loss for the HFT.

This creates adverse selection risk for the HFT, although one which is due to relative speed rather than an informational advantage about the fundamental value by the LFTs (unlike in Glosten and Milgrom (1985)). Stale quotes are attractive not only to patient and impatient LFTs but also to arbitrageurs, and it is plausible for λ_A to be high if the demand and supply functions of Figure 3 are very steep, and thus this risk can be substantial.

Figure 4 illustrates how the HFT's quotes can become stale after a price jump. When the fundamental price of the security jumps up by $J = 4C$, the HFT's earlier quote at the pre-jump $X_t + 3C$ becomes stale and leads to an arbitrage opportunity for LFTs. Patient, impatient and as well as arbitrageurs would certainly like to buy the asset at the post-jump $X_t - C$. If they submit an order during this stale-quote period, a trade occurs, and the HFT would lose C to the LFT that submitted the order. Note that the HFT's quote to buy the asset at the pre-jump $X_t - 5C$ is not attractive to any LFT, patient, impatient or arbitrageur, due to the demand functions in (2.4)-(2.5), and thus no trade would take place at that price and so the potential loss from staleness is not compensated by a potential gain.

2.5. HFT's Informational Advantage

We endow the HFT with a microstructure-driven informational advantage by letting him make quoting decisions based on two independent signals which are respectively informative about the direction of the next incoming market order (buy or sell), and about the type of the trader submitting the market order (patient or impatient). The signals are not predictive about arbitrageurs orders as they will only trade with the HFT if he is stuck with a stale quote due to a price jump (see Section 2.4), a situation which the HFT will never voluntarily face.

The motivation behind the existence of these signals is the empirical and anecdotal evidence that HFTs are able to identify and exploit fleeting opportunities arising from the "plumbing" of the trading process, derived for instance from collocating their trading engine near the exchange matching engine, obtaining and exploiting pricing information faster than other traders, extracting information from the current state of the limit order book, such as real time order book imbalances, observing the arrival of orders on one exchange before they hit other exchanges, running a securities information processor (SIP) that is faster than the publicly-available one, observing recent trading patterns that may be predictive of the direction of the future orders, etc. This assumption of an HFT informational advantage derived from the microstructure of the market is consistent with the empirical evidence presented in, e.g., Hirschey (2011) and Sağlam (2016).

2.5.1 First Signal: Direction of the Order Flow

The first signal is an i.i.d. Bernoulli random variable, $S^{\mathrm{dir}} \in \{B, S\}$, imperfectly predicting the direction of the next incoming market order: B corresponds to a LFT order to buy (i.e., on the ask side of the quotes from the perspective of the HFT) and S refers to a LFT order to sell (on the bid side of the HFT's quotes). The accuracy level of the signal is $1/2 \leq p < 1$. If $p = 1/2$, the signal is uninformative. If $M^{\mathrm{dir}} \in \{B, S\}$ denotes the direction of the next LFT order, and s denote the most recent directional signal received, then by definition

$$\mathbb{P}(S^{\mathrm{dir}} = B | M^{\mathrm{dir}} = B) = p \quad \text{and} \quad \mathbb{P}(S^{\mathrm{dir}} = S | M^{\mathrm{dir}} = B) = (1 - p). \tag{2.6}$$

Given that buy and sell LFT orders are equally likely, the unconditional probabilities for $\mathbb{P}\left(S^{\text{dir}}=B\right)$ and $\mathbb{P}\left(S^{\text{dir}}=S\right)$ are

$$\begin{aligned}\mathbb{P}\left(S^{\text{dir}}=B\right) &= \mathbb{P}(M^{\text{dir}}=B)\mathbb{P}(S^{\text{dir}}=B|M^{\text{dir}}=B)+\mathbb{P}(M^{\text{dir}}=S)\mathbb{P}(S^{\text{dir}}=B|M^{\text{dir}}=S)\\ &= 0.5p+0.5(1-p)\\ &= 0.5\end{aligned}$$

and similarly $\mathbb{P}\left(S^{\text{dir}}=S\right)=0.5$. The conditional probabilities, $\mathbb{P}(M^{\text{dir}}=B|S^{\text{dir}}=B)$ and $\mathbb{P}(M^{\text{dir}}=S|S^{\text{dir}}=B)$, using Bayes' rule, behave as expected:

$$\mathbb{P}(M^{\text{dir}}=B|S^{\text{dir}}=B)=\frac{\mathbb{P}(M^{\text{dir}}=B)\mathbb{P}(S^{\text{dir}}=B|M^{\text{dir}}=B)}{\mathbb{P}\left(S^{\text{dir}}=B\right)}=\frac{0.5p}{0.5p+0.5(1-p)}=p.$$
$$\mathbb{P}(M^{\text{dir}}=S|S^{\text{dir}}=B)=\frac{\mathbb{P}(M^{\text{dir}}=S)\mathbb{P}(S^{\text{dir}}=B|M^{\text{dir}}=S)}{\mathbb{P}\left(S^{\text{dir}}=B\right)}=\frac{0.5(1-p)}{0.5p+0.5(1-p)}=1-p.$$

The conditional probabilities for $S^{\text{dir}}=S$ are obtained by symmetry.

In order to avoid introducing any autocorrelation in the order flow from the signal due to (2.6), we assume that each previous signal is cancelled by either the arrival of a new signal (on a Poisson scale θ) or by the arrival of a LFT order (on a Poisson scale λ). It is natural to assume that each transaction cancels the previous signal and replaces it with a fresh one since the fact that the transaction occurred will often itself be informative. Both events lead to the replacement of the previous signal by a new signal, drawn from its unconditional distribution.

2.5.2 Second Signal: Type of Trader

We assume that the HFT is able, imperfectly again, to predict the type of LFT behind the next incoming order. This introduces the potential for price discrimination by the HFT: if his signal indicates that the next incoming order is more likely to come from an impatient investor, the HFT could potentially widen his quotes to benefit from the impatient LFTs willing to buy at a higher price and sell at a lower one (recall Figure 3). Institutional investors often criticize HFTs for providing "phantom liquidity", arguing that quotes at the best bid and ask suddenly disappear when they try to act on them, leaving them with quotes away from the best bid and ask prices. Quote widening by the HFT may not necessarily happen, though, as the HFT trades off the potential for a higher gain with the lower probability of a trade, since the signal could be incorrect and a patient LFT may materialize, as well as the HFT's inventory concerns.

Like the first, this signal is an i.i.d. Bernoulli random variable, $S^{\text{type}}\in\{P,I\}$, with P and I refer to patient and impatient LFTs respectively. The accuracy level of the second signal is $1/2 \leq q < 1$. As before, if $q = 1/2$, the second signal becomes uninformative. Letting $M^{\text{type}} \in \{P, I\}$ denote the type of the next LFT, we have

$$\mathbb{P}(S^{\text{type}} = P | M^{\text{type}} = P) = q \ \text{ and } \ \mathbb{P}(S^{\text{type}} = I | M^{\text{type}} = P) = (1 - q). \tag{2.7}$$

From 2.3, the unconditional distributions of patient and impatient market orders are given by

$$\mathbb{P}(M^{\text{type}} = P) = \frac{\lambda_P}{\lambda_P + \lambda_I} \ \text{ and } \ \mathbb{P}(M^{\text{type}} = I) = \frac{\lambda_I}{\lambda_P + \lambda_I}. \tag{2.8}$$

We can then compute the unconditional probabilities $\mathbb{P}\left(S^{\text{type}} = P\right)$ and $\mathbb{P}\left(S^{\text{type}} = I\right)$:

$$\begin{aligned}
\mathbb{P}\left(S^{\text{type}} = P\right) &= \mathbb{P}(M^{\text{type}} = P)\mathbb{P}(S^{\text{type}} = P | M^{\text{type}} = P) + \mathbb{P}(M^{\text{type}} = I)\mathbb{P}(S^{\text{type}} = P | M^{\text{type}} = I) \\
&= \frac{q\lambda_P + (1 - q)\lambda_I}{\lambda_P + \lambda_I}
\end{aligned}$$

Similarly, we have $\mathbb{P}\left(S^{\text{type}} = I\right) = \frac{q\lambda_I + (1-q)(\lambda_P)}{\lambda_P + \lambda_I}$. We can then compute the conditional probabilities, $\mathbb{P}(M^{\text{type}} = P | S^{\text{type}} = P)$ and $\mathbb{P}(M^{\text{type}} = I | S^{\text{type}} = P)$ using Bayes' rule:

$$\begin{aligned}
\mathbb{P}(M^{\text{type}} = P | S^{\text{type}} = P) &= \frac{\mathbb{P}(M^{\text{type}} = P)\mathbb{P}(S^{\text{type}} = P | M^{\text{type}} = P)}{\mathbb{P}\left(S^{\text{type}} = P\right)} = \frac{q\lambda_P}{q\lambda_P + (1 - q)\lambda_I}. \\
\mathbb{P}(M^{\text{type}} = I | S^{\text{type}} = P) &= \frac{\mathbb{P}(M^{\text{type}} = I)\mathbb{P}(S^{\text{type}} = P | M^{\text{type}} = I)}{\mathbb{P}\left(S^{\text{type}} = P\right)} = \frac{(1 - q)\lambda_I}{q\lambda_P + (1 - q)\lambda_I}.
\end{aligned}$$

The conditional probabilities for $y = I$ are obtained symmetrically. We assume again that each previous signal is cancelled by either the arrival of a new signal or by the arrival of a LFT order, so that each order arrival is preceded by a minimum of one new signal, in order to preclude any correlation between the type of incoming LFTs.

We can aggregate the conditional probabilities based on both signals $s = \left(S^{\text{dir}}, S^{\text{type}}\right)$ by independence of the two signals. For example,

$$\begin{aligned}
\mathbb{P}(M^{\text{dir}} = S, M^{\text{type}} = P | S^{\text{dir}} = S, S^{\text{type}} = P) &= \mathbb{P}(M^{\text{dir}} = S | S^{\text{dir}} = S)\mathbb{P}(M^{\text{type}} = P | S^{\text{type}} = P) \\
&= \frac{pq\lambda_P}{q\lambda_P + (1 - q)\lambda_I}.
\end{aligned} \tag{2.9}$$

In terms of timing, recall that the HFT is able to make quoting decisions upon the arrival of a Poisson process with parameter μ. He makes his quoting decisions on the basis of the most recent signals received before the arrival of his quote decision time. Once his quotes are in place, it is possible for a new set of signals to arrive before the next incoming market order, drawn from the signals' unconditional distribution and the new set of signals may be different from the existing one. If the HFT cannot update his quotes (no arrival of his quote decision time in between) then he will get stuck quoting of the basis of an out-of-date signal when the next market order arrives. The new order will always be in accordance with the latest signal, whether the HFT was able to act upon it or not. Thus a high arrival rate of the signals (θ) relative to the HFT decision speed (μ) is a risk factor from the HFT's perspective. In practice though, since signal acquisition is part of the HFT technology, it is likely that μ has the largest value.

Figure 5 illustrates the quoting decisions of a relatively slower and faster HFTs in response to two sequences of signal and order arrivals. In the upper sequence, the relatively slower HFT gets stuck with a quote determined on the basis of an out-of-date signal. Of course, the HFT internalizes the possibility of this happening when he computes his optimal quoting strategy. In the lower sequence, a relatively faster HFT is able to update his quotes in response to the latest signal and take advantage of it before the LFT order arrives.

2.6. *HFT Inventory Aversion and Objective Function*

The HFT's position in the asset is denoted by x_t. This position can be positive or negative. In the case of a stock, this means that we impose no restrictions on short selling, while in the case of a futures contract a positive (resp. negative) value of x_t denotes a long (resp. short) net position in the contract. We assume that the HFT is risk-neutral, but penalizes itself for holding excess inventory at a rate of $\Gamma|x_t|$ where Γ is a constant parameter of inventory aversion. In practice, limiting or penalizing inventory is one of the primary sources of risk mitigation by HFTs and is often what is coded into the actual algorithms that perform market making. Inventory aversion is also the reason why, even without competition and no jumps in the fundamental value, so no possibility of adverse selection, a monopolistic HFT may not systematically quote on both sides of the market and attempt to systematically capture every spread.

The HFT's objective is to maximize his expected discounted rewards earned from transacting against the incoming order flow from LFTs, which earns him the bid-ask spread, minus the amount of price jump if the quote is stale, and the potential penalty costs from holding an inventory. The discount rate $D > 0$ is assumed to be constant. Let π denote any feasible policy that chooses ℓ^b and ℓ^a at decision times, T_k^q, and T_i^a be the ith sell order submitted by LFT type $y_i^a \in \{P, I\}$ and T_j^b is the jth market buy order submitted by the LFT type $y_j^b \in \{P, I\}$ where i,j and k are positive integers.

To track the most recent decision time by the HFT before the arrival of market orders, define

$$\begin{cases} \tau_i = \max\left\{k : T_k^q < T_i^a\right\} \\ \tau_j = \max\left\{k : T_k^q < T_j^b\right\}. \end{cases} \tag{2.10}$$

There are three potential outcomes in terms of the HFT's reward function when an LFT submits a buy or sell order. By symmetry, we focus on the bid side of the HFT's quotes, i.e., the ith sell order. The first case refers to a zero payoff due to no trade. If the HFT is not quoting at the bid side ($\ell_{\tau_i}^b = 0$) or the fundamental price has a positive jump after the most recent decision time ($X_{T_i^a} - X_{\tau_i} > 3C$), there will be no trade. If a patient LFT sends the ith sell order ($y_i^a = P$) and the HFT's quote is at the second-best bid ($\ell_{\tau_i}^b = 2$), then there is again no trade.

In the second case, the HFT gains C by trading with either a patient or an impatient LFT ($y_i^a \in \{P, I\}$) when he is quoting at the best bid ($\ell_{\tau_i}^b = 1$) and the fundamental price has not changed since the HFT's most recent decision time ($X_{T_i^a} = X_{\tau_i}$). If there has been a negative jump during this period, the HFT would lose the jump amount mJ despite his spread gain where $X_{T_i^a} - X_{\tau_i} = -mJ$ with $m = 0, 1, 2 \ldots$.

In the third case, the HFT gains $3C$ if the order is submitted by an impatient LFT ($y_i^a = I$) when he is quoting at the second-best bid ($\ell_{\tau_i}^b = 2$) and the fundamental price has not changed since the HFT's most recent decision time ($X_{T_i^a} = X_{\tau_i}$). If there has been a negative jump during this period, the HFT would again lose the jump amount mJ ($X_{T_i^a} - X_{\tau_i} = -mJ$).

We can summarize these cases with the following HFT gains function from LFTs' sell orders:

$$G^-(\ell, y, X_T, X) = \begin{cases} C - mJ & \text{if } \ell = 1, X_T - X = -mJ \\ 3C - mJ & \text{if } \ell = 2, y = I, X_T - X = -mJ \\ 0 & \text{otherwise.} \end{cases} \tag{2.11}$$

Similarly, we can obtain the symmetric function to accommodate the gains from LFTs' buy orders:

$$G^+(\ell, y, X_T, X) = \begin{cases} C - mJ & \text{if } \ell = 1, X_T - X = mJ \\ 3C - mJ & \text{if } \ell = 2, y = I, X_T - X = mJ \\ 0 & \text{otherwise.} \end{cases} \tag{2.12}$$

Putting it all together, the HFT has the following formal objective function:

$$\max_{\pi} \mathbb{E}^{\pi}\left[\sum_{i=1}^{\infty} e^{-DT_i^a} G^-(\ell^b_{\tau_i}, y^a_i, X_{T^a_i}, X_{\tau_i}) + \sum_{j=1}^{\infty} e^{-DT_j^b} G^+(\ell^a_{\tau_j}, y^b_j, X_{T^b_j}, X_{\tau_j}) \right.$$
$$\left. - \Gamma \int_0^{\infty} e^{-Dt}|x_t|dt\right] \tag{2.13}$$

The first (resp. second) term is the HFT's gain from an incoming sell (resp. buy) order crossed against his existing limit order net of adverse selection costs due to stale quotes. The last term captures the HFT's inventory penalty costs.

2.7. Limitations of the Model

To keep the model tractable, a number of elements are left out and simplifying assumptions are made.

First, this paper is only about high frequency market making, not every possible trading strategy that a HFT might employ: in particular, the HFT in our model does not make a strategic choice between limit and market orders, but employs limit orders only. We refer to it as a "HFT" rather than a "HFMM" for ease of exposition, but this paper is only about HFMM. In Rosu (2009), traders dynamically choose between limit and market orders; however they are classical in the sense that there is no speed differential unlike the issue we are focusing on here.

Second, the HFT does not place orders larger than for one contract, so the quantity is not a strategic variable chosen by the HFT. This implies that the relevant notion of market depth at a given price tick in our model is the probability that a quote will be placed by the HFT at that price tick rather than the number of shares available.

Third, HFTs' limit orders are always placed at the best bid and ask price or at the next prices one tick away from the best ones, rather than at every possible price level. Extending the model to allow for quotes at levels further removed from X would needlessly complicate the model. But it means that we exclude order placement strategies known as "quote stuffing" that place large numbers of quotes far away from the best prices to falsely give the impression to other traders of an incoming imbalance, presumably without the intent of ever executing these orders. This sort of strategy is arguably not part of market making. It also means that we preclude the use by HFTs of the now-banned "stub quotes", which are place-holding quotes far from the current market price, employed by market makers to post quotes without any desire to trade, but which may become relevant in a flash crash event for instance. Nevertheless, the HFT decision to quote at the best bid and ask prices vs. quoting at one level removed from them is strategic, changes the spread faced by LFTs, and consequently the spread is determined endogenously in the model.

Fourth, we do not allow for parameter uncertainty: the HFT knows the parameters of the model, including the arrival rates of the various types of LFTs. Cartea et al. (2013) study market making strategies that are robust to model misspecification.

Finally, we do not model explicitly any rebate that might be provided by the exchange to market makers (or payment for order flow in the form of a rebate to liquidity takers). But from the perspective of the HFT, a rebate to liquidity makers can be easily incorporated into the G^- and G^+ functions.

3. Optimal Market Making in a Simplified Model

The key advantage of the model's Poisson-based setup is that we can convert the existing continuous-time model to a discrete-time equivalent one in terms of the arrival times of the LFTs' orders, signals and the HFT's decision time. This effectively merges the different Poisson clocks into a single chronologically-ordered one (with arrival rate equal to the sum of the individual ones), and then performing a time change from the (random) Poisson clock to the corresponding discrete-time event clock.

To demonstrate, we first analyze a simplified version of the model where the fundamental price is constant ($\sigma = 0$), and the arrival rate of impatient LFTs is zero ($\lambda_I = 0$). In this case, the second signal about the likely type of LFT submitting the next order becomes irrelevant, and the only signal will be $s = S^{\text{dir}}$, predicting the likely direction of the next order submitted by LFTs. In the next section, we build on the results of this simplified analysis and solve the full model.

We describe the main elements of the solution method and leave the technical details to the Appendix, where we also prove the results stated in this section. The state space after the discrete-time transformation is (x, s, l, e) where $x \in \{\ldots, -2, -1, 0, 1, 2, \ldots\}$ denotes the current holdings or inventory of the HFT. The second state variable is the order direction signal, $s \in \{B, S\}$. The third state variable, $l \in \{00, 10, 01, 11\}$, denotes the quotes of the HFT currently in the limit order book. $l = 10$ denotes that the HFT is quoting at the best available bid price and is not quoting at the best ask price, and similarly for the other three possible values. In the simplified model, there is no point for the HFT to quote away from the best bid and ask prices in the absence of impatient LFTs.

Recall that the HFT can only make quoting decisions after μ-events in the merged Poisson clock. In the remaining Poisson arrivals, the HFT does not have the ability to change his existing quotes but can only maintain his quotes initiated at the most recent decision time. The fourth state variable, $e \in \{0, 1\}$, is therefore a binary state variable denoting when $e = 1$ that the discrete date in the merged clock corresponds to a $\mu-$event , in which case the HFT can revise his quotes; whereas $e = 0$ refers to the arrival of either a LFT order (λ) or a signal event (θ), in which cases the HFT cannot revise his quotes in this state.

The action taken by the HFT at each state is whether to quote a limit order or not at the best bid and the best ask, i.e., $d^b(x, s, l, e) \in \{0, 1\}$ and $d^a(x, s, l, e) \in \{0, 1\}$. When $e = 0$, the corresponding action taken by the HFT is determined by l so we can actually consider these states as fake decision epochs that force the HFT to continue with existing active quotes. For example, if $l = 10$, then $d^{b*}(x, s, 10, 0) = 1$ and $d^{a*}(x, s, 10, 0) = 0$. Thus, the relevant states for the HFT's decision making are the ones when $e = 1$ and in this case the existing active quotes given by l are no longer binding. For this purpose, we will suppress the last two states when we refer to the optimal market-making policy writing $\ell^{b*}(x, s) = d^{b*}(x, s, l, 1)$ and $\ell^{a*}(x, s) = d^{a*}(x, s, l, 1)$.

3.1. Optimal Market Making Policy by the HFT

In the simplified model, the optimal policy of the HFT reflects only the economic trade-off between quoting in order to capture the spread as often as possible and the risk of quoting too much and increasing his inventory. We let $V(x, s, l, e)$ be the value function of the HFT at state (x, s, l, e). Since the model is symmetric around the bid and ask side of the market, we can first eliminate s from the state space:

Lemma 1.

$$V(-x, S, l, e) = \begin{cases} V(x, B, l, e) & \textit{when } l \in \{00, 11\}, \\ V(x, B, 01, e) & \textit{when } l = 10, \\ V(x, B, 10, e) & \textit{when } l = 01. \end{cases}$$

Second, the value function at a decision time is independent of the existing quotes of the HFT, as at this state ($e = 1$), the HFT can revise his quotes without any restriction. Formally, we have the following result:

Lemma 2.

$$V(x, B, 00, 1) = V(x, B, 01, 1) = V(x, B, 10, 1) = V(x, B, 11, 1).$$

We show in Section A in the Appendix how to express the transition probabilities of the system, and the HFT's reward function, as a function of the HFT's actions. Using these results, we can concisely state the Hamilton-Jacobi-Bellman optimality equations for the HFT's value functions:

Proposition 1. *Let $v(x, l) \equiv V(x, S, l, 0)$ and $h(x) \equiv V(x, S, 00, 1)$. Then, $v(x, l)$ and $h(x)$ jointly* *satisfy*

$$
\begin{aligned}
v(x,00) &= -\gamma|x| + \delta\left(\frac{\mu}{r}h(x) + \frac{\theta+\lambda}{2r}v(x,00) + \frac{\theta+\lambda}{2r}v(-x,00)\right)\\
v(x,10) &= -\gamma|x| + \delta\Big(\frac{\mu}{r}h(x) + \frac{\theta+(1-p)\lambda}{2r}v(x,10) + \frac{\theta+(1-p)\lambda}{2r}v(-x,01) + \frac{p\lambda}{2r}v(x+1,00)\\
&\quad + \frac{p\lambda}{2r}v(-x-1,00) + \frac{pc\lambda}{r}\Big)\\
v(x,11) &= -\gamma|x| + \delta\Big(\frac{\mu}{r}h(x) + \frac{\theta}{2r}(v(x,11)+v(-x,11)) + \frac{p\lambda}{2r}v(x+1,01) + \frac{p\lambda}{2r}v(-x-1,10)\\
&\quad + \frac{(1-p)\lambda}{2r}v(x-1,10) + \frac{(1-p)\lambda}{2r}v(-x+1,01) + \frac{c\lambda}{r}\Big)\\
v(x,01) &= -\gamma|x| + \delta\Big(\frac{\mu}{r}h(x) + \frac{\theta+p\lambda}{2r}v(x,01) + \frac{\theta+p\lambda}{2r}v(-x,10) + \frac{(1-p)\lambda}{2r}v(x-1,00)\\
&\quad + \frac{(1-p)\lambda}{2r}v(-x+1,00) + \frac{(1-p)c\lambda}{r}\Big)
\end{aligned}
\tag{3.1}
$$

where $h(x) = \max\{v(x,00), v(x,01), v(x,10), v(x,11)\}$ *and*

$$
\lambda \equiv 2\lambda_P, \quad \delta \equiv \frac{\lambda+\mu+\theta}{\lambda+\mu+\theta+D}, \quad c \equiv \delta C \quad \text{and} \quad \gamma \equiv \frac{\Gamma}{\lambda+\mu+\theta+D}. \tag{3.2}
$$

Proposition 1 shows that the HFT aims to choose the optimal action when the state is at $e = 1$ by maximizing over all possible quoting actions and the corresponding value is stored in $h(x)$. On the other hand, $v(x,l)$ computes the expected one-step reward resulting from possible transitions determined by the active quotes in state l.

We can now characterize the optimal quoting policy of the HFT. The following result shows that it is based on (endogenously determined) thresholds:

Theorem 1. *The optimal quoting policy* π^* *of the HFT consists in quoting at the best bid and the best ask according to a threshold policy, i.e., there exist* $L^* < 0 < U^*$ *such that*

$$
\ell^{b*}(x,S) = \begin{cases} 1 & \text{when } x < U^* \\ 0 & \text{when } x \geq U^* \end{cases} \qquad \ell^{a*}(x,B) = \begin{cases} 1 & \text{when } x > -U^* \\ 0 & \text{when } x \leq -U^* \end{cases}
$$

$$
\ell^{a*}(x,S) = \begin{cases} 1 & \text{when } x > L^* \\ 0 & \text{when } x \leq L^* \end{cases} \qquad \ell^{b*}(x,B) = \begin{cases} 1 & \text{when } x < -L^* \\ 0 & \text{when } x \geq -L^* \end{cases}
$$

The limits L^* *and* U^* *are functions of the model parameters, but not of the state.*

We can interpret the result of Theorem 1 as follows. U^* and $-U^*$ are the limits that matter for quoting in the direction of the anticipated sign of the next LFT order to arrive. Suppose that the

HFT holds a positive inventory $x > 0$. If he receives the signal $s = S$, he is going to act upon it by placing or keeping a limit order to buy, i.e., on the bid side of the book ($\ell^b = 1$), as long as his current long inventory is not already too high ($x < U^*$); if $x \geq U^*$ the HFT will not quote to buy and risk increasing his already positive inventory even further. Symmetrically, if he receives the signal B, he will quote to sell ($\ell^a = 1$) as long as his inventory is not already too negative ($x > -U^*$).

The HFT wishes to capture the spread, and his signal is imperfect, so the HFT may often quote on the side of the book that is opposite to what the signal predicts. For example, even if he receives the signal $s = S$, he may place or keep a limit order of the ask side of the book, that is, offer to sell from his inventory. L^* and $-L^*$ are the limits that matter for quoting in that case. If the HFT receives the signal $s = S$, he is going to quote to sell ($\ell^a = 1$) as long as his current long inventory is not already too negative ($x > L^*$). Symmetrically, if he receives the signal B, he will quote to buy ($\ell^b = 1$) as long as his inventory is not too positive ($x < -L^*$).

These limits are of course such that the HFT will always quote to buy when his inventory is negative, and quote to sell when his inventory is positive, irrespective of his signal. The limits only bind for quoting in the direction that would increase his inventory further.

A final remark is that the HFT's inventory will be contained in the inventory region $[-N, N]$ where $N \equiv \max(|L|, |U|)$ since at the widest inventory levels, the HFT will have either buy or sell quotes.

3.2. *Analytical Computation of L^* and U^**

We can use the structure of the optimal market making policy to derive a system of equations that will let us solve the threshold limits and the corresponding value functions. The following proposition provides a sufficient condition for threshold limits L and U to be optimal using the structure of the optimal policy:

Proposition 2. *The limits L and U are optimal if L satisfies $v(L, 10) > v(L, 11)$ and $v(L+1, 10) \leq v(L+1, 11)$, and U satisfies $v(U, 01) > v(U, 11)$ and $v(U-1, 01) \leq v(U-1, 11)$, and the value functions solve the system of equations (3.1) where the function $h(x)$ is replaced by the function*

$$m(x) = \mathbf{1}(x \leq L)v(x, 10) + \mathbf{1}(L < x < U)v(x, 11) + \mathbf{1}(x \geq U)v(x, 01). \tag{3.3}$$

We prove in Section A.5 in the Appendix that L^* and U^* must be finite. Combining this result with the sufficient conditions, we propose a simple and efficient algorithm to solve for the optimal thresholds in the form of Algorithm 1 in Section A.5.

4. Optimal Market Making in the Complete Model

We now solve the complete model as described in Section 2, adding back price volatility, impatient traders and a signal about trader type. Compared to the simplified model, the optimal policy of the HFT now reflects two additional trade-offs: quoting at the aggressive price level and transacting with every incoming LFTs vs. attempting to price-discriminate by quoting at a price one tick removed from the best one and trading only with impatient LFTs, but earning more from these rarer trades; and quoting in a volatile environment to earn the spread vs. running the adverse selection risk of stale quotes.

4.1. *Action and State Space*

In its full generality, the state space in our model can be represented by (x, s, l, e, j) where x is as before, and s is the most recent signal with now $s = \left(S^{\text{dir}}, S^{\text{type}}\right) \in \{BP, SP, BI, SI\}$. The third state variable, l, denotes the active quotes of the HFT in the limit order book at the best prices or one tick away from them, with $l \in \{00, 10, 20, 01, 02, 11, 22, 12, 21\}$. For example, $l = 12$ denotes that HFT is quoting at the best price level on the bid side and at the second level on the ask side. The fourth state variable, $e \in \{0, 1\}$, is as before. The fifth and last state variable, j, keeps track of the jumps realized in the asset fundamental value since the HFT's last quoting action, with $j \in \{\ldots, -2, -1, 0, 1, 2, \ldots\}$.

Using the symmetry result in Lemma 1, we can again reduce our signal state by only focusing on $s \in \{SP, SI\}$ as the remaining cases will be given by symmetric states. The action taken by the HFT at each state is whether to quote at the second-best or best available prices, i.e., $d^b(x, s, l, e, j) \in \{0, 1, 2\}$ and $d^a(x, s, l, e, j) \in \{0, 1, 2\}$. When $e = 0$, the corresponding action taken by the HFT is determined by l so we can, as in the simplified model, consider these states as fake decision epochs that force the HFT to continue with existing active quotes. For example, if $l = 10$, then $d^{b*}(x, s, 10, 0, j) = 1$ and $d^{a*}(x, s, 10, 0, j) = 0$. Thus, the relevant state for the HFT is when $e = 1$ and in this case the existing active quotes given by l is no longer binding. At this state, j also reverts to 0 as the HFT can now peg his quotes around the new fundamental price. For this purpose, we will suppress the last three states when we refer to the optimal market-making policy using $\ell^{b*}(x, s) = d^{b*}(x, s, l, 1, 0)$ and $\ell^{a*}(x, s) = d^{a*}(x, s, l, 1, 0)$.

4.2. *Optimal Market Making*

We show in the Appendix how to express the transition probabilities of the system, and the HFT's reward function, as a function of his actions. We again use two value functions to reflect the quoting decision times of the HFT. Let $V(x, s, l, e, j)$ denote the value function of the HFT. Using Lemma 2, we can again suppress some of the states in the case of $e = 1$ with $h(x, s) = V(x, s, 00, 1, 0)$, as the HFT can revise his quotes at this time. Similarly, let $v(x, s, l, j)$ be the value function in the no-decision event types, i.e., a LFT order, signal or volatility event arrives (i.e., when $e = 0$).

We can now characterize the optimal quoting policy of the HFT, showing that, as in the simplified model, it is based on thresholds but now with additional levels that dictate when the HFT should not quote ($\ell = 0$), quote at the best prices ($\ell = 1$) or quote at the second best prices ($\ell = 2$):

Theorem 2. *The optimal quoting policy π^* of the HFT consists in quoting according to a threshold policy, i.e., there exist U_P^0, U_P^1, U_I^0, U_I^1 and L_P^0, L_P^1, L_I^0, L_I^1 which are functions of the model parameters, but not of the state, such that*

$$\ell^{b*}(x, SP) = \begin{cases} 1 & x \leq U_P^1, \\ 2 & U_P^1 < x < U_P^0, \\ 0 & x \geq U_P^0. \end{cases} \quad \ell^{a*}(x, BP) = \begin{cases} 1 & x \geq -U_P^1, \\ 2 & -U_P^1 > x > -U_P^0, \\ 0 & x \leq -U_P^0. \end{cases}$$

$$\ell^{b*}(x, SI) = \begin{cases} 1 & x \leq U_I^1, \\ 2 & U_I^1 < x < U_I^0, \\ 0 & x \geq U_I^0. \end{cases} \quad \ell^{a*}(x, BI) = \begin{cases} 1 & x \geq -U_I^1, \\ 2 & -U_I^1 > x > -U_I^0, \\ 0 & x \leq -U_I^0. \end{cases}$$

$$\ell^{b*}(x, BP) = \begin{cases} 1 & x \leq L_P^1, \\ 2 & L_P^1 < x < L_P^0, \\ 0 & x \geq L_P^0. \end{cases} \quad \ell^{a*}(x, SP) = \begin{cases} 1 & x \geq -L_P^1, \\ 2 & -L_P^1 > x > -L_P^0, \\ 0 & x \leq -L_P^0. \end{cases}$$

$$\ell^{b*}(x, BI) = \begin{cases} 1 & x \leq L_I^1, \\ 2 & L_I^1 < x < L_I^0, \\ 0 & x \geq L_I^0. \end{cases} \quad \ell^{a*}(x, SI) = \begin{cases} 1 & x \geq -L_I^1, \\ 2 & -L_I^1 > x > -L_I^0, \\ 0 & x \leq -L_I^0. \end{cases}$$

The optimal quoting policy in the complete model can interpreted similarly as in the simplified model with obvious nuances. The limits U_P^0 and U_I^0 are the relevant limits for the HFT to stop quoting for inventory considerations when the direction signal is aligned with the direction of quoting. The limits L_P^0 and L_I^0 are the relevant limits for the HFT to stop quoting for inventory considerations when the direction signal is opposite to the direction of quoting. These limits now depend on the anticipated type of the LFT (patient or impatient) submitting the order.

The second main difference is the additional set of limits, U_P^1, U_I^1, L_P^1 and U_I^1, that lets the HFT to decide between quoting at the best bid or at the second best bid. These limits are driven by the economic consideration based on the likelihood of trade given the accuracy of the signals and the relative ratio of patient LFTs to impatient LFTs, and the additional gain from transacting away from the best prices with an impatient LFT. The presence of the volatility and the arrival rate of arbitrageurs (which increases the risk caused by price jumps) also affect the optimal policy by reducing or expanding the quoting regions.

4.3. Example

For illustrative purposes, we use the following realistic parameter values: the HFT will be able to make decisions in every 100 milliseconds which implies that $\mu = 600$ per minute; $C = \$0.005$ which makes the tick-size and the minimum spread to be \$0.01; the arrival rate of impatient LFTs on each side of the market is set to be $\lambda_I = 7.5$ per minute while the arrival rate of patient LFTs is given $\lambda_P = 22.5$ per minute which then implies that total arrival rate of LFTs on both sides of the market is 1 order per second; the arrival rate of arbitrageurs is given by, $\lambda_A = 300$ per minute on each side, which implies the same rate of arrivals as the HFT's decision events in aggregate. The discount rate is 10% per year which corresponds to roughly 10^{-6} per minute. For the accuracy of the signals, we set $p = 0.7$ and $q = 0.6$, i.e., the signal will predict the correct sign of the next market order with 70% chance and the type of the LFT submitting the next order with 60% chance. The signal will be subject to change $\theta = 30$ times per minute on average. The fundamental price will be subject to a jump occurring 10 times per minute, i.e., $\sigma = 10$. Each jump will be in the amount \$0.04 in either positive or negative direction with equal probability. Given $\mu \gg \sigma$, we use a truncated state space for the number of jumps since the HFT's last quoting action with letting $j \in \{-1, 0, 1\}$. Finally, $\Gamma = 0.05$ so that HFT is paying \$0.05 per minute for each non-zero inventory he is holding. This is of course a high value, necessary only to obtain tight inventory limits in order to illustrate the role these limits play.

With these parameters, the quoting limits are computed as $U_P^0 = 4$, $U_P^1 = 0$, $U_I^0 = 5$, $U_I^1 = -1$, and $L_P^0 = 2$, $L_P^1 = 0$, $L_I^0 = 4$, $U_I^1 = -1$. The HFT optimal quoting policy is given in (4.1). When the optimal policy consists in quoting $\ell^{b*} = 0$ and $\ell^{a*} = 1$, we write 01 below:

Inventory (x)	$s = SP$	$s = SI$	$s = BP$	$s = BI$
-5	10	10	10	10
-4	10	10	10	12
-3	10	12	12	12
-2	10	12	12	12
-1	12	12	12	12
0	11	22	11	22
1	21	21	21	21
2	21	21	01	21
3	21	21	01	21
4	01	21	01	01
5	01	01	01	01

(4.1)

This example illustrates the following phenomena. First, when the HFT is expecting to buy, i.e., the signal is $S^{\text{dir}} = S$, he may avoid quoting to sell the asset, to hedge himself against a price jump. Of course, the strength of this hedging motive compared to the trading motive on the side opposite to the signal is dependent upon the accuracy of the signal (p), the level of arbitrageur activity (λ_A) and inventory concerns. Second, when the order is anticipated to come from an impatient LFT, i.e., the signal is $S^{\text{type}} = I$, we see that the HFT quotes at the wider bid and ask prices, hoping to price-discriminate against the impatient LFT. Again, whether the HFT does this or not depends upon the accuracy of the signal (q); and if he needs to lower his inventory, he may forgo the higher spread in order to increase the probability of a transaction in a direction that lowers his inventory.

4.4. *Comparative statics*

We now turn to a comparative statics analysis of the model, examining the sensitivity of the HFT's objective value as a function of the model's parameters.

Proposition 3. *$h(0, s)$ is increasing in λ_I, λ_P, μ, p, q, C and decreasing in λ_A, θ, Γ, δ, and σ.*

This proposition shows the role that each parameter of the model plays in determining the overall value achieved by the HFT. The HFT achieves higher value when: he is faster ($\mu \uparrow$) since he is more likely to be able replenish his quotes between the arrival of successive LFT orders, thereby transacting with two LFT orders in close succession, more likely to be able to take advantage of the latest signal, and less likely to be caught with stale quotes when the asset value is volatile; he receives more LFT orders that he benefits from ($\lambda_I, \lambda_P \uparrow$) since he will capture the spread more often over time; his signal(s) are more accurate ($p, q \uparrow$) since he can price-discriminate more effectively and his inventory risk is lower, allowing him to adjust his quotes asymmetrically when close to his inventory limits; the tick size is higher ($C \uparrow$) since that is what he gains on each transaction. Conversely, the HFT achieves lower value when: more arbitrageurs arrive ($\lambda_A \uparrow$) and more fundamental price volatility occur ($\sigma \uparrow$) since both increase his risk of losing to arbitrageurs from stale quotes; his signals become more frequent ($\theta \uparrow$) since a signal update without the ability to update his quotes ($\mu-$event) is detrimental to the HFT as it makes it more likely that he will continue quoting on the basis of an outdated signal; he is more inventory averse ($\Gamma \uparrow$) since this is from his perspective a pure penalty with no benefits; his discount rate is higher ($\delta \uparrow$) since it lowers the present value of his rewards, i.e., his objective function.

5. Duopolistic Market Making with Queuing Priority

The model so far includes only a single HFT, who is enjoying monopolistic rents as the sole provider of liquidity in the market. We now study a more realistic situation where the HFT is competing with a second liquidity provider in order to capture the bid-ask spread earned by executing LFTs' orders. The limit order book sorts orders by price (first) and time (second) priority. For simplicity, this second liquidity provider is a nonstrategic trader who places limit orders into the book at a Poisson rate of β; we call him a medium-frequency trader (MFT). The HFT remains fully strategic, taking now into account not only the arrival of the LFTs but now also the presence of the MFT he is competing with. Both the MFT and the HFT trade with the LFTs' incoming orders. Allocation of the incoming LFT orders to either the HFT or the MFT is done according to price and time priority.

There is now competition for priority and the fully strategic HFT needs to take into account the state of the limit order book and the priority or lack thereof that he would achieve when computing his best response to the MFT in order to place his own quotes. The MFT, like the HFT, offers at most one share to buy or sell. For the sake of simplicity, to limit the number of possible states, we assume that the MFT is only placing quotes at the best bid and ask prices. The HFT can still choose to quote at the best bid/ask prices, or one tick away from them. When the HFT decides to quote away from the best prices, he obviously runs the risk that he will lose a transaction to the MFT even if the HFT had time priority, should the MFT enter the order book before the next impatient LFT's order arrival.

Suppose that order book events attributed to the MFT, which occur with a rate 2β, arrive symmetrically in the best bid and ask queues. Consider the best bid queue. The MFT sends a limit order to the best bid queue at rate β and waits for his order to be executed by a market order. His order stays in the queue until it is met by an incoming LFT order. If the HFT has a limit order in the best bid queue ahead of him, the MFT will wait for the HFT's order to be either executed or cancelled by the HFT; the MFT will then be at the front of the queue and his order will stay there until it is met by an incoming LFT order. Once the MFT order is executed, he waits another Poisson time with rate β before placing his next limit order in the best bid queue.

The analysis of the HFT's optimal strategy in this context proceeds by recognizing the presence of another Poisson time clock in the model, and taking into account his position in the limit order queue relative to the MFT's orders. So a new state variable, which tracks his position in the queue, becomes relevant for the HFT's optimal decisions process. Each side of the limit order book, bid or ask queues, can now be in one of seven states, represented in Figure 6: $b = 00$, denoting that the bid queue is empty, with neither HFT nor MFT currently quoting; $b = 10$, denoting that the HFT is quoting at the best bid and MFT is not quoting; $b = 20$, denoting that the HFT is quoting at the second-best bid and MFT is not quoting; $b = m2$, denoting that the MFT is quoting at the best bid and the HFT is quoting at the second-best bid, but the MFT's order has price priority; $b = m1$, denoting that both the MFT and the HFT are quoting at the best bid, but the MFT's order has time priority; $b = 1m$, denoting that both the MFT and the HFT are quoting at the best bid, but the HFT's order has time priority; $b = m0$, denoting that there is only an order by the MFT at the best bid, which has priority.

Consequently, in this framework, the main difference is again in state $l = b \times a$ which can be in the set of 49 states obtained by interacting the 7 states for the bid queue with the 7 states for the ask queue. We write the value function corresponding to each state. In each case, we now consider the possibility that the MFT may submit a limit order at the best bid or best ask. If the HFT chooses to quote an order on either side of the market, he gains priority for execution as the book is empty. However, if he does not quote, then in the next period he may find an MFT limit order in the book and his order will have lower time priority.

We then characterize the optimal quoting policy of the HFT, showing that it is again based on thresholds, which are determined by the queuing states.

Theorem 3. *The optimal quoting policy π^* of the HFT consists in quoting at the available bid and the ask prices according to a threshold policy, i.e., there exists state-dependent limits for both the bid and the ask side such that the optimal policy is monotonic with respect to inventory:*

$$\ell^{b*}(x,s,l) = \begin{cases} 1 & x \le U^1(s,l), \\ 2 & U^1(s,l) < x < U^0(s,l), \\ 0 & x \ge U^0(s,l), \end{cases} \quad \ell^{a*}(-x, \mathsf{sym}(s,l)) = \begin{cases} 1 & x \ge -U^1(s,l), \\ 2 & -U^1(s,l) > x > -U^0(s,l), \\ 0 & x \le -U^0(s,l). \end{cases}$$

where $\mathsf{sym}(s, l)$ *denotes the symmetric state of* (s, l)*.*

Theorem 3 shows that the quoting actions retain their monotonicity even in this complicated framework with a competing market maker. The intrinsic economic trade-offs faced by the HFT with respect to quoting at different prices remain in place in the presence of the MFT.

6. Implications of the Model for Market Structure

In this section, we consider the implications of the model for the HFT's liquidity provision, cancellation or modification of orders, and how liquidity provision is affected by fundamental price volatility. The figures in this section are computed using the same parameter values as those of 4.3.

Proposition 3 allows us to anticipate how a change in the environment will affect the optimal strategy of the HFT. We first compute the optimal inventory limits and characterize the optimal policy of the trader according to Theorem 2 (or Theorem 3). Under this optimal policy, the inventory of the trader will be in a finite set, $[-N^*, N^*]$. Therefore, under the optimal policy, the model is governed by a finite-state Markov Chain. Let P_{opt} be the probability transition matrix defined on this finite state space under the optimal policy. Since the Markov Chain is aperiodic and irreducible, a stationary distribution π exists for this Markov Chain, which solves $\pi P_{\mathsf{opt}} = \pi$.

We then compute the long-run liquidity metrics resulting from the equilibrium between the HFT's (or HFT's plus MFT's) quotes and the LFTs' orders under the Markov stationary distribution, π. We focus on the following three liquidity measures: the expected bid-offer spread, the fill rate of LFT orders and the probability that the HFT will quote at the best bid and ask prices, i.e., the probability of a one-tick market, which given the fixed quantity of one in the model captures the market depth provided by the HFT at the best prices.

6.1. *HFT's Liquidity Provision and HFT's Speed*

Does the HFT's speed have a positive impact on the market liquidity? Figure 7 illustrates that as the HFT gets faster (μ), market liquidity as measured by the three liquidity metrics is higher. This result is consistent with the empirical literature on HFT, which tends to suggest that HFTs have had a generally positive effect on spreads and depths (see, e.g., Hendershott et al. (2011), Hasbrouck and Saar (2013), Chaboud et al. (2010) and Menkveld (2013).)

In the model, as the HFT becomes faster, the various risks he is facing (getting caught with an out-of-date signal, getting caught with stale quotes, etc.) are reduced. Furthermore, getting faster also improves the likelihood that the HFT will be able to exploit signals that allow him to price-discriminate. Figure 8 illustrates that if the HFT's second signal accuracy (q) improves, then the average bid-offer spread in long-run equilibrium increases. This reflects the fact that the HFT is better able to price discriminate among types of LFTs and exploits this ability to engage in predatory trading against impatient LFTs. This result is consistent with the often-expressed complaints from LFTs regarding “phantom or fleeting liquidity”, i.e., liquidity provided at the best bid or ask prices that suddenly disappears, leading to an execution one tick away from the best prices.

6.2. Quote Cancellations by the HFT

Limit order cancellations are widely observed empirically in high frequency data. Hasbrouck and Saar (2009) note that over one third of limit orders are cancelled within two seconds and term those “fleeting orders”. Baruch and Glosten (2013) show that quote cancellations can emerge as an equilibrium strategy in a trading game. Angel et al. (2015) document that the ratio of quotes to trades, which was relatively stable at about 2 between 1993 and 2001, started increasing and is over 25 by 2013. Our model makes predictions that are consistent with this empirical fact: a high number of cancellations will occur as the HFT changes his quotes in responses to changes in the fundamental value, his inventory and the signals he receives.

Figure 9 and Figure 10 illustrate the HFT’s quote revisions and cancellations while following his optimal policy (recall (4.1)). Since the optimal limits may be different for each possible signal, endogenous quote cancellations may occur for instance when the HFT receives a different signal than the one that was previously in effect. For example, when his inventory is zero or close to it, the HFT makes his quoting decision depending upon the signal predicting the type of incoming LFT. Figure 9 shows that if the HFT is expecting a patient LFT, he quotes at the best bid and ask but if the signal suggests that an impatient LFT is next to arrive, he quotes at the wider price levels. From the perspective of the LFTs, these lead to cancellations and a widening of the spread. Similarly, the signal providing information about the likely direction of the next LFT order may induce the HFT to cancel his quotes to minimize the probability of a stale quote. The left panel in Figure 10 illustrates that if the signal changes to BP when the HFT is expecting to buy from a patient LFT, the HFT decides to cancel his quote at the second-best bid in order to lower the risk of leaving a stale quote. Similar withdrawal of an ask quote happens in the right panel of Figure 10.

6.3. Implications of Asset Price Volatility for the Provision of Liquidity

One fundamental issue regarding the HFT’s provision of liquidity concerns not its quantity but its “quality”. Possible definitions of that quality vary, but most include the notion that this liquidity is to be provided in a stable manner over time and over different market environments, consistent with the requirement to provide a "fair and orderly market" that is imposed on regulated specialists and market makers. Are unregulated HFTs fair weather liquidity providers ready to provide plenty of liquidity when the market is calm and doesn't really need it, only to remove it whenever the market becomes turbulent (and it would be needed)? This question is of central importance for market stability, and to understand the potential for systemic risk should HFTs suddenly suspend their provision of liquidity in response to a market shock, contributing to an amplification of that shock.

We quantify the change in the equilibrium bid-offer spread as a function of asset price volatility, measured by σ: see Figure 11. We find that the rate of liquidity provision by the HFT decreases as the fundamental price becomes more volatile, resulting in a higher spread in equilibrium. The model predicts that the HFT protects himself against unanticipated price jumps (over which he has no informational advantage) by quoting less frequently. The conclusion that it is optimal for the HFT to decrease his provision of liquidity when the price volatility risk increases has obvious consequences for the way markets can be expected to operate in times of stress, and for the potential need to regulate the provision of liquidity by market makers.

On the other hand, since λ^B and λ^S described in Section 2.3 aggregate the demand for and supply of the asset emanating from many different traders, who are likely to have heterogenous beliefs, trading motives and propensities, it is possible and even plausible for λ^B and λ^S to depend on the fundamental price volatility σ; if volatility leads to a higher intensity of trading, then $\partial\lambda^B/\partial\sigma \geq 0$ and $\partial\lambda^S/\partial\sigma \geq 0$. The resulting effect on the HFT's optimal quoting, and hence the equilibrium spread, is illustrated in Figure 11 in the case where $\lambda^B(Y-X,\sigma)$ and $\lambda^S(Y-X,\sigma)$ are both linearly increasing in σ, corresponding to a linear positive correlation between volatility and LFTs' supply and demand. Other cases are possible depending upon the assumptions made outside the model regarding the LFTs' motives for trading (portfolio rebalancing, limits reached, directional bets, etc.) that ultimately determine the dependence of λ^B and λ^S on σ.

Any positive dependence of λ^B and λ^S on σ makes it more attractive for the HFT to continue quoting, and in some cases increase his quoting in times of market stress, although as we see in Figure 11 the increase in the provision of liquidity is short-lived: as σ begins to increase, the HFT starts by providing more liquidity, resulting in a lower spread, but as σ keeps increasing the HFT starts to withdraw liquidity. The HFT's objective value is increasing, due both to the increase in trading opportunities with LFTs owing to the increased volatility, and to the higher profitability of each trade even when the HFT quotes less.

This U-shaped dependence of the equilibrium spread (or inverse U-shaped dependence of liquidity) on volatility is consistent with recent empirical evidence. Brogaard et al. (2014) and Anand and Venkataraman (2016) find that HFT participation is higher on more volatile days. However, in the presence of extreme volatility events, such as flash crashes, Easley et al. (2011) and Kirilenko et al. (2010) show that some HFTs may withdraw from their market-making roles.

6.4. Competition for Order Flow and Liquidity Provision

We next investigate the impact on liquidity of competition between the HFT and MFT, as they queue in the limit order book. We solve for the HFT's optimal policy taking into account the MFT's presence or absence in the form of an additional state variable, compute the optimal value the HFT is able to achieve, and compare it to the baseline model where the HFT is a monopolistic provider of liquidity. This duopoly model, for which we set a value of the MFT quote arrival rate parameter $\beta = 60$, subsumes the base model by setting $\beta = 0$. This value reflects the presumption that the MFT (β) is slower than the HFT (μ), but the MFT's quoting actions are faster than the arrival rate of market orders by the LFTs (λ_P and λ_I).

Figure 12 illustrates the impact of queueing on the HFT's value due to the presence of the MFT, as a function of the HFT speed in the monopoly and duopoly models. In the presence of the MFT, the HFT splits the rent by losing some of his profits to the MFT, as some of the LFTs' orders are now executed by the MFT. Comparing the slopes in the value function in both regimes, we see that lower speed is more detrimental when the HFT is extracting monopolistic rents in the absence of the MFT: with the MFT present, not being as fast is less costly for the HFT. This suggests that with competition, the HFT's incentive to increase his speed is lower.

We next study how the presence of the MFT affects market liquidity. Figure 13 shows that LFTs are better off when market makers compete, compared to the monopolistic HFT situation: the equilibrium bid-offer spread decreases. In the duopoly case, the HFT quotes less: after making a trade, either a sale or a buy, the HFT is now less likely to trade in the opposite side of the market immediately due to the presence of the MFT. Although the HFT provides less liquidity than he did as a monopolist, the total provision of liquidity facing the LFTs is higher in the duopoly situation. Also, competition with the MFT decreases the HFT's ability to price-discriminate, since removing his quotes at the best bid and ask prices to quote one tick away from them causes the HFT's to potentially lose his time priority in the best bid and ask queues to the MFT. This induces the HFT to quote more often at the best bid and ask, reducing the equilibrium spread and increasing the probability of a one-tick market.

7. Implications of The Model for HFT Regulations

Regulators worldwide are paying increasing attention to the impact of HFTs on the markets they supervise, and have started debating, developing and in some cases implementing policies designed primarily to limit any negative consequences stemming from the rise of HFTs. The basic premise is that markets should remain platforms to trade risk among investors with different needs and beliefs, and not become primarily race tracks with the potential to ultimately discourage slower users from participating, due to the perception of a lack of fairness.

In this section, we use the model to make predictions about the impact of some frequently proposed, and in some cases already enacted, regulatory policies. We analyze their impact on the HFT's objective value and provision of liquidity to the market. Although welcome, additional liquidity in good times when markets are already very liquid may not be particularly useful, and may even be counterproductive if it lulls agents into complacency only to dry up when markets experience volatility. So we view as a desirable outcome of a policy a reduced provision of liquidity by HFTs in good (low volatility) times, in exchange for an increased provision of liquidity in bad (high volatility) times. Such a countercyclical impact would be desirable for some of the same reasons that advocate for banks' capital requirements to increase in booms, in order to limit the provision of speculative credit, and decrease in recessions, in order to speed up the economic recovery.

We specifically examine the impact of four widely discussed HFT policies through the prism of our model: imposing a transaction tax on each trade, setting minimum-time limits before orders can be cancelled, taxing the cancellations of limit orders, and replacing time priority with a pro rata allocation of shares to all quotes in the book at that price level.

These policies, or combinations thereof, capture the main elements that have been proposed in various countries, and in some cases already implemented. In 2012, France introduced a 0.2% tax on transactions in large stocks, and a 0.01% tax on HFTs penalizing them for a high rate of order cancellations within a half-second[2],[3]. Similarly, in Italy, a tax of 0.02% on orders issued and then cancelled within half a second, once above a threshold, has been introduced[4]. The Deutsche Börse introduced a tax in 2012 that charges HFTs for high "order-to-trade" ratios[5] as does the London Stock Exchange[6]. Norwegian regulators too consider taxing traders who submit a large number of orders relative to their actual executions[7]. The CME Group, the world's largest futures exchange, has had for a number of years message volume caps, designed to prevent excessive numbers of orders from being placed[8], while Nasdaq and DirectEdge, two of the largest US stock exchanges have introduced fines to

[2] "Paris traders brace for financial transactions tax", Reuters, July 31, 2012.

[3] Somewhat predictably, many investors in France have avoided the tax by trading "contracts for difference" which allow them to profit from an asset's gain or loss without actuallys owning the shares.

[4] "All eyes on Italy's high-frequency rules", The Financial Times, February 19, 2013.

[5] "D Börse to charge for 'stupid algos'", The Financial Times, February 28, 2012.

[6] "Bourses play nice cop to head off speed-trade rules", Reuters, April 10, 2012.

[7] "Oslo Bors to charge for excessive orders", The Financial Times, May 24, 2012.

[8] CME Messaging Efficiency Program: http://www.cmegroup.com/globex/resources/cme-globex-messaging-efficiency-program.html

discourage excessive order placement[9]. Canadian regulators too began increasing the fees charged to HFTs that flood the market with orders[10], while Indian regulators are studying ways to curb HFTs[11]. On the other hand, Brazil appears more open to the influx of HFTs[12]. Australian regulators want HFTs to implement a "kill switch" to prevent future flash crashes, and are considering a tax charge, although they appear to take a more benevolent view of HFTs than some of their counterparts in Europe[13].

In January 2013, European Union finance ministers approved a transaction tax in Germany, France, Italy, Spain and seven other Eurozone countries[14]; the UK, concerned about the impact on the City, is opposed[15]. It seems unlikely at present that the initially far-reaching package will get implemented as proposed, if ever[16]. The German government has advanced legislation that would, among other things, force HFTs to register as such with the government[17] and limit their ability to rapidly place and cancel orders[18]. The European Parliament has voted to require HFTs to honor the quotes they submit for at least half a second; imposes a minimum half-second delay on executing orders in a bid; possible use of circuit breakers to interrupt a sudden market plunge; and fee structures that would discourage excessive algorithmic trades[19]. These rules could potentially apply to all 27 member states of the European Union if governments were to give their approval. In the US, the SEC and CFTC are discussing similar kinds of regulatory actions[20], while transaction tax legislation has been introduced in the Senate, although with little prospects of passage. Not surprisingly, many trade associations representing trading firms are opposing these proposals[21]. Other alternative policies involve bunching together incoming orders every few milliseconds, or randomizing their allocation ("scrambling"), so a HFT would face queuing risk, as well as a minimum rest time before a cancellation. EBS, one of the major trading platforms in the foreign exchange market, has discussed such a proposal with its users.

Pro rata matching is employed for instance at the Chicago Mercantile Exchange for Eurodollar futures contracts as an alternative to time priority. The objective is to reduce the value of being at the top of the queue for execution, an objective shared with the proposal for batch auctions (see, e.g.,

[9] "US bourses to fine HFT data-cloggers", The Financial Times, March 7, 2012.
[10] "Canada's 'hot' traders attract regulatory heat", The Financial Times, October 16, 2012.
[11] "India takes steps to rein in algos", The Financial Times, May 22, 2013.
[12] "Despite Risks, Brazil Courts the Millisecond Investor", The New York Times, May 22, 2013.
[13] "Australia finds HFT fears 'overstated'", The Financial Times, March 18, 2013.
[14] "Brussels proposes 30bn 'Tobin tax'", The Financial Times, February 14, 2013.
[15] "Britain challenges EU over 'Tobin tax'", The Financial Times, April 19, 2013.
[16] "Europe plans major scaling back of financial trading tax", Reuters, May 30, 2013.
[17] "Berlin forges ahead with trading controls", The Financial Times, September 25, 2012.
[18] "German Bundestag Passes Bill to Regulate High-frequency Trading", The Wall Street Journal, February 28, 2013.
[19] "EU Lawmakers Call for Enforced Delay on High-Frequency Trades", The Wall Street Journal, September 26, 2012.
[20] "High speed trading a stiff challenge for U.S. regulators", Reuters, May 19, 2013.
[21] "Traders see Europe's Tobin tax hurting savers", Reuters, October 11, 2012 and "Trading Clamps Spur Lobby Effort", The Wall Street Journal, March 24, 2013.

Budish et al. (2015)).

In all cases, we assume that the HFT continues to optimize taking into account the full regulatory environment, and we extend the model accordingly. This analysis remains dependent on the assumptions of the model, and excludes alternative responses by HFTs, such as simply moving their trading to an alternative non-regulated venue, instead of optimizing under the new constraints imposed by the regulation.

7.1. Tobin Tax: Taxing High Frequency Trades

The first policy we consider consists in taxing each trade that a HFT executes. Leaving aside the question of identifying HFT trades (perhaps by requiring HFT firms to register with the regulators, as has been proposed in Germany), a financial transactions tax is nothing new. Originally known as a "stamp duty," it was first implemented at the London Stock Exchange in the 17th century, was later advocated by Keynes on the grounds that speculation by uninformed traders increased volatility, and then by Tobin as a means of reducing currency fluctuations.

The argument in favor of a transactions tax is that financial trading is under-taxed relative to the rest of the economy; this encourages excessive trading, by HFTs in particular, which in turn undermines financial stability as the ability of HFTs to get out of the market quickly undermines the market's liquidity when it is most needed.

Of course, sophisticated traders may simply move their trading to financial instruments or jurisdictions not subject to the tax. Sweden for instance introduced a tax on the purchase or sale of stocks in 1984; the tax was repealed in 1990 after the country experienced a large displacement of trades.[22] A second argument often made against the tax is that it will depress economic activity by imposing a large burden on the financial sector. These two arguments are somewhat self-contradictory: either the tax is easily avoided so as to be inconsequential, or it imposes a large economic penalty, but not both together[23].

In the framework of our model, a transactions tax is straightforward to analyze. Suppose that the HFT pays κ dollars each time an LFT order crosses one of his limit orders. From the perspective of the HFT, the transaction tax, κ, merely reduces the gain that the HFT earns from each trade. We can analyze the impact of the tax policy on the objective value of the HFT and liquidity provision using our complete model.

Figure 14 displays three graphs illustrating objective value, equilibrium bid-offer spread with re-

[22] "Financial Transactions Taxes: The International Experience and the Lessons for Canada", by Marion G. Wrobel, Parliament of Canada Report, June 1996.

[23] "Europe should embrace a financial transaction tax", The Financial Times, May 28, 2013.

spect to the Tobin tax rate and equilibrium bid-offer spread as a function of volatility in the presence and absence of tax regulation. We use the same parameter values as in Section 6 to facilitate the comparison with earlier applications. Figure 14 illustrates that the HFT's objective value is decreasing with higher taxes, i.e., with a maximum κ considered of 25 bps in line with the proposals being considered or already implemented in Europe. We observe that in the long-run transaction taxes do not incentivize the HFT to quote more on both sides of the market. If the taxes are high enough, in fact, the HFT's liquidity provision actually decreases, as seen in the top panel of Figure 14.

These predictions of the model are consistent with what has been observed in Italy following the introduction of the Tobin tax: the average daily trading volume for Italian-domiciled stocks has fallen by nearly 40% in March compared to January and February 2013.[24]. Lastly, we investigate how market liquidity, measured by the equilibrium bid-offer spread, changes with respect to volatility when the Tobin tax is implemented. The HFT's quoting has higher sensitivity to volatility compared to the absence of a Tobin tax, so the tax produces no improvement on that front either.

7.2. Speed Bumps for HFTs: Imposing a Minimum Rest Time Before a Quote Can Be Cancelled

Another possible policy consists in imposing a minimum time before a quote can be cancelled by the HFT. This minimum "rest time" is a widely discussed policy among regulators and exchanges, including European and Australian ones. The objective is to improve the provision of liquidity by HFTs by effectively forcing them to stand behind their quotes for at least a brief period. One concern about the reported higher liquidity due to HFT activity is that the provided liquidity is very short-lived, or "phantom", i.e., HFTs cancel many of their quotes before LFTs get a chance to trade with it.

We analyze the effect of minimum rest limits in our model as follows. Although the policy proposals all suggest a fixed waiting time, typically 500 milliseconds, it will come as no surprise in the context of the model that it is actually more convenient to analyze a version of the policy where the waiting time before a cancellation is random, itself derived from a Poisson process, with an expected value of 500 milliseconds. The arrival rate of that Poisson process controls the expected amount of waiting time before a quote can be cancelled. That is, we suppose that each active quotes cannot be cancelled before a random time amount, τ^{cancel}, which is exponentially distributed with mean duration $1/\theta$.

Figure 15 illustrates the objective value of the HFT as a function of the average rest time, $1/\theta$, expressed in milliseconds. The limiting case $\theta \to \infty$, meaning that the minimum rest time is zero for active quotes, reverts to the base model. As the rest time increases, we observe in the upper left panel

[24] "Tax blow to Italian stock trading", The Financial Times, March 13, 2013.

that the objective value of the HFT decreases: like any other constraint, a rest time does reduce the maximum achievable value.

Importantly, we find that this form of regulation is good for the overall liquidity as measured by equilibrium bid-offer spread. When the HFT has active quotes on both sides of the market, the minimum rest time comes into effect and forces the HFT not to cancel as frequently as before due to signal changes, including signals about the next incoming LFT's type. As he internalizes this constraint when optimizing his quoting policy, the HFT ends up initiating fewer quotes but they stay active sufficiently longer to result in a higher probability that he will have a quote active at any given point in time, and in a higher fill rate for the incoming LFTs' orders.

This policy fails however to result in a countercyclical improvement to liquidity: as shown in the bottom-panel graph in Figure 15, with rest times in effect, the HFT becomes more sensitive to market volatility, i.e., when a rest time is imposed, equilibrium bid-offer spread is lower when volatility is low, but higher when volatility is high, compared to the situation without one. The reason is that as σ increases, the potential risk from being caught with stale quotes increases, and is compounded by the imposition of a rest time which forces the HFT to honor a quote that he would rather not. The HFT then optimally reduces his quoting. In other words, rest times lead HFTs to provide more liquidity when the market does not really need it (more rain in a monsoon) but less when the market would really benefit from it (less rain in a drought).

7.3. *Taxing Limit Order Cancellations*

The third policy we consider consists in taxing the HFT whenever he cancels an existing quote. Unlike rest times, which make cancellations impossible within a certain time interval, this policy simply makes them costlier. The two are not directly comparable in that rest times are effectively an infinite tax on cancellations but only for a brief period, whereas a cancellation tax is a small tax that is in effect permanently. We assume that the HFT must pay ε percent as a penalty for each cancelled limit order. For example, if the HFT chooses action $\ell^b = 0$ or $\ell^a = 1$ when his existing quotes are $l = 11$, he will pay a cancellation tax of ε as he cancelled his existing quote on the bid side.

Figure 16 considers the impact of such cancellation taxes. We consider cases where ε is as high as 10 bps for each quote cancellation. We observe that the objective value of the HFT decreases in the presence of cancellation taxes. The HFT tries to quote more at the best bid and the ask to avoid cancellation fees, which lowers the equilibrium bid-offer spread. A cancellation tax shares the same drawback as a minimum rest time: it improves market liquidity when volatility is low but when there is high volatility, the cancellation tax causes the HFT to quote less compared to the benchmark case.

7.4. *Pro Rata or Random Allocation*

The last policy we consider is pro rata allocation: as an alternative to price and time priority, or first in first out, pro rata allocation matches incoming orders against quotes by allocating shares proportionately to all quotes at the best price according to their size. This form of matching, employed for instance at the Chicago Mercantile Exchange for Eurodollar futures contracts, enables all liquidity providers to join the queue at a particular price level and have an opportunity to compete for the next fill at that price level, independently of their order's time priority. A closely-related alternative matching algorithm is one where, among the quotes in the book at the prevailing best price, the one to be executed is selected at random.

In this section, we use the duopoly model of Section 5 to address the impact of pro rata/random matching on the optimal strategy of the HFT, and the resulting equilibrium liquidity measures. So far in the model, when the HFT and the MFT have limit orders at the best price, the execution priority belongs to the liquidity provider who first submitted the order. Since quantities in the model are set at one share per transaction, we model pro rata/random allocation as follows. Instead of the first quote to enter the book being automatically selected for execution, an independent Bernoulli random variable is drawn and the order with the highest time-priority is selected with probability $\eta \in [0.5, 1]$. When $\eta = 1$, this reverts to the original duopoly model, i.e., pure time priority, whereas $\eta = 0.5$ completely eradicates time priority and corresponds to a fully random allocation. Intermediary values are possible.

Figure 17 considers the impact this policy on liquidity metrics. We find that the objective value of the HFT decreases in the presence of pro rata allocation. In the absence of this policy, HFT's orders often enjoy higher priority over the MFT's orders due to the HFT's speed advantage, which results in more trades for the HFT. However, once the pro rata policy in place, the HFT's speed advantage loses (part of) its significance and his objective value drops. Although this is indeed the objective of the policy, it results in less liquidity provision. Figure 17 illustrates that the HFT lowers his total liquidity provision under pro rata allocation, as it is now less likely for the HFT to trade with an LFT even if his order had time-priority. Since the HFT's objective is to make round-trip trades in minimum time, keeping the same quoting policy under pro rata allocation is now costly due to inventory risk. Consequently, the HFT is less willing to quote under a policy that diminishes his speed advantage. Specifically, the HFT will now find quoting at the second-best prices more attractive relative to the time-priority case as the HFT's probability of trading with a LFT at the best price is lower. A higher incentive to quote at the second-best prices leads to higher equilibrium bid-offer spreads as shown in Figure 17. Finally, the bottom panel in Figure 17 illustrates that the HFT's quoting, although lower across the volatility range, has similar sensitivity to volatility compared to time priority, so pro rata matching produces no improvement in terms of making the provision of liquidity countercyclical with respect to volatility.

Aside from the model, note that if quantities were not fixed, we would expect the HFT to quote larger quantities under pro rata matching; this incentive creates a separate risk, that of liquidity providers quoting larger quantities than what they can safely absorb since pro rata allocation makes it less likely that they would be forced to do so. In the case of extreme market moves, such book padding strategies, although harmless in normal times, could result in large losses (see McPartland (2015)).

8. Conclusions

We propose a theoretical model of high frequency market making. We superpose different Poisson processes running on different time clocks to represent the arrival of different elements of market information and orders, resulting in a tractable and flexible framework where the optimal market making strategy of the HFT and the equilibrium between the HFT quotes and incoming LFTs' orders is fully characterized. The model reproduces many important stylized facts about HFTs. We find that the HFT's liquidity provision increases when he gets faster. We find that the optimal quoting policy of the HFT also leads to cancellation rates of his orders in the presence of informative signals about the order flow. The model also quantifies the impact of higher volatility on liquidity provision. We find that the HFT will provide less liquidity when volatility increases if the arrival rate of LFTs is independent of price volatility. In the presence of positive correlation between the arrival rates of the LFTs and fundamental price volatility, the model predicts an inverse U-shaped pattern of liquidity provision: a small increase in volatility leads to more liquidity, followed by a decrease as volatility increases further. When analyzing competition for order flow with a second market maker, we find that competing market makers split the rent extracted from LFTs, liquidity provision increases and LFTs tend to be better off.

Finally, we provide the first model-based analysis of the impact of four widely discussed HFT policies: imposing a transactions tax, setting minimum-time limits before orders can be cancelled, taxing the cancellations of limit orders, and replacing time priority with pro rata or random allocation. We assess these regulatory policies on the basis of their potential to induce the HFT to provide liquidity that is more resilient to increases in volatility, i.e., countercyclical with respect to volatility. We find that none of the four policies result in an improvement. A transactions tax and pro rata matching result in less liquidity altogether. Minimum rest times and a cancellation tax result in more liquidity in good (low volatility) environments but less in bad (high volatility) environments, the opposite of the desired effect. Ultimately, our conclusion is that the microstructure-based speed and informational advantages of the HFT are difficult to even out.

References

Amihud, Y., Mendelson, H., 1980. Dealership market: Market-making and inventory. Journal of Financial Economics 8, 31–53.

Anand, A., Venkataraman, K., 2016. Market conditions, fragility and the economics of market making. Journal of Financial Economics 121, 327–349.

Angel, J., Harris, L., Spatt, C. S., 2015. Equity trading in the 21st century: An update. Quarterly Journal of Finance 5, 1–39.

Avellaneda, M., Stoikov, S., 2008. High-frequency trading in a limit order book. Quantitative Finance 8, 217–224.

Baruch, S., Glosten, L. R., 2013. Flickering quotes. Tech. rep., Working paper, Columbia University Graduate School of Business.

Biais, B., Foucault, T., Moinas, S., 2015. Equilibrium fast trading. Journal of Financial Economics , 292–313.

Biais, B., Woolley, P., 2011. High frequency trading. Tech. rep., Toulouse School of Economics and London School of Economics.

Brogaard, J. A., Hendershott, T., Riordan, R., 2014. High frequency trading and price discovery. Review of Financial Studies 27, 2267–2306.

Brunetti, C., Kirilenko, A., Mankad, S., 2011. Identifing high-frequency traders in electronic markets: Properties and forecasting. Tech. rep., Johns Hopkins University.

Budish, E., Cramton, P., Shim, J., 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. Quarterly Journal of Economics 130, 1547–1621.

Cartea, A., Donnelly, R. F., Jaimungal, S., 2013. Robust market making. Tech. rep., University of Oxford.

Cartea, A., Jaimungal, S., Ricci, J., 2014. Buy low, sell high: A high frequency trading perspective. SIAM Journal on Financial Mathematics 5, 415–444.

Cartea, A., Penalva, J., 2012. Where is the value in high frequency trading? Quarterly Journal of Finance 2, 1–46.

Chaboud, A., Chiquoine, B., Hjalmarsson, E., Vega, C., 2010. Rise of the machines: Algorithmic trading in the foreign exchange market. Tech. rep., IMF.

Cvitanić, J., Kirilenko, A., 2010. High frequency traders and asset prices. Tech. rep., Caltech.

Easley, D., de Prado, M. M. L., O'Hara, M., 2011. The microstructure of the "flash crash": Flow toxicity, liquidity crashes and the probability of informed trading. Journal of Portfolio Management 37, 118–128.

Foucault, T., Hombert, J., Rosu, I., 2016. News trading and speed. The Journal of Finance 71, 335–382.

Foucault, T., Kadan, O., Kandel, E., 2013. Liquidity cycles and make/take fees in electronic markets. The Journal of Finance 68, 299–341.

Garman, M. B., 1976. Market microstructure. Journal of Financial Economics 3, 257–275.

Glosten, L. R., Milgrom, P. R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.

Guéant, O., Lehalle, C.-A., Fernandez-Tapia, J., 2013. Dealing with the inventory risk: A solution to the market making problem. Mathematics and Financial Economics 7, 477–507.

Guilbaud, F., Pham, H., 2013. Optimal high-frequency trading with limit and market orders. Quantitative Finance 13, 79–94.

Hasbrouck, J., Saar, G., 2009. Technology and liquidity provision: The blurring of traditional definitions. Journal of Financial Markets 12, 143–172.

Hasbrouck, J., Saar, G., 2013. Low-latency trading. Journal of Financial Markets 16, 646–679.

Hendershott, T., Jones, C. M., Menkveld, A. J., 2011. Does algorithmic trading improve liquidity? The Journal of Finance 66, 1–33.

Hendershott, T., Menkveld, A. J., 2014. Price pressures. Journal of Financial Economics 114, 405–423.

Hirschey, N. H., 2011. Do high-frequency traders anticipate buying and selling pressure? Tech. rep., London Business School.

Ho, T. S. Y., Stoll, H. R., 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47–73.

Jarrow, R. A., Protter, P., 2012. A dysfunctional role of high frequency trading in electronic markets. International Journal of Theoretical and Applied Finance 15 (3), 1–15.

Jovanovic, B., Menkveld, A. J., 2010. Middlemen in limit-order markets. Tech. rep., New York University and VU University Amsterdam.

Kirilenko, A., Kyle, A. P., Samadi, M., Tuzun, T., 2010. The flash crash: The impact of high frequency trading on an electronic market. Tech. rep., University of Maryland.

McPartland, J. W., 2015. Recommendations for equitable allocation of trades in high frequency trading environments. Journal of Trading 10, 81–100.

Menkveld, A. J., 2013. High frequency trading and the new-market makers. Journal of Financial Markets 16, 712–740.

Menkveld, A. J., 2016. The economics of high-frequency trading: Taking stock. Annual Review of Financial Economics, forthcoming .

Moallemi, C., Sağlam, M., 2013. The cost of latency in high-frequency trading. Operations Research 61, 1070–1086.

Pagnotta, E., Philippon, T., 2011. Competing on speed. Tech. rep., NBER.

Puterman, M. L., 1994. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons.

Rosu, I., 2009. A dynamic model of the limit oder book. Review of Financial Studies 22, 4601–4641.

Sağlam, M., 2016. Order anticipation in large order executions. Available at SSRN .

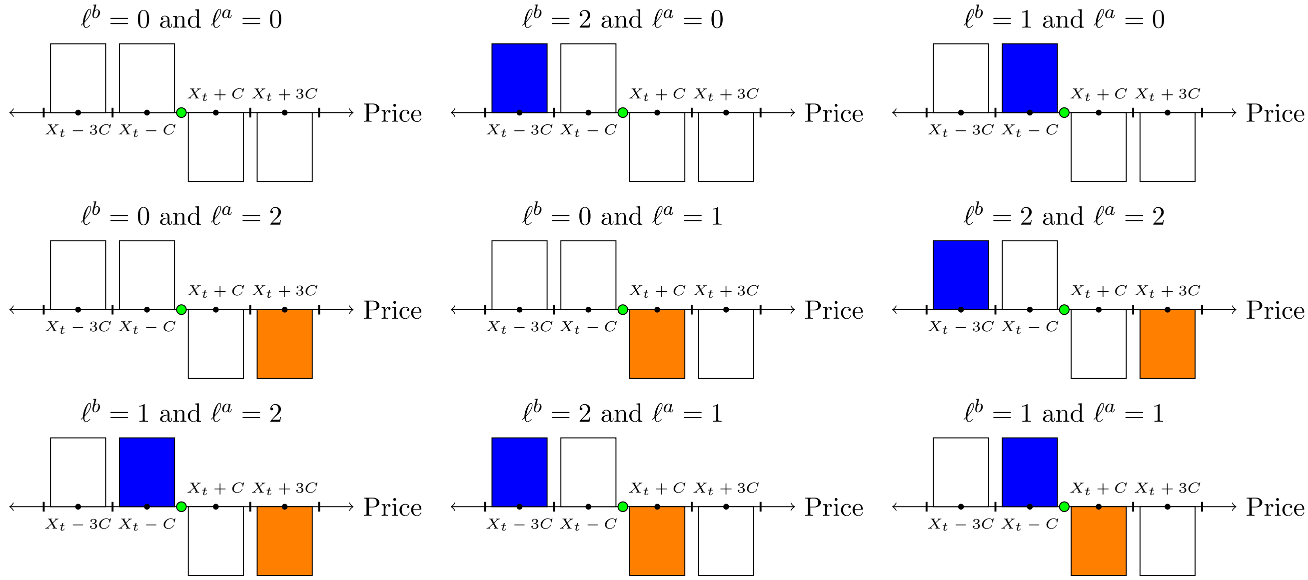

Fig. 1. HFT possible quotes.

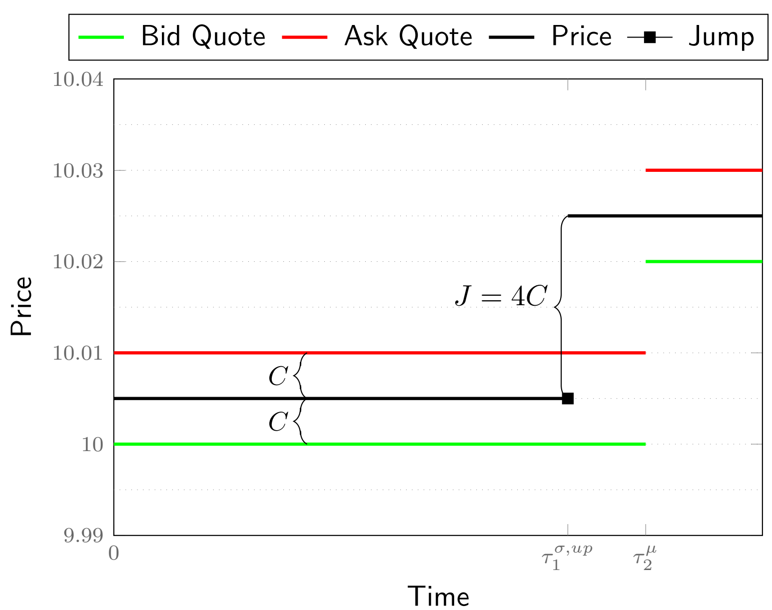

Fig. 2. Dynamics of the fundamental value: Example of a price jump.

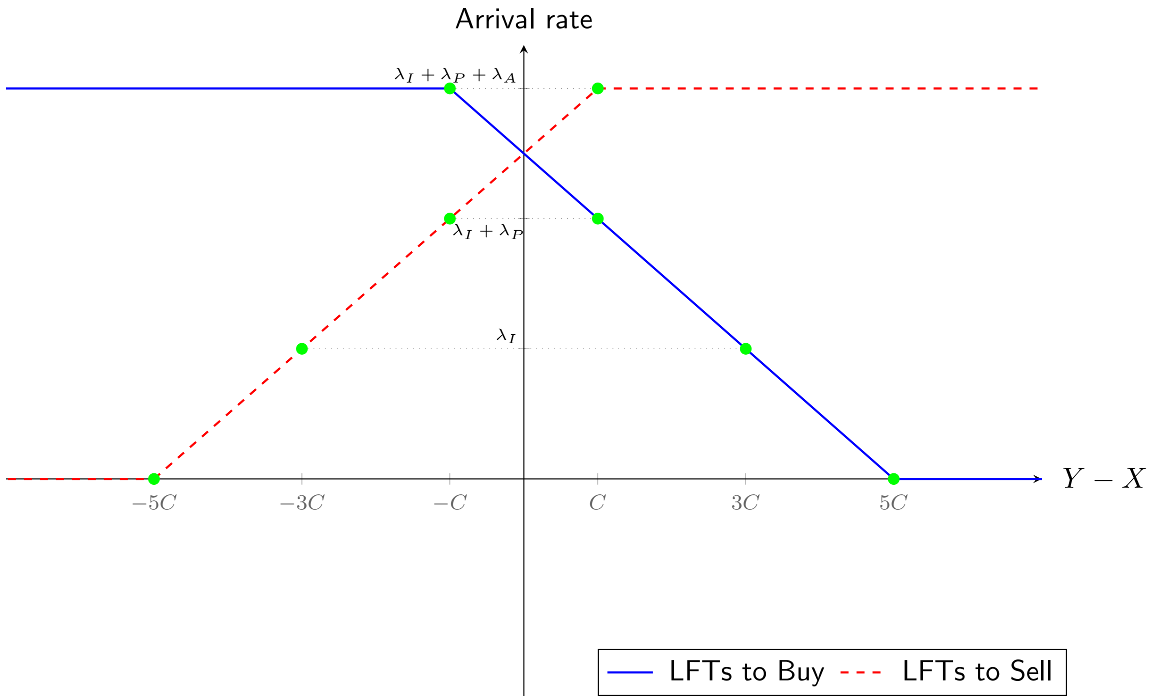

Fig. 3. Demand and supply curves from LFTs.

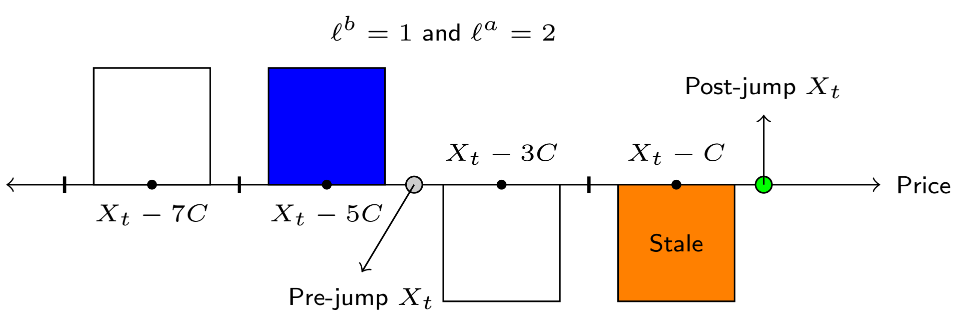

Fig. 4. Potential stale quotes of the HFT in the presence of price jumps.

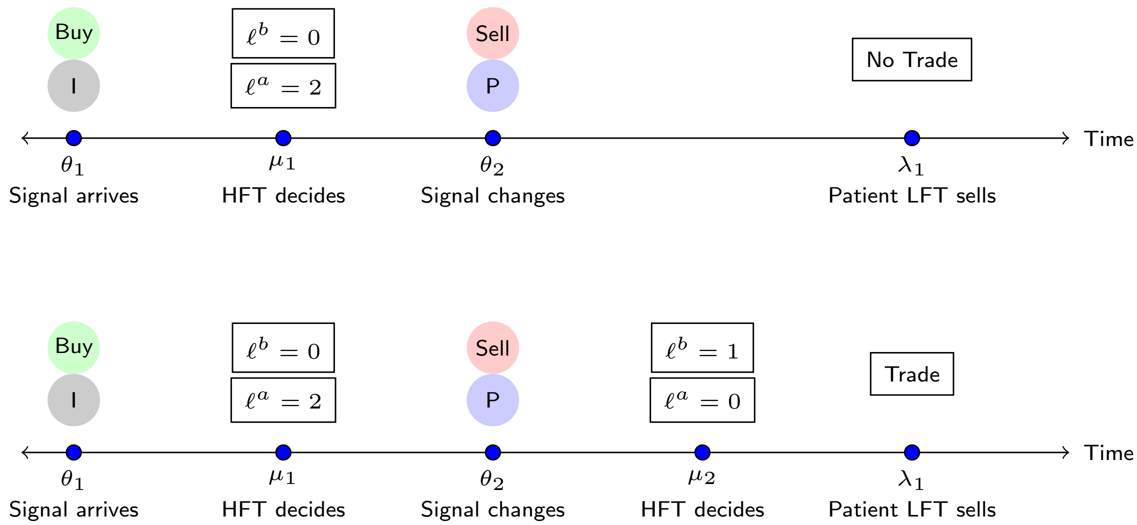

Fig. 5. Potential sequence of signal and order arrivals and HFT's corresponding actions.

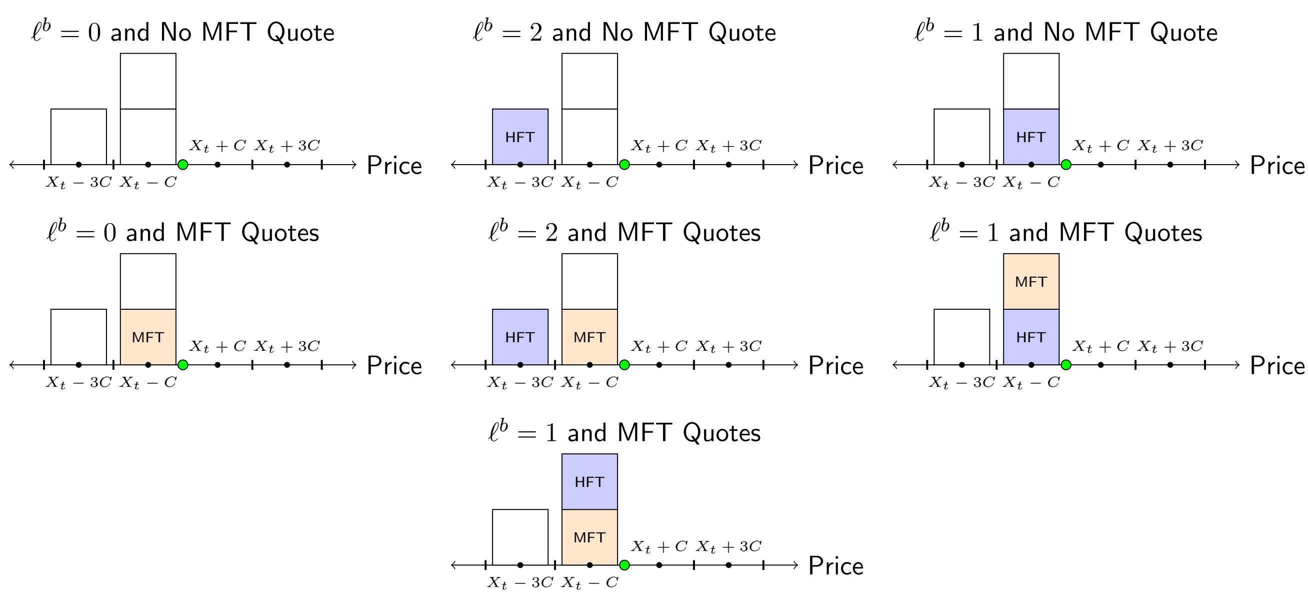

Fig. 6. The 7 possible states of the bid side of the book.

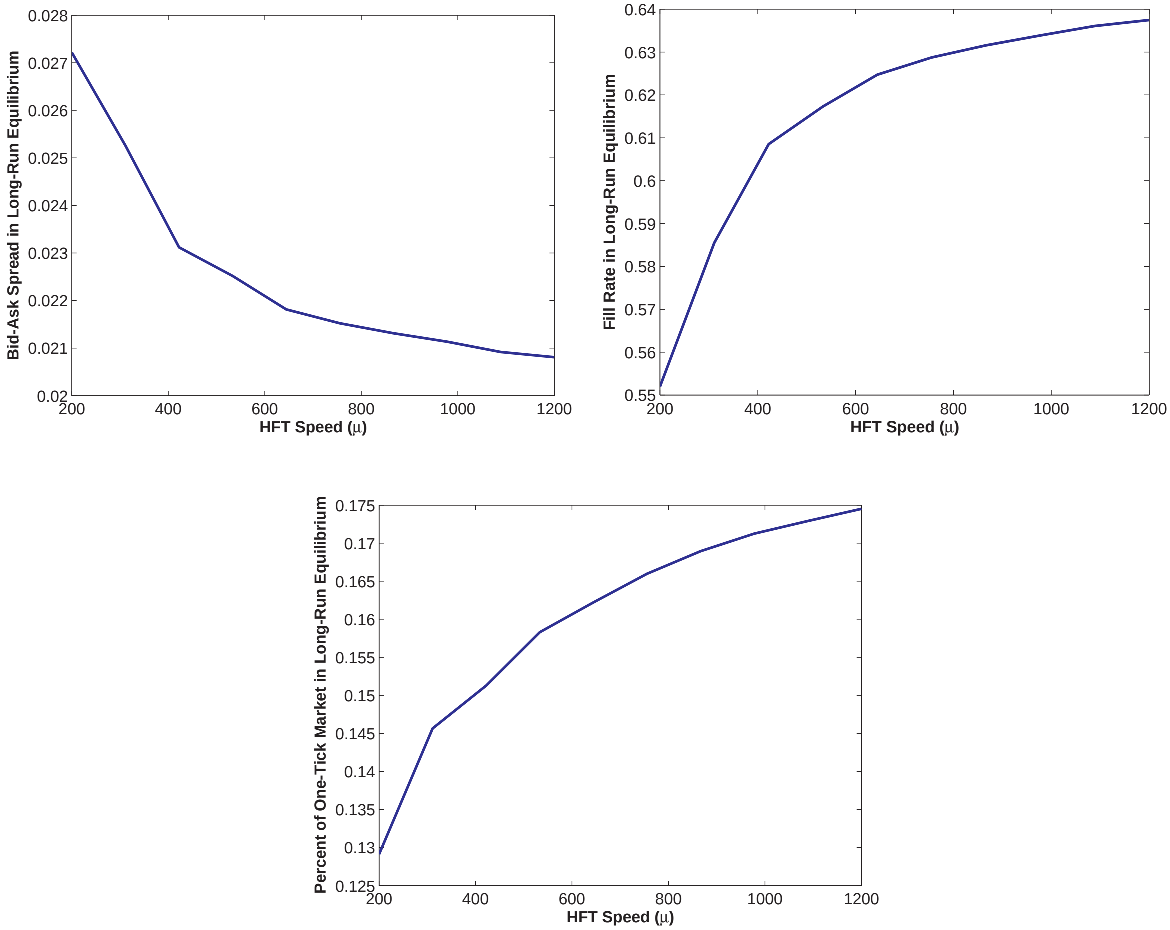

Fig. 7. Liquidity measures in long-run equilibrium as a function of the HFT's speed.

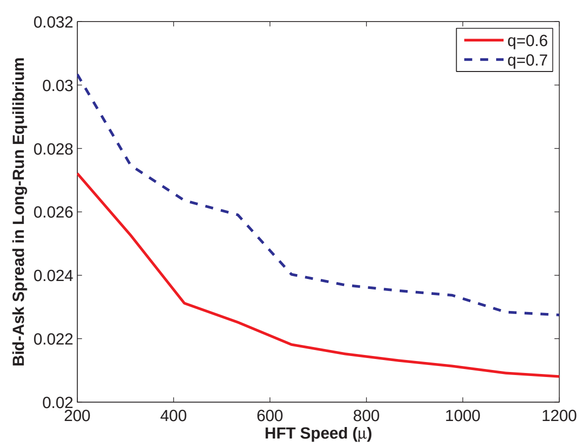

Fig. 8. Bid-offer spread in long-run equilibrium as a function of HFT's speed for two regimes of HFT's predictive ability regarding the LFT types.

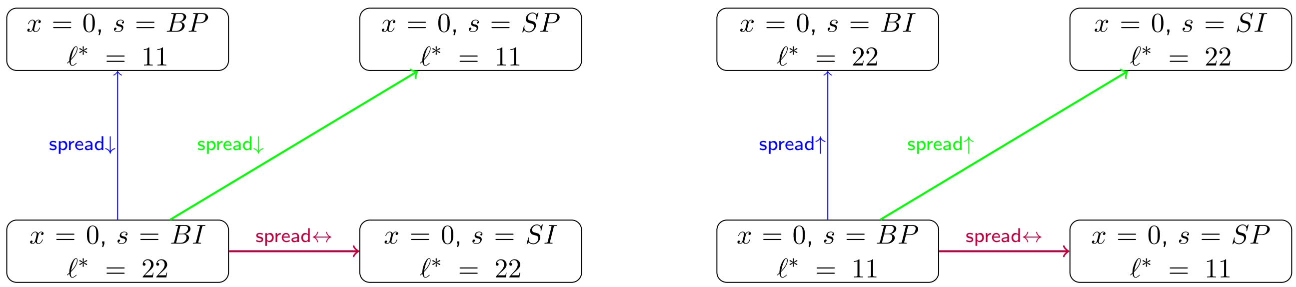

Fig. 9. HFT's quote changes under different inventory and signal states and their impact on the spread.

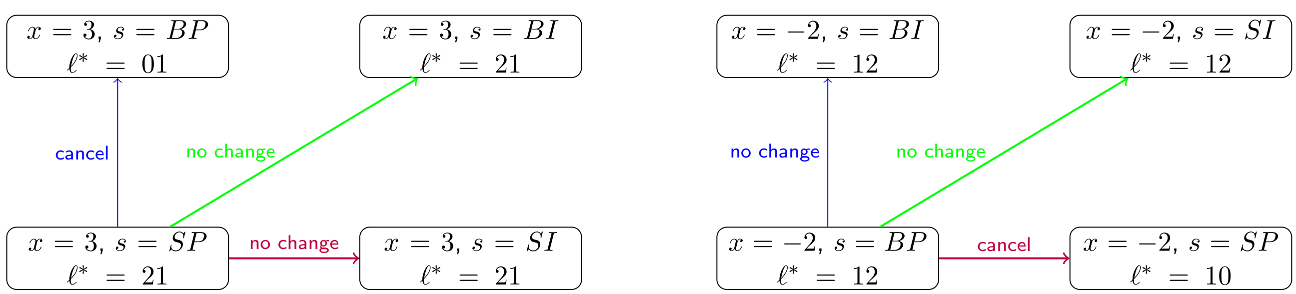

Fig. 10. HFT's order revisions and cancellations under different inventory and signal states.

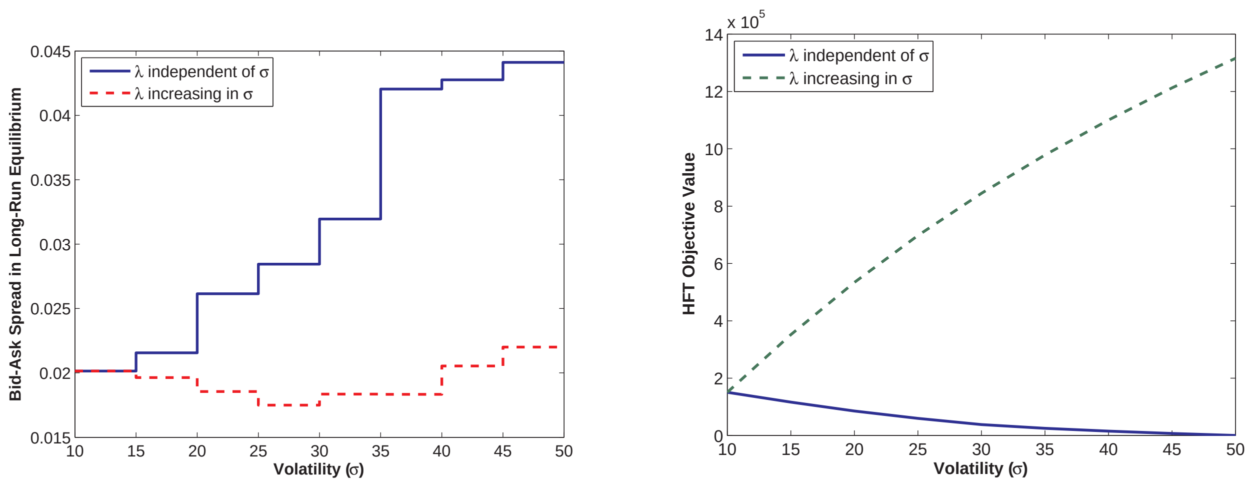

Fig. 11. Long-run equilibrium bid-offer spread as a function of price volatility.

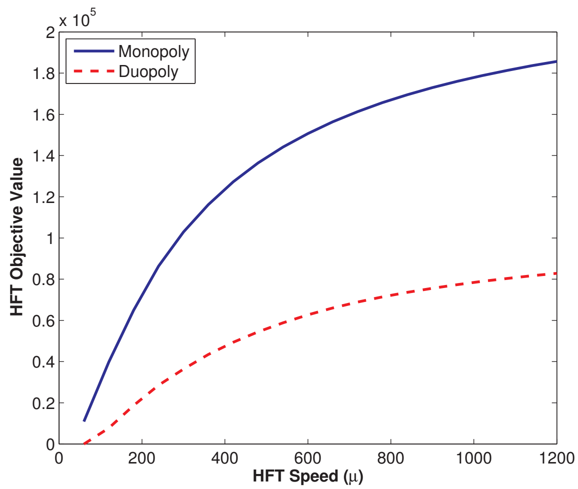

Fig. 12. HFT Optimal Values in the Monopoly (No Queuing) and Duopoly (Queuing) Situations.

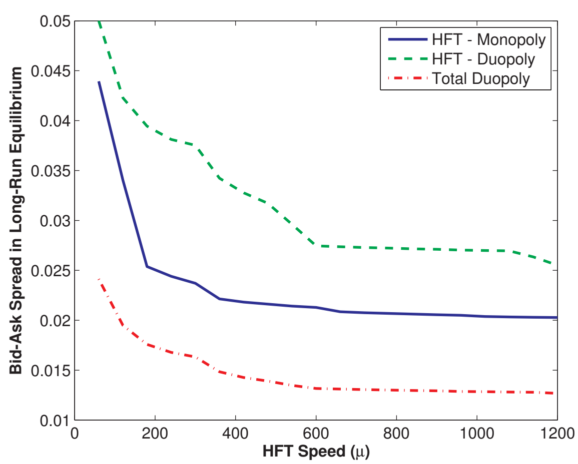

Fig. 13. Long-run equilibrium spread in the presence of the MFT as a function of the HFT's speed.

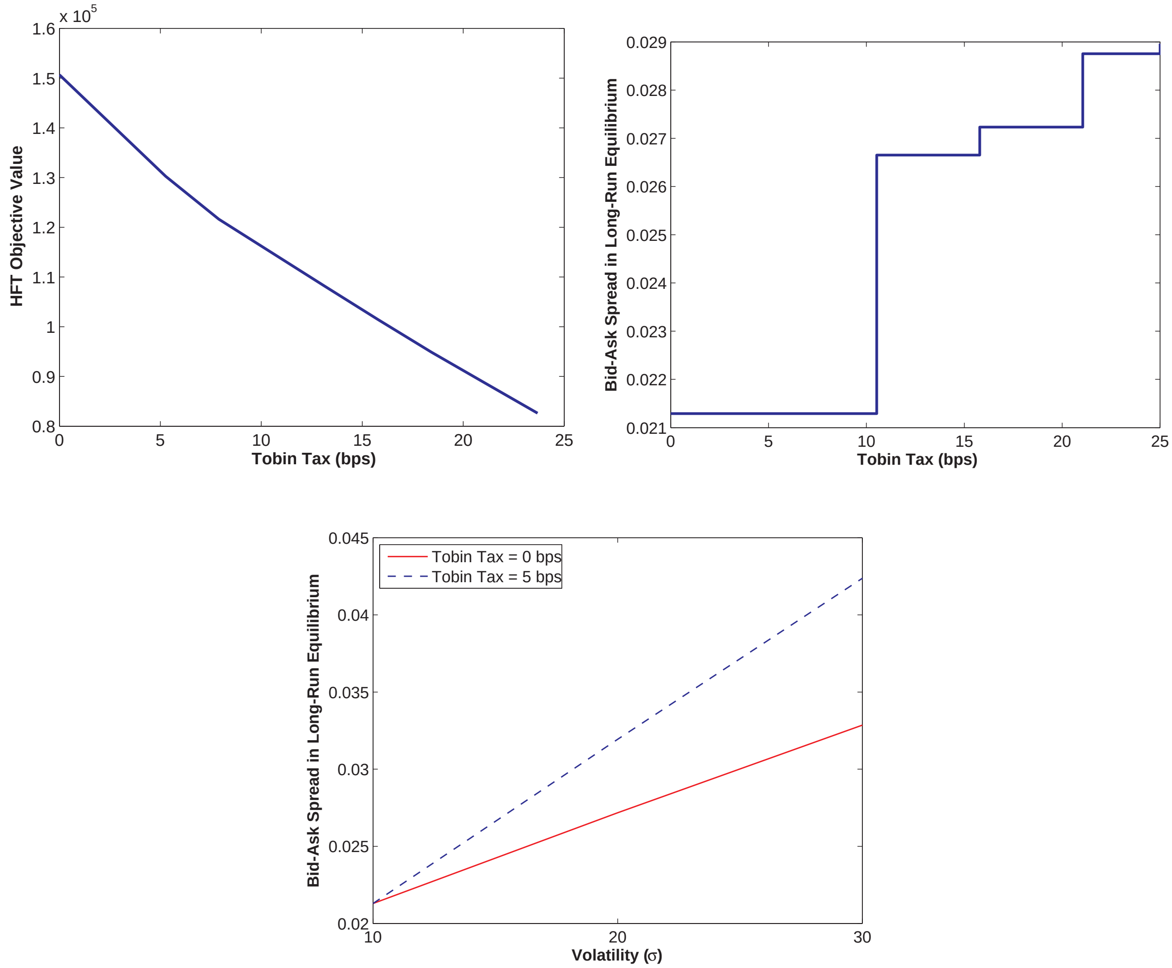

Fig. 14. Impact of a Tobin Tax

Notes: The top panels plot the effect of taxing transactions on the HFTs' value and his provision of liquidity. The bottom panel displays the sensitivity of the equilibrium bid-offer spread to volatility (in the form of price jumps) before and after the Tobin tax.

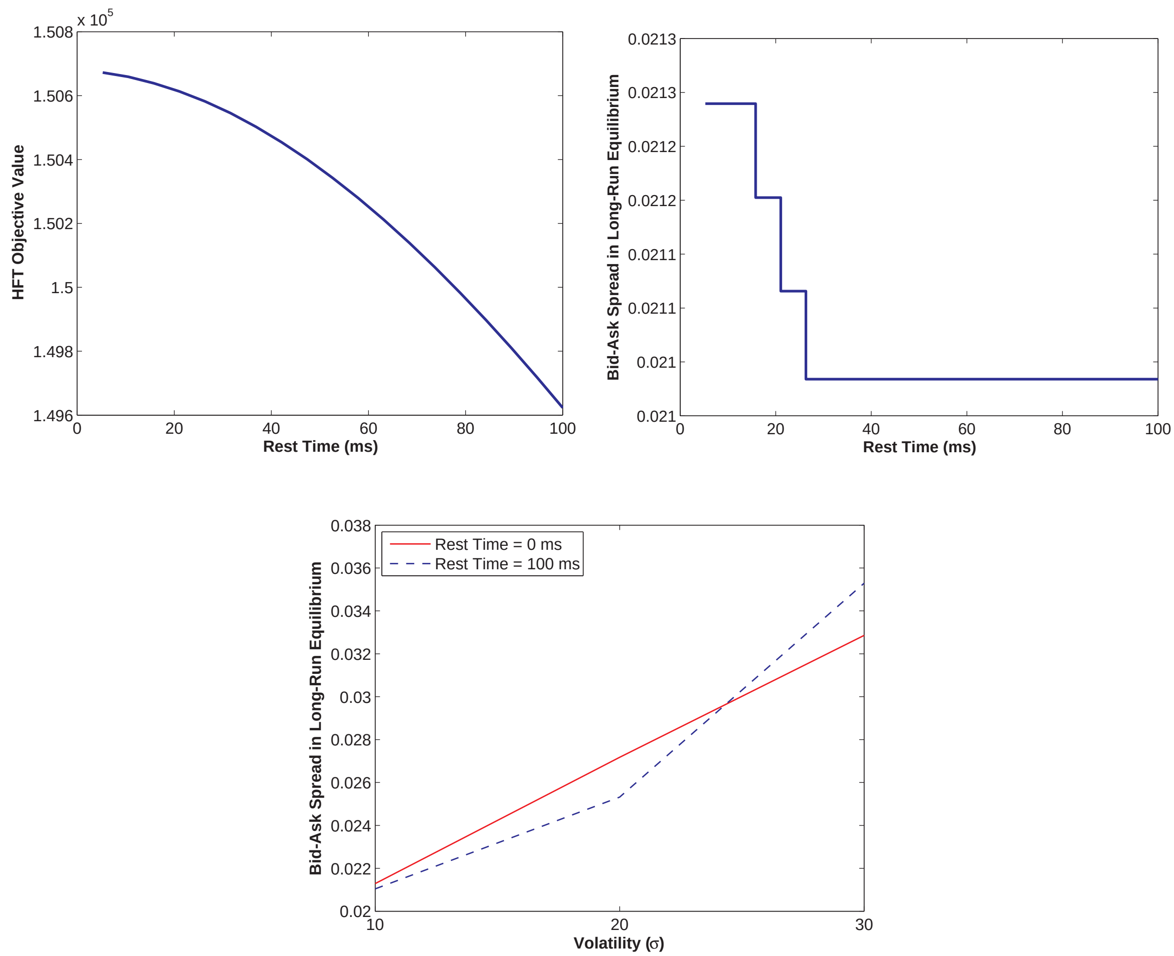

Fig. 15. Impact of a Mandatory Minimum Rest Time on HFT Quotes

Notes: The top panels plot the effect of minimum rest times on the HFTs' value and his provision of liquidity. The bottom panel displays the sensitivity of the equilibrium bid-offer spread to volatility (in the form of price jumps) before and after the minimum rest time.

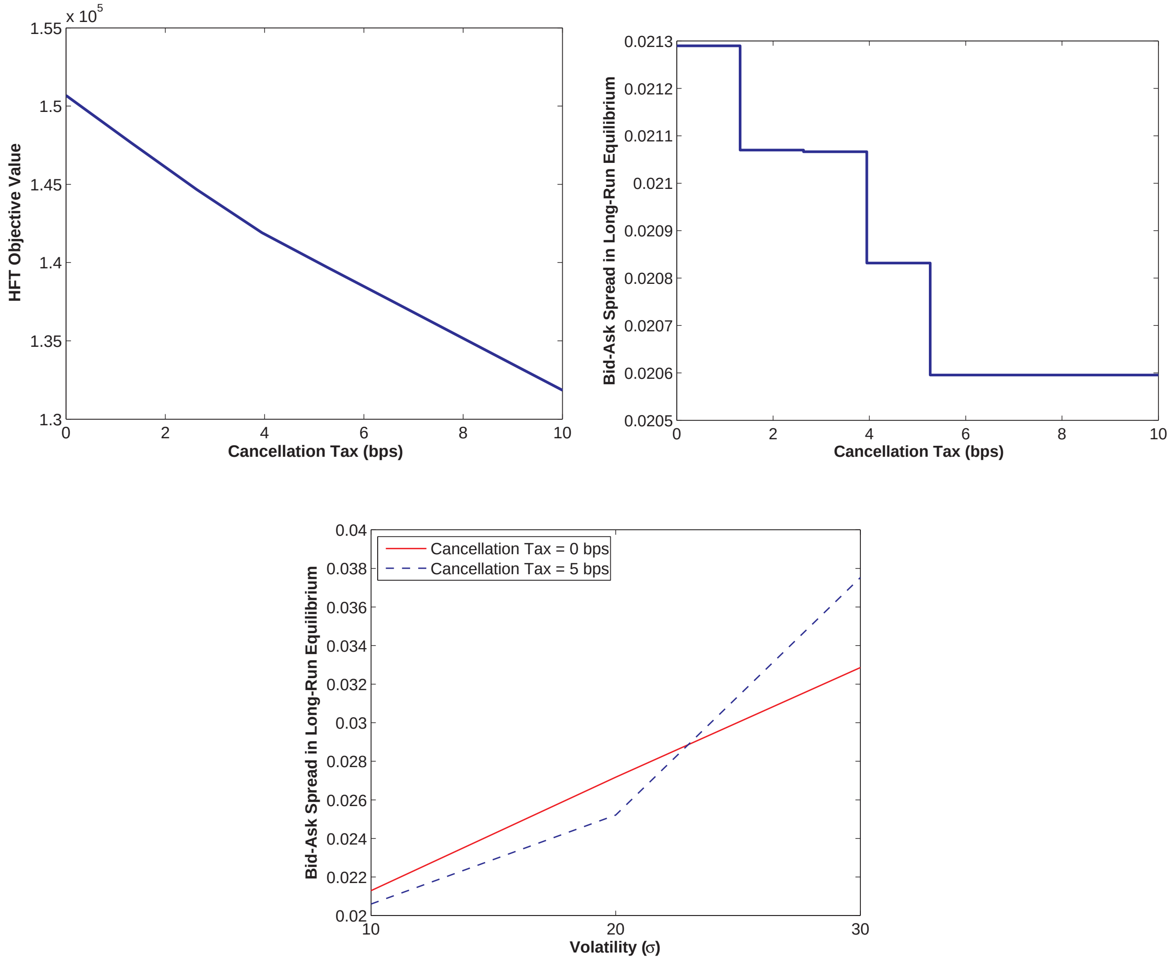

Fig. 16. Impact of a Cancellation Tax

Notes: The top panels plot the effect of taxing cancellations on the HFTs' value and his provision of liquidity. The bottom panel displays the sensitivity of the equilibrium bid-offer spread to volatility (in the form of price jumps) before and after the cancellation tax.

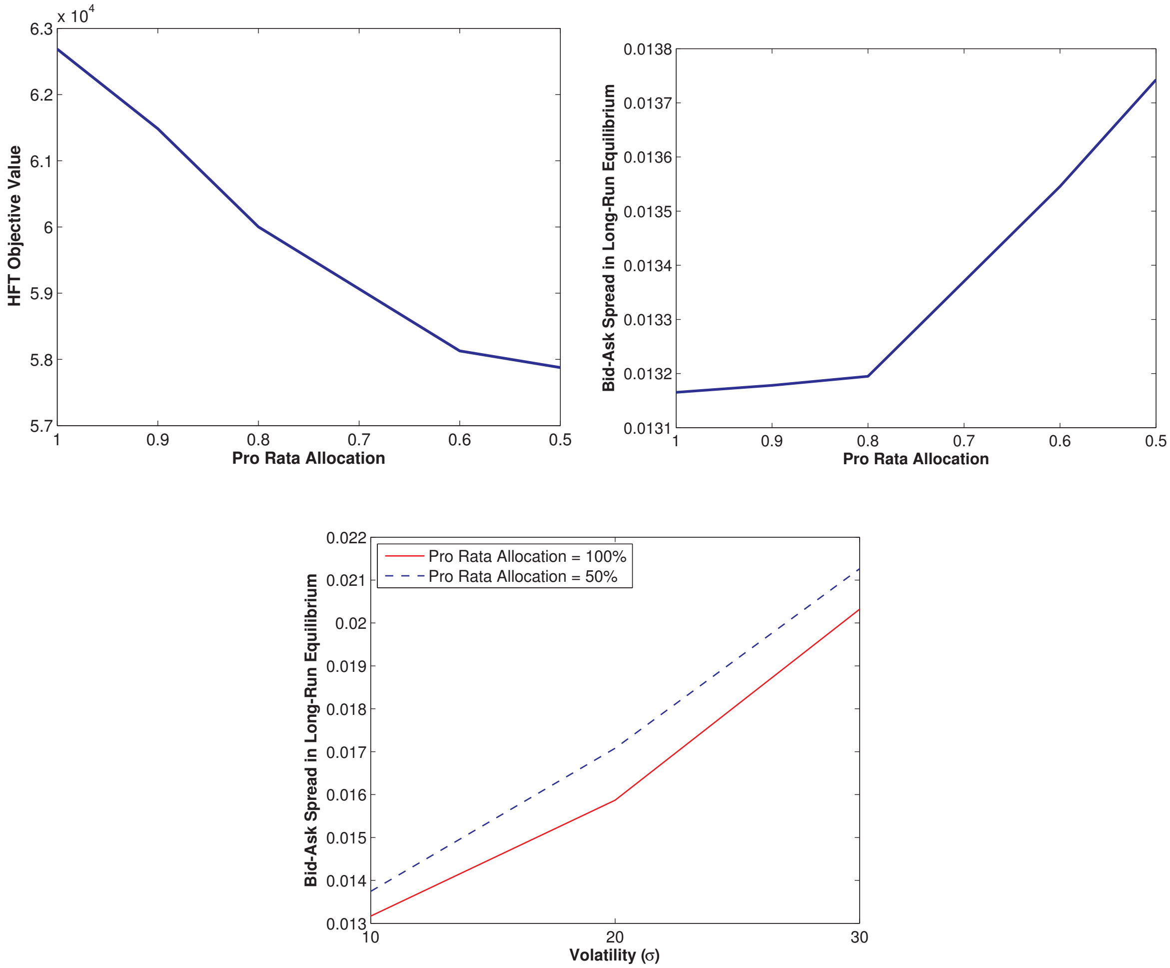

Fig. 17. Impact of Pro Rata Allocation

Notes: The top panels plot the effect of various pro rata levels on the HFTs' value and his provision of liquidity. The bottom panel displays the sensitivity of the equilibrium bid-offer spread to volatility (in the form of price jumps) before and after the pro rata policy with 50%.

Appendix for "High Frequency Market Making"
Technical Results and Proofs

A. Discrete-time Embedding in the Simplified Model

We start by recalling the definition of a discounted infinite horizon Markov Decision Process (MDP), before showing that our continuous-time HFT optimization problem can be represented as such. A MDP is defined by a 4-tuple, $(I, A_i, \mathbb{P}(.|i,a), \mathbb{R}(.|i,a))$, in which I is the state space, A_i is the action space, i.e., the set of possible actions that a decision maker can take when the state is $i \in I$, $\mathbb{P}(.|i,a)$ is the probability transition matrix determining the state of the system in the next decision time, and finally $\mathbb{R}(.|i,a)$ is the reward matrix, specifying the reward obtained using action a when the state is in i. The HFT seeks a quoting policy that maximizes the expected discounted reward

$$v(i) = \max_{\pi} \mathbb{E}\left[\sum_{t=0}^{\infty} \alpha^t \mathbb{R}(i_{t+1}|i_t, \pi(i_t))|i_0 = i\right], \tag{A.1}$$

where α is the discount rate. An admissible stationary policy π maps each state $i \in I$ to an action in A_i. Under mild technical conditions, we can guarantee the existence of optimal stationary policies (see Puterman (1994)). Conditioning on the first transition from i to i', we obtain the Hamilton-Jacobi-Bellman optimality equation

$$\begin{aligned} v(i) &= \max_{\pi} \left\{ \sum_{i'} \mathbb{P}(i'|i,\pi(i)) \left(\mathbb{R}(i'|i,\pi(i)) + \alpha \mathbb{E}\left[\sum_{t=1}^{\infty} \alpha^{t-1} \mathbb{R}(i_{t+1}|i_t,\pi(i_t))|i_1 = i' \right] \right) \right\} \\ &= \max_{\pi} \left\{ \sum_{i'} \mathbb{P}(i'|i,\pi(i)) \left(\mathbb{R}(i'|i,\pi(i)) + \alpha \mathbb{E}\left[\sum_{k=0}^{\infty} \alpha^{k} \mathbb{R}(i_{k+1}|i_k,\pi(i_k))|i_k = i' \right] \right) \right\} \\ &= \max_{a \in A_i} \left\{ \sum_{i'} \mathbb{P}(i'|i,a) \left(\mathbb{R}(i'|i,a) + \alpha v(i') \right) \right\}. \end{aligned} \tag{A.2}$$

A.1. Transition Probabilities

We now calculate the transition probabilities at each state of the HFT in the simplified model. First, note that the state transitions occur at a rate of $\mu + \lambda + \theta$ where $\lambda \equiv 2\lambda_P$ and the transition rate is the same for all states and actions. Let $\mathbb{P}((x', s', l', e')|(x, s, l, e), (\ell^b, \ell^a))$ be the probability of reaching state (x', s', l', e') when the system is in state (x, s, l, e) and the trader takes the actions of ℓ^b and ℓ^a.

First, we define

$$\mathsf{tr}(s) = \begin{cases} p & \text{if } s = S, \\ 1-p & \text{if } s = B. \end{cases}$$

Suppose that the current state of the HFT is (x, s, l, e). Let $r = (\lambda + \mu + \theta)$. We provide the transition probabilities with respect to four possible actions that can be employed by the HFT at decision epochs with $\ell^b \in \{0, 1\}$ and $\ell^a \in \{0, 1\}$.

First, if the HFT does not quote in either side of the market, the inventory level in the next state cannot change. If a new decision event arrives before the arrival of a market order or a signal, the state remains the same. Otherwise, we know that either a market order or a signal arrived before HFT makes a new quoting decision. Formally, we obtain

$$\mathbb{P}\left((x', s', l', e')|(x, s, l, 1), (0, 0)\right) = \begin{cases} \frac{(\lambda+\theta)}{2r} & \text{if } x = x', s' \in \{B, S\}, e' = 0, l' = 00, \\ \frac{\mu}{r} & \text{if } x = x', s = s', e' = 1, l' = 00, \\ 0 & \text{otherwise.} \end{cases}$$

When the HFT takes the action $(1, 0)$, he may increase his inventory by trading with the incoming market-sell order submitted by a patient LFT, which occurs with probability $\frac{\mathsf{tr}(s)\lambda}{r}$.

$$\mathbb{P}\left((x', s', l', e')|(x, s, l, 1), (1, 0)\right) = \begin{cases} \frac{(\lambda(1-\mathsf{tr}(s))+\theta)}{2r} & \text{if } x = x', s' \in \{B, S\}, e' = 0, l' = 10 \\ \frac{\lambda \mathsf{tr}(s)}{2r} & \text{if } x + 1 = x', s' \in \{B, S\}, e' = 0, l' = 00 \\ \frac{\mu}{r} & \text{if } x = x', e' = 1, s = s', l' = 10 \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, if the HFT takes the action $(0, 1)$, he may increase his inventory by trading with the incoming market-sell order submitted by a patient LFT, which occurs with probability $\frac{\mathsf{tr}(s)\lambda}{r}$. Similarly, he may decrease his inventory by only trading with the incoming market-buy order submitted by a patient

LFT, which occurs with probability $\frac{(1-\mathsf{tr}(s))\lambda}{r}$.

$$\mathbb{P}\left((x',s',l',e')|(x,s,l,1),(1,1)\right) = \begin{cases} \frac{\theta}{2r} & \text{if } x=x', s'\in\{B,S\}, e'=0, l'=11 \\ \frac{\lambda \mathsf{tr}(s)}{2r} & \text{if } x+1=x', s'\in\{B,S\}, e'=0, l'=01 \\ \frac{\lambda(1-\mathsf{tr}(s))}{2r} & \text{if } x-1=x', s'\in\{B,S\}, e'=0, l'=10 \\ \frac{\mu}{r} & \text{if } x=x', e'=1, s=s', l'=11 \\ 0 & \text{otherwise.} \end{cases}$$

If the system is observed at the arrival time of a market order or signal event, that is $e=0$, the HFT cannot revise his quotes. We can accommodate these states in our model using fake decisions that merely sets the action to the existing quotes tracked by $l=(l^b,l^a)$. In this case,

$$\mathbb{P}\Big((x',s',l',e')|(x,s,l,0),(\ell^b,\ell^a)\Big) = \begin{cases} \mathbb{P}\left((x',s',l',e')|(x,s,l,1),(\ell^b,\ell^a)\right) & \text{if } \ell^b=l^b, \ell^a=l^a \\ 0 & \text{otherwise.} \end{cases}$$

A.2. HFT's Reward Function

Let $\mathbb{R}\left((x',s',l',e')|(x,s,l,e),(\ell^b,\ell^a)\right)$ be the total reward achieved by the HFT when the system is in state (x,s,l,e), the HFT chooses quoting actions ℓ^b and ℓ^a and the system reaches the state (x',s',l',e'). In the simplified model, the HFT will be able to make C from trade events and will not lose anything in the absence of price jumps.

We would like to write the HFT's objective in (2.13) in the form of an MDP objective function as in (A.1). We first introduce the following notation. Let t_k be the time of the kth state transition due to a decision, signal or market order arrival (by convention $t_0=0$) and let τ_k be the length of this cycle, i.e., $\tau_k=t_k-t_{k-1}$. We start with the positive reward terms in $G^-(.)$ and $G^+(.)$ that measure the spreads earned by the HFT when there is a trade. We can track this sum of the discounted rewards in our MDP framework with

$$\sum_{k=1}^{\infty} e^{-Dt_k}\mathbb{R}^+\Big((x'_{t_k},s'_{t_k},l'_{t_k},e'_{t_k})|(x_{t_{k-1}},s_{t_{k-1}},l_{t_{k-1}},e_{t_{k-1}}),(\ell^b_{t_{k-1}},\ell^a_{t_{k-1}})\Big),$$

where

$$\mathbb{R}^+\Big((x',s',l',e')|(x,s,l,e),(\ell^b,\ell^a)\Big) = C\mathbb{1}\left(x+1=x'\right)\mathbb{1}\left(\ell^b=1\right)$$

We can take the expectation of the HFT's discounted earnings using the independence of each cycle length, τ_i, which is an exponentially distributed random variable with mean $1/r$:

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{k=1}^{\infty} e^{-Dt_k}\mathbb{R}^+\left((x'_{t_k}, s'_{t_k}, b'_{t_k}, e'_{t_k})|(x_{t_{k-1}}, s_{t_{k-1}}, b_{t_{k-1}}, e_{t_{k-1}}), (\ell^b_{t_{k-1}}, \ell^a_{t_{k-1}})\right)\right] \\
&= \sum_{k=1}^{\infty} \mathbb{E}\left[e^{-D\sum_{i=1}^{k}\tau_i}\right] \mathbb{E}\left[\mathbb{R}^+\left((x'_{t_k}, s'_{t_k}, b'_{t_k}, e'_{t_k})|(x_{t_{k-1}}, s_{t_{k-1}}, b_{t_{k-1}}, e_{t_{k-1}}), (\ell^b_{t_{k-1}}, \ell^a_{t_{k-1}})\right)\right] \\
&= \sum_{k=1}^{\infty} \mathbb{E}\left[e^{-D\tau_1}\right]^k \mathbb{E}\left[\mathbb{R}^+\left((x'_{t_k}, s'_{t_k}, b'_{t_k}, e'_{t_k})|(x_{t_{k-1}}, s_{t_{k-1}}, b_{t_{k-1}}, e_{t_{k-1}}), (\ell^b_{t_{k-1}}, \ell^a_{t_{k-1}})\right)\right] \\
&= \sum_{k=1}^{\infty} \left(\int_0^{\infty} re^{-(r+D)t}dt\right)^k \mathbb{E}\left[\mathbb{R}^+\left((x'_{t_k}, s'_{t_k}, b'_{t_k}, e'_{t_k})|(x_{t_{k-1}}, s_{t_{k-1}}, b_{t_{k-1}}, e_{t_{k-1}}), (\ell^b_{t_{k-1}}, \ell^a_{t_{k-1}})\right)\right] \\
&= \sum_{k=1}^{\infty} \left(\frac{r}{r+D}\right)^k \mathbb{E}\left[\mathbb{R}^+\left((x'_{t_k}, s'_{t_k}, b'_{t_k}, e'_{t_k})|(x_{t_{k-1}}, s_{t_{k-1}}, b_{t_{k-1}}, e_{t_{k-1}}), (\ell^b_{t_{k-1}}, \ell^a_{t_{k-1}})\right)\right] \\
&= \delta \sum_{k=0}^{\infty} \delta^k \mathbb{E}\left[\mathbb{R}^+\left((x'_{t_{k+1}}, s'_{t_{k+1}}, b'_{t_{k+1}}, e'_{t_{k+1}})|(x_{t_k}, s_{t_k}, b_{t_k}, e_{t_k}), (\ell^b_{t_k}, \ell^a_{t_k})\right)\right].
\end{aligned}
$$

where δ is the "adjusted discount factor," defined in (3.2). Inventory costs in the third term of (2.13) can be simplified as

$$
\begin{aligned}
\mathbb{E}\left[\Gamma \int_0^{\infty} e^{-Dt}|x_t|dt\right] &= \Gamma \sum_{k=0}^{\infty} \mathbb{E}\left[\int_{t_k}^{t_{k+1}} e^{-Dt}|x_t|dt\right] \\
&= \Gamma \sum_{k=0}^{\infty} \mathbb{E}\left[\int_{t_k}^{t_{k+1}} e^{-Dt}dt\right] \mathbb{E}\left[|x_{t_k}|\right] \\
&= \frac{\Gamma}{D} \sum_{k=0}^{\infty} \mathbb{E}\left[e^{-Dt_k}\right]\left(1 - \mathbb{E}\left[e^{-D\tau_{k+1}}\right]\right) \mathbb{E}\left[|x_{t_k}|\right] \\
&= \frac{\Gamma}{D} \sum_{k=0}^{\infty} \delta^k \left(\frac{D}{\lambda+\mu+D}\right) \mathbb{E}\left[|x_{t_k}|\right] \\
&= \frac{\Gamma}{r+D} \sum_{k=0}^{\infty} \left(\frac{r}{r+D}\right)^k \mathbb{E}\left[|x_{t_k}|\right].
\end{aligned}
$$

Let $\mathbb{R}\left((x', s', b', e')|(x, s, l, e), (\ell^b, \ell^a)\right)$ be the probability of reaching state (x', s', l', e') when the system is in state (x, s, l, e) and the trader takes the actions of ℓ^b and ℓ^a. We are now ready to define the total reward matrix. Let

$$
\mathbb{R}\left((x', s', l', e')|(x, s, l, e), (\ell^b, \ell^a)\right) = c\mathbb{1}\left(x + 1 = x'\right)\mathbb{1}\left(\ell^b = 1\right) + c\mathbb{1}\left(x - 1 = x'\right)\mathbb{1}\left(\ell^a = 1\right) - \gamma|x| \tag{A.3}
$$

where c and γ, defined in (3.2) are the "adjusted spread" and "adjusted inventory aversion" parameters for the discrete-time formulation. Then, the HFT maximizes

$$
V(x, s, l, e) = \max_{\pi} \mathbb{E}^{\pi}\left[\sum_{k=0}^{\infty} \delta^k \mathbb{R}\left((x'_{t_{k+1}}, s'_{t_{k+1}}, b'_{t_{k+1}}, e'_{t_{k+1}})|(x_{t_k}, s_{t_k}, b_{t_k}, e_{t_k}), (\ell^b_{t_k}, \ell^a_{t_k})\right)\right], \tag{A.4}
$$

starting from his initial state, (x, s, l, e), which is in the requisite MDP form.

A.3. *HFT's Value Function*

We have now transformed our continuous-time problem into an equivalent discrete-time MDP. Using the Hamilton-Jacobi-Bellman optimality equations, $V(x, s, l, e)$ in (A.4) can be computed by solving the following set of equations:

$$V(x,s,l,e) = \max_{\ell^b,\ \ell^a} \Bigg\{ \sum_{(x',s',l',e')} \mathbb{P}\left((x',s',l',e')|(x,s,l,e),(\ell^b,\ell^a)\right) \Big[\mathbb{R}\left((x',s',l',e')|(x,s,l,e),(\ell^b,\ell^a)\right) + \delta V(x',s',l',e')\Big]\Bigg\}. \tag{A.5}$$

By substituting the expressions for $\mathbb{P}\left((x',s',l',e')|(x,s,l,e),(\ell^b,\ell^a)\right)$ and $\mathbb{R}\left((x',s',l',e')|(x,s,l,e),(\ell^b,\ell^a)\right)$, we can obtain the implicit equations for the value functions of each state.

Proof of Proposition 1. Since the model is symmetric around the bid and ask side of the market, we can first eliminate s from our state space. We have that:

$$V(-x,S,l,e) = \begin{cases} V(x,B,l,e) & \text{when } l \in \{00,11\}, \\ V(x,B,01,e) & \text{when } l = 10, \\ V(x,B,10,e) & \text{when } l = 01. \end{cases}$$

Using this result, we let $v(x,l) \equiv V(x,S,l,0)$ and $h(x) \equiv V(x,S,l,1)$ for ease in notation and obtain the following set of equations for the value functions:

$$\begin{aligned}
v(x,00) &= -\gamma|x| + \delta\left(\frac{\mu}{r}h(x,00) + \frac{\theta+\lambda}{2r}v(x,00) + \frac{\theta+\lambda}{2r}v(-x,00)\right) \\
v(x,10) &= -\gamma|x| + \delta\Big(\frac{\mu}{r}h(x,10) + \frac{\theta+(1-p)\lambda}{2r}v(x,10) + \frac{\theta+(1-p)\lambda}{2r}v(-x,01) + \frac{p\lambda}{2r}v(x+1,00) \\
&\quad + \frac{p\lambda}{2r}v(-x-1,00) + \frac{pc\lambda}{r}\Big) \\
v(x,01) &= -\gamma|x| + \delta\Big(\frac{\mu}{r}h(x,01) + \frac{\theta+p\lambda}{2r}v(x,01) + \frac{\theta+p\lambda}{2r}v(-x,10) + \frac{(1-p)\lambda}{2r}v(x-1,00) \\
&\quad + \frac{(1-p)\lambda}{2r}v(-x+1,00) + \frac{(1-p)c\lambda}{r}\Big) \\
v(x,11) &= -\gamma|x| + \delta\Big(\frac{\mu}{r}h(x,11) + \frac{\theta}{2r}(v(x,11)+v(-x,11)) + \frac{p\lambda}{2r}v(x+1,01) + \frac{p\lambda}{2r}v(-x-1,10) \\
&\quad + \frac{(1-p)\lambda}{2r}v(x-1,10) + \frac{(1-p)\lambda}{2r}v(-x+1,01) + \frac{c\lambda}{r}\Big)
\end{aligned}$$

and

$$h(x) = \max\left\{v(x,00), v(x,01), v(x,10), v(x,11)\right\}.$$

$\square$

A.4. *Optimal Market Making Solution*

Proof of Theorem 1. First, we prove some auxiliary results.

Lemma 3. *If* $x \leq -\frac{c}{\gamma(1-\delta)}$ *(*$x \geq \frac{c}{\gamma(1-\delta)}$*) then* $\ell^{b*} = 1$ *and* $\ell^{a*} = 0$ *(*$\ell^{b*} = 0$ *and* $\ell^{a*} = 1$*)*

Proof. We know that the discounted expected cost between decision epochs is $\gamma|x|$. We know that the maximum discounted revenue from earning spreads is less than $\frac{c}{1-\delta}$. Thus, you would not quote to sell (buy) if $x \leq -\frac{c}{\gamma(1-\delta)}$ ($x \geq \frac{c}{\gamma(1-\delta)}$). □

We establish by induction that $v(x,l)$ is concave in x, $v(x,11) - v(x,01)$, $v(x,10) - v(x,00)$ is non-increasing in x and $v(x,10) - v(x,11)$, $v(x,01) - v(x,00)$ is non-decreasing in x.

We use induction on the steps of the dynamic programming operator. We focus on $v(x,l)$ for simplicity. Let $v^{(0)}(x,l) = 0$ for all x and l. Then, in the base case, we obtain

$$\begin{aligned}
v^{(1)}(x,00) &= -\gamma|x|,\\
v^{(1)}(x,10) &= -\gamma|x| + \frac{\delta p\lambda c}{r},\\
v^{(1)}(x,01) &= -\gamma|x| + \frac{\delta(1-p)\lambda c}{r},\\
v^{(1)}(x,11) &= -\gamma|x| + \frac{\delta\lambda c}{r},
\end{aligned}$$

which shows that $v^{(1)}(x,l)$ is concave. Assume that $v^{(n)}(x,l)$ satisfies the induction hypothesis. Then, $v^{(n+1)}(x,l)$ satisfies the following set of equations:

$$\begin{aligned}
v^{(n+1)}(x,00) &= -\gamma|x| + \delta\left(\tfrac{\mu}{r}h^{(n)}(x) + \tfrac{\theta+\lambda}{2r}v^{(n)}(x,00) + \tfrac{\theta+\lambda}{2r}v^{(n)}(-x,00)\right)\\
v^{(n+1)}(x,10) &= -\gamma|x| + \delta\Big(\tfrac{\mu}{r}h^{(n)}(x) + \tfrac{\theta+(1-p)\lambda}{2r}v^{(n)}(x,10) + \tfrac{\theta+(1-p)\lambda}{2r}v^{(n)}(-x,01) + \tfrac{p\lambda}{2r}v^{(n)}(x+1,00)\\
&\quad + \tfrac{p\lambda}{2r}v^{(n)}(-x-1,00) + \tfrac{pc\lambda}{r}\Big)\\
v^{(n+1)}(x,01) &= -\gamma|x| + \delta\Big(\tfrac{\mu}{r}h^{(n)}(x) + \tfrac{\theta+p\lambda}{2r}v^{(n)}(x,01) + \tfrac{\theta+p\lambda}{2r}v^{(n)}(-x,10) + \tfrac{(1-p)\lambda}{2r}v^{(n)}(x-1,00)\\
&\quad + \tfrac{(1-p)\lambda}{2r}v^{(n)}(-x-1,00) + \tfrac{(1-p)c\lambda}{r}\Big)\\
v^{(n+1)}(x,11) &= -\gamma|x| + \delta\Big(\tfrac{\mu}{r}h^{(n)}(x) + \tfrac{\theta}{2r}(v^{(n)}(x,11) + v^{(n)}(-x,11)) + \tfrac{p\lambda}{2r}v^{(n)}(x+1,01)\\
&\quad + \tfrac{p\lambda}{2r}v^{(n)}(-x-1,10) + \tfrac{(1-p)\lambda}{2r}v^{(n)}(x-1,10) + \tfrac{(1-p)\lambda}{2r}v^{(n)}(-x+1,01) + \tfrac{c\lambda}{r}\Big)
\end{aligned}$$

Using this set of equations, $v^{(n+1)}(x,11) - v^{(n+1)}(x,01)$ is non-increasing as it is a positive sum of non-increasing functions:

$$\begin{aligned}
v^{(n+1)}(x,11)-v^{(n+1)}(x,01) &= \tfrac{\theta\delta}{2r}\left(v^{(n)}(x,11)-v^{(n)}(x,01)\right)+\tfrac{\theta\delta}{2r}\left(z^{(n)}(x,11)-z^{(n)}(x,01)\right)\\
&+\tfrac{p\lambda\delta}{2r}\left(v^{(n)}(x+1,01)-v^{(n)}(x,01)\right)+\tfrac{p\lambda\delta}{2r}\left(z^{(n)}(x+1,01)-z^{(n)}(x,01)\right)\\
&+\tfrac{(1-p)\lambda\delta}{2r}\left(v^{(n)}(x-1,10)-v^{(n)}(x-1,00)\right)+\tfrac{(1-p)\lambda\delta}{2r}\left(z^{(n)}(x-1,10)-z^{(n)}(x-1,00)\right)+\tfrac{pc\lambda\delta}{r},
\end{aligned}$$

where $z(x,l)\equiv v(-x,\mathsf{sym}(l))$. Similarly, $v^{(n+1)}(x,10)-v^{(n+1)}(x,00)$ is also non-increasing:

$$\begin{aligned}
v^{(n+1)}(x,10)-v^{(n+1)}(x,00) &= \tfrac{(\theta+(1-p)\lambda)\delta}{2r}\left(v^{(n)}(x,10)-v^{(n)}(x,00)\right)+\tfrac{(\theta+(1-p)\lambda)\delta}{2r}\left(z^{(n)}(x,10)-z^{(n)}(x,00)\right)\\
&+\tfrac{p\lambda\delta}{2r}\left(v^{(n)}(x+1,00)-v^{(n)}(x,00)\right)+\tfrac{p\lambda\delta}{2r}\left(z^{(n)}(x+1,00)-z^{(n)}(x,00)\right)+\tfrac{pc\lambda\delta}{r}.
\end{aligned}$$

Using the same reasoning, $v^{(n+1)}(x,11)-v^{(n+1)}(x,10)$ and $v^{(n+1)}(x,01)-v^{(n+1)}(x,00)$ are non-decreasing in x as they are positive sum of non-decreasing functions:

$$\begin{aligned}
v^{(n+1)}(x,11)-v^{(n+1)}(x,10) &= \tfrac{\theta\delta}{2r}\left(v^{(n)}(x,11)-v^{(n)}(x,10)\right)+\tfrac{\theta\delta}{2r}\left(z^{(n)}(x,11)-z^{(n)}(x,10)\right)\\
&+\tfrac{p\lambda\delta}{2r}\left(v^{(n)}(x+1,01)-v^{(n)}(x+1,00)\right)+\tfrac{p\lambda\delta}{2r}\left(z^{(n)}(x+1,01)-z^{(n)}(x+1,00)\right)\\
&+\tfrac{(1-p)\lambda\delta}{2r}\left(v^{(n)}(x-1,10)-v^{(n)}(x,00)\right)+\tfrac{(1-p)\lambda\delta}{2r}\left(z^{(n)}(x-1,10)-z^{(n)}(x,00)\right)+\tfrac{(1-p)c\lambda\delta}{r},\\
v^{(n+1)}(x,01)-v^{(n+1)}(x,00) &= \tfrac{(\theta+p\lambda)\delta}{2r}\left(v^{(n)}(x,01)-v^{(n)}(x,00)\right)+\tfrac{(\theta+p\lambda)\delta}{2r}\left(z^{(n)}(x,01)-z^{(n)}(x,00)\right)\\
&+\tfrac{(1-p)\lambda\delta}{2r}\left(v^{(n)}(x-1,00)-v^{(n)}(x,00)\right)+\tfrac{(1-p)\lambda\delta}{2r}\left(z^{(n)}(x-1,00)-z^{(n)}(x,00)\right)+\tfrac{(1-p)c\lambda\delta}{r}.
\end{aligned}$$

Finally, $v^{(n+1)}(x,l)$ is concave for all l if $h^{(n)}(x)$ is concave. Let L and U satisfy

$$L\triangleq\max\{x:v(x,11)<v(x,10)\}\ \text{ and } U\triangleq\min\{x:v(x,11)<v(x,01)\}.$$

We now show that for all inventory regions $h^{(n)}(x)$ is concave. For $x\leq L-1$,

$$h^{(n)}(x)-h^{(n)}(x+1)=v^{(n)}(x,10)-v^{(n)}(x+1,10),$$

which is nondecreasing as $v^{(n)}(x,10)$ is concave. For $x\geq U$,

$$h^{(n)}(x)-h^{(n)}(x+1)=v^{(n)}(x,01)-v^{(n)}(x+1,01),$$

which is nondecreasing as $v^{(n)}(x,01)$ is concave. If $x\geq L$ and $x+1<U$,

$$h^{(n)}(x)-h^{(n)}(x+1)=v^{(n)}(x,10)-v^{(n)}(x+1,11)=(v^{(n)}(x,10)-v^{(n)}(x,11))+(v^{(n)}(x,11)-v^{(n)}(x+1,11)),$$

which is again nondecreasing as it is the sum of nondecreasing functions. If $x \geq L$ and $x+1=U$,

$$h^{(n)}(x) - h^{(n)}(x+1) = v^{(n)}(x,10) - v^{(n)}(x+1,01)$$
$$= (v^{(n)}(x,10) - v^{(n)}(x+1,10)) + (v^{(n)}(x+1,10) - v^{(n)}(x+1,11)) + (v^{(n)}(x+1,11) - v^{(n)}(x+1,01)),$$

which is again nondecreasing as it is the sum of nondecreasing functions. If $x > L$ and $x+1 < U$,

$$h^{(n)}(x) - h^{(n)}(x+1) = v^{(n)}(x,11) - v^{(n)}(x+1,11),$$

which is nondecreasing as $v^{(n)}(x,11)$ is concave. Finally, if $x > L$ and $x+1=U$,

$$h^{(n)}(x) - h^{(n)}(x+1) = v^{(n)}(x,11) - v^{(n)}(x+1,01) = (v^{(n)}(x,11) - v^{(n)}(x+1,11)) + (v^{(n)}(x+1,11) - v^{(n)}(x+1,01)),$$

which is sum of nondecreasing functions. □

A.5. *Computation of the HFT's Threshold Quoting Policy*

We proved that the optimal market making policy involves thresholds. In this section, we exploit this solution structure and provide an efficient algorithm to solve for the threshold limits L^* and U^*, and the value functions v. Using the linear system of equations in Proposition 1 and the optimality conditions in Proposition 2, we obtain the following algorithm to compute the optimal limits.

Algorithm 1: Efficient algorithm to compute L^* and U^*.

```
Output: L* and U*
Initialize L = −1, K = c/(γ(1−δ)) and flag = 0;
while flag = 0 do
    U ⟵ 1;
    while U ≤ K do
        Solve for v(L,l), v(L+1,l), ..., v(U,l) using the linear system in Proposition 1;
        if v(L,10) > v(L,11), v(L+1,10) ≤ v(L,11), v(U,01) > v(U,11),
        v(U−1,01) ≤ v(U−1,11) then
            flag ⟵ 1, L* ⟵ L and U* ⟵ U ;
            break ;
        U ⟵ U + 1;
    L ⟵ L − 1;
```

B. Analysis of the Complete Model

B.1. Transition Probabilities

We now calculate the transition probabilities at each state of the HFT. First, note that the state transitions occur at a rate of $\mu+\lambda+\theta+\sigma$ where $\lambda \equiv 2(\lambda_P+\lambda_I+\lambda_A)$ and using uniformization we can have the same transition rate for all states and actions. Let $\mathbb{P}((x',s',l',e',j')|(x,s,l,e,j),(\ell^b,\ell^a))$ be the probability of reaching state (x',s',l',e',j') when the system is in state (x,s,l,e,j) and the trader takes the actions of ℓ^b and ℓ^a. First, we define our auxiliary variables.

Let $\mathsf{pr}(s)$ denote the unconditional probability of receiving signal $s=\left(S^{\text{dir}},S^{\text{type}}\right)$ right after the arrival of a market order or a signal.

$$
\begin{aligned}
\mathsf{pr}(s) &= \sum_{i\in\{P,I\}}\sum_{j\in\{B,S\}} \mathbb{P}(M^{\text{type}}=i,M^{\text{dir}}=j)\mathbb{P}(s|M^{\text{type}}=i,M^{\text{dir}}=j)\\
&= \sum_{i\in\{P,I\}}\sum_{j\in\{B,S\}} (0.5/\lambda)\left(p\mathbb{1}\{s_1=i\}+(1-p)\mathbb{1}\{s_1\neq i\}\right)\left(q\mathbb{1}\{s_2=j\}+(1-q)\mathbb{1}\{s_2\neq j\}\right)\\
&\qquad\qquad \left(\lambda^I\mathbb{1}\{i=I\}+\lambda^P\mathbb{1}\{i=P\}\right)\\
&= \begin{cases} 0.5p\lambda^I/\lambda+0.5(1-p)\lambda^P/\lambda & \text{if } s_1=I,\\ 0.5(1-p)\lambda^I/\lambda+0.5p\lambda^P/\lambda & \text{if } s_1=P. \end{cases}
\end{aligned}
$$

Let $\mathsf{m}_s^{s'}$ denote the conditional probability of receiving a market order with type s' (e.g., buy order submitted by an impatient LFT will be denoted by $s'=IB$) when the last signal appeared is s.

$$
\begin{aligned}
\mathsf{m}_s^{s'} &= \frac{\mathbb{P}(M^{\text{type}}=s'_1,M^{\text{dir}}=s'_2)\mathbb{P}(S^{\text{type}}=s_1,S^{\text{dir}}=s_2|M^{\text{type}}=s'_1,M^{\text{dir}}=s'_2)}{\sum_{i\in\{P,I\}}\sum_{j\in\{B,S\}}\mathbb{P}(M^{\text{type}}=i,M^{\text{dir}}=j)\mathbb{P}(S^{\text{type}}=s_1,S^{\text{dir}}=s_2|M^{\text{type}}=i,M^{\text{dir}}=j)}\\
&= \left(p\mathbb{1}\{s_1=s'_1\}+(1-p)\mathbb{1}\{s_1\neq s'_1\}\right)\left(q\mathbb{1}\{s_2=s'_2\}+(1-q)\mathbb{1}\{s_2\neq s'_2\}\right)\\
&\quad \left(\lambda^I\mathbb{1}\{s'_1=I\}+\lambda^P\mathbb{1}\{s'_1=P\}\right)/(\lambda\mathsf{pr}(s))
\end{aligned}
$$

Suppose that the current state of the HFT is (x,s,l,e,j). Let $r=(\lambda+\mu+\theta+\sigma)$. First, we provide the transition probabilities with respect to each action taken at decision epochs.

If the HFT does not quote in either side of the market, the inventory level cannot change. If a new decision event arrives before the arrival of an LFT order, a signal or jump, the state for tracking jumps, j, reverts to zero. Since the arrival of arbitrageurs will not change the state of the system, we also have the additional self-transition at the rate of $2\lambda_A/r$ to uniformize the model. Formally, we obtain

$$\mathbb{P}\left((x',s',l',e',j')|(x,s,b,1,j),(0,0)\right)=\begin{cases}\frac{(\lambda^{PI}+\theta)\mathsf{pr}(s')}{r} & \text{if } x=x', e'=0, l'=00\\ \frac{\sigma}{2r} & \text{if } x=x', e'=0, l'=00, j'=j+1\\ \frac{\sigma}{2r} & \text{if } x=x', e'=0, l'=00, j'=j-1\\ \frac{\mu}{r} & \text{if } x=x', e'=1, s=s', l'=00, j'=0\\ \frac{2\lambda_A}{r} & \text{if } x=x', e'=e, s=s', l'=l, j'=j\\ 0 & \text{otherwise,}\end{cases}$$

where $\lambda^{PI}\equiv 2(\lambda_P+\lambda_I)$.

We now analyze the trade scenarios for one-sided quoting. We are going to state the transition probabilities in which the HFT takes the action $(1,0)$. The remaining one-sided HFT actions are also very similar. When the HFT's action is $(1,0)$ he may increase his inventory by trading with the incoming market-sell order submitted by a patient or an impatient LFT, which occurs with probability $m_s^{IS}+m_s^{PS}$. We also need to account for the existence of stale quotes. Since the HFT is quoting at the bid side in this case, the stale quotes can only appear if $j<0$. Formally, we have the following transition probabilities:

$$\mathbb{P}\left((x',s',l',e',j')|(x,s,b,1,j),(1,0)\right)=\begin{cases}\frac{(\lambda^{\mathsf{sell}}+\theta)\mathsf{pr}(s')}{r} & \text{if } x=x', e'=0, l'=10, j\le 0\\ \frac{(\lambda_P+\lambda_I+\theta)\mathsf{pr}(s')}{r} & \text{if } x=x', e'=0, l'=10, j>0\\ \frac{\lambda^{\mathsf{buy}}\mathsf{pr}(s')}{r} & \text{if } x+1=x', e'=0, l'=00, j=0\\ \frac{\left(\lambda_A+\lambda^{\mathsf{buy}}\right)\mathsf{pr}(s')}{r} & \text{if } x+1=x', e'=0, l'=00, j<0\\ \frac{\sigma}{2r} & \text{if } x=x', e'=0, l'=10, j'=j+1\\ \frac{\sigma}{2r} & \text{if } x=x', e'=0, l'=10, j'=j-1\\ \frac{\mu}{r} & \text{if } x=x', e'=1, s=s', l'=10, j'=0\\ \frac{2\lambda_A}{r} & \text{if } x=x', e'=e, s=s', l'=l, j'=j\ge 0\\ \frac{\lambda_A}{r} & \text{if } x=x', e'=e, s=s', l'=l, j'=j<0\\ 0 & \text{otherwise,}\end{cases}$$

where $\lambda^{\mathsf{buy}}=\lambda^{PI}(m_s^{IS}+m_s^{PS})$ and $\lambda^{\mathsf{sell}}=\lambda^{PI}(m_s^{IB}+m_s^{PB})$ denotes the corresponding intensity for HFT's buy and sell trade in absence of any jumps, respectively.

We now analyze the trade scenarios for both-sided quoting. We are going to state the transition probabilities in which the HFT takes the action $(2,2)$. The remaining both-sided HFT actions are also very similar. When the HFT's action is $(2,2)$ he may increase his inventory by trading with the incoming market-sell order submitted by an impatient LFT, which occurs with probability m_s^{IS}. He may decrease his inventory by trading with the incoming market-buy order submitted by a patient LFT, which occurs with probability m_s^{IB}. We also need to account for the existence of stale quotes. Since the HFT is quoting at both sides of the market, the stale quotes will emerge when $j \neq 0$. Formally, we have the following transition probabilities:

$$\mathbb{P}\left((x',s',l',e',j')|(x,s,b,1,j),(2,2)\right) = \begin{cases} \frac{(\lambda_P+\theta)\mathsf{pr}(s')}{r} & \text{if } x=x', e'=0, l'=22, j=0 \\ \frac{(\lambda^{\mathsf{buy}}+\theta)\mathsf{pr}(s')}{r} & \text{if } x=x', e'=0, l'=22, j>0 \\ \frac{(\lambda^{\mathsf{sell}}+\theta)\mathsf{pr}(s')}{r} & \text{if } x=x', e'=0, l'=22, j<0 \\ \frac{(\lambda_P+\lambda_I)m_s^{IS}\mathsf{pr}(s')}{r} & \text{if } x+1=x', e'=0, l'=02, j=0 \\ \frac{\left(\lambda_A+\lambda^{\mathsf{buy}}\right)\mathsf{pr}(s')}{r} & \text{if } x+1=x', e'=0, l'=02, j<0 \\ \frac{(\lambda_P+\lambda_I)m_s^{IB}\mathsf{pr}(s')}{r} & \text{if } x-1=x', e'=0, l'=20, j=0 \\ \frac{\left(\lambda_A+\lambda^{\mathsf{sell}}\right)\mathsf{pr}(s')}{r} & \text{if } x-1=x', e'=0, l'=20, j<0 \\ \frac{\sigma}{2r} & \text{if } x=x', e'=0, l'=22, j'=j+1 \\ \frac{\sigma}{2r} & \text{if } x=x', e'=0, l'=22, j'=j-1 \\ \frac{\mu}{r} & \text{if } x=x', e'=1, s=s', l'=22, j'=0 \\ \frac{2\lambda_A}{r} & \text{if } x=x', e'=e, s=s', l'=l, j'=j=0 \\ \frac{\lambda_A}{r} & \text{if } x=x', e'=e, s=s', l'=l, j'=j\neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

If the system is observed at the arrival time of a market order or signal event, that is $e=0$, the HFT cannot revise his quotes. We can accommodate these states in our model using fake decisions that merely sets the action to the existing quotes. In this case,

$$\mathbb{P}\left((x',s',l',e',j')|(x,s,b,0,j),(\ell^b,\ell^a)\right) = \begin{cases} \mathbb{P}\left((x',s',l',e',j')|(x,s,b,1,j),(\ell^b,\ell^a)\right) & \text{if } \ell^b=b_1, \ell^a=b_2 \\ 0 & \text{otherwise.} \end{cases}$$

B.2. HFT's Reward Function

Let $\mathbb{R}\left((x',s',l',e',j')|(x,s,l,e,j),(\ell^b,\ell^a)\right)$ be the probability of reaching state (x',s',l',e',j') when the system is in state (x,s,l,e,j) and the trader takes the actions of ℓ^b and ℓ^a. We can define the total reward matrix using a similar analysis as in Section A.2. The main difference in the complete model is the possibility of earning $3C$ and losing the jump amount to the LFT in the presence of a stale quote.

Formally, let

$$\mathbb{R}\left((x',s',l',e',j')|(x,s,l,e,j),(\ell^b,\ell^a)\right) = \mathbb{1}\left(x+1=x'\right)\left(c(1-kj^-)\mathbb{1}\left(\ell^b=1\right)+c(3-kj^-)\mathbb{1}\left(\ell^b=2\right)\right)$$
$$\mathbb{1}\left(x-1=x'\right)\left(c(1-kj^+)\mathbb{1}\left(\ell^a=1\right)+c(3-kj^+)\mathbb{1}\left(\ell^a=2\right)\right)-\gamma|x|,$$

where $j^+ \equiv \max(j,0)$, $j^- \equiv \max(-j,0)$, c and γ, are defined as the "adjusted spread" and "adjusted inventory aversion" parameters for the discrete-time formulation of the complete model and given by

$$\lambda \equiv 2(\lambda_P+\lambda_I+\lambda_A), \quad r \equiv \lambda+\mu+\theta+\sigma, \quad \delta \equiv \frac{r}{r+D}, \quad c \equiv \delta C \;\text{ and }\; \gamma \equiv \frac{\Gamma}{r+D}. \tag{B.1}$$

Using this definition, the HFT maximizes

$$V(x,s,l,e,j) = \max_{\pi}\mathbb{E}^{\pi}\left[\sum_{k=0}^{\infty}\delta^k\mathbb{R}\left((x'_{t_{k+1}},s'_{t_{k+1}},b'_{t_{k+1}},e'_{t_{k+1}},j'_{t_{k+1}})|(x_{t_k},s_{t_k},b_{t_k},e_{t_k},j_{t_k}),(\ell^b_{t_k},\ell^a_{t_k})\right)\right], \tag{B.2}$$

starting from his initial state, (x,s,l,e,j), which is in the requisite MDP form.

B.3. HFT's Value Function

We have now transformed our continuous-time problem into an equivalent discrete-time MDP. We can now solve for $V(x,s,l,e,j)$ in (B.2) using the discrete-time HJB equations:

$$V(x,s,l,e,j) = \max_{\ell^b,\ \ell^a}\left\{\sum_{(x',s',l',e',j')}\mathbb{P}\left((x',s',l',e',j')|(x,s,l,e,j),(\ell^b,\ell^a)\right)\Big[\mathbb{R}\left((x',s',l',e',j')|(x,s,l,e,j),(\ell^b,\ell^a)\right)\right.$$
$$\left.+\,\delta V(x',s',l',e',j')\Big]\right\}. \tag{B.3}$$

Since the model is symmetric around the bid and ask side of the market, we can first eliminate the order direction signal from our state space. We have the following reduction for each $s_1 \in \{P, I\}$.

$$V(-x, s_1S, l, e, j) = \begin{cases} V(x, s_1B, l, e, -j) & \text{when } l \in \{00, 11, 22\}, \\ V(x, s_1B, 01, e, -j) & \text{when } l = 10, \\ V(x, s_1B, 10, e, -j) & \text{when } l = 01, \\ V(x, s_1B, 21, e, -j) & \text{when } l = 12, \\ V(x, s_1B, 12, e, -j) & \text{when } l = 21, \\ V(x, s_1B, 02, e, -j) & \text{when } l = 20, \\ V(x, s_1B, 20, e, -j) & \text{when } l = 02. \end{cases}$$

Using this result, we let $v(x, P, l, j) \equiv V(x, PS, l, 0, j)$, $v(x, I, l, j) \equiv V(x, IS, l, 0, j)$ and $h(x, P) \equiv V(x, PS, l, 1, 0)$, $h(x, I) \equiv V(x, IS, l, 1, 0)$ for ease in notation.

By substituting the expressions for $\mathbb{P}$ and $\mathbb{R}$, we can obtain the implicit equations for the value functions of each state. In the following proposition, we state the HJB equations for our value function. We provide value functions for the bid side of the market by excluding the symmetric states.

Proposition 4. *$v(x, P, 00, j)$ satisfies the following set of equations. Since there is no active quote, there is no risk of stale quotes.*

$$\begin{aligned} v(x, P, 00, j) = -\gamma|x| + \delta\Big\{ \frac{\mu}{r} h(x, P) + \frac{\lambda^{PI} + \theta}{r} \Big(\mathsf{pr}(PS) v(x, P, 10, j) + \mathsf{pr}(IS) v(x, I, 10, j) + \mathsf{pr}(PB) z(x, P, 10, j) \\ + \mathsf{pr}(IB) z(x, I, 10, j) \Big) + \frac{\sigma}{2r} \Big(v(x, P, 10, j-1) + v(x, P, 10, j+1) \Big) \\ + \frac{2\lambda_A}{r} v(x, P, 10, j) \Big\} \end{aligned}$$

$v(x, P, 10, j)$ satisfies the following set of equations. If $j < 0$, the HFT is subject to the risk of leaving stale quotes:

$$\begin{aligned} v(x, P, 10, j) = -\gamma|x| + \delta\Big\{ \frac{\mu}{r} h(x, P) + \frac{\lambda^{\mathsf{sell}} + \theta}{r} \Big(\mathsf{pr}(PS) v(x, P, 10, j) + \mathsf{pr}(IS) v(x, I, 10, j) + \mathsf{pr}(PB) z(x, P, 10, j) \\ + \mathsf{pr}(IB) z(x, I, 10, j) \Big) + \frac{\lambda^{\mathsf{buy}} + \lambda_A}{r} \Big(c + jJ + \mathsf{pr}(PS) v(x+1, P, 00, j) + \mathsf{pr}(IS) v(x+1, I, 00, j) \\ + \mathsf{pr}(PB) z(x+1, P, 00, j) + \mathsf{pr}(IB) z(x+1, I, 00, j) \Big) + \frac{\sigma}{2r} \Big(v(x, P, 10, j-1) + v(x, P, 10, j+1) \Big) \\ + \frac{\lambda_A}{r} v(x, P, 10, j) \Big\} \end{aligned}$$

For $j = 0$, the HFT does not suffer from stale quotes:

$$\begin{aligned}
v(x,P,10,0) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\text{sell}}+\theta}{r}\Big(&\text{pr}(PS)v(x,P,10,0) + \text{pr}(IS)v(x,I,10,0) + \text{pr}(PB)z(x,P,10,0) \\
&+ \text{pr}(IB)z(x,I,10,0)\Big) + \frac{\lambda^{\text{buy}}}{r}\Big(c + \text{pr}(PS)v(x+1,P,00,0) + \text{pr}(IS)v(x+1,I,00,0) \\
&+ \text{pr}(PB)z(x+1,P,00,0) + \text{pr}(IB)z(x+1,I,00,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,10,-1) + v(x,P,10,1)\Big) \\
&+ \frac{2\lambda_A}{r}v(x,P,10,0)\Big\}
\end{aligned}$$

If $j > 0$, the HFT's bid quote is not attractive to any LFT so there is no possibility of a trade:

$$\begin{aligned}
v(x,P,10,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{PI}+\theta}{r}\Big(&\text{pr}(PS)v(x,P,00,j) + \text{pr}(IS)v(x,I,00,j) + \text{pr}(PB)z(x,P,00,j) \\
&+ \text{pr}(IB)z(x,I,00,j)\Big) + \frac{\sigma}{2r}\Big(v(x,P,10,j+1) + v(x,P,00,j-1)\Big) \\
&+ \frac{2\lambda_A}{r}v(x,P,00,j)\Big\}
\end{aligned}$$

$v(x,P,20,j)$ *satisfies the following set of equations. If $j < 0$, the HFT is again subject to the risk of leaving stale quotes at the bid side:*

$$\begin{aligned}
v(x,P,20,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\text{sell}}+\theta}{r}\Big(&\text{pr}(PS)v(x,P,20,j) + \text{pr}(IS)v(x,I,20,j) + \text{pr}(PB)z(x,P,20,j) \\
&+ \text{pr}(IB)z(x,I,20,j)\Big) + \frac{\lambda^{\text{buy}}+\lambda_A}{r}\Big(c + jJ + \text{pr}(PS)v(x+1,P,00,j) + \text{pr}(IS)v(x+1,I,00,j) \\
&+ \text{pr}(PB)z(x+1,P,00,j) + \text{pr}(IB)z(x+1,I,00,j)\Big) + \frac{\sigma}{2r}\Big(v(x,P,20,j-1) + v(x,P,20,j+1)\Big) \\
&+ \frac{\lambda_A}{r}v(x,P,20,j)\Big\}
\end{aligned}$$

For $j = 0$, the HFT does not suffer from stale quotes:

$$\begin{aligned}
v(x,P,20,0) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{PI}(1-m_s^{IS})+\theta}{r}\Big(&\text{pr}(PS)v(x,P,20,0) + \text{pr}(IS)v(x,I,20,0) + \text{pr}(PB)z(x,P,20,0) \\
&+ \text{pr}(IB)z(x,I,20,0)\Big) + \frac{\lambda^{PI}m_s^{IS}}{r}\Big(c + \text{pr}(PS)v(x+1,P,00,0) + \text{pr}(IS)v(x+1,I,00,0) \\
&+ \text{pr}(PB)z(x+1,P,00,0) + \text{pr}(IB)z(x+1,I,00,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,20,-1) + v(x,P,20,1)\Big) \\
&+ \frac{2\lambda_A}{r}v(x,P,20,0)\Big\}
\end{aligned}$$

If $j > 0$, the HFT's bid quote is not attractive to any LFT so there is no possibility of a trade:

$$\begin{aligned}
v(x,P,20,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{PI}+\theta}{r}\Big(&\text{pr}(PS)v(x,P,20,j) + \text{pr}(IS)v(x,I,20,j) + \text{pr}(PB)z(x,P,20,j) \\
&+ \text{pr}(IB)z(x,I,20,j)\Big) + \frac{\sigma}{2r}\Big(v(x,P,20,j+1) + v(x,P,20,j-1)\Big) \\
&+ \frac{2\lambda_A}{r}v(x,P,20,j)\Big\}
\end{aligned}$$

$v(x,P,11,j)$ *satisfies the following set of equations. In this case, the HFT is subject to the risk of* *leaving stale quotes at the both sides of the market. If $j < 0$,*

$$v(x,P,11,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\mathsf{sell}}+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,11,j) + \mathsf{pr}(IS)v(x,I,11,j) + \mathsf{pr}(PB)z(x,P,11,j) + \mathsf{pr}(IB)z(x,I,11,j)\Big) + \frac{\lambda^{\mathsf{buy}}+\lambda_A}{r}\Big(c + jJ + \mathsf{pr}(PS)v(x+1,P,01,j) + \mathsf{pr}(IS)v(x+1,I,01,j) + \mathsf{pr}(PB)z(x+1,P,01,j) + \mathsf{pr}(IB)z(x+1,I,01,j)\Big) + \frac{\sigma}{2r}\Big(v(x,P,11,j-1) + v(x,P,11,j+1)\Big) + \frac{\lambda_A}{r}v(x,P,11,j)\Big\}$$

For $j = 0$, the HFT does not suffer from stale quotes:

$$v(x,P,11,0) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,11,0) + \mathsf{pr}(IS)v(x,I,11,0) + \mathsf{pr}(PB)z(x,P,11,0) + \mathsf{pr}(IB)z(x,I,11,0)\Big) + \frac{\lambda^{\mathsf{buy}}}{r}\Big(c + \mathsf{pr}(PS)v(x+1,P,01,0) + \mathsf{pr}(IS)v(x+1,I,01,0) + \mathsf{pr}(PB)z(x+1,P,01,0) + \mathsf{pr}(IB)z(x+1,I,01,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,11,-1) + v(x,P,11,1)\Big) + \frac{2\lambda_A}{r}v(x,P,11,0) + \frac{\lambda^{\mathsf{sell}}}{r}\Big(c + \mathsf{pr}(PS)v(x-1,P,10,0) + \mathsf{pr}(IS)v(x-1,I,10,0) + \mathsf{pr}(PB)z(x-1,P,10,0) + \mathsf{pr}(IB)z(x-1,I,10,0)\Big)\Big\}$$

If $j > 0$,

$$v(x,P,11,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\mathsf{buy}}+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,11,j) + \mathsf{pr}(IS)v(x,I,11,j) + \mathsf{pr}(PB)z(x,P,11,j) + \mathsf{pr}(IB)z(x,I,11,j)\Big) + \frac{\lambda^{\mathsf{sell}}+\lambda_A}{r}\Big(c - jJ + \mathsf{pr}(PS)v(x-1,P,10,0) + \mathsf{pr}(IS)v(x-1,I,10,0) + \mathsf{pr}(PB)z(x-1,P,10,0) + \mathsf{pr}(IB)z(x-1,I,10,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,11,j+1) + v(x,P,11,j-1)\Big) + \frac{\lambda_A}{r}v(x,P,11,j)\Big\}$$

$v(x,P,22,j)$ satisfies the following set of equations. In this case, the HFT is subject to the risk of leaving stale quotes at the both sides of the market. If $j < 0$,

$$v(x,P,22,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\mathsf{sell}}+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,22,j) + \mathsf{pr}(IS)v(x,I,22,j) + \mathsf{pr}(PB)z(x,P,22,j) + \mathsf{pr}(IB)z(x,I,22,j)\Big) + \frac{\lambda^{\mathsf{buy}}+\lambda_A}{r}\Big(c + jJ + \mathsf{pr}(PS)v(x+1,P,02,j) + \mathsf{pr}(IS)v(x+1,I,02,j) + \mathsf{pr}(PB)z(x+1,P,02,j) + \mathsf{pr}(IB)z(x+1,I,02,j)\Big) + \frac{\sigma}{2r}\Big(v(x,P,22,j-1) + v(x,P,22,j+1)\Big) + \frac{\lambda_A}{r}v(x,P,22,j)\Big\}$$

For $j = 0$, the HFT does not suffer from stale quotes:

$$v(x,P,22,0) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{PI}(1-m_s^{IB}-m_s^{IS})+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,22,0) + \mathsf{pr}(IS)v(x,I,22,0) + \mathsf{pr}(PB)z(x,P,22,0)$$
$$+ \mathsf{pr}(IB)z(x,I,22,0)\Big) + \frac{\lambda^{PI}m_s^{IS}}{r}\Big(c + \mathsf{pr}(PS)v(x+1,P,02,0) + \mathsf{pr}(IS)v(x+1,I,02,0)$$
$$+ \mathsf{pr}(PB)z(x+1,P,02,0) + \mathsf{pr}(IB)z(x+1,I,02,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,22,-1) + v(x,P,22,1)\Big)$$
$$+ \frac{2\lambda_A}{r}v(x,P,22,0) + \frac{\lambda^{PI}m_s^{IB}}{r}\Big(c + \mathsf{pr}(PS)v(x-1,P,20,0) + \mathsf{pr}(IS)v(x-1,I,20,0)$$
$$+ \mathsf{pr}(PB)z(x-1,P,20,0) + \mathsf{pr}(IB)z(x-1,I,20,0)\Big)\Big\}$$

If $j > 0$,

$$v(x,P,22,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\mathsf{buy}}+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,22,j) + \mathsf{pr}(IS)v(x,I,22,j) + \mathsf{pr}(PB)z(x,P,22,j)$$
$$+ \mathsf{pr}(IB)z(x,I,22,j)\Big) + \frac{\lambda^{\mathsf{sell}}+\lambda_A}{r}\Big(c - jJ\mathsf{pr}(PS)v(x-1,P,20,0) + \mathsf{pr}(IS)v(x-1,I,20,0)$$
$$+ \mathsf{pr}(PB)z(x-1,P,20,0) + \mathsf{pr}(IB)z(x-1,I,20,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,22,j+1) + v(x,P,22,j-1)\Big)$$
$$+ \frac{\lambda_A}{r}v(x,P,22,j)\Big\}$$

$v(x,P,21,j)$ satisfies the following set of equations. In this case, the HFT is subject to the risk of leaving stale quotes at the both sides of the market. If $j < 0$,

$$v(x,P,21,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\mathsf{sell}}+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,21,j) + \mathsf{pr}(IS)v(x,I,21,j) + \mathsf{pr}(PB)z(x,P,21,j)$$
$$+ \mathsf{pr}(IB)z(x,I,21,j)\Big) + \frac{\lambda^{\mathsf{buy}}+\lambda_A}{r}\Big(c + jJ + \mathsf{pr}(PS)v(x+1,P,01,j) + \mathsf{pr}(IS)v(x+1,I,01,j)$$
$$+ \mathsf{pr}(PB)z(x+1,P,01,j) + \mathsf{pr}(IB)z(x+1,I,01,j)\Big) + \frac{\sigma}{2r}\Big(v(x,P,21,j-1) + v(x,P,21,j+1)\Big)$$
$$+ \frac{\lambda_A}{r}v(x,P,21,j)\Big\}$$

For $j = 0$, the HFT does not suffer from stale quotes:

$$v(x,P,21,0) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{PI}(m_s^{PS})+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,21,0) + \mathsf{pr}(IS)v(x,I,21,0) + \mathsf{pr}(PB)z(x,P,21,0)$$
$$+ \mathsf{pr}(IB)z(x,I,21,0)\Big) + \frac{\lambda^{PI}m_s^{IS}}{r}\Big(c + \mathsf{pr}(PS)v(x+1,P,01,0) + \mathsf{pr}(IS)v(x+1,I,01,0)$$
$$+ \mathsf{pr}(PB)z(x+1,P,01,0) + \mathsf{pr}(IB)z(x+1,I,01,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,21,-1) + v(x,P,21,1)\Big)$$
$$+ \frac{2\lambda_A}{r}v(x,P,21,0) + \frac{\lambda^{\mathsf{sell}}}{r}\Big(c + \mathsf{pr}(PS)v(x-1,P,20,0) + \mathsf{pr}(IS)v(x-1,I,20,0)$$
$$+ \mathsf{pr}(PB)z(x-1,P,20,0) + \mathsf{pr}(IB)z(x-1,I,20,0)\Big)\Big\}$$

If $j > 0$,

$$v(x,P,21,j) = -\gamma|x| + \delta\Big\{\frac{\mu}{r}h(x,P) + \frac{\lambda^{\text{buy}}+\theta}{r}\Big(\mathsf{pr}(PS)v(x,P,21,j) + \mathsf{pr}(IS)v(x,I,21,j) + \mathsf{pr}(PB)z(x,P,21,j) + \mathsf{pr}(IB)z(x,I,21,j)\Big) + \frac{\lambda^{\text{sell}}+\lambda_A}{r}\Big(c - jJ\mathsf{pr}(PS)v(x-1,P,20,0) + \mathsf{pr}(IS)v(x-1,I,20,0) + \mathsf{pr}(PB)z(x-1,P,20,0) + \mathsf{pr}(IB)z(x-1,I,20,0)\Big) + \frac{\sigma}{2r}\Big(v(x,P,21,j+1) + v(x,P,21,j-1)\Big) + \frac{\lambda_A}{r}v(x,P,21,j)\Big\}$$

Finally, as in the simplified model, HFT's value function at the decision epoch is given by

$$h(x,P) = \max\Big\{v(x,P,00,0), v(x,P,10,0), v(x,P,01,0), v(x,P,11,0), v(x,P,20,0), v(x,P,02,0), v(x,P,22,0), v(x,P,12,0), v(x,P,21,0)\Big\}$$

Proposition 4 shows that the HFT aims to choose the optimal action when the state is in $e = 1$ by maximizing over all possible quoting actions and the corresponding value is stored in $h(x,s)$. On the other hand, $v(x,s,l,j)$ computes the expected one-step reward resulting from possible transitions determined by the active quotes in state l.

B.4. Optimal Market Making Solution

Proof of Theorem 2. The proof is very similar to that of Theorem 1. Using the value iteration algorithm, we first establish by induction that value functions are concave in x. We also need to show that as inventory gets larger (smaller), less quoting at the bid (ask) side will be more and more attractive. Formally, we need to show that $v(x,s,0a,j) - v(x,s,2a,j)$, $v(x,s,2a,j) - v(x,s,1a,j)$ will be nondecreasing in x for any fixed policy $a \in \{0,1,2\}$ at the ask side and similarly, $v(x,s,b2,j) - v(x,s,b0,j)$, $v(x,s,b1,j) - v(x,s,b2,j)$ will be nondecreasing in x for any fixed policy $b \in \{0,1,2\}$ at the bid side. We will refer to this condition as the "nondecreasing property."

Let $v^{(0)}(x,s,l,j) = 0$ for all (x,s,l,j). Then, in the base case, all value functions will include $\gamma|x|$ and an appropriate constant term. Therefore, the base case will satisfy the concavity and nondecreasing property. Assume that $v^{(n)}(x,s,l,j)$ satisfies the induction hypothesis. Then, we first illustrate that the nondecreasing property holds using $v^{(n+1)}(x,P,20,j) - v^{(n+1)}(x,P,10,j)$. If $j \neq 0$,

$$v^{(n+1)}(x,P,20,j) - v^{(n+1)}(x,P,10,j) = \delta\Big\{\frac{\lambda^{\text{sell}}+\theta}{r}\Big(\mathsf{pr}(PS)\Big(v^{(n)}(x,P,20,j) - v^{(n)}(x,P,10,j)\Big) + \mathsf{pr}(IS)\Big(v^{(n)}(x,I,20,j) - v^{(n)}(x,I,10,j)\Big) + \mathsf{pr}(PB)\Big(z^{(n)}(x,I,20,j) - z^{(n)}(x,I,10,j)\Big) + \mathsf{pr}(IB)\Big(z^{(n)}(x,I,20,j) - z^{(n)}(x,I,10,j)\Big)\Big) + \frac{\sigma}{2r}\Big(v^{(n)}(x,P,20,j-1) - v^{(n)}(x,P,10,j-1)\Big) + \frac{\sigma}{2r}\Big(v^{(n)}(x,P,20,j+1) - v^{(n)}(x,P,10,j+1)\Big) + \frac{\lambda_A}{r}\Big(v^{(n)}(x,P,20,j) + v^{(n)}(x,P,10,j)\Big)\Big\}$$

is also nondecreasing in x as each term satisfies the non-decreasing property via the induction hypothesis. If $j = 0$, we have the following additional terms:

$$\frac{\lambda^{PI} m_s^{PS}}{r} \Big(\mathsf{pr}(PS) \left(v^{(n)}(x, P, 20, j) - v^{(n)}(x+1, P, 00, j) \right) + \mathsf{pr}(IS) \left(v^{(n)}(x, I, 20, j) - v^{(n)}(x+1, I, 00, j) \right) \\ + \mathsf{pr}(PB) \left(z^{(n)}(x, I, 20, j) - z^{(n)}(x+1, I, 00, j) \right) + \mathsf{pr}(IB) \left(z^{(n)}(x, I, 20, j) - z^{(n)}(x+1, I, 00, j) \right) \Big)$$

which equals to

$$\frac{\lambda^{PI} m_s^{PS}}{r} \Big(\mathsf{pr}(PS) \left(v^{(n)}(x, P, 20, j) - v^{(n)}(x, P, 00, j) \right) + \mathsf{pr}(IS) \left(v^{(n)}(x, I, 20, j) - v^{(n)}(x, I, 00, j) \right) \\ + \mathsf{pr}(PB) \left(z^{(n)}(x, I, 20, j) - z^{(n)}(x, I, 00, j) \right) + \mathsf{pr}(IB) \left(z^{(n)}(x, I, 20, j) - z^{(n)}(x, I, 00, j) \right) \\ + \mathsf{pr}(PS) \left(v^{(n)}(x, P, 00, j) - v^{(n)}(x+1, P, 00, j) \right) + \mathsf{pr}(IS) \left(v^{(n)}(x, I, 00, j) - v^{(n)}(x+1, I, 00, j) \right) \\ + \mathsf{pr}(PB) \left(z^{(n)}(x, I, 00, j) - z^{(n)}(x+1, I, 00, j) \right) + \mathsf{pr}(IB) \left(z^{(n)}(x, I, 00, j) - z^{(n)}(x+1, I, 00, j) \right) \Big)$$

which is also nondecreasing in x as the first four terms satisfies the non-decreasing property via the induction hypothesis, and the last four terms satisfy the concavity in x via the induction hypothesis.

□

B.5. Comparative statics

Proof of Proposition 3. Suppose that we increase μ by $\Delta\mu$. Assume that π^* is the optimal quoting policy when the arrival rate of decision epochs is μ. Using the similar thinning argument, if the HFT does not quote at the new decision times due to $\Delta\mu$, he is going to get the same objective value as in the original model. However, he can quote to buy when his inventory is negative and quote to sell when his inventory is positive and thus he can achieve a higher value.

The sensitivity with respect to λ_I and λ_P is very similar. We prove for λ_I. Suppose that we increase λ_I by $\Delta\lambda_I$. Assume that π^* is the optimal quoting policy when the arrival rate of impatient orders is λ_I. Consider any sample path under this model. For any sample path, we can construct a model with an arrival rate of impatient orders of $\lambda_I + \Delta\lambda_I$ in which via thinning Poisson processes, we can obtain the original sequence of orders plus the new arrivals due to $\Delta\lambda_I$. Now at these new arrivals, there is a positive probability that the HFT will reduce his inventory and achieve higher value.

Suppose that we decrease θ by $\Delta\theta$. Assume that π^* is the optimal quoting policy when the arrival rate of signal is θ. In the new model π^* is still feasible, and the probability of receiving an order as predicted by the signal is higher and thus the HFT's value increases as long as the signals are informative. The relationship with respect to p and q are similar as by construction increasing the accuracy of the signals will increase the probability of receiving the order that the signal predicts at the decision epoch. Thus, the value of the HFT will be nondecreasing in p and q.

The relationship with respect to C and Γ can be obtained using induction on the steps of the value iteration algorithm. We use the simplified model for ease in notation. Let $v^{(0)}(x,l)=0$ for all x and l. Then, in the base case, we obtain

$$\begin{aligned}
v^{(1)}(x,00) &= -\gamma|x|,\\
v^{(1)}(x,10) &= -\gamma|x|+\frac{\delta p\lambda c}{r},\\
v^{(1)}(x,01) &= -\gamma|x|+\frac{\delta(1-p)\lambda c}{r},\\
v^{(1)}(x,11) &= -\gamma|x|+\frac{\delta\lambda c}{r},
\end{aligned}$$

which shows that $v^{(1)}(x,l)$ is increasing in C and Γ as $c\equiv C\delta$ and $\gamma\equiv\Gamma/(r+D)$. Assume that $v^{(n)}(x,l)$ satisfies the induction hypothesis. Then, $v^{(n+1)}(x,l)$ satisfies the following set of equations:

$$\begin{aligned}
v^{(n+1)}(x,00) &= -\gamma|x|+\delta\left(\tfrac{\mu}{r}h^{(n)}(x)+\tfrac{\theta+\lambda}{2r}v^{(n)}(x,00)+\tfrac{\theta+\lambda}{2r}v^{(n)}(-x,00)\right)\\
v^{(n+1)}(x,10) &= -\gamma|x|+\delta\Big(\tfrac{\mu}{r}h^{(n)}(x)+\tfrac{\theta+(1-p)\lambda}{2r}v^{(n)}(x,10)+\tfrac{\theta+(1-p)\lambda}{2r}v^{(n)}(-x,01)+\tfrac{p\lambda}{2r}v^{(n)}(x+1,00)\\
&\quad+\tfrac{p\lambda}{2r}v^{(n)}(-x-1,00)+\tfrac{pc\lambda}{r}\Big)\\
v^{(n+1)}(x,01) &= -\gamma|x|+\delta\Big(\tfrac{\mu}{r}h^{(n)}(x)+\tfrac{\theta+p\lambda}{2r}v^{(n)}(x,01)+\tfrac{\theta+p\lambda}{2r}v^{(n)}(-x,10)+\tfrac{(1-p)\lambda}{2r}v^{(n)}(x-1,00)\\
&\quad+\tfrac{(1-p)\lambda}{2r}v^{(n)}(-x-1,00)+\tfrac{(1-p)c\lambda}{r}\Big)\\
v^{(n+1)}(x,11) &= -\gamma|x|+\delta\Big(\tfrac{\mu}{r}h^{(n)}(x)+\tfrac{\theta}{2r}(v^{(n)}(x,11)+v^{(n)}(-x,11))+\tfrac{p\lambda}{2r}v^{(n)}(x+1,01)\\
&\quad+\tfrac{p\lambda}{2r}v^{(n)}(-x-1,10)+\tfrac{(1-p)\lambda}{2r}v^{(n)}(x-1,10)+\tfrac{(1-p)\lambda}{2r}v^{(n)}(-x+1,01)+\tfrac{c\lambda}{r}\Big)
\end{aligned}$$

We observe that each v is increasing in C and decreasing in Γ as the positive sum of these functions preserve the property.

The objective value is decreasing in λ_A and J as under any fixed policy these two parameters control the expected loss of the HFT in the presence of stale quotes. The value is also decreasing in σ as the probability of HFT's stale quote increases with σ. Finally, the value is also decreasing in D as the expected reward from each period is discounted more under any fixed policy. □

C. Analysis of the Duopoly Model

C.1. Transition Probabilities

We will calculate the transition probabilities at each state of the HFT in the presence of a competing MFT. First, note that the state transitions occur at a rate of $\mu+\lambda+\theta+\sigma+\beta$ where $\lambda\equiv\lambda_P+\lambda_I+\lambda_A$ and using uniformization we can have the same transition rate for all states and actions. Let

$\mathbb{P}((x', s', b', a', e', j')|(x, s, a, b, e, j), (\ell^b, \ell^a))$ be the probability of reaching state (x', s', b', a', e', j') when the system is in state (x, s, a, b, e, j) and the trader takes the actions of ℓ^b and ℓ^a. Here b (a) is the current state of the bid (ask) market.

Suppose that the current state of the HFT is (x, s, a, b, e, j). Let $r = (\lambda + \mu + \theta + \sigma + \beta)$. In order to illustrate our methodology, we will provide a few examples of the transition probabilities.

If the only quote at the bid side is from the MFT, the MFT will have priority for execution at the bid side. For example, suppose that the state is $(x, s, m0, 00, 1, j)$ and the HFT decides to take action $(1, 0)$. Since the MFT order has a priority, in the next epoch, this order can be executed and the HFT's order will now have the highest priority. Thus, there will be no possibility of an inventory increase compared to the monopoly model. Since the ask side is empty, the MFT can submit a new order at the ask side as well. In this case the stale quote risk still applies and will be applicable if $j < 0$. Formally, we have the following transition probabilities for $\mathbb{P}\left((x, s', b', a', e', j')|(x, s, m0, 00, 1, j), (1, 0)\right)$:

$$
\begin{cases}
\frac{(\lambda^{\mathsf{sell}}+\theta)\mathsf{pr}(s')}{r} & \text{if } e' = 0, b' = m1, a' = 00, j \leq 0 \\
\frac{(\lambda^{\mathsf{sell}}+\lambda^{\mathsf{buy}}+\theta)\mathsf{pr}(s')}{r} & \text{if } e' = 0, b' = m1, a' = 00, j > 0 \\
\frac{\lambda^{\mathsf{buy}}\mathsf{pr}(s')}{r} & \text{if } e' = 0, b' = 10, a' = 00, j = 0 \\
\frac{\left(\frac{\lambda_A}{2}+\lambda^{\mathsf{buy}}\right)\mathsf{pr}(s')}{r} & \text{if } e' = 0, b' = 10, a' = 00, j < 0 \\
\frac{\sigma}{2r} & \text{if } e' = 0, s = s', b' = m1, a' = 00, j' = j + 1 \\
\frac{\sigma}{2r} & \text{if } e' = 0, s = s', b' = m1, a' = 00, j' = j - 1 \\
\frac{\mu}{r} & \text{if } e' = 1, s = s', b' = m1, a' = 00, j' = 0 \\
\frac{2\lambda_A+\beta}{r} & \text{if } e' = e, s = s', b' = b, a' = a, j' = j \geq 0 \\
\frac{\lambda_A+\beta}{r} & \text{if } e' = e, s = s', b' = b, a' = a, j' = j < 0 \\
0 & \text{otherwise,}
\end{cases}
$$

If the system is observed at the arrival time of a market order or signal event, that is $e = 0$, the HFT cannot revise his quotes. We can accommodate these states in our model using fake decisions that merely sets the action to the existing quotes. In this case,

$$
\begin{aligned}
&\mathbb{P}\left((x', s', b', a', e', j')|(x, s, b, a, 0, j), (\ell^b, \ell^a)\right) \\
&= \begin{cases}
\mathbb{P}\left((x', s', b', a', e', j')|(x, s, b, a, 1, j), (\ell^b, \ell^a)\right) & \text{if } \ell^b = b, \ell^a = a \\
0 & \text{otherwise.}
\end{cases}
\end{aligned}
$$

C.2. HFT's Reward Function

We can define the total reward matrix using a similar analysis as in Section B.2. The main difference in the duopoly model is the priority of the HFT over the MFT. Let the priority states of the HFT given by $\mathcal{P} = \{10, 20, 1m\}$. Formally, let

$$
\begin{aligned}
&\mathbb{R}\left((x', s', b', a', e', j')|(x, s, b, a, e, j), (\ell^b, \ell^a)\right) \\
&= \mathbb{1}\left(x + 1 = x'\right)\left(c(1 - kj^-)\mathbb{1}\left(\ell^b = 1, b \in \mathcal{P}\right) + c(3 - kj^-)\mathbb{1}\left(\ell^b = 2, b \in \mathcal{P}\right)\right) \\
&+ \mathbb{1}\left(x - 1 = x'\right)\left(c(1 - kj^+)\mathbb{1}\left(\ell^a = 1, b \in \mathcal{P}\right) + c(3 - kj^+)\mathbb{1}\left(\ell^a = 2, b \in \mathcal{P}\right)\right) \\
&- \gamma|x|
\end{aligned}
$$

where c and γ, are defined as the "adjusted spread" and "adjusted inventory aversion" parameters for the resulting discrete-time formulation of the duopoly model and given by

$$
r \equiv \lambda + \mu + \theta + \sigma + 2\beta, \quad \delta \equiv \frac{r}{r + D}, \quad c \equiv \delta C \text{ and } \gamma \equiv \frac{\Gamma}{r + D}. \tag{C.1}
$$

By substituting the expressions for $\mathbb{P}$ and $\mathbb{R}$, we can obtain the implicit equations for the value functions of each state. The proof of Theorem 3 follows steps identical to the proof of Theorem 2, and is thus omitted.