Discrete or continuous trading?

HFT competition and liquidity on batch auction markets

Marlene D. Haas and Marius A. Zoican[*]

February 26, 2016

Abstract

A batch auction market does not necessarily improve liquidity relative to continuous-time trading. HFTs submit quotes that become stale if the market clears before they process new information. Such stale quotes are adversely selected by informed HFTs. In equilibrium, HFTs supply excess liquidity in the auction market. Consequently, arbitrage profits are larger and the spread increases to compensate. On the other hand, price competition between arbitrageurs reduces adverse selection costs: the spread decreases. Except for particularly high or low auction frequencies, the batch auction market can hurt liquidity. The HFT "arms' race" stimulates price competition between arbitrageurs, generating a lower spread.

Discrete or continuous trading?

HFT competition and liquidity on batch auction markets

Abstract

A batch auction market does not necessarily improve liquidity relative to continuous-time trading. HFTs submit quotes that become stale if the market clears before they process new information. Such stale quotes are adversely selected by informed HFTs. In equilibrium, HFTs supply excess liquidity in the auction market. Consequently, arbitrage profits are larger and the spread increases to compensate. On the other hand, price competition between arbitrageurs reduces adverse selection costs: the spread decreases. Except for particularly high or low auction frequencies, the batch auction market can hurt liquidity. The HFT "arms' race" stimulates price competition between arbitrageurs, generating a lower spread.

1 Introduction

Is continuous-time trading inherently flawed? Modern exchanges are largely organized as continuous-time limit order books: One can buy and sell assets at any given time. Traders who react first to profitable opportunities have a comparative advantage. Consequently, continuous-time trading generates incentives for each trader to become *marginally* faster than her competitors. As a result, an "arms' race" emerged between high-frequency traders: The round-trip trading times between New York and Chicago dropped from 16 milliseconds in 2010 to 8.02 milliseconds in July 2015. A London-based trader can buy stocks in Frankfurt within just 2.21 milliseconds. As a benchmark, light needs 2.12 milliseconds to travel the same distance.[1]

Such arms' race is not necessarily benign: Ever higher trading speeds come at a non-trivial social cost. In 2010, Spread Networks spent USD 300 million building a straight-line fiber optic cable between New York and Chicago for a 3 millisecond latency gain (Laughlin, Aguirre, and Grundfest, 2014). Moreover, faster trading does not necessarily improve market quality. Budish, Cramton, and Shim (2015) argue that, while socially costly, the HFT "arms race" has no impact on spreads. Ye, Yao, and Gai (2013) document supporting evidence: a drop in exchange latency on NASDAQ from the microsecond to the nanosecond level did not have an impact on the bid-ask spread.

The proposed alternative to the prevailing continuos-time market design is *discrete-time trading*, i.e., a frequent batch auction market. While traders can submit orders at any time, the batch auction market clears at discrete intervals (e.g., one second) through an uniform auction. Budish, Cramton, and Shim (2015) and McPartland (2015) argue that a batch auction market eliminates the "inherent flaw" of the continuous-time limit order book: the discrete advantage from being marginally (e.g., one nanosecond) faster than competitors. Consequently, the scope for an HFT "arms' race" is limited.

There is an active ongoing debate over the relative merits of continuous- and discrete-time trading mechanisms. The Securities and Exchanges Commission (SEC) chair indicated in June 2014 her interest in batch auction markets as a "*more flexible, competitive*" exchange design. In October 2015, Chicago Stock Exchange received approval from the SEC to launch a batch-auction platform, CHX SNAP. In March 2016,

---

[1]Sources for this paragraph: McKay Brothers Microwave Latencies Table and Bloomberg: "Wall Street Grabs NATO Towers in Traders Speed-of-Light Quest."

London Stock Exchange is expected to launch a midday auction for the most liquid securities.[2]

This paper builds a model of batch auction markets to analyze the costs and benefits of a transition from continuous- to discrete-time trading, in the presence of high-frequency traders. In the model, high-frequency traders choose between two canonical strategies (see, e.g., Hagstromer and Norden, 2013; SEC, 2010): to provide liquidity, to speculate on short-lived arbitrage opportunities, or both. HFTs' reaction to new information about an asset crucially depends on the market structure. On a limit-order market, the first HFT to react to an arbitrage opportunity earns maximum rents: HFTs compete primarily in speed rather than prices. On a batch auction market, however, both price and speed competition are important.

To fix intuition, consider the following setup. Alice and her brother Bob would like to buy their mother a gold necklace for Mother's Day. They each have 200 dollars to spend. Alice finds the necklace auctioned off on eBay for 100 dollars; the auction, however, closes in 5 minutes. If Alice sees her brother sleeping in the next room, she can bid 100 dollars, the seller's reservation value, and buy the necklace – that is, Alice acts as a monopolist. On the other hand, if Alice sees Bob focused at his computer, she will bid 200 dollars, her own reservation price – that is, Alice and Bob engage in Bertrand competition.

What is auction outcome if the door between the siblings' room is closed? Alice does not know with certainty whether Bob is also bidding. In equilibrium, she chooses a random price between one and two hundred dollars, a function on the probability Bob is trying to buy the necklace. This probability crucially depends on the time left until the auction and Bob's diligence in monitoring eBay: *How likely is it that Bob checks the eBay offers in the next five minutes?* Further, if Alice had more brothers, the number of competing siblings would also influences her bid.

If, rather than the five minutes lag, the item is immediately auctioned off to the first person offering 100 dollars or more, then both Alice and Bob will bid exactly 100 dollars; Therefore, the sibling with the fastest computer wins the auction.

The two siblings are a metaphor for high-frequency traders chasing a profit opportunity (the necklace). The "immediate sale" scenario corresponds to the limit order market: HFTs aggressively trade against stale

---

quotes and the winner earns the maximum rent. On a batch auction market, the first HFT to react to a new profit opportunity does not necessarily capture it, as other HFTs may also react before the market clears.

On batch auction markets, HFTs engage in price competition over arbitrage opportunities and consequently earn lower rents. Such price competition is stronger as the expected number of informed HFTs is larger. In turn, the expected number of "active" HFTs depends on the *batch auction interval* (the five minutes before auction time), *the total number of HFTs* trading in that market (number of siblings), and on the *frequency* with which HFTs monitor the market (the siblings' diligence) – a measure of trading speed. These three factors determine the expected size of the arbitrage profits for high-frequency traders.

In the model, HFTs choose whether to supply quotes to liquidity traders or not (as in, e.g., Menkveld and Zoican, 2015). If they do, HFTs can earn the bid-ask spread. At the same time, they are exposed to adverse selection risk if they do not process new information before the auction takes place. The alternative for HFTs is to become "pure" arbitrageurs who only react on news and try to snipe stale quotes. In equilibrium, HFTs are indifferent between providing liquidity or not. Consequently, a positive equilibrium spread arises due to adverse selection costs.

The paper's main result is that the transition to a batch auction market can reduce liquidity. We identify two channels that distinguish limit order from frequent auction markets. On the one hand, a batch auction market generates *excessive liquidity supply*, which reduces each HFT's expected profit from providing liquidity. As a consequence, the equilibrium spread increases. On the other hand, a batch auction market promotes *price competition* on arbitrage trades, reducing arbitrage rents. Consequently, the adverse selection costs for liquidity providers are lower and the equilibrium spread decreases. Depending on which effect dominates, the bid-ask spread on batch auction markets can be either higher, equal, or lower than on an equivalent limit order book.

Excessive liquidity supply emerges as batch auction markets do not feature time priority. On a limit order book, if expected liquidity demand exceeds the outstanding market depth, a marginal HFT does not want to supply additional liquidity. Being at the back of the queue, such HFT does not trade with liquidity traders in expectation. However, he is still exposed to adverse selection risk.[3] In contrast, on batch auction

---

[3] The result carries through both if HFTs submit limit orders sequentially and can observe the state of the book and if HFTs submit limit orders simultaneously, but can condition on the market depth – such as the "Top of Book" order on Xetra.

markets, all HFTs with orders at the best price share the liquidity demand, and earn the same expected profit. In equilibrium, *all* HFTs choose to provide liquidity on the batch auction market and snipe competitors' quotes if they observe an arbitrage opportunity. Therefore, the liquidity supply exceeds the expected liquidity demand. As a result, traded quantities are rationed, reducing the expected profit from market-making. The equilibrium spread increases to compensate for the rationing effect. The economic channel is related to Yueshen (2014), who also finds that liquidity overshoots on a limit order market if the HFTs' position in the queue is uncertain.

Further, the batch auction market stimulates price competition. As in the siblings metaphor, HFT arbitrageurs compete in prices on the batch auction market. Consequently, their expected rents are lower than on the limit order market.

The price competition effect depends non-linearly on both the number of HFTs and the frequency of batch auctions. If there are many HFTs on the market, competition is stronger and per-unit arbitrageur rents decrease. At the same time, however, the expected number of stale quotes increases in the number of HFTs. It follows that arbitrageur rents are lowest for markets with either very few or very many HFTs. A similar channel emerges with respect to batch auction frequency. If batch auctions are very frequent, it is unlikely an HFT learns new information before the market clears. Consequently, arbitrage profits are low. At the other end of the spectrum, if the market clears at long intervals, it is very likely that two or more HFTs observe the news. Therefore, if the batch auction frequency decreases below a certain threshold, HFT competition on arbitrage opportunities converges to the Bertrand case, depressing arbitrage profits and the opportunity cost for providing liquidity.

The paper offers several novel policy-relevant insights. First, the "arms' race" on a batch auction market intensifies competition between arbitrageurs and decreases the spread. One way to promote HFT competition is to set a low batch auction frequency, i.e., to allow HFTs more time to process news. This is costly in terms of waiting times for liquidity traders. If HFTs process information faster (i.e., a more intense "arms' race"), such constraint on the batch auction frequency can be relaxed at no cost for liquidity. In the context of batch auction market, speed competition stimulates price competition. Discrete-time trading can thus align private and social incentives.

4

Second, whether a batch auction market improves liquidity only depends on the number of HFT, their speed, and the batch auction frequency. In particular, stock-specific measures such as liquidity demand and volatility do not influence the ranking between batch auction and limit order market liquidity. A natural consequence is that the batch auction can be implemented uniformly across markets with a given HFT activity, without the need to adjust auction frequency for each stock.

Our paper contributes to a growing literature on HFT and market design. A closely related paper is Budish, Cramton, and Shim (2015), who study a model of batch auction markets. The authors assume HFTs react to new information with no delay. As a consequence, arbitrageurs always compete à la Bertrand. In terms of our metaphor, Alice and Bob are constantly monitoring eBay offers. Therefore, both the arbitrageur expected profit and the equilibrium spread are zero. We introduce adverse selection risk, as HFTs may not be able to observe or react to news before the auction takes place. The richer model we propose features a positive bid-ask spread and unveils new economic channels: excess liquidity supply on batch auction markets, and imperfect price competition between arbitrageurs. We can therefore establish necessary and sufficient conditions for a batch auction market to improve liquidity relative to the current setup. Our model nests the Budish, Cramton, and Shim (2015) setup if HFTs monitor news with infinite intensity, i.e., the limit of the "arms' race."

Fricke and Gerig (2015) calibrate a batch auction trading model with risk-averse traders to U.S. data and find the optimal batch length is between 0.2 and 0.9 seconds. However, the authors focus on liquidity risk rather than an adverse selection channel.

In a policy paper, Farmer and Skouras (2012) estimate the worldwide benefits of the transition from limit order to frequent auction markets to be around USD 500 billion per year. Their paper is similar to ours as it models the batch auction time as a Poisson process. However, the authors do not consider the effects of the batch auction on arbitrageur competition, nor the endogenous order choice for HFTs. Wah and Wellman (2013) and Wah, Hurd, and Wellman (2015) develop agent-based models to showcase the benefits of batch auction markets. Batch auctions improve welfare as they better aggregate supply and demand. If trader can choose between a limit order and a batch auction market, HFTs will always follow the choice of slow traders to increase arbitrage profits.

Madhavan (1992) and Economides and Schwartz (1995) argue that batch auction markets improve price efficiency as they aggregate disparate information from traders for a longer interval of time. In a model where "fast" traders act as intermediaries rather than arbitrageurs, Du and Zhu (2015) find that the socially optimal trading frequency corresponds to the information arrival frequency.

The paper is also related to the literature on auctions in financial markets. Janssen and Rasmusen (2002) and Jovanovic and Menkveld (2015) study auction mechanisms where the number of competing bids (in our case, the number of informed HFTs) is not common knowledge. The bidding equilibrium is always symmetric and in mixed strategies. Our paper proposes asymmetric information as a rationale for the uncertain number of auction participants. Kremer and Nyborg (2004) study various allocation rules in uniform price auctions and find that a discrete tick size or uniform rationing at infra-marginal prices eliminates arbitrarily large underpricing.

Finally, the paper relates to a growing literature on high-frequency trading. Several theoretical papers argue that the benefits from a speed "arms' race" between HFTs are limited. Biais, Foucault, and Moinas (2015) find socially excessive investment in fast trading technology. According to Menkveld (2014), the HFT "arms' race" can hurt market liquidity. In the same spirit, Menkveld and Zoican (2015) argue that ever faster exchanges promote a higher frequency of inter-HFT trades, increasing the adverse selection cost and consequently the spread. Empirical evidence suggests high-frequency traders use strategies to snipe stale quotes. Hendershott and Moulton (2011), Baron, Brogaard, and Kirilenko (2012), and Brogaard, Hendershott, and Riordan (2014) find HFT market orders have a larger price impact.

The rest of the paper is structured as follows. Section 2 describes the model of the batch auction market. Section 3 compares the equilibrium spread on the batch auction and limit order markets and establishes necessary and sufficient conditions for a batch auction market to improve liquidity. Section 4 discusses the role of the HFT arms' race in batch auction market. Section 5 extends the baseline model by studying the case of impatient liquidity traders. Section 6 concludes.

2 A model of the batch auction market

2.1 Primitives

Trading environment. A single risky asset is traded on a batch auction market as in Budish, Cramton, and Shim (2015). There is no time priority as in a limit order market. Orders are processed in batches at discrete time intervals, using a uniform price auction. The market clearing time is random and follows a Poisson process with intensity $\tau > 0$. Consequently, the expected *batch interval length* is $1/\tau$.

All traders can post orders at any time between $t = 0$ and the moment of the batch auction. At $t = 0$, there are no outstanding un-matched orders. Traders can submit both limit and market orders: Limit orders specify an offer to buy or sell a certain quantity at a given price, whereas market orders specify only the quantity to trade. At the end of the batch interval, the buy and sell orders are matched; an unique market-clearing price and quantity are determined.

Market clearing. The market clearing mechanism is a *uniform price* auction with *uniform rationing*, as defined in, e.g., Kremer and Nyborg (2004). First, there is an unique price for all units traded. Second, agents with both marginal and infra-marginal bids (i.e., at the clearing price or at a better price) trade equal amounts.[4]

Agents. The risky asset is traded by two types of agents: high-frequency traders (HFTs) and liquidity traders (LTs). There are $N > 2$ HFTs and an infinite number of LTs. Liquidity traders submit only market orders. High-frequency traders submit both limit and marketable orders. All orders are submitted at zero cost.

All players are risk-neutral. However, HFTs have a limit on risky positions: they can submit only one limit order to buy and one limit order to sell the asset.[5] Liquidity traders are impatient; conditional on a

---

[4]Budish, Cramton, and Shim (2014) also suggest an uniform price frequent batch auction. The additional uniform rationing assumption guarantees the existence of a Nash equilibrium in pure strategies. It eliminates the HFT's incentives to undercut each other's quotes by infinitesimal amounts. In an alternative, albeit less tractable setup, a positive tick size generates the same qualitative results.

[5]This assumption is equivalent to a certain risk-aversion coefficient for HFTs. Note that the position limit does not apply to arbitrage triggered orders, as the payoff on such orders is non-negative.

liquidity shock, the probability of submitting a market order is:[6]

$$Prob \text{ (LT initiates trade)} = \frac{1}{1 + \xi \times \text{expected execution time}} \tag{1}$$

where $\xi > 0$ is an "impatience" coefficient. For $\xi \to \infty$, the LT only initiates a trade if the expected execution time is zero. Sections 2 through 4 consider the case of patient liquidity traders, i.e., $\xi = 0$. The impatient LT traders equilibrium is discussed in Section 5.

Events. Two types of exogenous events might occur, as in Menkveld and Zoican (2015). First, there can be *news*: common value innovations are described by a compound Poisson with intensity $\eta > 0$. The common value at time t is v_t; Conditional on news, it either jumps to $v_t + \sigma$ for "good" news or $v_t - \sigma$ for "bad" news. The common value of the asset is a martingale; hence, good and bad news are equally likely.

Second, an LT might receive a private value shock. Private value shocks arrive as a Poisson process with intensity $\mu > 0$. The size of the private value shock is either $+\sigma'$ or $-\sigma'$, with equal probabilities. Further, we assume $\sigma' > \sigma$. Conditional on a private value shock, with the probability in equation (1), the LT submits an order for Q units, where $Q \in [2, N)$, a deterministic quantity.[7]

At most one event arrival is possible before the market clears (as in, e.g., Dugast, 2015). The three possible states (news arrival, LT arrival, or no arrival) are tabulated below together with their probabilities.

Event	Probability
News arrival	$\frac{\eta}{\eta+\mu+\tau}$
LT arrival	$\frac{\mu}{\eta+\mu+\tau}$
No event	$\frac{\tau}{\eta+\mu+\tau}$

Information structure. HFTs learn about common value shocks with a delay. For all HFTs, the learning delay is exponentially distributed with parameter ϕ. Conditional on news, each HFT independently learns the

---

[6]Equivalently, the private value shocks of liquidity traders have a finite and random life.

[7]The assumption $Q < N$ guarantees price competition between HFTs. Further, $Q \geq 2$ implies no single HFT can unilaterally set the clearing price.

new common value before market clearing with probability

$$p \equiv \frac{\phi}{\phi + \tau}. \tag{2}$$

If an HFT learns the news before market clears, we refer to him as "informed," or HFI. Otherwise, we denote an "uninformed" HFT by HFU. LTs are uninformed and only motivated to trade by private value shocks. All model parameters are summarized in Appendix A.

2.2 Adverse selection on batch auction markets

2.2.1 Informed HFT sniping profits

Due to the risky position limit, at $t = 0$ HFTs post at most one limit order to buy and one limit order to sell the risky asset. Let s_l denote the half-spread on the limit orders around the common value v. If there is news, HFIs can "snipe" outstanding quotes of HFUs.

To build intuition, consider first the case where the number of informed competitors is common knowledge to HFIs. Assume a good news arrival: the value of the asset is $v + \sigma$ and the outstanding ask quote is $v + s_l$. If a single HFT learns about the good news, then he can post a marketable buy order at $v + s_l$ for $N - 1$ units and earn $(N - 1)(\sigma - s_l)$. In this case, the single HFI has a monopoly over the information and can maximize his rents. If two or more HFTs learn about the good news, Bertrand competition emerges.[8] Consequently, they post $N - 1$ marketable buy orders at $v + \sigma$ and earn zero profit.

If the number of HFIs is unknown, informed high-frequency traders are not aware whether they are "monopolists" or "Bertrand competitors." As a consequence, there is no equilibrium in pure strategies (Janssen and Rasmusen, 2002). High-frequency informed traders submit orders at $v + s_m$, with $s_m \in [s_l, \sigma]$

---

[8]Budish, Cramton, and Shim (2015) assume Betrand competition across HFT snipers. Our paper nests their model as a special case when $\phi \to \infty$.

drawn from a distribution $F(s_m)$. Their expected profit conditional on news arrival and s_l is π_{snipe}, where

$$\pi_{snipe}(s_m, F(s_m)|s_l) = p \underbrace{\sum_{k=0}^{N-1} \binom{N-1}{k} (1-p)^{N-1-k} p^k}_{\text{Probability of } k \text{ HFIs}} \underbrace{F(s_m)^k}_{\text{Probability winning bid}} \underbrace{(N-1-k)}_{\text{Sniped quotes}} (\sigma - s_m)$$

$$= p(1-p)(\sigma - s_m)(N-1)[1 - p(1 - F(s_m))]^{N-2}. \tag{3}$$

Figure 1 illustrates how HFI sniping orders are matched with HFU stale quotes for the risky asset. The HFI sniping orders are not exposed to adverse selection risk and, conditional on execution, yield a guaranteed profit. Therefore, there is no position limit on arbitrage orders.

[insert Figure 1 here]

With probability p, an HFT is informed. Out of the remaining $N-1$ HFTs, exactly k are informed with probability $\binom{N-1}{k}(1-p)^{N-1-k}p^k$ and submit bids from the same distribution F. An informed HFT has the highest bid of all k HFIs with probability $F(s_m)^k$. The HFT with the highest bid trades against all the outstanding quotes of $N-k-1$ HFUs and earns $\sigma - s_m$ for each "sniped" quote.

Lemma 1 states the partial equilibrium distribution for prices on HFI sniping orders.

Lemma 1. (Distribution of sniping orders) *If they observe a common value innovation, all HFTs withdraw all outstanding quotes and submit marketable limit orders for $N-1$ units at price $v + s_m$ (buy orders, for good news) or $v - s_m$ (sell orders, for bad news), where s_m is drawn from the distribution:*

$$F(s_m) = \begin{cases} 0 & , \text{if } s_m < s_l \\ \frac{1-p}{p}\left(\sqrt[N-2]{\frac{\sigma - s_l}{\sigma - s_m}} - 1 \right) & , \text{if } s_m \in \left[s_l, \sigma - (\sigma - s_l)(1-p)^{N-2} \right) \\ 1 & , \text{if } s_m \geq \sigma - (\sigma - s_l)(1-p)^{N-2}. \end{cases} \tag{4}$$

In a mixed strategy equilibrium, the HFI is indifferent between all pure strategies in the support.

10

Consequently, all pure strategies in the support have equal expected profits. If an HFT learns the news (probability p) and submits a marketable order at $v \pm s_l$ (lowest half spread in the support) then he is only successful if no other HFT is active, with probability $(1 - p)^{N-1}$. The sniping expected profit is therefore

$$\pi_{\text{snipe}}(s_l) = (N - 1)\, p\, (1 - p)^{N-1}\, (\sigma - s_l). \tag{5}$$

Figure 2 illustrates the distribution of prices on HFI sniping orders for different values of τ and N.

[insert Figure 2 here]

The conditional expected HFI per-unit sniping profit, $\sigma - s_m$, decreases in both the expected length of the batch interval $1/\tau$ and the number of HFTs N. First, as the expected batch interval length increases, more HFTs are likely to become informed. Second, as N becomes larger, more HFTs (in absolute terms) become informed. The two effects strengthen the competition between HFIs. With stronger competition, HFIs post sniping orders closer to the efficient price than to the stale outstanding quotes. Consequently, per-unit HFI sniping rents decrease with both the batch interval length and the number of HFTs on the market.

2.2.2 Uninformed HFT adverse selection costs

Next, we compute the expected adverse selection cost on limit orders for uninformed HFTs. Consider an HFT submits who submits limit orders with half-spread s_l. Conditional on news arrival, with probability $1 - p$, the HFT is uninformed (HFU). Consequently, he faces adverse selection risk if there is at least one HFI on the market.

The conditional loss is a function of the maximum bid submitted by HFIs, which depends in turn on the number k of HFIs. The expected loss is $\ell(s_l)$, where

$$\ell(s_l) = (1 - p) \sum_{k=0}^{N-1} \underbrace{\binom{N-1}{k} (1 - p)^{N-1-k}\, p^k}_{\text{Probability of } k \text{ HFIs}} \underbrace{\mathbb{E}_s\left(\sigma - \max_{i=1}^{k} s_{m,i} \middle| k\right)}_{\text{Expected loss conditional on k HFIs}}. \tag{6}$$

11

If the HFT learns the new common value before market clearing, he updates the stale quote and faces no adverse selection risk. With probability $1 - p$, the HFT is uninformed. He faces adverse selection risk from k HFIs; the probability of exactly k HFIs is again $\binom{N-1}{k}(1-p)^{N-1-k}p^k$. The HFU expected loss conditional on trading is given by the absolute expected difference between the new common value and the closest price to it across the k HFIs' sniping orders, that is $\mathbb{E}_s\left(\sigma - \max_{i=1}^k s_{m,i}|k\right)$

The cumulative distribution function of $\max_{i=1}^k s_{m,i}$ is $F^k(s_m)$. It follows that

$$\mathbb{E}_s\left(\sigma - \max_{i=1}^k s_{m,i}|k\right) = \int_{s_l}^{\sigma-(\sigma-s_l)(1-p)^{N-2}} (\sigma - s_m)kF^{k-1}(s_m)\frac{\partial F(s_m)}{\partial s_m}ds_m \tag{7}$$

From equations (6) and (7) it follows that:

$$\ell(s_l) = (1-p)\int_{s_l}^{\sigma-(\sigma-s_l)(1-p)^{N-2}} \sum_{k=0}^{N-1}\binom{N-1}{k}(1-p)^{N-1-k}p^k(\sigma - s_m)kF^{k-1}(s_m)\frac{\partial F(s_m)}{\partial s_m}ds_m$$

$$= (1-p)\int_{s_l}^{\sigma-(\sigma-s_l)(1-p)^{N-2}} (N-1)p\left[1-p(1-F(s_m))\right]^{N-2}\frac{\partial F(s_m)}{\partial s_m}ds_m$$

$$= \int_{s_l}^{\sigma-(\sigma-s_l)(1-p)^{N-2}} \pi_{\text{snipe}}(s_m|s_l)\frac{\partial F(s_m)}{\partial s_m}ds \tag{8}$$

The expected adverse selection cost on limit orders for HFUs is equal to the expected sniping profit of HFIs over all potential sniping order half-spreads s_m. Since HFIs are in equilibrium indifferent between all half-spreads s_m in the mixed strategy support, the right hand side of equation (8) is simply $\pi_{\text{snipe}}(s_l)$. Therefore, the HFI sniping profits are equal to the HFU sniping losses.

$$\ell(s_l) = \pi_{\text{snipe}}(s_l). \tag{9}$$

2.3   Equilibrium

We search for HFT-symmetric Nash equilibria in pure and mixed strategies. In particular, at any point in time between $t = 0$ and the batch auction, an equilibrium consists of the HFT orders to trade a specific price and quantity of the risky asset. Since the equilibrium is symmetric, all HFTs take the same actions.

The HFT expected profit from providing liquidity, i.e., submitting a buy and a sell order at $v \pm s_l$ is π_{liq}, where

$$\pi_{\text{liqudity}} = \underbrace{\frac{\mu}{\mu + \eta + \tau}}_{\text{LT arrives}} \frac{Q}{N} s_l - \underbrace{\frac{\eta}{\mu + \eta + \tau}}_{\text{News arrives}} \left[\underbrace{\ell(s_l) - \pi_{\text{snipe}}(s_l)}_{=0} \right] \tag{10}$$

With probability $\frac{\mu}{\mu+\eta+\tau}$, a liquidity trader arrives with a demand for Q units of the asset. Each HFT trades an equal share $\frac{Q}{N}$ and earns the half-spread s_l. Alternatively, with probability $\frac{\eta}{\mu+\eta+\tau}$, news arrives. In expectation, the HFT loses $\ell(s_l)$ on the stale quote if uninformed and earns $\pi_{\text{snipe}}(s_l)$ if informed. With probability $\frac{\tau}{\mu+\eta+\tau}$ the market clears before any of the events.

From equation (9), if an HFT submits limit orders at $t = 0$, its expected payoff conditional on news arrival is zero. If liquidity provision were mandatory, then price competition on limit orders would push the spread to zero, as in Budish, Cramton, and Shim (2015).

However, HFTs may choose not to provide liquidity and act instead as arbitrageurs. As an arbitrageur, an HFT does not expose himself to adverse selection risk by posting quotes at $t = 0$. Rather, he only submits sniping orders after news. In this case his expected profit is $\pi_{\text{arbitrageur}}$, where:

$$\pi_{\text{arbitrageur}} = \underbrace{\frac{\eta}{\mu + \eta + \tau}}_{\text{News arrives}} \pi_{\text{snipe}}(s_l). \tag{11}$$

In equilibrium, $\pi_{\text{liquidity}} = \pi_{\text{arbitrageur}}$. Consequently the equilibrium half-spread s^* is pinned down by equation (12):

$$\frac{\mu}{\mu + \eta + \tau} \frac{Q}{N} s^* = \frac{\eta}{\mu + \eta + \tau} (N - 1) p (1 - p)^{N-1} (\sigma - s^*). \tag{12}$$

Proposition 1 describes equilibrium strategies for HFTs.

Proposition 1. (Equilibrium) *The following HFT strategies form a Nash equilibrium in the trading game:*

(i) At $t = 0$, all HFTs submit a buy limit order for one unit at $v_0 - s^*$ and a sell limit order for one unit at $v_0 + s^*$, where $p = \frac{\phi}{\phi+\tau}$ and

$$s^* = \sigma \frac{\eta (N - 1) p (1 - p)^{N-1}}{\mu \frac{Q}{N} + \eta (N - 1) p (1 - p)^{N-1}}. \tag{13}$$

13

(ii) If an HFT observes a good (bad) news item, then he immediately cancels the outstanding ask (bid) limit order and submits a marketable buy (sell) order for $(N-1)$ units at price $v_0 + s_m$, where s_m has distribution $F(s_m)$, as defined in equation (4).

In particular, no HFT has unilateral incentive to deviate from the half-spread s^*. Figure 3 illustrates the equilibrium mechanism for the ask side. Suppose one HFT, e.g., HFT_1 posts a sell order at $v + s^* - \varepsilon$. Since $Q \geq 2$, the clearing price is still s^*. Due to uniform rationing, HFT_1 traded volume also remains unchanged, i.e., $\frac{Q}{N}$. Such a deviation is not strictly profitable. Otherwise, suppose HFT_1 posts a sell order at $v + s^* + \varepsilon$. Such deviation is strictly dominated by the equilibrium strategy, as HFT_1 never trades; the other $N - 1$ high-frequency traders split the liquidity demand Q among themselves.

Lemma 2 describes the behavior of the equilibrium half-spread with respect to news intensity, news arrival size, liquidity shock intensity, and liquidity demand.

Lemma 2. *The equilibrium half-spread s^*,*

$$s^* = \sigma \frac{\eta(N-1)p(1-p)^{N-1}}{\mu\frac{Q}{N} + \eta(N-1)p(1-p)^{N-1}}, \tag{14}$$

increases in the size of value innovations (σ), news intensity (η), and decreases in the liquidity traders' arrival intensity (μ) and liquidity demand (Q).

Lemma 2 is consistent with existing results in the literature (for a detailed survey see, e.g., Biais, Glosten, and Spatt, 2005). The positive spread emerges as a compensation for the opportunity cost of being a pure HFT arbitrageur and is thus proportional to sniping profits. Sniping profits increase in the news intensity η and news size σ, and so does the equilibrium spread. A larger μ or Q increase the HFT payoff from providing liquidity relative to the foregone arbitrageur profits and therefore the spread decreases.

Proposition 2 describes the behaviour of the equilibrium spread s^* with respect to the batch auction frequency τ and the number of HFTs N.

Proposition 2. *The equilibrium spread on the batch auction market, s_l,*

(i) increases (decreases) in the batch auction frequency τ if $\tau < \tau^*(N)$ (and $\tau \geq \tau^*(N)$, respectively),

(ii) increases (decreases) in increases in the number of HFTs N if $N < N^*(\tau)$ (and $N \geq N^*(\tau)$, respectively),

where τ^* and N^* are defined as

$$\tau^*(N) \equiv (N-1)\phi \tag{15}$$

$$N^*(\tau) \equiv \frac{\log\frac{\tau}{\phi+\tau} - 2 - \sqrt{4 + \left(\log\frac{\tau}{\phi+\tau}\right)^2}}{2\log\frac{\tau}{\phi+\tau}} \tag{16}$$

Figure 4 illustrates the result. If batch auctions are very frequent, i.e., if the expected batch interval length $1/\tau$ approaches zero, the equilibrium spread becomes arbitrarily small: No HFT learns the common value innovation, so there is no sniping. As the batch frequency decreases, the probability of each HFT becoming informed increases and so do the sniping profits. However, as the batch frequency decreases even more, more HFTs become informed in expectation. Competition between arbitrageurs becomes stronger, pushing down the expected sniping rent. Therefore, the opportunity cost for providing liquidity is lower, and the equilibrium spread decreases. The competition effect dominates for long enough batch intervals, i.e., for $\tau > \tau^*$.

[insert Figure 4 here]

A similar trade-off emerges as the number of HFTs, N, is allowed to vary. For a low N, arbitrageur competition is weak. On the other hand, the expected absolute number of outstanding stale quotes, the "size of the prize," is proportional to N: the expected sniping profit is also low. As N increases, stronger competition between HFIs drives the expected-profit-per-stale-quote down. At the same time, the number of stale quotes is larger, generating an opposite channel. For $N \geq N^*$, the competition effect dominates, and profits decrease with N; for $N < N^*$, the "size of the prize" effect dominates, and profits increase with N.

15

3 Liquidity benchmark: the limit order market

A natural benchmark for batch auction market quality is the limit order market, the prevailing market design on modern exchanges. Does the transition from continuos trading to discrete auctions improve liquidity?

In this section, we compare the batch auction equilibrium spread in Section 2 to the outcome of a model where the risky asset is traded on a limit order market. The limit order market is modelled as in Budish, Cramton, and Shim (2015). Orders have price-time priority: they are executed in the order they arrive at the market. Since HFTs have equal monitoring intensities ϕ, each HFT has a probability $\frac{1}{N}$ of being first to the market.

First, no more than Q HFTs submit limit orders. Only the first Q HFT orders that reach the exchange are executed against the liquidity trader's market order. Unlike in the batch auction market, there is no rationing of (infra-)marginal HFT liquidity bids. Consequently, the $(Q + 1)$th HFT never trades with LT, but is still exposed to adverse selection risk. Hence, he would prefer not to submit a limit order.[9]

As in Section 2, HFTs are indifferent between being a liquidity provider and pure arbitrageur in equilibrium. Let s_{LOB} be the quoted half-spread on the limit order market. If exactly $Q' < Q$ HFTs submit one limit order on each side of the book, it follows that in equilibrium

$$\underbrace{\mu s_{\text{LOB}} - \frac{N-1}{N}\eta(\sigma - s_{\text{LOB}}) + \frac{1}{N}(Q' - 1)\eta(\sigma - s_{\text{LOB}})}_{\text{Liquidity provider payoff}} = \underbrace{\frac{1}{N}Q'\eta(\sigma - s_{\text{LOB}})}_{\text{Sniper payoff}}. \tag{17}$$

The left hand side of equation (17) represents the expected profit of an HFT who submits limit orders at $v \pm s_{\text{LOB}}$. Liquidity traders arrive to the market with intensity μ. Each of the Q' HFTs with limit orders in the book trade one unit and earn the spread s_{LOB}. Alternatively, with intensity η, there is news. An HFT with a quote in the book incurs the adverse selection cost $\sigma - s_{\text{LOB}}$ whenever he is not first to the market, with probability $\frac{N-1}{N}$. If the HFT is first to the market after news, with probability $\frac{1}{N}$, he cancels his own quote and consumes the remaining $Q' - 1$, earning $\sigma - s_{LOB}$ for each.

---

[9]An important question is how to achieve HFTs coordination on the subset of limit order submitters. An easily available solution is to use *top-of-the-book* limit orders. Top-of-the-book orders only become effective if the cumulative depth at the desired price is below a threshold. Any order that would have queued behind the threshold is automatically cancelled. Such orders are available on, for example, Xetra Deutsche Boerse. Importantly, they do not require HFTs to continuosly observe the state of the book.

The right hand side of (17) represents the expected profit of an HFT who does not submit limit orders. Conditional on news (intensity η) and being first to the market (probability $\frac{1}{N}$), he consumes all Q' stale quotes and earns a profit of $\sigma - s_{\text{LOB}}$ per traded unit.

The equilibrium spread on a limit order market is given by the unique solution to equation (17), i.e.,

$$s_{\text{LOB}}^* = \sigma \frac{\eta}{\eta + \mu}. \tag{18}$$

The result is qualitatively identical to the limit order market equilibrium spread in Budish, Cramton, and Shim (2015). In particular, the equilibrium spread s_{LOB}^* does not depend on the number of HFTs N. There is an equilibrium in which exactly Q HFTs submit limit orders at half-spread s_{LOB}^*.

The time priority rule generates two significant differences between the batch auction and limit order markets. On the one hand, there is an *excess supply of liquidity* in the batch auction market. All N HFTs submit limit orders to the batch auction, but only $Q < N$ do so on the limit order market. As a consequence, in equilibrium trades on the batch auction market are rationed, whereas on the limit order market they are not. Rationing decreases the profits from liquidity supply, and increases the spread.

On the other hand, there is no *price competition* between HFT arbitrageurs on the limit order market; the first HFT to the market trades against all stale quotes. Since there is no time priority on the batch auction market, there is price competition between all informed HFT arbitrageurs. As a consequence, the expected adverse selection cost and the equilibrium decrease.

Depending on whether the excess liquidity supply or the price competition effect dominates, the batch auction market spread can be higher or lower than for the limit order market. Proposition 3 establishes the necessary and sufficient condition for a batch auction market to improve liquidity.

Proposition 3. *A batch auction market improves liquidity relative to a limit order book if and only if*

$$\Gamma(Q, N, \phi, \tau) \equiv \frac{N}{Q}(N-1)\frac{\phi}{\phi + \tau}\left(\frac{\tau}{\phi + \tau}\right)^{N-1} < 1. \tag{19}$$

Channel decomposition. We can identify the liquidity supply and the competition channels in equation (19). To this end, we take the natural logarithm of both the left- and the right-hand sides:

$$\underbrace{\log \frac{N}{Q}}_{\text{Liquidity supply channel, >0}} + \underbrace{\log \left[(N-1) \frac{\phi}{\phi+\tau} \left(\frac{\tau}{\phi+\tau} \right)^{N-1} \right]}_{\text{Competition channel, <0}} < 0. \tag{20}$$

The first term, $\log \frac{N}{Q}$, is always positive. It is the log-ratio between available liquidity supply on batch auction and limit order markets, i.e., the excess liquidity supply channel. The second term, $\log \left[(N-1) \frac{\phi}{\phi+\tau} \left(\frac{\tau}{\phi+\tau} \right)^{N-1} \right]$, is always negative.[10] From equation (5), the second term is proportional to the expected sniping profit and therefore inversely related to price competition among arbitrageurs.

Consequently, a batch auction improves liquidity if and only if the competition channel dominates the liquidity supply channel, that is if

$$\underbrace{\left\| \log \left[(N-1) \frac{\phi}{\phi+\tau} \left(\frac{\tau}{\phi+\tau} \right)^{N-1} \right] \right\|}_{\text{Competition channel}} > \underbrace{\log \frac{N}{Q}}_{\text{Liquidity supply channel}}. \tag{21}$$

From Corollary 1, whether the batch auction improves liquidity is determined by market-wide parameters (batch frequency, HFT competition) rather than stock-specific ones (e.g., news intensity). An important policy implication is that if the transition from continuous trading to batch auctions is optimal, it can implemented market-wide without the need to fine tune batch frequency for each individual stock.

Corollary 1. *The ranking of the batch auction and limit order equilibrium spreads only depends on LT demand size, batch length, and HFT competition. In particular, it does not depend on news intensity and size, and liquidity trader arrival rate.*

Complements or substitutes? Figure 5 displays the contour plots for the ratio between the batch auction and the limit order book equilibrium spreads, as a function of the number of HFTs and the batch interval

---

[10]It immediately follows from the fact that $(N-1) \frac{\phi}{\phi+\tau} \left(\frac{\tau}{\phi+\tau} \right)^{N-1} < 1$.

expected length. The number of HFTs N and the batch interval $\frac{1}{\tau}$ have an ambiguous impact on the equilibrium spread on the batch auction market.

[insert Figure 5 here]

To determine whether a higher N and a longer batch interval $\frac{1}{\tau}$ are substitutes in promoting stronger arbitrageur competition, we apply the implicit function theorem to $\Gamma(Q, N, \phi, \tau)$:

$$\frac{d\tau}{dN} = -\left(\frac{\partial \Gamma}{\partial \tau}\right)^{-1} \frac{\partial \Gamma}{\partial N} \implies \frac{d\tau}{dN} = \frac{\tau(\phi + \tau)\left[1 - 2N - (N-1)N\log\frac{\tau}{\phi+\tau}\right]}{(N-1)N\left[\tau - (N-1)\phi\right]} \tag{22}$$

From the proof to Proposition 2, it follows that a longer batch length and higher HFT competition are (locally) substitutes if and only if

$$[\tau - \tau^*(N)]\,[N - N^*(\tau)] < 0, \tag{23}$$

and (locally) complements otherwise.

Equation (23) indicates the optimal policy with respect to τ following a change in HFT competition N. Suppose an HFT exogenously withdraws from the market. To preserve liquidity, the regulator should increase batch auction frequency if $\tau \geq \tau^*$ and decrease it otherwise. Equation (22) implies the necessary auction frequency change to keep the equilibrium spread unchanged.

4   The information "arms race'" in batch auction markets

In this section, we study the effect of more intense HFT market monitoring (a higher ϕ) on equilibrium liquidity. High-frequency traders have an incentive to have more information than their competitors: an arms' race emerges (Budish, Cramton, and Shim, 2015), characterized by an ever higher monitoring intensity ϕ.

On the limit order market, such an HFT "arms' race" does not influence liquidity. With time priority, the race objective is to be first to the market. If all HFTs invest equally in monitoring technology, then all of them are first to the market with probability $1/N$.

19

However, on a batch auction market, being first to the market is not so important as processing new information before the market clears. On the one hand, a higher ϕ increases the probability each individual HFT becomes informed; consequently it generates a higher sniping profits and the spread. On the other hand, a higher ϕ implies *all* HFTs are more likely to be informed. Therefore, price competition between HFT arbitrageurs is in expectation stronger, driving spreads down.

Proposition 4 formalizes the result. Liquidity improves with HFT monitoring in batch auction markets if the arms' race exceeds a certain intensity. Technology investment costs notwithstanding, the HFT arms race can improve liquidity on batch auction markets – as opposed to on limit order books. The result arises since the arms race promotes price competition on batch auctions markets, as opposed to speed competition.

Proposition 4. *The equilibrium spread on the batch auction market, s^*, decreases in the monitoring intensity ϕ if and only if $\phi > \frac{\tau}{N-1}$. In particular,*

$$\lim_{\phi \to \infty} s^* = 0. \tag{24}$$

The limiting result in Proposition 4 is an exact counterpart to the batch auction market equilibrium in Budish, Cramton, and Shim (2015). The authors assume "fast" HFTs act immediately on the information, i.e., infinite monitoring intensity generates a zero spread.

Figure 6 illustrates how a more intense arms race can change the optimal market structure from a limit order book to a batch auction market.

[insert Figure 6 here]

Therefore, policies that intensify HFT monitoring (e.g., allowing for colocation) can reduce the equilibrium spread and can facilitate the transition from limit order markets to batch auctions. The "arms race" implications strikingly differ with market structure.

5   Extension: Impatient liquidity traders

A salient implication of Proposition 2 is that a lower auction frequency improves liquidity, as in expectation more HFTs become informed and the resulting stronger competition decreases sniping profits. However, fewer auctions comes at the cost of longer waiting times for liquidity traders.

In this section, we relax the "infinitely patient LTs" assumption, i.e., we allow $\xi > 0$. If liquidity traders are impatient, then upon arrival LTs trade with probability $\frac{\tau}{\tau + \xi} < 1$. The liquidity provider expected profit in equation (10) becomes

$$\pi_{\text{liquidity}} = \underbrace{\frac{\mu}{\mu + \eta + \tau}}_{\text{LT arrives}} \frac{Q}{N} \frac{\tau}{\tau + \xi} s_l - \underbrace{\frac{\eta}{\mu + \eta + \tau}}_{\text{News arrives}} \left[\ell(s_l) + \pi_{\text{snipe}}(s_l) \right]. \tag{25}$$

Consequently, the equilibrium spread also becomes a function of ξ:

$$s_\xi^* = \sigma \frac{\eta (N-1) p (1-p)^{N-1}}{\mu \frac{Q}{N} \frac{\tau}{\tau + \xi} + \eta (N-1) p (1-p)^{N-1}}. \tag{26}$$

The equilibrium spread s_ξ^* increases with LT impatience ξ. LT transaction costs increase twofold: first due to the longer waiting time, and second due to the higher spread.

The equilibrium spread in Section 2 is a lower bound for spread if liquidity traders are impatient. Therefore, if a limit order market dominates a batch auction market with patient LTs, it also dominates a batch auction market with *impatient* LTs. Figure 7 illustrates the result.

[insert Figure 7 here]

Proposition 3 has the following corollary for the setup with impatient liquidity traders.

Corollary 2. *A batch equation market improves liquidity relative to a limit order market if*

$$\Gamma(Q, N, \phi, \tau) < \frac{\tau}{\tau + \xi} < 1. \tag{27}$$

For any $\tau > 0$, there exists a $\xi^*(\tau)$ such that the batch auction market is always worse relative to the limit order book if $\xi > \xi^*(\tau)$.

From Corollary 2 and equation (28), we identify a third channel, LT impatience, that determines whether the transition to a batch auction market is optimal. The equilibrium spread is lower on the auction market than on the limit order book if and only if

$$\underbrace{\log \frac{N}{Q}}_{\text{Liquidity supply channel}} + \underbrace{\log\left(1 + \frac{\xi}{\tau}\right)}_{\text{Impatient LT channel}} < \underbrace{\left\|\log\left[(N-1)\frac{\phi}{\phi+\tau}\left(\frac{\tau}{\phi+\tau}\right)^{N-1}\right]\right\|}_{\text{Competition channel}} \qquad (28)$$

From Proposition 2, both a very high and very low auction frequency have a positive impact of liquidity. However, if liquidity traders are significantly impatient, Corollary 2 implies more frequent batch auctions can be optimal.

6 Conclusions

This paper finds that the transition from continuous- to discrete-time trading can increase transaction costs for liquidity traders. Two offsetting effects emerge on a batch auction market. First, HFTs supply excess liquidity which leads to trade rationing. The profit from supplying liquidity decreases, and the equilibrium spread increases as a result. Second, HFTs with private information compete in prices over the arbitrage opportunities. Price competition reduces arbitrage rents and adverse selection costs: the equilibrium spread decreases. The two effects have opposite signs; depending on which one dominates, the batch auction can improve or hurt liquidity relative to the limit order market.

The paper's findings contribute to the public debate on alternative market design. It generates two important implications. First, on batch auction markets speed competition reinforces speed competition. Therefore, the HFT "arms' race" reduces expected arbitrage rents and has positive social value. Second, the optimal decision on whether to replace limit-order with batch auction market only depends on HFT competition, and not on stock-specific characteristics. Therefore, such market design changes can be

implemented exchange-wide.

Finally, the model is appealing as it can be solved in closed-form. The framework can be extended to analyse other important questions, such as the competition for order flow between batch auction and limit order markets, or the equilibrium arms' race intensity on auction markets.

References

Baron, Matthew, Jonathan Brogaard, and Andrei A. Kirilenko, 2012, The trading profits of high-frequency traders, *Working paper*.

Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium fast trading, *Journal of Financial Economics* 116, 292–313.

Biais, Bruno, Larry Glosten, and Chester Spatt, 2005, Market microstructure: A survey of microfoundations, empirical results, and policy implications, *Journal of Financial Markets* 8, 217–264.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014, High frequency trading and price discovery, *Review of Financial Studies* 27, 2267–2306.

Budish, Eric, Peter Cramton, and John Shim, 2014, Implementation details for frequent batch auctions, *American Economic Review: Papers and Proceedings* 104, 418–424.

——— , 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *The Quarterly Journal of Economics* 130, 1547–1600.

Du, Songzi, and Haoxiang Zhu, 2015, Welfare and optimal trading frequency in dynamic double auctions, *Working paper*.

Dugast, Jérôme, 2015, Unscheduled news and market dynamics, *Working paper*.

Economides, Nicholas, and Robert A. Schwartz, 1995, Electronic call market trading, *The Journal of Portfolio Management* 21, 10–18.

Farmer, J. Doyne, and Spyros Skouras, 2012, Review of the benefits of a continuous market vs. randomised stop auctions and of alternative priority rules, *Review of the Markets in Financial Instruments Directive*.

Fricke, Daniel, and Austin Gerig, 2015, Too fast or too slow? Determining the optimal speed of financial markets, *Working paper*.

Hagstromer, Björn, and Lars Norden, 2013, The diversity of high-frequency traders, *Journal of Financial Markets* 16, 741–770.

Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld, 2011, Does algorithmic trading improve liquidity?, *Journal of Finance* 66, 1–33.

Hendershott, Terrence, and Pamela C. Moulton, 2011, Automation, speed, and stock market quality: The nyse's hybrid, *Journal of Financial Markets* 14, 568–604.

Janssen, Maarten, and Eric Rasmusen, 2002, Bertrand competition under uncertainty, *The Journal of Industrial Economics* 50, 11–21.

Jovanovic, Boyan, and Albert J. Menkveld, 2015, Dispersion and skewness of bid prices, *Working paper*.

Kremer, Ilan, and Kjell G. Nyborg, 2004, Underpricing and market power in uniform price auctions, *The Review of Financial Studies* 17, 849–877.

Laughlin, Gregory, Anthony Aguirre, and Joseph Grundfest, 2014, Information transmission between financial markets in chicago and new york, *Financial Review* 49, 283–312.

Madhavan, Ananth, 1992, Trading mechanisms in securities markets, *The Journal of Finance* 47, 607–641.

Malinova, Katya, Andreas Park, and Ryan Riordan, 2013, Do retail traders suffer from high frequency traders?, *WFA 2013 paper*.

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

Menkveld, Albert, and Marius A. Zoican, 2015, Need for speed? Exchange latency and liquidity, *Tinbergen Institute Discussion Paper 14-097/IV/DSF78*.

Menkveld, Albert J., 2014, High frequency traders and market structure, *The Financial Review (Special Issue: Computerized and High-Frequency Trading)* 49, 333–344.

Riordan, Ryan, and Andreas Storkenmaier, 2012, Latency, liquidity and price discovery, *Journal of Financial Markets* 15, 416–437.

SEC, 2010, Concept release on equity market structure, Release No. 34-61358; File No. S7-02-10.

Wah, Elaine, Dylan Hurd, and Mic Wellman, 2015, Strategic market choice: Frequent call markets vs. continuous double auctions for fast and slow traders, *Working paper*.

Wah, Elaine, and Michael Wellman, 2013, Latency arbitrage, market fragmentation, and efficiency: A two-market model, *Proceedings of the fourteenth ACM Conference*.

Ye, Mao, Chen Yao, and Jiading Gai, 2013, The externalities of high frequency trading, *Working paper*.

Yueshen, Bart Zhou, 2014, Queuing uncertainty in limit order market, *Working paper*.

Appendix

A Notation summary

Model parameters and their interpretation.

Parameter	Definition
v_t	Common value of the risky asset at time t.
τ	Poisson intensity of batch auction times.
η	Poisson intensity of news arrival.
μ	Poisson intensity of LT arrival.
ϕ	Poisson intensity of HFT news monitoring.
σ	Size of news, i.e., common value innovations.
σ'	Size of liquidity shocks, i.e., private value innovations.
N	Number of high-frequency traders.
ξ	Liquidity traders' impatience coefficient.
Q	Liquidity demand size (units).

B Proofs

Lemma 1

Proof. First, there are no values of s_m such that HFTs submit marketable orders at $v \pm s_m$ with positive probability. If an HFT assigns positive probability to s_m, then either $s_m + \epsilon$ or $s_m - \epsilon$, with ϵ close to zero, is a profitable deviation. We refer the reader to the discussion in Janssen and Rasmusen (2002), on pages 12 and 13, for an in-depth discussion of this point.

Let $F(s_m)$ be the cumulative distribution function of HFT half-spread bids on marketable orders. The expected profit from a marketable order at $v \pm s_m$ targeted at a limit order at $v \pm s_l$, conditional on news and on the limit order submitter being unaware of news, is

$$\pi_{mo}(s_m, F(s_m)) = \sum_{k=0}^{N-1} \binom{N-1}{k} (1-p)^k \, p^{N-k-1} F(s_m)^{N-k-1} \, k \, (\sigma - s_m), \tag{B.1}$$

where $p = \frac{\phi}{\phi+\tau}$, the probability an HFT observes the news. From the binomial rule, it follows that

$$\pi_{mo}(s_m, F(s_m)) = [1 - p + pF(s_m)]^{N-2} (1-p)(N-1)(\sigma - s_m). \tag{B.2}$$

In a mixed strategy equilibrium, HFTs are indifferent between all s_m in the support. Hence, for some support of s_m, the following first-order condition is true in equilibrium:

$$\frac{\partial \pi_{mo}(s_m, F(s_m))}{\partial s_m} = 0. \tag{B.3}$$

From (B.3), the equilibrium cumulative distribution function $F(s_m)$ solves the differential equation

$$1 - p + pF(s_m) - (N-2)p(\sigma - s_m)\frac{\partial F(s_m)}{\partial s_m} = 0. \tag{B.4}$$

We set the boundary condition $F(s) = 0$. That is, no HFT submits a marketable order that never executes against the stale quote. From (B.4) and the boundary condition it follows that

$$F(s_m) = \begin{cases} 0 & \text{, if } s_m < s_l \\ \frac{1-p}{p} \sqrt[N-2]{\frac{\sigma - s_l}{\sigma - s_m}} & \text{, if } s_l \geq s_m < \sigma - (\sigma - s_l)(1-p)^{N-2} \\ 1 & \text{, if } s_m \geq \sigma - (\sigma - s_l)(1-p)^{N-2} \end{cases}, \tag{B.5}$$

which is what we intended to show. □

Proposition 1

Proof. First, at $t = 0$, no HFT would strictly prefer to deviate and post quotes with a lower half-spread, e.g., $s^* - \epsilon$. Suppose HFT$_1$ deviates as such. Since $Q \geq 2$, the clearing price is still s^*, determined the other HFTs. Due to uniform rationing, HFT$_1$ traded volume also remains unchanged, i.e., $\frac{Q}{N}$. Undercutting s^* is not a strictly profitable deviation.

Second, no HFT would prefer to deviate at $t = 0$ and post orders at a larger half-spread, e.g., $s^* + \varepsilon$. Suppose again HFT$_1$ deviates as such. It follows that HFT$_1$ never trades: the market clears at s^* and each of the remaining $N - 1$ HFTs split the liquidity demand Q.

Third, from equation (12) it follows that for s^*, HFTs are indifferent between posting quotes or not. Therefore, not submitting a quote is not a strictly profitable deviation either.

After a news event, it is optimal for HFTs to cancel the outstanding stale quote as he incurs thus zero adverse selection cost. If an HFT does not cancel the quote, he is exposed to lose $\pi_{\text{arbitrageur}}$ due to adverse selection.

Finally, HFTs submit $N - 1$ marketable orders against stale quotes. These orders have either a positive profit (if sniping is successful, or quotes have been cancelled) or a zero profit if another HFT submits a better price against the stale quotes. The expected profit scales linearly with the number of market orders submitted, so each HFT attempts to snipe all of his $N - 1$ competitors. Lemma 1 establishes the distribution of HFT sniping orders prices. □

Lemma 2

Proof. First, we compute the partial derivative of s^* with respect to σ:

$$\frac{\partial s^*}{\partial \sigma} = \frac{\eta(N-1)p(1-p)^{N-1}}{\mu\frac{Q}{N} + \eta(N-1)p(1-p)^{N-1}} > 0. \tag{B.6}$$

Since the partial derivative is positive, the equilibrium half-spread s^* increases in σ.

Next, we compute the partial derivative of s^* with respect to η:

$$\frac{\partial s^*}{\partial \eta} = \sigma \frac{\mu (N-1) N Q p (1-p)^{N-1}}{\left(\eta (N-1) N (1-p)^N p + \mu (1-p) Q \right)^2} > 0. \tag{B.7}$$

Since the partial derivative is positive, the equilibrium half-spread s^* increases in η.

Next, we compute the partial derivative of s^* with respect to μ:

$$\frac{\partial s^*}{\partial \mu} = -\sigma \frac{\eta (N-1) N Q p (1-p)^{N-1}}{\left(\eta (N-1) N (1-p)^{N+1} p + \mu (1-p) Q \right)^2} < 0. \tag{B.8}$$

Since the partial derivative is negative, the equilibrium half-spread s^* decreases in μ.

Finally, we compute the partial derivative of s^* with respect to Q:

$$\frac{\partial s^*}{\partial \mu} = -\sigma \frac{\eta \mu (N-1) N p (1-p)^{N-1}}{\left(\eta (N-1) N (1-p)^{N+1} p + \mu (1-p) Q \right)^2} < 0. \tag{B.9}$$

Since the partial derivative is negative, the equilibrium half-spread s^* decreases in Q. $\qquad\square$

Proposition 2

Proof. First, we compute the partial derivative of s^* with respect to τ:

$$\frac{\partial s^*}{\partial \tau} = \sigma \frac{\eta \mu (N-1) N Q \phi \left(\frac{\tau}{\tau+\phi} \right)^N [(N-1)\phi - \tau]}{(\tau + \phi) \left(\eta (N-1) N \phi \left(\frac{\tau}{\tau+\phi} \right)^N + \mu Q \tau \right)^2}. \tag{B.10}$$

It follows that $\frac{\partial s^*}{\partial \tau} \lesseqgtr 0$ if and only if $(N-1)\phi - \tau \lesseqgtr 0$. Let $\tau^* \equiv (N-1)\phi$. Consequently, the equilibrium half-spread s^* increases in τ for $\tau < \tau^*$ and decreases in τ for $\tau \geq \tau^*$.

Second, we compute the partial derivative of s^* with respect to N:

$$\frac{\partial s^*}{\partial N} = \frac{\eta \mu Q \sigma \tau \phi \left(\frac{\tau}{\tau+\phi} \right)^N \left((N-1) N \log \left(\frac{\tau}{\tau+\phi} \right) + 2N - 1 \right)}{\left(\eta (N-1) N \phi \left(\frac{\tau}{\tau+\phi} \right)^N + \mu Q \tau \right)^2}. \tag{B.11}$$

It follows that $\frac{\partial s^*}{\partial N} \lesseqgtr 0$ if and only if

$$(N-1) N \log \left(\frac{\tau}{\tau + \phi} \right) + 2N - 1 \lesseqgtr 0. \tag{B.12}$$

28

Equation (B.12) has two real roots,

$$N_{1,2} = \frac{\log\left(\frac{\tau}{\phi+\tau}\right) - 2 \pm \sqrt{\log^2\left(\frac{\tau}{\phi+\tau}\right) + 4}}{2\log\left(\frac{\tau}{\phi+\tau}\right)}. \tag{B.13}$$

Next, we prove that one of the roots $- N_2 -$ is always smaller than one, i.e., the smallest possible number of HFTs on the market. Denote $p \equiv \frac{\phi}{\phi+\tau}$, then N_2 can be written as

$$N_2 = \frac{\log p - 2 + \sqrt{\log^2 p + 4}}{2\log p}. \tag{B.14}$$

We observe N_2 increases in p,

$$\frac{\partial N_2}{\partial p} = \frac{1 - \frac{2}{\sqrt{\log^2(p)+4}}}{p\log^2(p)} > 0, \tag{B.15}$$

and further that $\lim_{p\to 1} N_2 = \frac{1}{2}$. It follows that for all p, $N_2 < \frac{1}{2} < 1$. We denote the other root of (B.12) as N^*, i.e.,

$$N^* \equiv N_1 = \frac{\log\left(\frac{\tau}{\phi+\tau}\right) - 2 - \sqrt{\log^2\left(\frac{\tau}{\phi+\tau}\right) + 4}}{2\log\left(\frac{\tau}{\phi+\tau}\right)}. \tag{B.16}$$

Consequently, it follows that the equilibrium half-spread s^* increases in N for $N < N^*$ and decreases in N for $N \geq N^*$. □

Proposition 3

Proof. From equations (13) and (18), we compute the ratio between the batch auction market and limit order market equilibrium half-spreads,

$$\frac{s_l^*}{s_{\text{LOB}}^*} = \frac{(N-1)N\phi(\eta+\mu)\left(\frac{\tau}{\tau+\phi}\right)^N}{\eta(N-1)N\phi\left(\frac{\tau}{\tau+\phi}\right)^N + \mu Q\tau}. \tag{B.17}$$

The batch auction market improves liquidity relative to the limit order book if and only if $\frac{s_l^*}{s_{\text{LOB}}^*} < 1$ or, equivalently, if

$$(N-1)N\phi(\eta+\mu)\left(\frac{\tau}{\tau+\phi}\right)^N < \eta(N-1)N\phi\left(\frac{\tau}{\tau+\phi}\right)^N + \mu Q\tau. \tag{B.18}$$

We subtract $\eta(N-1)N\phi\left(\frac{\tau}{\tau+\phi}\right)^N$ from each side,

$$(N-1)N\phi\mu\left(\frac{\tau}{\tau+\phi}\right)^N < \mu Q\tau, \tag{B.19}$$

and we divide by $\mu > 0$ to obtain

$$(N-1)N\phi \left(\frac{\tau}{\tau+\phi}\right)^N < +Q\tau. \tag{B.20}$$

Rearranging, we obtain

$$(N-1)\frac{N}{Q}\frac{\phi}{\tau+\phi}\left(\frac{\tau}{\tau+\phi}\right)^{N-1} < 1. \tag{B.21}$$

$\square$

Corollary 1

Proof. Follows immediately from inspection of equation (B.21). $\square$

Proposition 4

Proof. We compute the partial derivative of s^* with respect to ϕ:

$$\frac{\partial s^*}{\partial \phi} = \frac{\eta\mu(N-1)NQ\sigma\left(\frac{\tau}{\tau+\phi}\right)^{N+1}[\tau-(N-1)\phi]}{\left(\eta(N-1)N\phi\left(\frac{\tau}{\tau+\phi}\right)^N + \mu Q\tau\right)^2} \tag{B.22}$$

It follows that $\frac{\partial s^*}{\partial \phi} \lessgtr 0$ if and only if $\tau - (N-1)\phi \lessgtr 0$. Let $\phi^* \equiv \frac{\tau}{N-1}$. Consequently, the equilibrium half-spread s^* increases in ϕ for $\phi < \phi^*$ and decreases in ϕ for $\phi \geq \phi^*$.

As ϕ increases to infinity, $\left(\frac{\tau}{\tau+\phi}\right)^{N+1}$ decreases to zero faster than $\tau - (N-1)\phi$. Consequently, the equilibrium half-spread s^* converges to zero. $\square$

Corollary 2

Proof. Follows immediately from the proof of Proposition 3 if one adjust the expected liquidity demand to $Q' \equiv Q\frac{\tau}{\tau+\xi}$. Therefore, a batch equation market improves liquidity relative to a limit order market if

$$\Gamma(Q,N,\phi,\tau) < \frac{\tau}{\tau+\xi}. \tag{B.23}$$

Note that $\Gamma(Q,N,\phi,\tau) > 0$ and does not depend on ξ, whereas $\frac{\tau}{\tau+\xi} > 0$ decreases in ξ, with $\lim_{\xi\to\infty}\frac{\tau}{\tau+\xi} = 0$. Consequently, regardless of the values of Q,N,ϕ, or τ one can find an arbitrarily large impatience coefficient ξ^* such that

$$\Gamma(Q,N,\phi,\tau) > \frac{\tau}{\tau+\xi^*}. \tag{B.24}$$

$\square$

Figure 1: **Stale quote sniping on the batch auction market**

This figure illustrates the clearing mechanism on a batch auction market following news. We focus on the ask side of the market. Assume five HFTs: three informed (HFIs), two uninformed (HFUs). Initially, each HFT has a sell quote in the order book at $v + s_l$. The three HFIs cancel their quotes (yellow crossed quotes) and post buy orders for four units at random prices between $v + s_l$ and $v + \sigma$. The best buy price, in this case of HFI_3 trades against the two remaining stale quotes in the book.
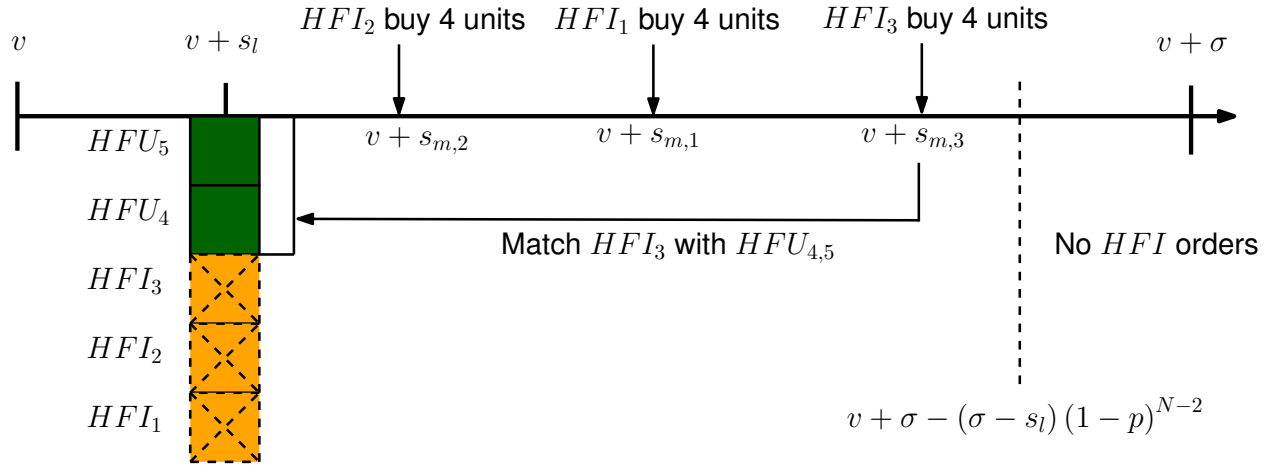
Figure 2: **Price distribution on HFI sniping orders**

This figure illustrates the distribution of s_m, i.e., the deviation of prices on HFI sniping orders from the stale quote midpoint. In Panel (a), the distribution converges to the new efficient price as the expected batch length increases. In Panel (b), the distribution converges to the new efficient price as there are more HFTs on the market.


(a) Expected batch interval length

(b) Number of HFTs

32
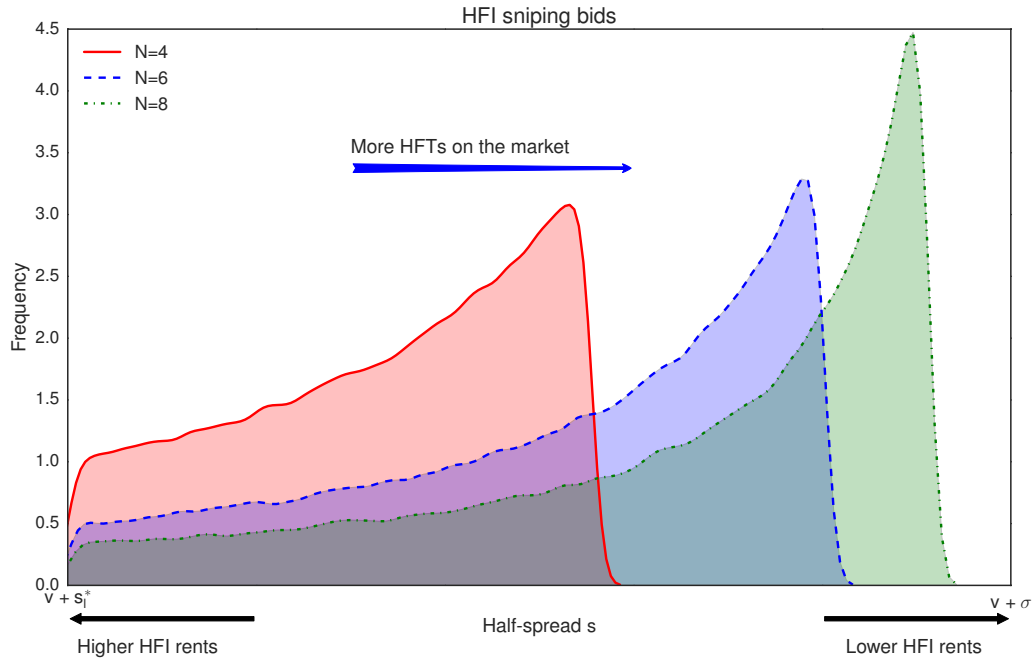
Figure 3: **Uniform-price, uniform-rationing batch auctions mechanism**

This figure illustrates the formation of the equilibrium price and quantity on the batch auction market with uniform price and uniform rationing across all marginal and infra-marginal bids. In particular, the figure illustrates HFT_1 has no incentive to deviate from the equilibrium spread s^* either by undercutting – Panel (a) – or by posting a wider spread – Panel (b).
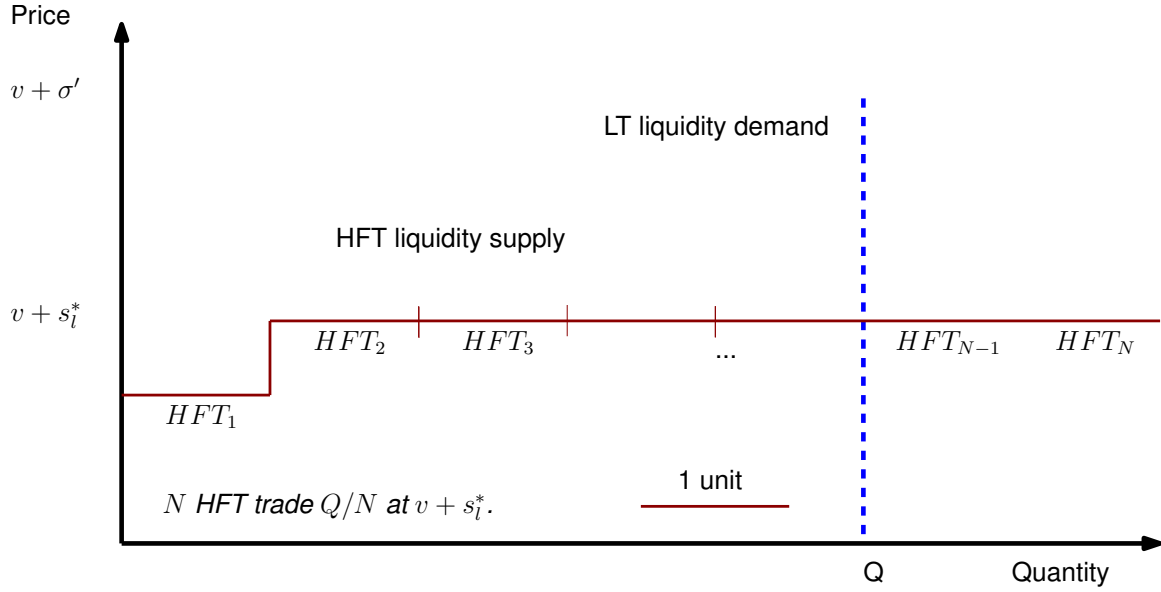

(a) One HFT undercuts the equilibrium price
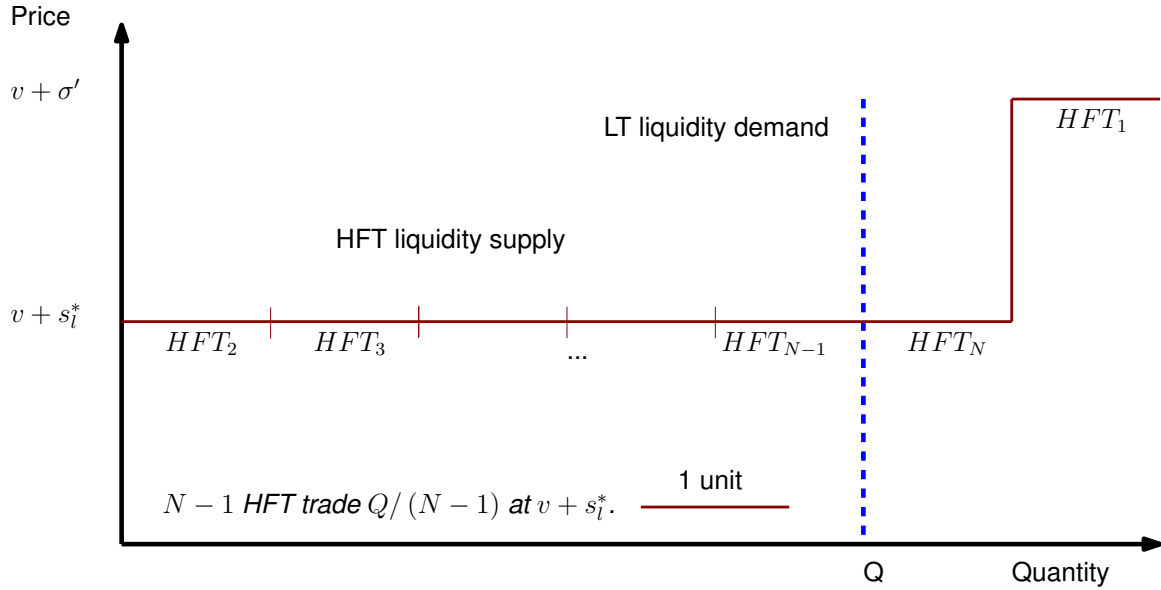
(b) One HFT widens the spread relative to the equilibrium price

Figure 4: **Equilibrium half-spread on batch auction markets**

This figure plots the equilibrium half-spread on batch auction markets as a function of the expected length of the batch interval (first panel) and the number of HFTs (second panel). The batch auction market half-spread is benchmarked against the half-spread in limit order markets, i.e., the red dashed line.
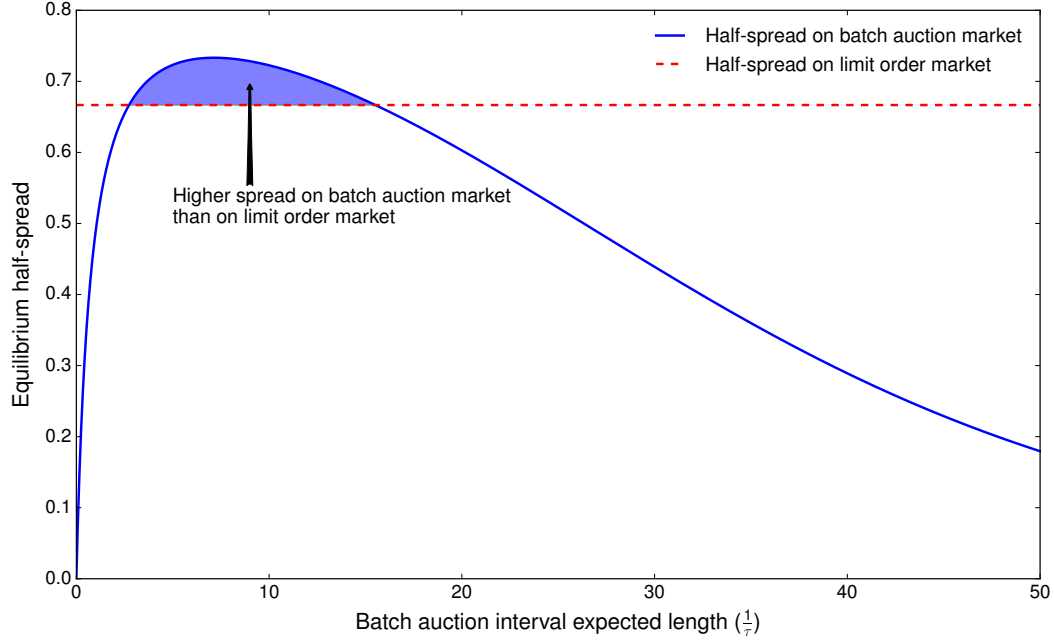
(a) Expected batch interval length
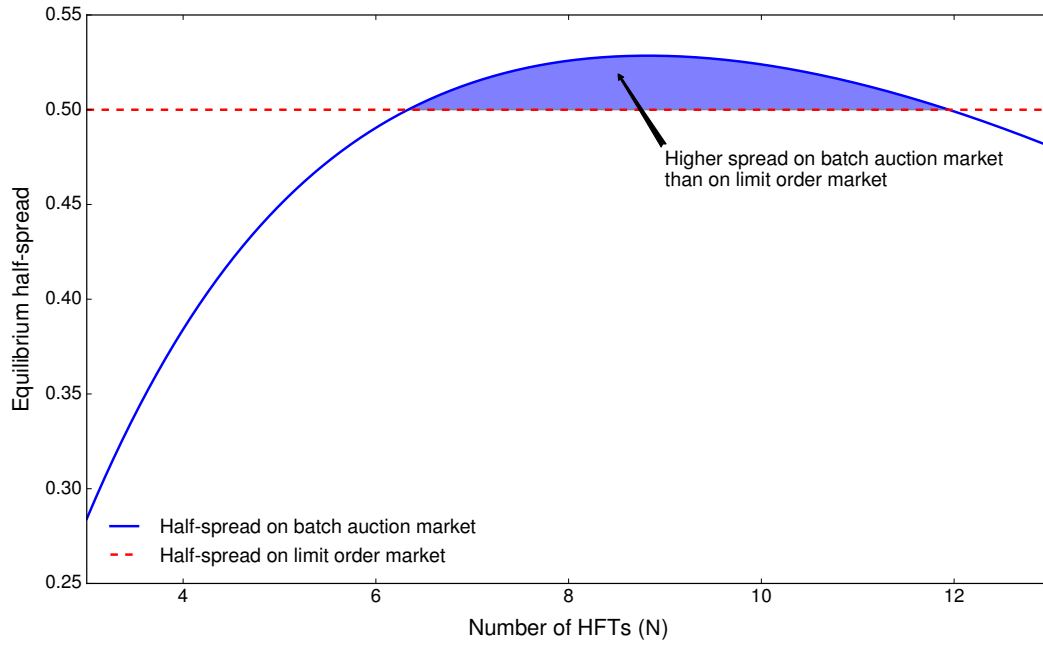
(b) Number of HFTs

Figure 5: **Relative equilibrium spread on the batch auction market.**
This figure plots the ratio between the batch auction and the limit order book equilibrium half-spreads, i.e., s_l^*/s_{LOB}^*, as a function of the expected batch length and the number of HFTs. A ratio higher of one indicates worse liquidity on the batch auction market.
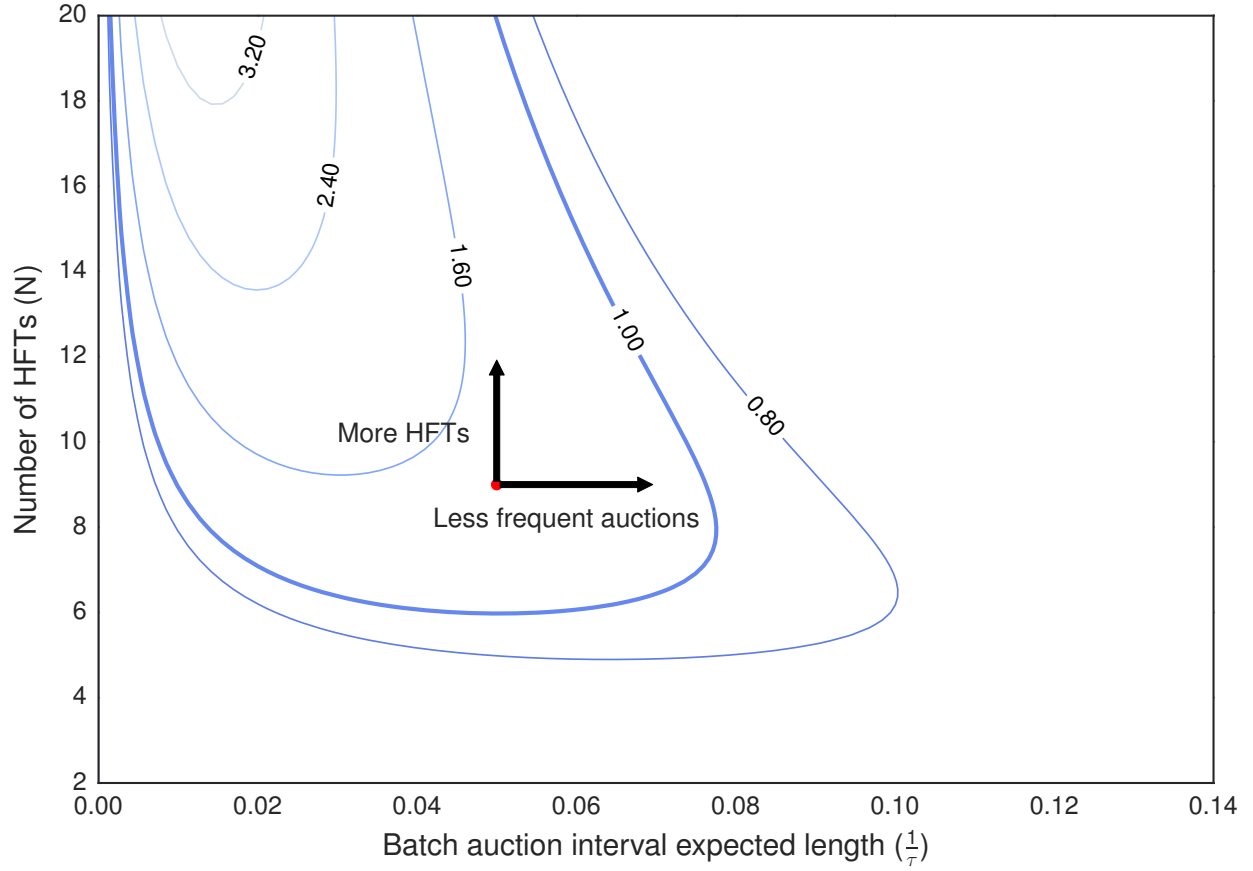
Figure 6: **HFT "arms race" on batch auction markets**

This figure plots, for two values of the monitoring intensity ϕ, the ratio between the batch auction and the limit order book equilibrium half-spreads, i.e., s_l^*/s_{LOB}^*. The red point emphasises a combination of N and τ for which the HFT arms' race (higher ϕ) improves liquidity on the batch auction market.
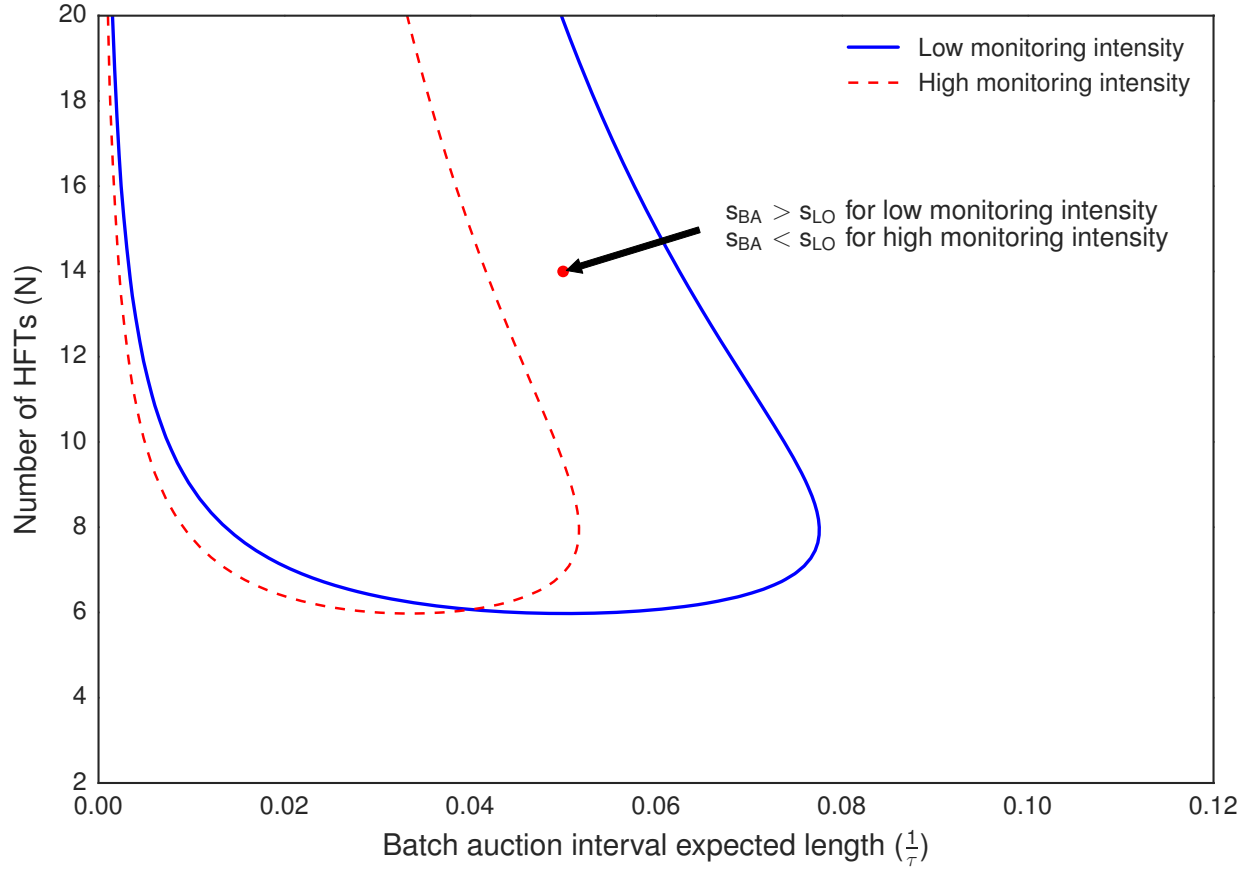
Figure 7: **Equilibrium spread with impatient liquidity traders**
This figure plots the equilibrium half-spread on batch auction markets as a function of the expected length of the batch interval. The continuous blue line corresponds to the case of infinitely patient LTs ($\xi = 0$). The dashed green line corresponds to the case of impatient LTs ($\xi > 0$). The batch auction market half-spreads are benchmarked against the half-spread in limit order markets, i.e., the red dashed line.