Beyond the Frequency Wall:
Speed and Liquidity on Batch Auction Markets

Marlene D. Haas and Marius A. Zoican*

August 23, 2016

Abstract

Frequent batch auction (FBA) markets do not necessarily improve liquidity relative to continuous-time trading. HFTs submit quotes that become stale if the market clears before new information becomes public, and may be sniped by other privately-informed HFTs. Less frequent auctions enhance learning by allowing for more public news: liquidity improves. For infrequent enough auctions, a frequency "wall" emerges as lower adverse selection is offset by lower liquidity demand by impatient traders. Beyond this "wall," longer auction intervals do not improve liquidity anymore. However, the HFT "arms' race" on speed stimulates price competition between arbitrageurs, allowing for a lower spread.

*Marlene Haas is affiliated with University of Lugano, Institute of Finance. Marius Zoican (corresponding author) is affiliated with Université Paris-Dauphine, PSL Research University, CNRS, UMR 7088, DRM Finance, 75016 Paris, France. Marlene Haas can be contacted at marlene.denise.haas@usi.ch. Marius Zoican can be contacted at marius-andrei.zoican@dauphine.fr. The paper received the 2016 Josseph de la Vega Prize from the Federation of European Securities Exchanges. We have greatly benefited from discussion on this research with Matthieu Bouvard, Eric Budish, Jérôme Dugast, Carole Gresse, Björn Hagströmer, Mario Milone, Sophie Moinas, Klaus Ritzberger, and Bart Yueshen Zhou. We are grateful to seminar participants at the Université Paris-Dauphine for insightful comments, as well as conference participants at the 2016 New York Stern Microstructure meeting, the 2016 IFABS meeting, and the 2016 Northern Finance Association conference.

Beyond the Frequency Wall:

Speed and Liquidity on Batch Auction Markets

Abstract

Frequent batch auction (FBA) markets do not necessarily improve liquidity relative to continuous-time trading. HFTs submit quotes that become stale if the market clears before new information becomes public, and may be sniped by other privately-informed HFTs. Less frequent auctions enhance learning by allowing for more public news: liquidity improves. For infrequent enough auctions, a frequency "wall" emerges as lower adverse selection is offset by lower liquidity demand by impatient traders. Beyond this "wall," longer auction intervals do not improve liquidity anymore. However, the HFT "arms' race" on speed stimulates price competition between arbitrageurs, allowing for a lower spread.

Keywords: Market design, high-frequency trading, batch auction markets, liquidity, adverse selection

JEL Codes: D43, D47, G10, G14

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 continuous-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 a 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 around 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

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

2016, London Stock Exchange launched 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 follow two canonical strategies (see, e.g., Hagstromer and Norden, 2013; SEC, 2010; Menkveld and Zoican, 2015): to provide liquidity or to speculate on short-lived arbitrage opportunities. 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 five 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 the 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 of the probability that Bob also bids for 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 influence 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

[2]Sources for this paragraph: SEC Chair Speech on June 5, 2014; the CHX SNAP Auction Market webpage; and the London Stock Exchange webpage.

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 might 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 our model, a competitive high-frequency market-maker (HFM) supplies quotes to liquidity traders (as in, e.g., Kyle, 1985). Therefore, the HFM stands to earn the bid-ask spread from liquidity traders. Other HFTs are "pure" arbitrageurs who only react to news and try to snipe stale quotes, i.e., high-frequency snipers. Consequently, the HFM is exposed to adverse selection risk in case she does not process new information before the auction takes place. A positive equilibrium spread arises due to adverse selection costs.

There are two economic frictions in the model. First, imperfect learning generates temporarily asymmetric information. In the model, news become public with a deterministic delay. High-frequency traders monitor the market and may privately learn news, at random, before it becomes public. In a batch auction that is short relative to the public news delay, the share of public-to-private learning is lower. Consequently, the information asymmetry is larger. Second, liquidity traders are impatient in the sense that they face private value shocks which subject them to urgent trading needs. If batch auctions are sufficiently long, some liquidity traders drop out of the market and liquidity demand is lower.

Given these frictions, a threshold result emerges, i.e., an auction length "wall" beyond which the equilibrium quoted spread reaches a low plateau. At first, consistent with Budish, Cramton, and Shim (2015), a longer auction interval improves liquidity. Less frequent auctions lead to an increase in the share of public-to-private news and consequently reduce the HFM's adverse selection cost. However, for long enough batch auctions, the impatience friction becomes binding. The HFM is less likely to trade with impatient liquidity traders, and she raises the spread to match the adverse selection cost. We find there exists a auction frequency threshold such that for any longer auction interval the two effects, i.e., the lower adverse selection and the lower liquidity demand, exactly offset each other. Beyond this "wall," a longer batch auction interval has no effect on the spread. Importantly, even beyond the auction frequency threshold, batch auction markets do not necessarily improve liquidity relative to limit order markets.

Beyond the auction frequency wall, stimulating (price) competition between HFT snipers improves liquidity. The price competition effect depends non-linearly on both the learning speed and the number of HFTs. If HFT speed increases, it is very likely that two or more HFTs privately observe the news before the market clears. Consequently, sniper competition on arbitrage opportunities converges to the Bertrand case, depressing arbitrage profits and the opportunity cost for providing liquidity. Conversely, if HFT speed is very low, then no HFT privately observes news before the market clears – snipers are very inefficient and the adverse selection cost for the market maker is low. A similar reasoning is valid if we vary the number of HFTs, rather than their speed. It follows that total arbitrageur rents are lowest for markets with either very few or very many HFTs.

Stronger arbitrageur competition also improves liquidity indirectly. Liquidity traders face two trading costs: the spread, which compensates for the adverse selection cost of the market-maker, and a non-pecuniary waiting cost. If arbitrageur price competition is stronger, then the adverse selection cost – and consequently the spread – drops. Therefore, the liquidity traders become relatively "more patient," in the sense they will accept longer waiting times in exchange for a lower spread. The auction frequency wall shifts until the impatience constraint is again binding. Therefore, stronger arbitrageur competition restores some of the auction length's efficacity as a tool to improve liquidity.

This paper offers an important policy-relevant insight: the HFT "arms' race" on a batch auction market intensifies competition between arbitrageurs and decreases the spread. In the context of batch auction markets, speed competition stimulates price competition. Stimulating HFT speed competition is effective in scenarios where longer auction intervals conflict with short-term trading needs. Discrete-time trading can thus align private and social incentives, since the arms' race is arguably privately optimal for HFTs (Budish, Cramton, and Shim, 2015).

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 to react to new information without a 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's 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, as well as impatient LTs. The richer model we propose features a positive bid-ask spread and unveils new economic channels, in particular the role of speed competition and the limits on batch auction length as a tool to improve liquidity (i.e., the existence of an auction frequency wall). 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 to be between 0.2 and 0.9 seconds. However, the authors focus on liquidity risk rather than on 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 batch auctions 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 traders 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 of double-auction markets where "fast" traders act as intermediaries rather than arbitrageurs, Du and Zhu (2015) find that, if news are unscheduled, the socially optimal trading frequency exceeds information arrival frequency. However, trading speed does not offer an informational advantage. In our model, stochastic learning leads to a novel channel: infrequent trading improves the public-to-private news ratio, enhances learning, and reduces the adverse selection cost.

This 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. Kandel, Rindi, and Bosetti (2012), Pagano and Schwartz (2003), and Pagano, Peng, and Schwartz (2013) document that the introduction of opening and closing call auctions on Paris Bourse in 1996 and Nasdaq in 2004, respectively, reduced spreads and volatility and improved price discovery.

Finally, this 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 to have a larger price impact.

There is also a growing literature on trader impatience. Here, our paper is most closely related to the theoretical models of impatience in Foucault, Kadan, and Kandel (2005) and Roşu (2009). Ellul, Holden, Jain, and Jennings (2007) document extreme impatience for orders submitted to the automatic execution system of NYSE. Demarchi and Thomas (2001) argue that quantitative managers such as index trackers are more likely to be impatient than long-horizon value traders. Several empirical papers, e.g., Ainsworth and Lee (2014) and Hagströmer and Nordén (2014), use order aggressiveness to measure trader impatience. Trader impatience becomes especially important under unusual circumstances, such as when traders or entire stock markets receive liquidity shocks. It is hence important to include this dimension in the framework of batch auctions in order to understand possible consequences of impatience on liquidity in such markets.

The rest of the paper is structured as follows. Section 2 formalizes the model of the batch auction market. Section 3 analyzes adverse selection on batch auction markets and its relationship with HFT competition. Section 4 discusses the trading game equilibrium and the determinants of liquidity on batch auction markets. Section 5 benchmarks the batch auction liquidity against the limit order market and discusses the different effects of the HFT "arms race" on the two market structures. Section 6 concludes.

2 A model of the Frequent Batch Auction (FBA) market

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 within the auction interval. Traders post orders from $t = 0$ until $t = \tau$, where τ is the batch auction length. 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 clearing price. Both limit and market orders can be cancelled at no cost and during the respective batch auction interval before the market clears.[3]

Market clearing. The market clearing mechanism is a *uniform-price* auction with price priority, as defined in, e.g., Budish, Cramton, and Shim (2014) or Kremer and Nyborg (2004). First, there is an unique price for all units traded. Second, quotes at infra-marginal prices are executed first. The assumption of uniform-price clearing follows the standard industry practice (see, e.g., Bats, 2015; Cinnober, 2010).

Agents. The risky asset is traded by two types of agents: N high-frequency traders (HFTs) and a large number of liquidity traders (LTs). Liquidity traders submit only market orders.

One of the N high-frequency traders, selected at random, is a competitive market-maker (HFM). She can submit at most one limit order to buy and one limit order to sell the asset. The other $N - 1$ HFTs trade only speculatively on arbitrage opportunities, i.e., they are high-frequency snipers (HFS).[4] The type of each high-frequency trader (i.e., HFM or HFS) is fixed for the duration of the auction.

[3]This is similar to the set-up of Roşu (2009). At $t = \tau$, the market clears using a uniform-price auction.

[4]This assumption is equivalent to a higher risk-aversion coefficient for the $N - 1$ HFSs than for HFM.

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 size of news is σ. The common value at time $t = 0$ is v: Conditional on news, it either jumps to $v + \sigma$ for "good" news or $v - \sigma$ 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 $+\sigma_P$ or $-\sigma_P$, with equal probabilities, where $\sigma_P > \sigma$.

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	$\left(1 - e^{-(\eta+\mu)\tau}\right) \frac{\eta}{\eta+\mu}$
LT arrival	$\left(1 - e^{-(\eta+\mu)\tau}\right) \frac{\mu}{\eta+\mu}$
No event	$e^{-(\eta+\mu)\tau}$

Trader impatience. Liquidity traders are impatient as in Foucault, Kadan, and Kandel (2005) and Roşu (2009). They incur a linear waiting cost between arrival at time t_l and order execution at time τ, that is $c\,(\tau - t_\ell)$, for each unit they desire to trade, where $c > 0$ is a measure of the waiting cost. Therefore, LTs experience utility losses which are proportional to their expected waiting time.

The waiting cost can be interpreted as a funding liquidity constraint. For example, if traders hedge outstanding risky positions with a delay, they need to post higher margins to compensate for additional risk. In this instance, impatience reflects the shadow cost of collateral. Roşu (2009) also mentions time-discounting and ambiguity aversion as potential sources of impatience.

Information structure and learning. A public signal on the common value of the asset arrives with delay Δ. Therefore, all news arriving before $\tau - \Delta$ are common knowledge by the auction time τ. Conversely, news that arrive after $\tau - \Delta$ do not become public by the auction time. However, HFTs may independently learn from a private signal that arrives as a Poisson process with intensity ϕ. The private signal contains information both on the new common value and on the timing of an innovation. Conditional on news at $t_n \in [\tau - \Delta, \tau)$, the probability an HFT becomes informed is

$$p(t_n) = 1 - \exp[-\phi(\tau - t_n)] \tag{1}$$

where $\phi > 0$. We interpret ϕ as HFT speed, i.e., the intensity with which an HFT is able to process a signal and make a trading decision. Two possible learning scenarios are illustrated below.

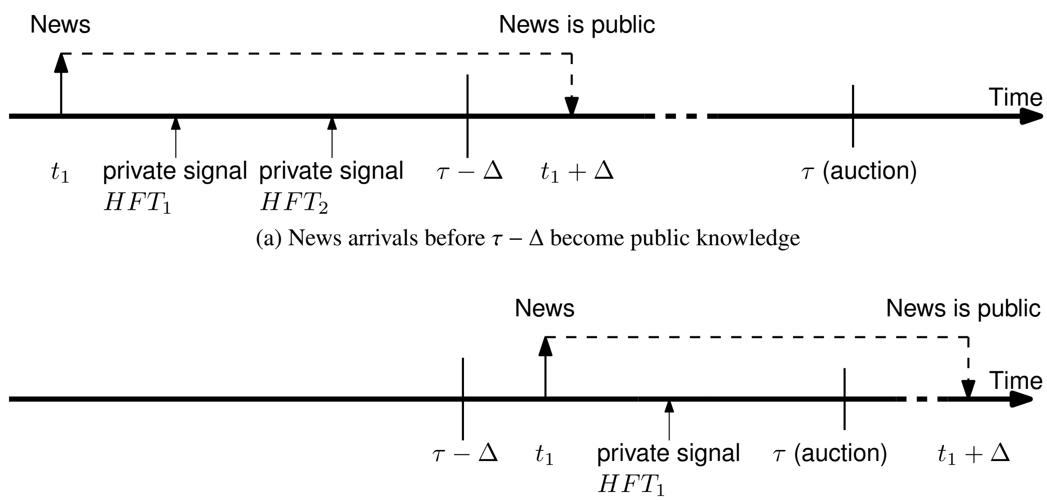

(a) News arrivals before $\tau - \Delta$ become public knowledge

(b) News arrivals after $\tau - \Delta$ revealed through private signals

Public and private signals

In the first panel, the HFTs receive a private signal before $\tau - \Delta$. In this case, the value of the signal is zero since the news becomes public knowledge before the auction. In the second panel, however, the news arrives after $\tau - \Delta$. Therefore, HFTs may only learn the new common value through private signals. A faster HFT (either a sniper or the market-maker) is more likely to learn and react to news before the auction takes place.

If an HFT learns the news before the market clears, we refer to her as "informed," or HFI. Otherwise, we denote an "uninformed" HFT by HFU.[5] LTs are uninformed and only motivated to trade by private value shocks.

[5]In equilibrium, as we show in Proposition 1, *all* traders learn the common value after the market clears from the unique, public clearing price.

All model parameters are summarized in Appendix A.

3 Adverse selection on FBA markets

Due to the learning friction, HFM quotes are exposed to adverse selection risk if news arrives between $(\tau - \Delta, \tau)$. In this section we focus on the market-maker's adverse selection costs and snipers' speculative rents on FBA markets, i.e., we restrict our attention to news arrivals between $(\tau - \Delta, \tau)$. We solve for general equilibrium in Section 4.

3.1 Informed HFS sniping profits

Let s_l denote the outstanding half-spread on the limit orders around the common value v. If there is news, informed HFS can "snipe" outstanding quotes of an uninformed HFM. Let $p(t_\mathrm{n})$ denote the probability of an HFT becoming informed by the auction time, conditional on the news arrival time t_n. For simplicity of notation, let $p \equiv p(t_\mathrm{n})$ throughout this section.

To build intuition, consider first the case where the number of informed competitors is common knowledge to each sniper. 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 sniper learns about the good news, then she can post a marketable buy order at $v + s_l$ and earn $(\sigma - s_l)$. In this case, the sniper has a monopoly over the information and can maximize his rents. If two or more snipers learn about the good news, Bertrand competition emerges.[6] Consequently, they post a buy order at $v + \sigma$ and earn zero profit.

If the number of HFIs is unknown to each sniper, snipers cannot infer 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 hence submit orders at $v + s_m$, with $s_m \in [s_l, \sigma]$ drawn

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

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

$$\pi_{snipe}(s_m, F(s_m)|s_l) = \underbrace{(1-p)}_{\text{HFM uninformed}} \quad p \underbrace{\sum_{k=0}^{N-2} \begin{pmatrix} N-2 \\ k \end{pmatrix} (1-p)^{N-2-k} p^k}_{\text{Probability of } k \text{ HFIs}} \underbrace{F(s_m)^k}_{\text{Probability winning bid}} (\sigma - s_m)$$
$$= p(1-p)(\sigma - s_m)\left[1 - p(1 - F(s_m))\right]^{N-2}. \tag{2}$$

Figure 1 illustrates how HFS sniping orders are matched with HFM stale quotes for the risky asset.

[insert Figure 1 here]

To illustrate, consider the ask side of the market. If good news arrives and the outstanding ask quote is $v + s_l$, then all informed HFS send buy orders at random prices between $v + s_l$ and $v + \sigma$. The highest sniping buy price is matched with the stale HFM quote, if the HFM remains uninformed. The sniping orders are not exposed to adverse selection risk and, conditional on execution, yield a guaranteed profit. Competition between snipers ensures that the market-maker stale quote is matched with the price closest to the efficient one, therefore limiting the adverse selection cost.

Formally, with probability $p(1-p)$, an HFS is informed and the HFM is not. Out of the remaining $N-2$ HFTs, exactly k are informed with probability $\begin{pmatrix} N-2 \\ k \end{pmatrix}(1-p)^{N-2-k} p^k$ 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 the stale quote and earns $\sigma - s_m$.

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

Lemma 1. (Conditional distribution of sniping orders) *If they observe a common value innovation at time t_n, all HFSs submit marketable limit orders for one unit 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:*

$$F(s_m) = \begin{cases} 0 & , \text{if } s_m < s_l \\ \frac{1-p(t_n)}{p(t_n)}\left(\sqrt[N-2]{\frac{\sigma - s_l}{\sigma - s_m}} - 1\right) & , \text{if } s_m \in \left[s_l, \sigma - (\sigma - s_l)(1 - p(t_n))^{N-2}\right) \\ 1 & , \text{if } s_m \geq \sigma - (\sigma - s_l)(1 - p(t_n))^{N-2}. \end{cases} \tag{3}$$

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

$$\pi_{\text{snipe}}(s_l) = p(1-p)^{N-1}(\sigma - s_l). \tag{4}$$

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

[insert Figure 2 here]

The conditional expected HFI per-unit sniping profit decreases in both the HFT speed ϕ and the number of HFTs N. First, as HFT speed increases, each is independently more likely to learn the private signal. 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.

3.2 HFM adverse selection costs

In this section, we compute the expected adverse selection cost on limit orders for an uninformed HFM. Let s_l denote the outstanding half-spread.

The conditional loss is a function of the maximum bid submitted by snipers, which depends in turn on the number k of informed snipers and on the news arrival time t_n. The expected loss is $\ell(s_l)$, where

$$\ell(s_l) = \int_{\tau-\Delta}^{\tau} \underbrace{[1-p(t_n)]}_{\text{Uninformed HFM}} \sum_{k=0}^{N-1} \underbrace{\binom{N-1}{k}[1-p(t_n)]^{N-1-k}\, p(t_n)^k}_{\text{Probability of } k \text{ informed snipers}} \underbrace{\mathbb{E}_s\left(\sigma - \max_{i=1}^{k} s_{m,i} | k\right)}_{\text{Expected loss conditional on k HFSs}} \frac{1}{\tau}\, dt_n. \tag{5}$$

All news arriving between $t = 0$ and $t = \tau - \Delta$ become public by the time of the auction and consequently do not generate adverse selection costs. On the other hand, news arrivals between $\tau - \Delta$ and τ generate positive adverse selection costs.

If the HFT learns the new common value before market clearing, she updates the stale quote and faces no adverse selection risk. With probability $1 - p(t_{\mathrm{n}})$, the HFM is uninformed and faces adverse selection risk from k informed snipers; the probability of exactly k informed HFSs is again $\begin{pmatrix} N-1 \\ k \end{pmatrix} [1 - p(t_{\mathrm{n}})]^{N-1-k} p(t_{\mathrm{n}})^k$. The market-makers' 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 HFS' 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 $F^k(s_m)$. From this, it follows that

$$\mathbb{E}_s \left(\sigma - \max_{i=1}^{k} s_{m,i} | k \right) = \int_{s_l}^{\sigma - (\sigma - s_l)(1-p)^{N-2}} (\sigma - s_m) k F^{k-1}(s_m) \frac{\partial F(s_m)}{\partial s_m} ds_m \tag{6}$$

From equations (5) and (6) it follows that the expected HFM loss conditional on the news arrival time is:

$$\begin{aligned} \ell(s_l | t_{\mathrm{n}}) &= (1-p) \int_{s_l}^{\sigma - (\sigma - s_l)(1-p)^{N-2}} \sum_{k=0}^{N-1} \begin{pmatrix} N-1 \\ k \end{pmatrix} (1-p)^{N-1-k} p^k (\sigma - s_m) k F^{k-1}(s_m) \frac{\partial F(s_m)}{\partial s_m} ds_m \\ &= (1-p) \int_{s_l}^{\sigma - (\sigma - s_l)(1-p)^{N-2}} (N-1) p [1 - p(1 - F(s_m))]^{N-2} \frac{\partial F(s_m)}{\partial s_m} ds_m \\ &= (N-1) \int_{s_l}^{\sigma - (\sigma - s_l)(1-p)^{N-2}} \pi_{\mathrm{snipe}}(s_m | s_l, t_{\mathrm{n}}) \frac{\partial F(s_m)}{\partial s_m} ds_m \end{aligned} \tag{7}$$

The expected adverse selection cost on limit orders for HFM is equal to the total expected sniping profit of the $(N-1)$ snipers over all potential sniping order half-spreads s_m. Since snipers are in equilibrium indifferent between all half-spreads s_m in the mixed strategy support, the right hand side of equation (7) is simply $\pi_{\mathrm{snipe}}(s_l)$. Therefore, conditional on the news arrival time, the HFS sniping profits are equal to the HFM sniping losses,

$$\ell(s_l | t_{\mathrm{n}}) = (N-1) \pi_{\mathrm{snipe}}(s_l | t_{\mathrm{n}}). \tag{8}$$

From equations (4), (5) and (8), it follows that

$$\ell(s_l) = (N-1)(\sigma - s_l)\int_{\tau-\Delta}^{\tau} p(t_n)[1-p(t_n)]^{N-1}\frac{1}{\tau}\,\mathrm{d}t_n, \tag{9}$$

since news arrival times are uniformly distributed over $(\tau - \Delta, \tau]$. From the definition of $p(t_n)$ in (1),

$$\ell(s_l) = \frac{1}{\tau}(\sigma - s_l)\frac{1 + (N-1)\exp(-N\Delta\phi) - N\exp[-(N-1)\Delta\phi]}{N\phi} = \frac{1}{\tau}L \times (\sigma - s_l), \tag{10}$$

where

$$L = \frac{1 + (N-1)\exp(-N\Delta\phi) - N\exp[-(N-1)\Delta\phi]}{N\phi}, \tag{11}$$

the HFM *snipe loss* function.

Lemma 2. (Adverse selection loss) *The* snipe loss *function L, as defined in equation* (11)*,*

(i) increases in Δ, the latency of the public signal;

(ii) there is an unique $\bar{\phi} > 0$ such that L increases in ϕ for $\phi \leq \bar{\phi}$ and decreases in ϕ for $\phi > \bar{\phi}$;

(iii) there is an unique $\bar{N} > 1$ such that L increases in N for $N \leq \bar{N}$ and decreases in N for $\phi > \bar{N}$;

The latency of the public signal is positively related to the HFM snipe loss. A higher delay for public news decreases the public-to-private news ratio, i.e., leads to a higher share of common value innovations exposed to information assymetry.

If HFT speed approaches zero, the equilibrium spread becomes arbitrarily small: No HFT learns the common value innovation, so there is no sniping. As the HFT speed increases, each HFT is independently more likely to receive a private signal; consequently, the snipe loss increases. However, as speed increases even further, 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 snipe loss decreases again. The competition effect dominates for $\phi > \bar{\phi}$, with $L = 0$ for $\phi \to \infty$.

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 probability that at least one sniper is informed is also small. As

N increases, stronger competition between HFIs drives the expected-profit-per-stale-quote down. At the same time, the probability that at least one sniper is informed also increases, generating an opposite channel. For $N \geq \bar{N}$, the competition effect dominates and the snipe loss L decreases with N.

The concave effect of HFT competition and speed on expected adverse selection costs follows from the opposite effects on the *extensive* and the *intensive* margins. Increasing the number or speed of HFTs leads to a larger probability of a sniping trade (the extensive margin), but to lower costs conditional on sniping due to enhanced competition (the intensive margin).

4 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 at $t = \tau$, an equilibrium consists of HFT and LT orders to trade a specific price and quantity of the risky asset.

4.1 LT trading decision

From Section 2, with probability $e^{-(\mu+\eta)\tau}\frac{\mu}{\mu+\eta}$, during the auction interval a single liquidity trader receives a private value shock of $\pm\sigma_P$.

Liquidity traders only trade if the private value minus the spread exceeds the waiting cost (see, e.g., Admati and Pfleiderer, 1988). This reflects the LT's impatience.

Let u_{LT} be the LT's utility from trading one unit of asset. For tractability, we assume the private value shock σ_P to be one order of magnitude larger than the common value shock σ, and therefore one order of magnitude larger than the spread. It follows that LTs submit orders only if $u_{\text{LT}} > 0$, where

$$u_{\text{LT}} = \sigma_P - s - c \max\{\tau - t_\ell, 0\}. \tag{12}$$

It follows that an LT submits an order only if she receives a private value shock close enough to the auction, that is if

$$t_\ell > \tau - \frac{\sigma_P - s}{c}. \tag{13}$$

From the properties of a Poisson process, the endowment shock timing t_ℓ has a conditional uniform distribution between $t = 0$ and the auction time $t = \tau$. Therefore, the ex ante probability that an LT submits an order, conditional on receiving a private value shock, is

$$\text{Prob}\left(t_\ell > \tau - \frac{\sigma_P - s}{c}\right) = \begin{cases} 1, \text{ if } \tau \leq \frac{\sigma_P - s}{c} \text{ and} \\ \frac{\sigma_P - s}{c\tau}, \text{ if } \tau > \frac{\sigma_P - s}{c}. \end{cases} \tag{14}$$

Lemma 3. *In continuos-time markets ($\tau = 0$), an LT always submits orders to trade on its private value. In discrete-time markets ($\tau > 0$), the ex-ante probability of an LT order weakly decreases in the length of the auction interval, τ.*

Lemma 3 describes the difference in LT behavior across continuous-time markets and discrete-time markets. As an LT faces time pressure due to the private value shocks and is hence to a certain degree impatient, he will only trade in discrete-time markets if his private value shock is close enough to the auction such that his waiting costs are smaller than the private value minus the spread. If the waiting time is too large, he will refrain from trading. This is different in continuous-time markets in which an LT can trade at every instant of time and hence does not occur any waiting costs.

4.2 HFM liquidity provision

Next, we turn to the liquidity provision of the HFM. The HFM expected profit from providing liquidity, i.e., submitting a buy and a sell order at $v \pm s_l$ is π_{liqudity}, where

$$\pi_{\text{liqudity}} = \underbrace{\left[1 - \exp\left(-\left(\eta + \mu\right)\tau\right)\right]\frac{\mu}{\eta + \mu}}_{\text{Probability of LT arrival}} \underbrace{\text{Prob}\left(t_\ell > \tau - \frac{\sigma_P - s}{c}\right)}_{\substack{\text{Probability of LT trade} \\ \text{conditional on arrival}}} s_l \tag{15}$$

With probability $\left[1 - \exp\left(-\left(\eta + \mu\right)\tau\right)\right]\frac{\mu}{\eta+\mu}$, a liquidity trader receives a liquidity shock and wants to trade. Conditional on arrival, he submits a market order to buy or to sell the asset with probability $\text{Prob}\left(t_\ell > \tau - \frac{\sigma_P - s}{c}\right)$, and the HFM earns s_l.

The HFM is competitive (as in, e.g., Kyle, 1985), and therefore its expected net profit (i.e., is the profit from providing liquidity minus the adverse selection cost) is zero:

$$\pi_{\text{HFM}} = \pi_{\text{liquidity}} - \underbrace{\left[1 - \exp\left(-\left(\eta + \mu\right)\tau\right)\right] \frac{\eta}{\eta + \mu}}_{\text{Probability of news arrival}} \underbrace{\frac{1}{\tau} L \times (\sigma - s_l)}_{\text{Conditional HFM loss}} = 0. \tag{16}$$

The competitive spread s^* is the solution to equation (16). From equations (14) and (15), it follows that s^* solves

$$s^* = \sigma \frac{\eta L}{\eta L + \mu\tau \min\left\{\frac{\sigma_P - s^*}{c\tau}, 1\right\}}. \tag{17}$$

Lemma 4. (Competitive spread) *Let*

$$\tau^* \equiv \frac{\sqrt{(c\eta L - \mu\sigma_P)^2 + 4c\eta L\mu(\sigma_P - \sigma)} - (c\eta L - \mu\sigma_P)}{2c\mu} > 0. \tag{18}$$

The competitive spread $s^*(\tau)$ is given by the continuos function

$$s^*(\tau) = \begin{cases} \sigma\frac{\eta L}{\eta L + \mu\tau}, \text{ if } \tau \leq \tau^*, \\ \frac{c\eta L + \mu\sigma_P - \sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}}{2\mu}, \text{ if } \tau > \tau^*. \end{cases} \tag{19}$$

For $\tau < \tau^*$, the impatience constraint is not binding even for liquidity traders who receive a shock at $t = 0$. Therefore, it follows that, if $\tau < \tau^*$,

$$\min\left\{\frac{\sigma_P - s^*}{c\tau}, 1\right\} = 1. \tag{20}$$

The competitive spread s^* – which we show in Proposition 1 to be the equilibrium spread – decreases in τ as the ratio of public-to-private news increases.

For $\tau \geq \tau^*$, the impatience constraint starts to bind, i.e., some liquidity traders choose not to trade on the batch auction market. The adverse selection cost and the profit from liquidity trades change in the auction length at the same rate, but in opposite directions. Therefore, the competitive spread is no longer a function of the auction length if $\tau \geq \tau^*$, i.e., the *frequency wall*. Figure 3 illustrates the result.

[insert Figure 3 here]

4.3 Equilibrium strategies

Proposition 1 describes equilibrium strategies in the trading game.

Proposition 1. (Equilibrium) *The following HFT strategies form a Nash equilibrium in the trading game. First, if there is any public signal update in* $(0, \tau)$*, HFTs do not submit any orders. Second, if there is no public signal update in* $(0, \tau)$*, then:*

(i) If the HFM does not observe any private signal, she submits at $t = \tau$ *a buy limit order for one unit at* $v_0 - s^*$ *and a sell limit order for one unit at* $v_0 + s^*$*, where* s^* *is defined in equation* (19)*. If the HFM receives any private signal, she does not submit any order.*

(ii) If an HFS receives a good (bad) private signal, then she submits a marketable buy (sell) order for one unit at price $v_0 + s_m$*, where* s_m *has distribution* $F(s_m)$*, as defined in equation* (3)*.*

(iii) An LT submits a marketable order to buy (sell) one unit of the asset at the clearing price if and only if he receives a private value shock at $t_\ell > \max\left\{0, \tau - \frac{\sigma_P - s^*}{c}\right\}$.

If there is any public news before the auction time, then all HFTs correctly infer there is no LT arrival. Consequently, the HFM does not post a quote since she never earns a positive payoff. High-frequency snipers only race to the market conditional on learning a private signal, standing to earn a strictly positive profit if the HFM is uninformed. Finally, liquidity traders only submit orders if they receive a private value shock close enough to the auction time.

Lemma 5 describes the behavior of the equilibrium half-spread with respect to news intensity, news arrival size, liquidity shock intensity, and liquidity demand.

Lemma 5. *The equilibrium half-spread s^* increases in the size of value innovations (σ), news intensity (η), and decreases in the liquidity traders' arrival intensity (μ).*

Lemma 5 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, the number of HFTs N, and the HFT speed ϕ.

Proposition 2. *The equilibrium half-spread on the batch auction market, s^*,*

(i) decreases in the auction interval length τ if $\tau < \tau^*$ and does not depend on τ for $\tau > \tau^*$.

(ii) increases in the latency of the public signal, Δ.

(iii) increases (decreases) in the number of HFTs N if $N \leq \bar{N}$ (and $N > \bar{N}$, respectively).

(iv) increases (decreases) in the HFT speed ϕ if $\phi \leq \bar{\phi}$ (and $\phi > \bar{\phi}$, respectively).

The first result of Proposition 2 follows directly from Lemma 4 and is illustrated in Figure 3. A higher auction length (i.e., a lower auction frequency) improves liquidity only below a threshold τ^*, the *auction frequency wall*, beyond which lower liquidity demand generates an opposite effect.

For $\tau > \tau^*$, competition-based instruments can be used to improve liquidity. Proposition 2 shows how the number and speed of HFTs impact the equilibrium spread. From Lemma 2, the number of HFTs N and their speed ϕ have a concave effect on the adverse selection loss of the HFT. This effect is transfered to the equilibrium spread. In particular, high HFT speed and a large number of HFTs stimulate price competition between snipers, reducing information rents. Figure 4 illustrates this result.

[insert Figure 4 here]

Complements or substitutes? Figure 5 displays the contour plot for the batch auction equilibrium spread, as a function of the number of HFTs and their speed. The number of HFTs N and the HFT speed ϕ have an ambiguous impact on the equilibrium spread on the batch auction market.

[insert Figure 5 here]

From the proof to Lemma 2, it follows that the number of HFTs and HFT speed are (locally) substitutes if and only if

$$(\phi - \bar{\phi})\left(N - \bar{N}\right) < 0, \tag{21}$$

and (locally) complements otherwise.

Equation (21) indicates the optimal policy with respect to ϕ following a change in HFT competition N. Suppose an HFT exogenously withdraws from the market. To preserve liquidity, the regulator should encourage higher speed investment if $\phi \geq \phi^*$ and discourage it otherwise.

Proposition 3. *The auction length "wall" τ^*,*

(i) decreases in the snipe loss L, news intensity η, news size σ, LT waiting cost c;

(ii) increases in LT arrival intensity μ and LT private value σ_P.

The auction length "wall" τ^* is the maximum interval between two consecutive auctions such that *all* liquidity traders are willing to trade. Conditional on a trade, LTs gain their private value σ_P and pay both the spread and waiting cost.

First, a lower LT waiting cost c or a higher private value σ_P increase LTs' willingness to wait until the auction time. Second, for a fixed σ_P, a lower spread also translates into a higher capacity for LTs to bear the waiting cost, i.e., LTs endogenously become more patient. If LTs' are relatively more patient, the auction length "wall" shifts to the right.

It follows that, if regulators always choose the optimal batch length τ^* that minimizes the spread, then a change in HFT competition has both a direct and an indirect effect on the spread.

The direct effect follows immediately from Proposition 2. The indirect effect is more subtle. If sniper competition is stronger, then the adverse selection loss for HFM decreases. Consequently, the competitive HFM charges a lower spread to liquidity traders. Since crossing the spread is cheaper, liquidity traders are willing to support a higher waiting cost. Therefore, the regulator can decrease the auction frequency with no adverse effect on liquidity demand. Finally, since at a lower auction frequency the public-to-private news ratio is higher, the adverse selection cost and the spread decrease further. Formally, the total effect of moving from a speed ϕ_L to ϕ_H where $\phi_H > \phi_L > \bar{\phi}$ is

$$\underbrace{s^*\left(\tau^*_{\phi_H}, \phi_H\right) - s^*\left(\tau^*_{\phi_L}, \phi_L\right)}_{\text{Total effect}} = \underbrace{s^*\left(\tau^*_{\phi_L}, \phi_H\right) - s^*\left(\tau^*_{\phi_L}, \phi_L\right)}_{\text{<0, Direct effect}} + \underbrace{s^*\left(\tau^*_{\phi_H}, \phi_H\right) - s^*\left(\tau^*_{\phi_L}, \phi_H\right)}_{\text{<0, Indirect effect}}. \tag{22}$$

Figure 6 illustrates the direct and indirect effects on liquidity of stronger speed competition between HFT snipers.

[insert Figure 6 here]

5 Benchmark: the limit order market

A natural benchmark for batch auction market quality is the limit order market. Since limit order books represent the dominant market structure, it is important to evaluate whether the transition from continuous trading to discrete auctions improves liquidity.

Let $\underline{s}^*$ denote the equilibrium spread on the batch auction market, conditional on the optimal auction length ($\tau \geq \tau^*$), i.e.,

$$\underline{s}^* = \inf_{\tau} s^* = s^*\left(\tau^*\right) \tag{23}$$

In this section, we compare the batch auction equilibrium spread in equation (23) to the outcome of a model where the risky asset is traded on a limit order market. If the equilibrium spread $s^*\left(\tau^*\right)$ is larger than the equilibrium spread on the limit order market s^*_{LOB}, it follows from Proposition 2 that $s^*\left(\tau\right) > s^*_{\text{LOB}}$ for any $\tau < \tau^*$.

The limit order market is modelled as in Budish, Cramton, and Shim (2015), with a competitive market-maker. 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 an ex ante identical probability $\frac{1}{N}$ of being first to the market. Since the waiting time is zero, all LTs trade and realize their private values.

As in Section 4, the HFM is competitive. Let s_{LOB} be the quoted half-spread on the limit order market. It follows that in equilibrium

$$\pi_{\text{HFM}}^{\text{LOB}} = \mu s_{\text{LOB}} - \frac{N-1}{N}\eta\left(\sigma - s_{\text{LOB}}\right) = 0 \tag{24}$$

The left hand side of equation (24) 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. The HFM trades one unit and earn the spread s_{LOB}. Alternatively, with intensity η, there is news. The HFM with a quote in the book incurs the adverse selection cost $\sigma - s_{\text{LOB}}$ whenever she is not first to the market, with probability $\frac{N-1}{N}$. If the HFT is first to the market after news, with probability $\frac{1}{N}$, she cancels his own quote and consumes the remaining $Q' - 1$, earning $\sigma - s_{LOB}$ for each.

The equilibrium spread on a limit order market is given by the unique solution to equation (24), i.e.,

$$s_{\text{LOB}}^{*} = \sigma\frac{(N-1)\eta}{(N-1)\eta + N\mu}. \tag{25}$$

Proposition 4. *For $\tau \geq \tau^*$, a batch auction market improves liquidity relative to a limit order book only if*

$$L \times c \leq \frac{N-1}{N}\frac{\eta(\sigma_P - \sigma) + \mu\sigma_P}{(\eta + \mu}. \tag{26}$$

Proposition 4 establishes a sufficient condition for a batch auction market to improve liquidity. Intuitively, the batch auction market yields a lower equilibrium spread if the product of the two frictions, the snipe loss L and the waiting cost c, is low enough relative to the adverse selection cost on the limit order market.

A batch auction market may worsen liquidity, even if the auction frequency is optimally chosen. Figure 7 illustrate how HFT competition and speed can change the optimal market structure from a limit order book to a batch auction market. The batch auction market is only optimal either for sufficiently low or sufficiently high level of competition.

[insert Figure 7 here]

5.1 The information "arms race"' in batch auction and limit order markets

Our model reveals a fundamental difference between the effects of the speed "arms' race" on frequent batch auction, on the one hand, and limit order markets, on the other hand.

High-frequency have an incentive to invest in relative speed, i.e., be faster and more informed than their competitors, leading to a costly arms' race (Budish, Cramton, and Shim, 2015). In the context of our model, the arms' race translates into an ever higher monitoring intensity ϕ.

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

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.

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.

Corollary 1. *The equilibrium spread on the batch auction market, s^* converges to zero if HFTs receive private signals infinitely often, i.e.,*

$$\lim_{\phi \to \infty} s^* = 0. \tag{27}$$

The limiting result in Corollary 1 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.

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.

6 Conclusions

We find that a batch auction market could potentially hurt liquidity. Strong HFT speed competition, in contrast to a limit order market, stimulates price competition and improves liquidity on frequent batch auction markets even when other instruments, such as auction frequency, become ineffective.

First, fine-tuning the batch auction frequency is an imperfect instrument to improve liquidity. If batch auctions are very frequent, the market may clear before new information becomes publicly available. High-frequency traders with the ability to privately process and act fast on new information benefit as they are able to snipe stale quotes. An obvious solution is to decrease the frequency of batch auctions and allow all traders to catch up with news. On the one hand, longer batch auctions increase the public-to-private news ratio and reduces adverse selection. On the other hand, longer batch auction discourages impatient liquidity traders to participate in the market. We find liquidity improves with the auction interval length until a threshold, an *auction frequency wall*.

How can liquidity be improved once the auction frequency wall is reached? On a batch auction market, HFTs with private information compete in prices over the arbitrage opportunities. Price competition reduces arbitrage rents and adverse selection costs: the equilibrium spread decreases. We find that higher

HFT speed in monitoring and acting on information stimulates price competition, as it reduces the probability each individual arbitrageur is a monopolist on information in any given auction.

The paper's findings contribute to the public debate on alternative market design. It generates an important implication: Policies that aim at intensifying HFT monitoring (e.g., allowing for colocation) can reduce the equilibrium spread and can facilitate the transition from limit order markets to batch auctions. This in turn implies that the "arms race" implication strikingly differs on batch auction markets when compared to limit order markets. Whereas in the latter case it is viewed to be socially costly, it turns out to be beneficial in the case of frequent batch auction markets.

Finally, the model is appealing as it can be solved in closed-form. The framework can be extended to analyze 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

Admati, Anat R., and Paul Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3–40.

Ainsworth, Andrew, and Adrian D. Lee, 2014, Waiting costs and limit order book liquidity: Evidence from the ex-dividend deadline in australia, *Journal of Financial Markets* 20, 101–128.

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

Bats, 2015, Guidance note: Periodic auctions book, Manuscript Bats Europe.

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.

Cinnober, 2010, Using adaptive micro auctions to provide efficient price discovery when access in terms of latency is differentiated among market participants, Manuscript Cinnober.

Demarchi, Marianne, and Solenn Thomas, 2001, French institutional investors: Investment process, trading practices, and expectations, in *The Electronic Call Auction: Market Mechanism and Trading: Building A Better Stock Market* . pp. 229–254 (Springer US: Boston, MA).

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.

Ellul, Andrew, Craig W. Holden, Pankaj Jain, and Robert Jennings, 2007, Order dynamics: Recent evidence from the nyse, *Journal of Empirical Finance* 14, 636–661.

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*.

Foucault, Thierry, Ohad Kadan, and Eugene Kandel, 2005, Limit order book as a market for liquidity, *Review of Financial Studies* 18, 1171–1217.

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.

Hagströmer, Björn, and Lars Nordén, 2014, Closing call auctions at the index futures market, *Journal of Futures Markets* 34, 299–319.

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*.

Kandel, Eugene, Barbara Rindi, and Luisella Bosetti, 2012, The effect of a closing call auction on market quality and trading strategies, *Journal of Financial Intermediation* 21, 23–49.

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

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

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.

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.

Pagano, Michael S., Lin Peng, and Robert A. Schwartz, 2013, A call auction's impact on price formation and order routing: Evidence from the nasdaq stock market, *Journal of Financial Markets* 16, 331–361.

Pagano, Michael S., and Robert A. Schwartz, 2003, A closing call's impact on market quality at euronext paris, *Journal of Financial Economics* 68, 439–484.

Roşu, Ioanid, 2009, A dynamic model of the limit order book, *Review of Financial Studies* 22, 4601–4641.

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*.

Appendix

A Notation summary

Model parameters and their interpretation.

Parameter	Definition
v_t	Common value of the risky asset at time t.
τ	Length of auction interval
Δ	Latency of public signal
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_P	Size of liquidity shocks, i.e., private value innovations.
N	Number of high-frequency traders.
c	Liquidity traders' waiting cost.

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

$$\pi_{mo}(s_m, F(s_m)) = (1-p)\sum_{k=0}^{N-2}\binom{N-2}{k}(1-p)^k p^{N-k-2} F(s_m)^{N-k-2}(\sigma - s_m), \tag{B.1}$$

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

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

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. □

Lemma 2

Proof. The function L is defined as in equation (11), i.e.,

$$L = \frac{1 + (N-1)\exp(-N\Delta\phi) - N\exp[-(N-1)\Delta\phi]}{N\phi}. \tag{B.6}$$

Comparative statics with respect to Δ. First, we compute the partial derivative of L with respect to Δ:

$$\frac{\partial L}{\partial \Delta} = (N-1)\exp(-N\Delta\phi)\left[\exp(\Delta\phi) - 1\right] > 0. \tag{B.7}$$

Since the partial derivative is positive, L increases in the public signal latency Δ. Further, since $\Delta > 0$ it follows that L is positive, as

$$L > L_{\Delta=0} = 0. \tag{B.8}$$

Comparative statics with respect to ϕ. To begin with, we compute the partial derivative of L with respect to ϕ, that is

$$\frac{\partial L}{\partial \phi} = \frac{e^{-\Delta N\phi}}{N\phi^2}\underbrace{\left(Ne^{\Delta\phi}(\Delta(N-1)\phi + 1) - e^{\Delta N\phi} - (N-1)(\Delta N\phi + 1)\right)}_{L_\phi}. \tag{B.9}$$

From equation (11), it follows that

$$\lim_{\phi\to 0} L = \lim_{\phi\to\infty} L = 0. \tag{B.10}$$

and, from (B.8), that $L > 0$. Let L_ϕ be defined as

$$L_\phi \equiv Ne^{\Delta\phi}(\Delta(N-1)\phi + 1) - e^{\Delta N\phi} - (N-1)(\Delta N\phi + 1). \tag{B.11}$$

To prove that L has an unique extremum in ϕ, it is enough to show that L_ϕ has an unique solution for $\phi > 0$. We first compute the limits of L_ϕ at $\phi = 0$ and $\phi \to \infty$,

$$\lim_{\phi \to 0} L_\phi = N - 1 - (N - 1) = 0 \tag{B.12}$$

and

$$\lim_{\phi \to \infty} L_\phi = \lim_{\phi \to \infty} e^{\Delta\phi} \left(N + \Delta\phi (N - 1)\right) \underbrace{\left(1 - \frac{(N - 1)(1 + \Delta N \phi)}{N + \Delta\phi (N - 1)}\right)}_{\to 1 - N < 0} = -\infty. \tag{B.13}$$

The second derivative of L_ϕ with respect to ϕ is

$$L''_\phi = \frac{\partial^2 L_\phi}{\partial \phi^2} = N\Delta^2 \left(-e^{N\Delta\phi} + e^{\Delta\phi} \left(\Delta\phi (N - 1) + 2N - 1\right)\right). \tag{B.14}$$

Since $N \exp(1 - 2N) \in \left(0, e^{-1}\right)$ for $N \geq 1$, the equation $L''_\phi = 0$ has two real solutions ϕ_1^* and ϕ_2^* – corresponding to the two real branches of the Lambert function W.[7] The solutions are expressions of the Lambert function, that is

$$\phi_1^* = \frac{1}{\Delta (N - 1)} \left(1 - 2N - W_{-1}\left(-Ne^{1-2N}\right)\right) > 0 \text{ and} \tag{B.15}$$

$$\phi_2^* = \frac{1}{\Delta (N - 1)} \left(1 - 2N - W_0\left(-Ne^{1-2N}\right)\right) < 0, \tag{B.16}$$

where $W_{-1}(\cdot)$ and $W_0(\cdot)$ denote the lower, respectively upper branches of the Lambert function. To prove only ϕ_1^* is positive, we recall that the upper branch of the Lambert function is increasing and the lower branch is decreasing. Therefore,

$$\begin{aligned} 1 - 2N - W_{-1}\left(-Ne^{1-2N}\right) &> 1 - 2N - \underbrace{W_{-1}\left((1 - 2N)e^{1-2N}\right)}_{1-2N} = 0, \text{ and} \\ 1 - 2N - W_0\left(-Ne^{1-2N}\right) &< 1 - 2N - \underbrace{W_0\left((1 - 2N)e^{1-2N}\right)}_{1-2N} = 0. \end{aligned} \tag{B.17}$$

It directly follows that $\phi_1^* > 0$ and $\phi_2^* < 0$. Since

$$\lim_{\phi \to 0} L''_\phi = 2N\Delta^2 (N - 1) > 0 \text{ and } \lim_{\phi \to \infty} L''_\phi = -\infty, \tag{B.18}$$

it follows that $L''_\phi \geq 0$ for $\phi \in \left(0, \phi_1^*\right]$ and $L''_\phi < 0$ for $\phi \in \left(\phi_1^*, \infty\right)$. The table below describes the behaviour and limits of L_ϕ and its first two derivatives.

[7] It is enough to see $N \exp(1 - 2N)$ decreases in N and takes the value one for $N = 1$.

Function	$\phi = 0$	$\phi \in \left(0, \phi_1^*\right)$	$\phi = \phi^*$	$\phi \in \left(\phi_1^*, \infty\right)$	$\phi \to \infty$
L''_ϕ	>0	Positive	0	Negative	$-\infty$
L'_ϕ	0	Increasing	Maximum	Decreasing	$-\infty$
L_ϕ	0	Positive	Positive	Undetermined sign	$-\infty$

From the table above, it follows that there is an unique solution for $L'_\phi = 0$, in the interval $\left(\phi_1^*, \infty\right)$. Let this solution be $\phi^{**} > \phi_1^*$. Therefore, L_ϕ is zero for $\phi = 0$, increases for $\phi < \phi^{**}$ and decreases for $\phi > \phi^{**}$. Consequently, L_ϕ only crosses zero once for $\phi \in (0, \infty)$. Therefore, there exists a $\bar{\phi} > 0$ such that the function L increases for $\phi \leq \bar{\phi}$ and decreases for $\phi > \bar{\phi}$.

Comparative statics with respect to N. To begin with, we compute the partial derivative of L with respect to N, that is

$$\frac{\partial L}{\partial N} = \frac{e^{-\Delta N \phi}}{N^2 \phi^2} \underbrace{\left(1 - e^{\Delta \phi N} - \Delta \phi N (N-1) + \Delta \phi N^2 e^{\Delta \phi}\right)}_{L_N}. \tag{B.19}$$

Again, from equation (11), it follows that

$$\lim_{N \to 1} L = \lim_{N \to \infty} L = 0. \tag{B.20}$$

Let L_N be defined as

$$L_N = 1 - e^{\Delta \phi N} - \Delta \phi N (N-1) + \Delta \phi N^2 e^{\Delta \phi}. \tag{B.21}$$

As in the previous part of the proof, it is enough to show that L_N has a unique solution in N for $N \geq 1$. We begin by taking the limits of L_N, that is

$$\lim_{N \to 1} L_N = 1 + e^{\phi \Delta} (\phi \Delta - 1) > 0 \text{ and } \lim_{N \to \infty} L_N = -\infty. \tag{B.22}$$

The first derivative of L_N with respect to N is

$$L'_N = \frac{\partial L_N}{\partial N} = \Delta \phi \left(1 - e^{N \Delta \phi} + 2N \left(e^{\Delta \phi} - 1\right)\right). \tag{B.23}$$

Since $\frac{\Delta \phi}{2(e^{\Delta \phi} - 1)} \exp\left(\frac{\Delta \phi}{2(e^{\Delta \phi} - 1)}\right) \in \left(0, e^{-1}\right)$ for $\Delta \phi > 0$[8], the equation $L'_N = 0$ has two real solutions N_1^* and N_2^*, that is

$$N_1^* = \frac{1}{2\left(1 - e^{\Delta \phi}\right)} - \frac{1}{\Delta \phi} W_0 \left[\frac{\Delta \phi}{2\left(1 - e^{\Delta \phi}\right)} \exp\left(\frac{\Delta \phi}{2\left(1 - e^{\Delta \phi}\right)}\right)\right] \text{ and}$$
$$N_2^* = \frac{1}{2\left(1 - e^{\Delta \phi}\right)} - \frac{1}{\Delta \phi} W_{-1} \left[\frac{\Delta \phi}{2\left(1 - e^{\Delta \phi}\right)} \exp\left(\frac{\Delta \phi}{2\left(1 - e^{\Delta \phi}\right)}\right)\right] \tag{B.24}$$

where $W_{-1}(\cdot)$ and $W_0(\cdot)$ denote the lower, respectively upper branches of the Lambert function.

[8] It is easy to show that $\frac{x}{2(e^x - 1)} \exp\left(\frac{x}{2(e^x - 1)}\right)$ decreases in x, is larger than zero for $x > 0$ and equal to $\frac{1}{2\sqrt{e}} < \frac{1}{e}$ for $x = 0$.

It follows that $N_1^* = 0$, since $W_0\left(xe^x\right) = x$, that is

$$N_1^* = \frac{1}{2\left(1-e^{\Delta\phi}\right)} - \frac{1}{\Delta\phi}\frac{\Delta\phi}{2\left(1-e^{\Delta\phi}\right)} = 0. \tag{B.25}$$

Since $W_{-1}(x) < W_0(x)$,

$$N_2^* = \frac{1}{2\left(1-e^{\Delta\phi}\right)} - \frac{1}{\Delta\phi}W_{-1}\left[\frac{\Delta\phi}{2\left(1-e^{\Delta\phi}\right)}\exp\left(\frac{\Delta\phi}{2\left(1-e^{\Delta\phi}\right)}\right)\right] > N_1^* = 0. \tag{B.26}$$

Finally, we need to show $N_2^* > 1$. We prove this by contradiction. First, assume $N_2^* < 1$. Since N_2^* is the largest solution to $L_N' = 0$, then L_N' has the same sign on $\left(N_2^*, \infty\right)$. For $\Delta\phi > 0$,

$$\lim_{N\to\infty} L_N' = \infty. \tag{B.27}$$

Further, for $N = 1$, $L_N' = \Delta\phi\left(e^{\Delta\phi} - 1\right) > 0$. It follows that L_N' changes sign on $(1, \infty)$. Consequently, it follows that $N_2^* > 1$.

The table below describes the behaviour and limits of L_N and its first derivative.

Function	$N = 1$	$N \in \left(1, N_2^*\right)$	$\phi = N_2^*$	$N \in \left(N_2^*, \infty\right)$	$N \to \infty$
L_N'	$\Delta\phi\left(e^{\Delta\phi} - 1\right) > 0$	Positive	0	Negative	$-\infty$
L_N	$1 + e^{\phi\Delta}(\phi\Delta - 1) > 0$	Increasing	Maximum	Decreasing	$-\infty$

From the table above, it follows that there is an unique solution for $L_N = 0$, in the interval $\left(N_2^*, \infty\right)$. Therefore, there exists a $\bar{N} > 0$ such that the function L increases for $N \leq \bar{N}$ and decreases for $N > \bar{N}$. □

Lemma 3

Proof. If $\tau = 0$, then the LT utility function becomes

$$u_{\mathrm{LT}} = \sigma_P - s - c\max\{0 - t_\ell, 0\} = \sigma_P - s, \tag{B.28}$$

which does not depend on τ and is strictly positive if $s < \sigma_P$. As we show in Lemma 4, in equilibrium $s^* < \sigma < \sigma_P$. Therefore the LT always submits a market order to trade on its private value.

The ex ante probability of an LT order does not depend on τ for $\tau \leq \frac{\sigma_P - s}{c}$ and strictly decreases in τ for $\tau > \frac{\sigma_P - s}{c}$ since

$$\frac{\partial\xi\tau^{-1}}{\partial\tau} = -\xi\tau^{-2} < 0. \tag{B.29}$$

□

Lemma 4

Proof. From equation (17), it follows that the competitive half-spread s^* is the solution to

$$s^* = \sigma \frac{\eta L}{\eta L + \mu\tau \min\left\{\frac{\sigma_P - s^*}{c\tau}, 1\right\}}. \tag{B.30}$$

If $\tau \le \frac{\sigma_P - s^*}{c}$, then $s^* = \sigma \frac{\eta L}{\eta L + \mu\tau}$, a decreasing function of τ. If $\tau > \frac{\sigma_P - s^*}{c}$, then s^* solves

$$s^* = \sigma \frac{\eta L}{\eta L + \mu \frac{\sigma_P - s^*}{c}}, \tag{B.31}$$

which is not a function of τ. It follows that there exists a threshold τ^* such that for $\tau > \tau^*$ the competitive half-spread is

$$s^* = \frac{c\eta L + \mu\sigma_P - \sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}}{2\mu} \tag{B.32}$$

The threshold τ^* is the positive solution of the quadratic equation

$$\tau^* = \frac{\sigma_P - \sigma \frac{\eta L}{\eta L + \mu\tau^*}}{c}, \tag{B.33}$$

that is

$$\tau^* = \frac{\sqrt{(c\eta L - \mu\sigma_P)^2 + 4c\eta L\mu(\sigma_P - \sigma)} - (c\eta L - \mu\sigma_P)}{2c\mu} > 0. \tag{B.34}$$

The competitive half-spread s^* is continuous in τ since

$$\lim_{\tau \uparrow \tau^*} s^* = \sigma \frac{\eta L}{\eta L + \mu\tau^*} = \frac{2c\eta L\sigma}{\sqrt{(c\eta L - \mu\sigma_P)^2 - 4c\eta L\mu(\sigma - \sigma_P)} + c\eta L + \mu\sigma_P} = \lim_{\tau \downarrow \tau^*} s^*, \tag{B.35}$$

which follows from the equality

$$\frac{2c\eta L\sigma}{\sqrt{(c\eta L - \mu\sigma_P)^2 - 4c\eta L\mu(\sigma - \sigma_P)} + c\eta L + \mu\sigma_P} = \frac{c\eta L + \mu\sigma_P - \sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}}{2\mu}. \tag{B.36}$$

□

Proposition 1

Proof. First, it is optimal for HFTs to submit all orders, if any, just before $t = \tau$. This strategy maximizes the information HFTs have, without the loss of any trading opportunity.

If the public signal is updated in $(0, \tau)$, that is, before market clearing, all HFTs are informed about a news arrival. Since there was a news arrival, it follows that no LT arrived between $t = 0$ and $t = \tau$. Hence, the HFM has no incentive to submit limit orders, regardless of the HFS strategy. Since the order book is empty, there are no sniping opportunities. Therefore, no HFS submits a limit order targeted at an existing quote.

Further, we analyze the equilibrium strategies if there is no public signal update in $(0, \tau)$.

If the HFM does not observe any private news, his exected utility is given by equation (16), i.e.,

$$\pi_{\text{HFM}} = \pi_{\text{liqudity}} - \underbrace{[1-\exp(-(\eta+\mu)\tau)]\frac{\eta}{\eta+\mu}}_{\text{Probability of news arrival}} \underbrace{\frac{1}{\tau}L\times(\sigma-s_l)}_{\text{Conditional HFM loss}} . \tag{B.37}$$

Since the HFM is competitive, its expected utility should be zero. Consequently, the half-spread is the competitive one, s^* as defined in Lemma 4.

Again, if the HFM observes any news privately, then she infers no LT arrival and has no incentive to submit limit orders.

If an HFS receives a private signal, then she submits a marketable order against the stale quote. This order has either a positive profit (if sniping is successful) or a zero profit if another HFS submits a better price against the stale quote, or the HFM is also informed. Lemma 1 establishes the distribution of HFT sniping orders prices.

Finally, an LT receiving a private value shock submits a marketable order if and only if his expected utility is positive, i.e., if

$$u_{\text{LT}} = \sigma_P - s - c\max\{\tau - t_\ell, 0\} > 0 \iff t_\ell > \max\left\{0, \tau - \frac{\sigma_P - s^*}{c}\right\}. \tag{B.38}$$

□

Lemma 5

Proof. From Lemma 4 and Proposition 1, the equilibrium half-spread is

$$s^*(\tau) = \begin{cases} \sigma\frac{\eta L}{\eta L+\mu\tau}, \text{ if } \tau \le \tau^*, \\ \frac{c\eta L+\mu\sigma_P - \sqrt{(c\eta L+\mu\sigma_P)^2 - 4c\eta L\mu\sigma}}{2\mu}, \text{ if } \tau > \tau^*. \end{cases} \tag{B.39}$$

Let s_1 and s_2 denote the upper, respectively lower branches of equation (B.39).

Equilibrium half-spread and σ. The partial derivative of s_1 with respect to σ is

$$\frac{\partial s_1}{\partial \sigma} = \frac{\eta L}{\eta L + \mu\tau} > 0. \tag{B.40}$$

The partial derivative of s_2 with respect to σ is

$$\frac{\partial s_1}{\partial \sigma} = \frac{c\eta L}{\sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}} > 0. \tag{B.41}$$

Since both the partial derivatives of s_1 and s_2 with respect to σ are positive, the equilibrium half-spread s^* increases in σ.

Equilibrium half-spread and μ. The partial derivative of s_1 with respect to μ is

$$\frac{\partial s_1}{\partial \mu} = -\frac{\eta L \tau}{(\eta L + \mu\tau)^2} < 0. \tag{B.42}$$

The partial derivative of s_2 with respect to μ is

$$\frac{\partial s_2}{\partial \mu} = \frac{c\eta L\,(c\eta L - 2\mu\sigma + \mu\sigma_P) - \sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}}{2\mu^2\sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}}. \tag{B.43}$$

It remains to show that

$$c\eta L - 2\mu\sigma + \mu\sigma_P - \sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma} < 0, \tag{B.44}$$

which is equivalent to showing that

$$c\eta L - 2\mu\sigma + \mu\sigma_P < \sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}. \tag{B.45}$$

Squaring both the left- and right-hand sides, equation (B.45) is equivalent to

$$4\mu^2\sigma\,(\sigma - \sigma_P) < 0, \tag{B.46}$$

which is true since $\sigma < \sigma_P$. Since both the partial derivatives of s_1 and s_2 with respect to μ are negative, the equilibrium half-spread s^* decreases in μ.

Equilibrium half-spread and η. The partial derivative of s_1 with respect to η is

$$\frac{\partial s_1}{\partial \eta} = \frac{L\mu\tau}{(\eta L + \mu\tau)^2} > 0. \tag{B.47}$$

The partial derivative of s_2 with respect to η is

$$\frac{\partial s_2}{\partial \eta} = \frac{cL\left(1 - \frac{c\eta L + \mu(\sigma_P - 2\sigma)}{\sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}}\right)}{2\mu} > 0, \tag{B.48}$$

since from equation (B.45), $\frac{c\eta L + \mu(\sigma_P - 2\sigma)}{\sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma}} < 1$. Since both the partial derivatives of s_1 and s_2 with respect to η are positive, the equilibrium half-spread s^* increases in η. □

Proposition 2

Proof.

Equilibrium half-spread and τ. The partial derivative of s_1 with respect to τ is

$$\frac{\partial s_1}{\partial \tau} = -\frac{\eta L \mu}{(\eta L + \mu\tau)^2} < 0. \tag{B.49}$$

Since the partial derivative of s_1 is negative, s^* decreases in τ for $\tau \leq \tau^*$. The partial derivative of s_2 with respect to τ is

$$\frac{\partial s_2}{\partial \tau} = 0. \tag{B.50}$$

Since the partial derivative of s_2 is zero, s^* does not depend on τ for $\tau > \tau^*$.

Equilibrium half-spread and L, Δ, N, and ϕ. The partial derivative of s_1 with respect to L is

$$\frac{\partial s_1}{\partial L} = \frac{\eta\mu\tau}{(\eta L + \mu\tau)^2} > 0. \tag{B.51}$$

The partial derivative of s_2 with respect to L is

$$\frac{\partial s_2}{\partial L} = \frac{c\eta\left(1 - \frac{c\eta L+\mu(\sigma_P-2\sigma)}{\sqrt{(c\eta L+\mu\sigma_P)^2-4c\eta L\mu\sigma}}\right)}{2\mu} > 0, \tag{B.52}$$

since from equation (B.45), $\frac{c\eta L+\mu(\sigma_P-2\sigma)}{\sqrt{(c\eta L+\mu\sigma_P)^2-4c\eta L\mu\sigma}} < 1$. Since both the partial derivatives of s_1 and s_2 with respect to L are positive, the equilibrium half-spread s^* increases in L.

For any parameter $\nu \in \{\Delta, N, \phi\}$, the equilibrium half-spread s^* depends on ν only through the snipe loss L, i.e.,

$$\frac{\partial s^*}{\partial \nu} = \underbrace{\frac{\partial s^*}{\partial L}}_{>0}\frac{\partial L}{\partial \nu}. \tag{B.53}$$

Lemma 2 establishes the sign of $\frac{\partial L}{\partial \nu}$ for all $\nu \in \{\Delta, N, \phi\}$ and therefore completes the proof. □

Proposition 3

Proof.

Auction length threshold τ^* and L. The partial derivative of τ^* with respect to L is

$$\frac{\partial \tau^*}{\partial L} = \frac{\eta\left(\frac{c\eta L+\mu(\sigma_P-2\sigma)}{\sqrt{(c\eta L-\mu\sigma_P)^2-4c\eta L\mu(\sigma-\sigma_P)}} - 1\right)}{2\mu} < 0, \tag{B.54}$$

since from equation (B.45), $\frac{c\eta L+\mu(\sigma_P-2\sigma)}{\sqrt{(c\eta L+\mu\sigma_P)^2-4c\eta L\mu\sigma}} < 1$. Since the partial derivative of τ^* with respect to L is negative, it follows that τ^* decreases in L.

Auction length threshold τ^* and η. The partial derivative of τ^* with respect to η is

$$\frac{\partial \tau^*}{\partial \eta} = \frac{L\left(\frac{c\eta L+\mu(\sigma_P-2\sigma)}{\sqrt{(c\eta L-\mu\sigma_P)^2-4c\eta L\mu(\sigma-\sigma_P)}} - 1\right)}{2\mu} < 0, \tag{B.55}$$

since from equation (B.45), $\frac{c\eta L+\mu(\sigma_P-2\sigma)}{\sqrt{(c\eta L+\mu\sigma_P)^2-4c\eta L\mu\sigma}} < 1$. Since the partial derivative of τ^* with respect to η is negative, it follows that τ^* decreases in η.

Auction length threshold τ^* and μ. The partial derivative of τ^* with respect to μ is

$$\frac{\partial \tau^*}{\partial \mu} = \frac{\eta L\left(\sqrt{(c\eta L-\mu\sigma_P)^2 - 4c\eta L\mu(\sigma-\sigma_P)} - c\eta L - \mu(\sigma_P - 2\sigma)\right)}{2\mu^2\sqrt{(c\eta L-\mu\sigma_P)^2 - 4c\eta L\mu(\sigma-\sigma_P)}} > 0. \tag{B.56}$$

since from equation (B.45), $\frac{c\eta L+\mu(\sigma_P-2\sigma)}{\sqrt{(c\eta L+\mu\sigma_P)^2-4c\eta L\mu\sigma}} < 1$. Since the partial derivative of τ^* with respect to μ is negative, it follows that τ^* increases in μ.

Auction length threshold τ^* and c. The partial derivative of τ^* with respect to c is

$$\frac{\partial \tau^*}{\partial c} = \frac{c\eta L(2\sigma-\sigma_P) - \sigma_P\left(\sqrt{c^2\eta^2L^2 + 2c\eta L\mu(\sigma_P-2\sigma) + \mu^2\sigma_P^2} + \mu\sigma_P\right)}{2c^2\sqrt{(c\eta L-\mu\sigma_P)^2 - 4c\eta L\mu(\sigma-\sigma_P)}}. \tag{B.57}$$

It remains to show that

$$c\eta L(\sigma_P - 2\sigma) + \sigma_P\left(\sqrt{c^2\eta^2L^2 + 2c\eta L\mu(\sigma_P-2\sigma) + \mu^2\sigma_P^2} + \mu\sigma_P\right) > 0, \tag{B.58}$$

which is equivalent to showing that

$$\sigma_P\left[c\eta L\frac{\sigma_P-2\sigma}{\sigma_P} + \mu\sigma_P + \sqrt{\left(c\eta L\frac{\sigma_P-2\sigma}{\sigma_P} + \mu\sigma_P\right)^2 + (c\eta L)^2\left(1-\left(\frac{\sigma_P-2\sigma}{\sigma_P}\right)^2\right)}\right] > 0. \tag{B.59}$$

The last equation follows directly from $\frac{\sigma_P-2\sigma}{\sigma_P} < 1$. Therefore, it follows that $\frac{\partial\tau^*}{\partial c} < 0$. Since the partial derivative of τ^* with respect to c is negative, it follows that τ^* decreases in c.

Auction length threshold τ^* and σ_P. The partial derivative of τ^* with respect to σ_P is

$$\frac{\partial \tau^*}{\partial \sigma_P} = \frac{\frac{c\eta L+\mu\sigma_P}{\sqrt{(c\eta L-\mu\sigma_P)^2-4c\eta L\mu(\sigma-\sigma_P)}} + 1}{2c} > 0, \tag{B.60}$$

Since the partial derivative of τ^* with respect to σ_P is positive, it follows that τ^* increases in σ_P.

Auction length threshold τ^* and σ. The partial derivative of τ^* with respect to σ_P is

$$\frac{\partial \tau^*}{\partial \sigma} = -\frac{\eta L}{\sqrt{(c\eta L - \mu\sigma_P)^2 - 4c\eta L\mu(\sigma - \sigma_P)}} < 0, \tag{B.61}$$

Since the partial derivative of τ^* with respect to σ is negative, it follows that τ^* decreases in σ.

□

Proposition 4

Proof. The batch auction market improves liquidity if and only if $\underline{s}^* < s^*_{\mathrm{LOB}}$. From Lemma 4 and equation (25),

$$s^* < s^*_{\mathrm{LOB}} \iff \sigma \frac{\eta L}{\eta L + \mu \frac{\sigma_P - s^*}{c}} < \sigma \frac{(N-1)\eta}{(N-1)\eta + N\mu}, \tag{B.62}$$

or, equivalently,

$$L < \frac{N-1}{N} \frac{\sigma_P - s^*}{c}. \tag{B.63}$$

Condition (B.63) is equivalent to

$$\begin{aligned} &L < \frac{N-1}{N} \frac{\sqrt{(c\eta L + \mu\sigma_P)^2 - 4c\eta L\mu\sigma} - c\eta L + \mu\sigma_P}{2c\mu} \iff \\ &L < \frac{N-1}{N} \frac{\eta(\sigma_P - \sigma) + \mu\sigma_P}{c(\eta + \mu)} \iff L \times c < \frac{N-1}{N} \frac{\eta(\sigma_P - \sigma) + \mu\sigma_P}{(\eta + \mu)} \end{aligned} \tag{B.64}$$

□

Corollary 1

Proof. From equation (B.10), it follows that

$$\lim_{\phi \to \infty} L = 0. \tag{B.65}$$

It immediately follows from equation (16) that

$$\lim_{\phi \to \infty} \pi_{\mathrm{HFM}} = \pi_{\mathrm{liqudity}}. \tag{B.66}$$

Consequently, for $\phi \to \infty$, s^* solves

$$\pi_{\mathrm{liqudity}} = \left[1 - \exp\left(-(\eta + \mu)\,\tau\right)\right] \frac{\mu}{\eta + \mu} \mathrm{Prob}\left(t_\ell > \tau - \frac{\sigma_P - s^*}{c}\right) s^* = 0, \tag{B.67}$$

that is,

$$\lim_{\phi \to \infty} s^* = 0. \tag{B.68}$$

□

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 and assume good news and an uninformed HFM. Initially, the HFM has a sell quote in the order book at $v + s_l$. Three informed HFS and post buy orders at random prices between $v + s_l$ and $v + \sigma$. The best buy price, in this case of HFS_3 trades against the stale quotes in the book.

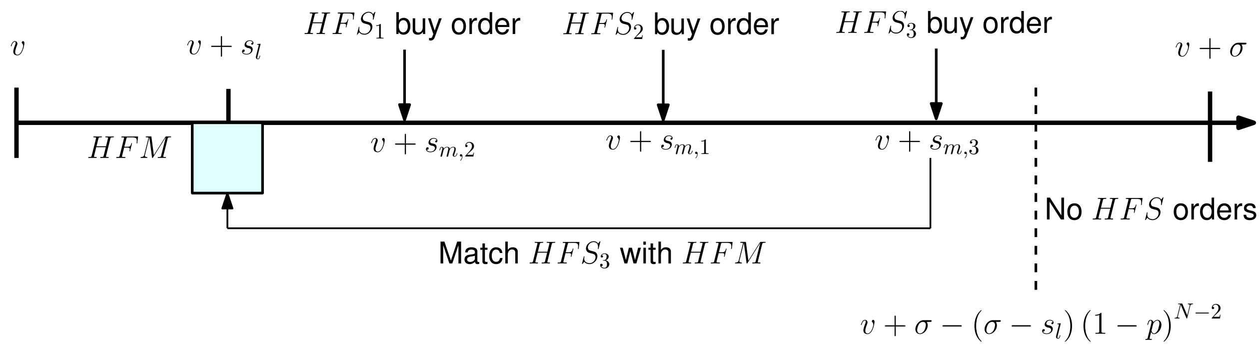

Figure 2: **Price distribution on HFI sniping orders**

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

(a) HFT speed

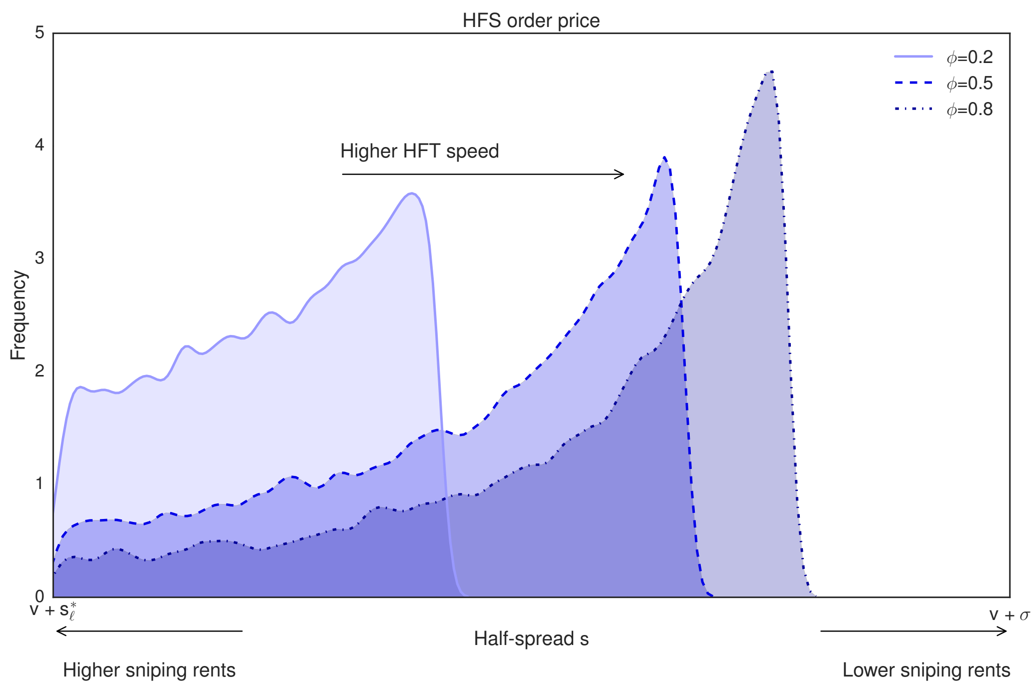

(b) Number of HFTs

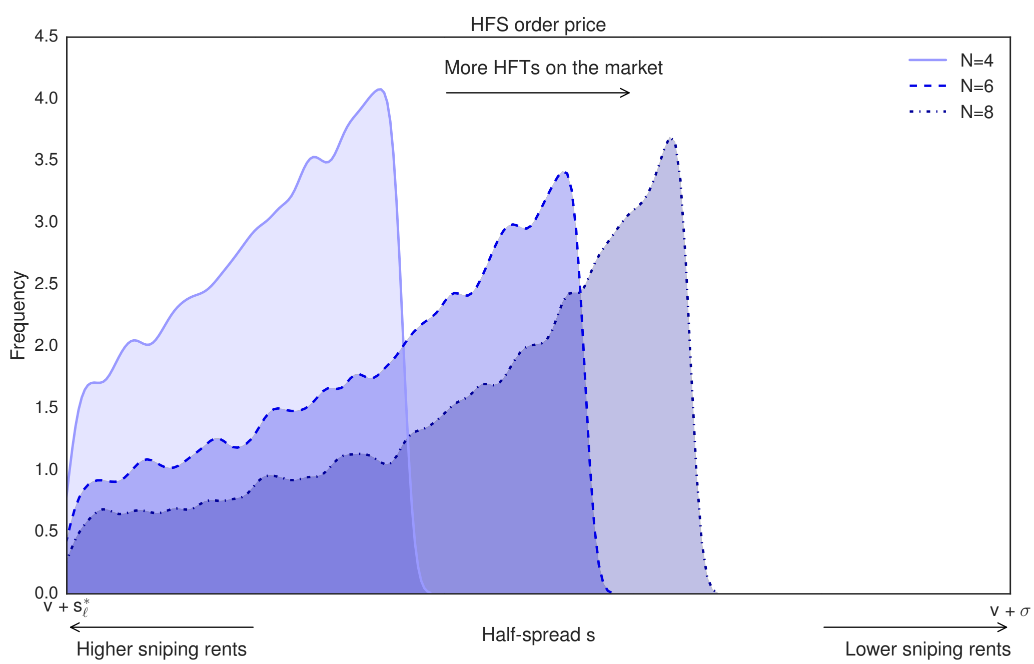

Figure 3: **Equilibrium spread and batch auction length**

This figure plots the equilibrium half-spread on batch auction markets as a function of the batch interval length τ.

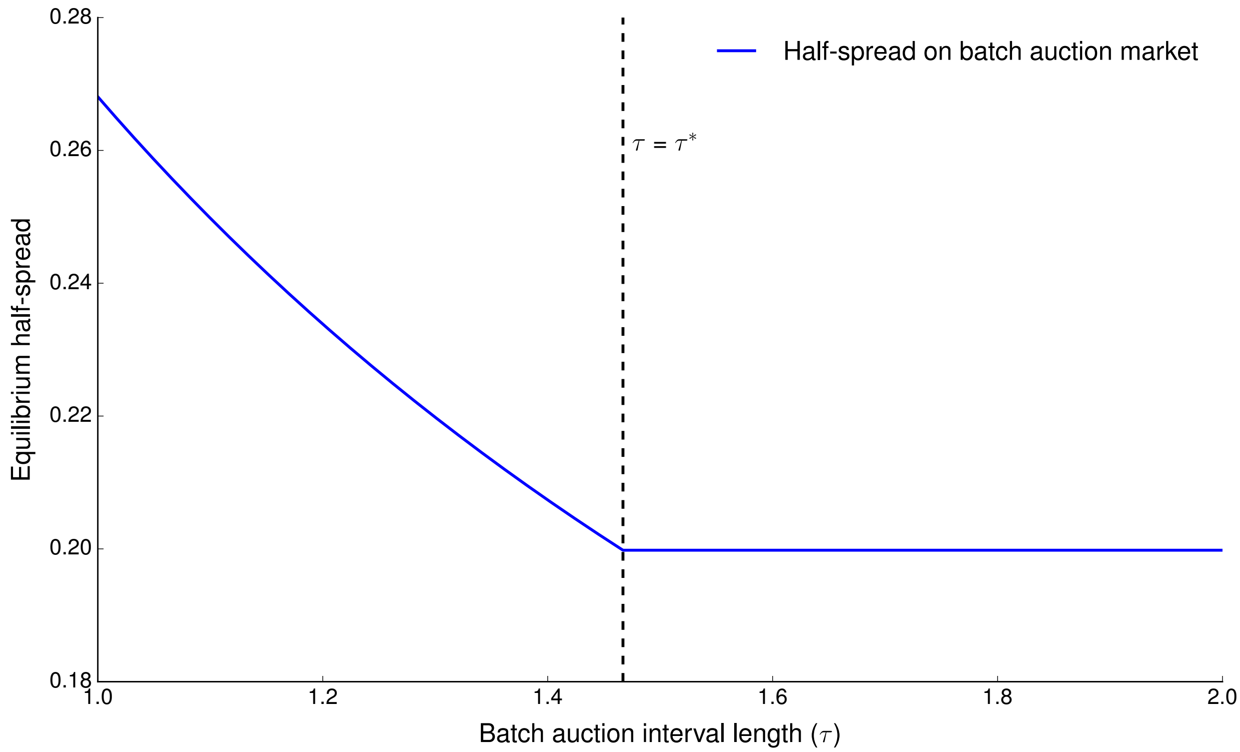

Figure 4: **Equilibrium half-spread on batch auction markets**

This figure plots the equilibrium half-spread on batch auction markets as a function of the HFT speed (first panel) and the number of HFTs (second panel). Parameter values are chosen such that $\tau > \tau^*$.

(a) HFT speed

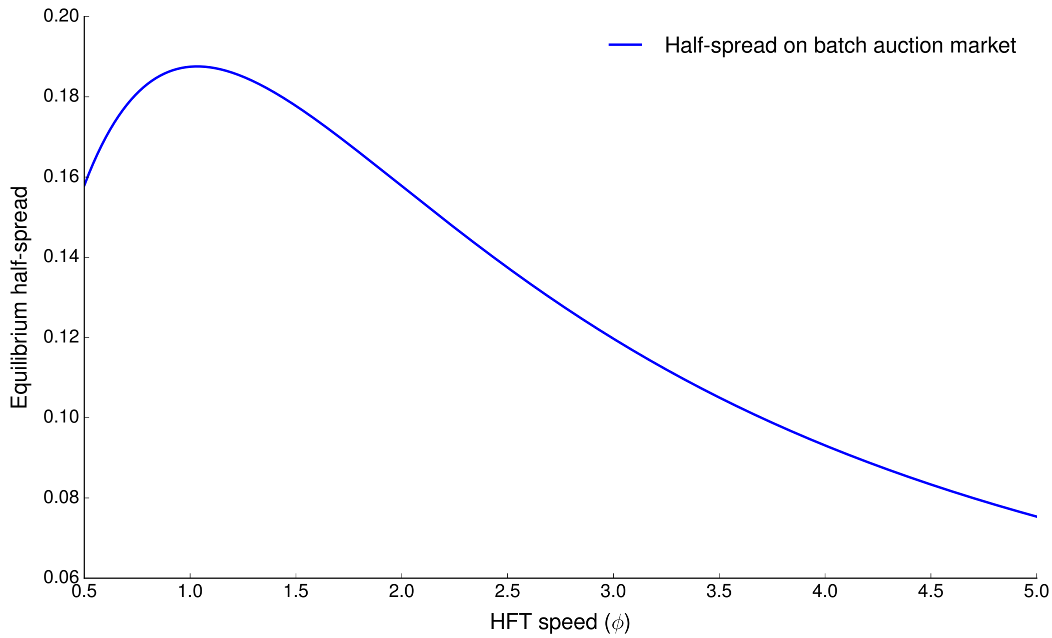

(b) Number of HFTs

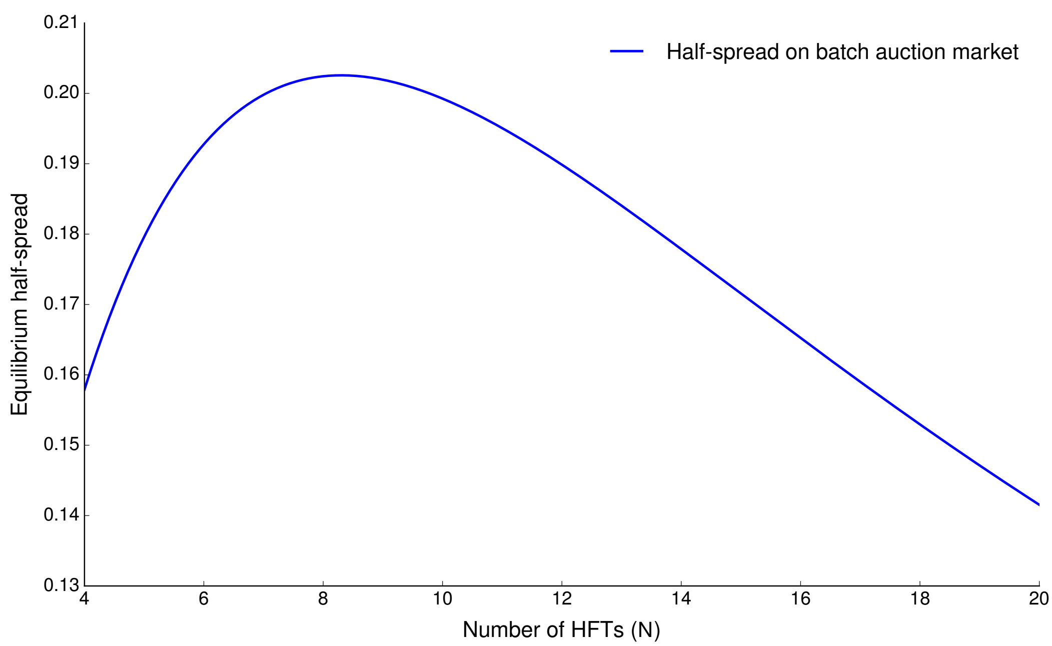

Figure 5: **Equilibrium spread, HFT competition and speed.**

This figure displays the contour plot of the equilibrium spread on a batch auction market as a function of both the number of HFTs and the HFT speed.

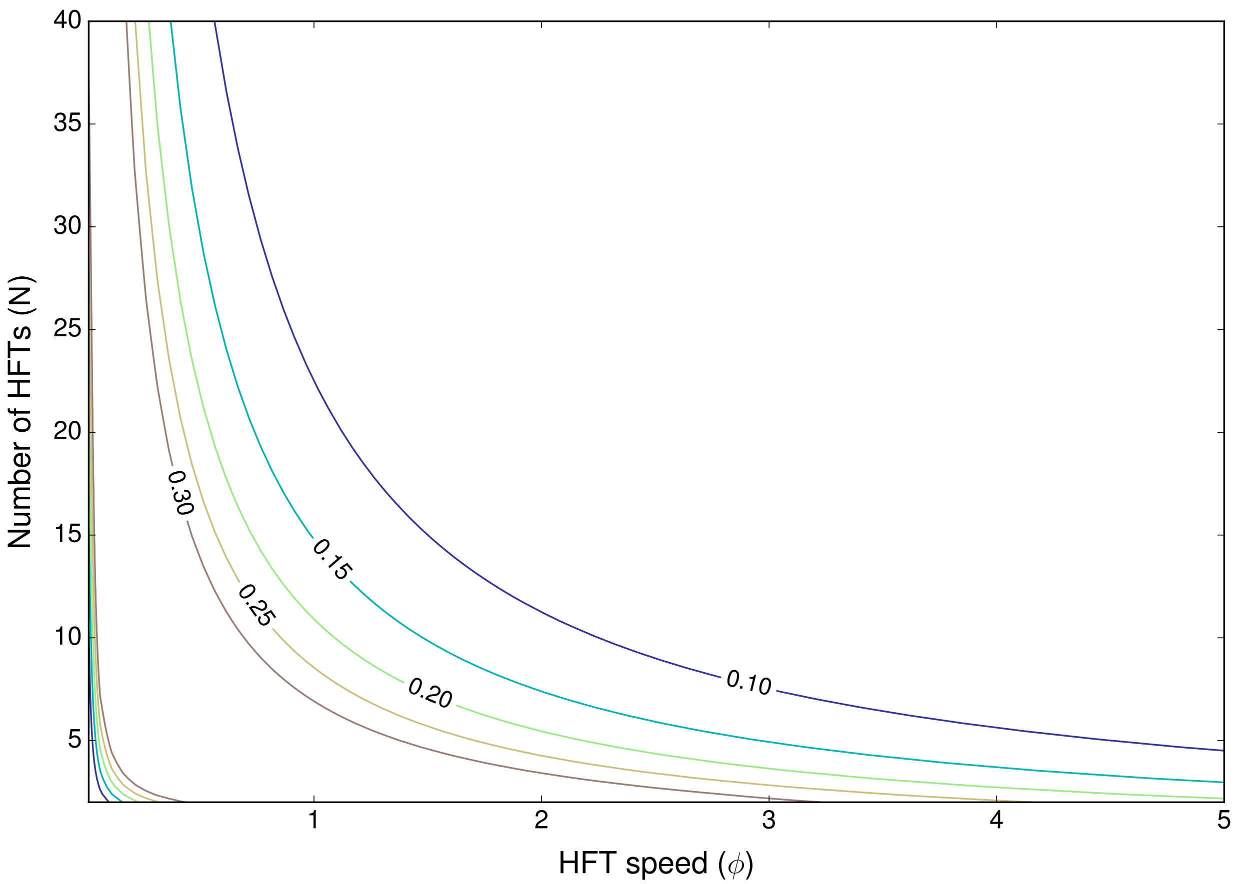

Figure 6: **Direct and indirect effects of speed on liquidity**

This figure plots the equilibrium half-spread on batch auction markets as a function of the batch interval length τ. It illustrates the direct effect of a speed increase on liquidity, and the indirect effect through a higher auction length threshold τ^*.

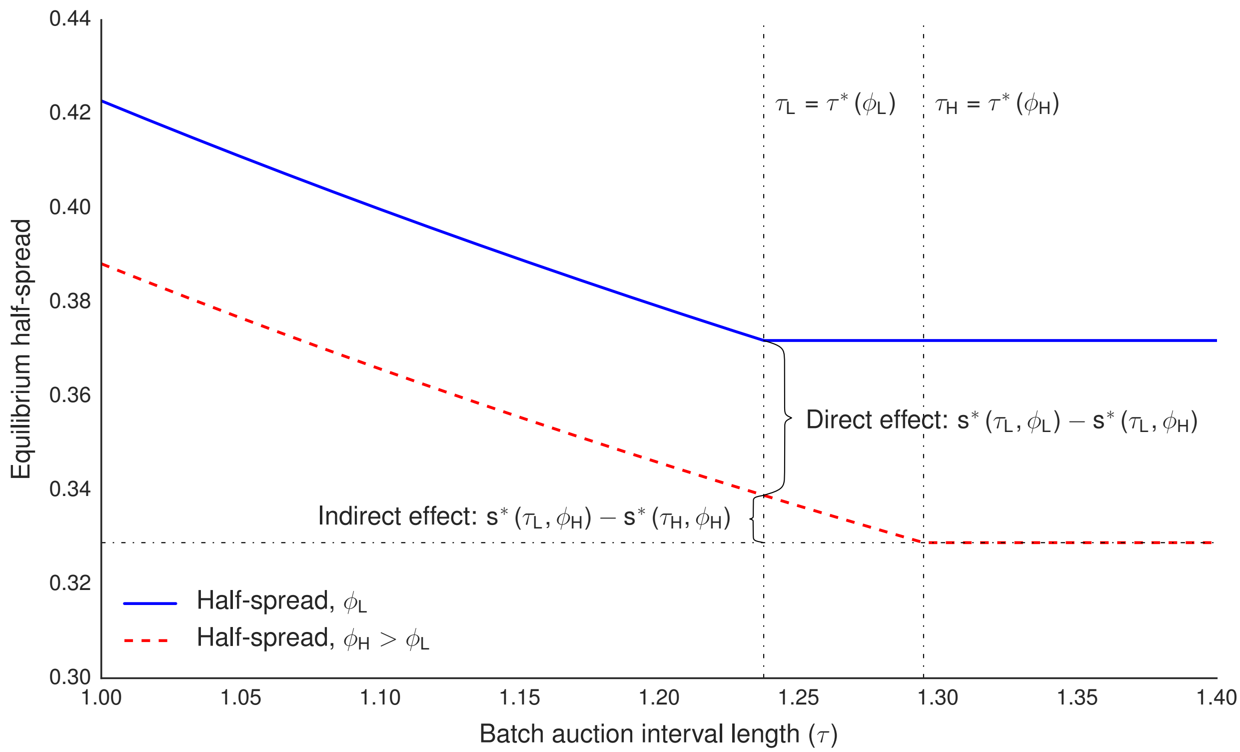

Figure 7: **Spread on benchmark limit order markets**

This figure plots the equilibrium half-spread on batch auction markets as a function of the HFT speed (first panel) and the number of HFTs (second panel). Parameter values are chosen such that $\tau > \tau^*$. The batch auction market half-spread is benchmarked against the half-spread in limit order markets, i.e., the red dashed line.

(a) HFT speed

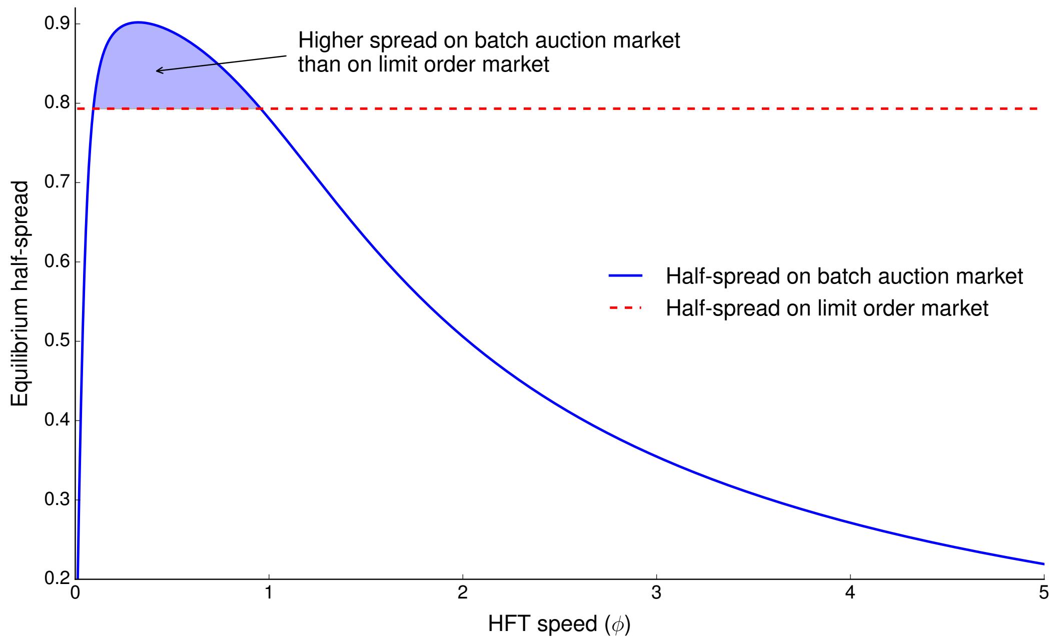

(b) Number of HFTs