

Hoffmann, Peter

Working Paper

A dynamic limit order market with fast and slow traders

ECB Working Paper, No. 1526

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Hoffmann, Peter (2013) : A dynamic limit order market with fast and slow traders, ECB Working Paper, No. 1526, European Central Bank (ECB), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/153959>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM



WORKING PAPER SERIES

NO 1526 / MARCH 2013

A DYNAMIC LIMIT ORDER MARKET WITH FAST AND SLOW TRADERS

Peter Hoffmann



In 2013 all ECB
publications
feature a motif
taken from
the €5 banknote.



NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Acknowledgements

This paper is a substantially revised version of the second chapter of my doctoral dissertation at Universitat Pompeu Fabra. Financial support from the Generalitat de Catalunya is gratefully acknowledged. I would like to thank Álvaro Cartea, Jean-Edouard Colliard, Hans Degryse, Başak Güneş, Roberto Pascual, Francisco Peñaranda, and Augusto Rupérez Micola for helpful discussions at various stages of the project. I am particularly grateful to Thierry Foucault, whose detailed comments on an earlier draft have greatly improved the paper. All remaining errors are mine. The views expressed in this paper are those of the author and do not necessarily reflect those of the European Central Bank or the Eurosystem.

Peter Hoffmann

European Central Bank; e-mail: peter.hoffmann@ecb.europa.eu

© European Central Bank, 2013

Address	Kaiserstrasse 29, 60311 Frankfurt am Main, Germany
Postal address	Postfach 16 03 19, 60066 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Internet	http://www.ecb.europa.eu
Fax	+49 69 1344 6000

All rights reserved.

ISSN	1725-2806 (online)
EU Catalogue No	QB-AR-13-023-EN-N (online)

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from <http://www.ecb.europa.eu> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=1969392.

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, <http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html>

Abstract

We study the role of high-frequency trading in a dynamic limit order market. Being fast is valuable because it enables traders to revise outstanding limit orders upon news arrivals when interacting with slow market participants. On the one hand, the existence of fast traders can help to reduce the inefficiency that is rooted in the risk of being "picked off" after unfavourable price movements and therefore allows more gains from trade to be realized. On the other hand, slow traders face a relative loss in bargaining power which leads them to strategically submit limit orders with a lower execution probability, thereby reducing trade. Due to this negative externality, the equilibrium level of investment is always welfare-reducing. The model generates additional testable implications regarding the effects of high-frequency trading on order flow statistics.

Keywords: High-frequency trading, Limit Order Market

JEL Classification: G19, C72, D62

Non-Technical Summary

During the past decade or so, trading in financial markets has become increasingly automated, with recent estimates suggesting that high-frequency trading (HFT) is now accounting for more than half of all transactions in U.S. equity markets. While the use of information technology for trading purposes is not a recent phenomenon per se, it has become the decisive factor in an industry where the proverb “time is money” is being taken literally and market participants are engaged in an arms race at the millisecond level.

This strong growth has ignited a heated debate among academics, practitioners, and regulators about the benefits and concerns associated with an increasing automation of the trading process. While proponents argue that HFT increases market efficiency via improved liquidity and price discovery, critiques claim that at least some of the associated trading strategies are making profits at the expense of other market participants and have the potential to destabilize markets.

This paper contributes to the on-going debate by presenting a stylized model of trading in a limit order market where agents differ in their trading speed, which is thought to capture the difference between (fast) HFTs and (slow) human market participants. Following Foucault (1999), we assume that agents’ limit orders are subject to the risk of being “picked off” by other traders upon the arrival of new value-relevant information, a known inefficiency of limit order markets that reduces the likelihood of realizing gains from trade.

In such a setting, HFT can have two opposing effects. First, fast traders’ may revise their limit orders sufficiently quickly upon news arrivals, and this ability reduces the inefficiency related to the risk of being “picked off” by allowing more trades to take place. Second, slow traders face a loss in bargaining power relative to fast traders, and this induces them to submit more cautious limit orders (i.e. with a lower execution probability). While being privately optimal, this strategy results in an efficiency loss because it reduces the likelihood of agents trading with each other.

We show that slow traders are always worse off in the presence of fast traders compared to a market with identical traders. While speed can allow for the realization of more gains from trade, it also allows fast traders to extract this additional surplus from slow traders. If the proportion of fast traders is determined endogenously, this directly implies that HFT leads to welfare loss in equilibrium. As in Biais et al. (2012), fast traders exert a negative externality on slow traders, but the nature of this effect is fundamentally different (shifts in “bargaining power” between fast and slow market participants). Despite this market failure, the socially

optimal level of HFT can be strictly positive and dominate a market with only slow traders. We argue that proposed regulatory changes such as limits on order cancellations or minimum order resting times are likely to lead to sub-optimal outcomes because they precisely address the source of potential efficiency gains derived from HFT. Direct taxation of HFT activity may improve upon rules that lead this type of activity to disappear completely. Furthermore, a number of testable implications regarding the effects of high-frequency trading on order flow statistics in limit order markets are derived. While some are novel to the literature, others are consistent with recent empirical findings.

1 Introduction

While the proverb "time is money" applies to virtually all economic activities, the accelerated proliferation of electronic trading has taken this wisdom to the extreme. High-frequency trading (HFT), a variant of algorithmic trading, relies on sophisticated computer programs for the implementation of trading strategies that involve a vast amount of orders in very small time intervals. Proprietary trading desks, hedge funds and so-called pure-play HFT outlets are investing large sums into human ("Quants") and physical (IT and data) capital in an effort to outpace the competition. Recent estimates suggest that HFTs are now responsible for more than 50% of trading in U.S. equities.¹

These developments are being accompanied by a heated debate among financial economists, practitioners, and regulators about the implications of an increasing computerization of the trading process. While proponents² argue that technology increases market efficiency via improved liquidity and price discovery, critiques³ claim that HFTs make profits at the expense of other (slow) market participants and have the potential to destabilize markets.

This paper contributes to this debate by presenting a stylized model of trading in a limit order market where agents differ in their trading speed, which is thought to capture the difference between (fast) HFTs and (slow) human market participants. We build on the model of Foucault (1999), in which limit orders cannot be revised after submission and thus may become stale due to the arrival of new value-relevant information. This risk of being "picked off" (Copeland and Galai (1983)) by the following trader is a source of inefficiency because a high level of asset price volatility leads agents to choose limit orders with a low execution probability, which implies that some potential gains from trade are not realized. We extend Foucault's model by endowing a proportion α of the trading crowd with a *relative* speed advantage that improves their ability to manage outstanding limit orders compared to the remaining market participants. More specifically, we assume that fast traders (FTs) are able to revise their limit orders after news releases, but only in case the next agent is a slow trader (ST).

We analyze the stationary equilibrium of this dynamic limit order market and compare it with the baseline case of identical traders studied by Foucault (1999). In order to convey the intuition it is convenient to interpret the model as a sequential bargaining process over a total surplus of $2L$, which is the difference in valuations across buyers and sellers. Upon arriving at the market, a trader makes a take-it-or-leave-it offer (limit order) to the following agent, who either accepts this offer (via a market order) or in turn makes another offer to the trader arriving after him. As usual in these situations, each agent's bargaining power is determined by his outside option which here is given by

¹See Financial Times, "High-frequency boom time hits slowdown", April 12, 2011.

²See e.g. Optiver, "High Frequency Trading", Position Paper, 2011, <http://fragmentation.fidessa.com/wp-content/uploads/High-Frequency-Trading-Optiver-Position-Paper.pdf>

³See SEC Chairman Mary Schapiro's speech in front of the Security Traders Association "Remarks Before the Security Traders Association", www.sec.gov/news/speech/2010/spch092210mls.htm.

the expected payoff earned from posting a limit order.

Now the fact that FTs do not face the risk of being picked off when interacting with STs has two opposing effects. First, the ability to revise quotes after news arrivals reduces the inefficiency that arises when agents submit limit orders with a low execution probability. This allows more gains from trade to be reaped. Second, the outside option of posting limit orders is more valuable to FTs than to STs and therefore they require a higher share of the surplus in order to accept an outstanding offer. Hence STs face a dilemma: They can either keep their chances of execution constant by posting more aggressive quotes that are attractive to both STs and FTs, or instead accept a decrease in the execution probability of their limit orders by only targeting STs. While their expected profits from posting limit orders decrease in either case, the latter choice (which is optimal for α sufficiently small) gives rise to an inefficiency because it reduces the chances that the surplus $2L$ is shared.

As the choice between market and limit orders is endogenous in this model, the decrease in STs' outside option indirectly also affects the profits they derive from market orders because they are willing to accept worse quotes than in a market without FTs. Additionally, they are too slow to "pick off" FTs and therefore lose profitable trading opportunities. Because agents' expected utility is a weighted average of profits due to submitting limit and market orders, it follows that STs are always worse off.

This has important implications for welfare, in particular once one drops the assumption that α is exogenous and instead assumes that agents may become fast by investing into HFT technology at a fixed cost. As an (interior) equilibrium requires equal payoffs across slow and fast agents, any investment into HFT leads to a social welfare loss in equilibrium.

We additionally derive several testable implications regarding the effect of HFT on the order flow composition. For example our model predicts that FTs are more likely to be makers (liquidity suppliers) than takers (liquidity consumers) and display a higher ratio of limit to market orders than STs. These implications are consistent with the widespread market making activities of large HFT firms and are also in line with recent empirical findings. We also link make-take and order submission decisions to the traded assets fundamental volatility. Finally, one additional testable prediction from our model is that STs react to HFT entry by submitting limit orders with lower execution probabilities.

Because HFT activity constitutes a negative externality, our results have important implications for policy makers. Currently, a number of different measures are being discussed among regulators that aim to curb HFT activity, including minimum order resting times and limits on the number of messages that individual traders can send.⁴ While these measures could eventually improve on the market outcome by eliminating HFT activity, the present analysis suggests that they would also

⁴See e.g. Financial Times, "Berlin forges ahead with trading controls", September 25th, 2012.

deprive the market of potential efficiency gains because it is precisely FTs' ability to revise their limit orders quickly that allows more gains from trade to be realized. It may be more effective to tax HFT activities directly in order to ensure that at least some of the potential benefits of speed are reaped.

The literature on algorithmic and high-frequency trading has grown substantially in the past couple of years (see e.g. the surveys by Biais and Woolley (2012) and Foucault (2012)). Most closely related to this work is the paper by Biais et al. (2012) which studies the impact of HFT in a Glosten and Milgrom (1985) framework. In their model FTs have a higher chance of finding trading opportunities than slow market participants which increases the likelihood that gains from trade are realized. But at the same time, FTs are a source of adverse selection because they possess private information, which raises the bid-ask spread payable by everyone and therefore reduces trade. If the proportion of FTs is endogenized, this negative externality implies that investment in HFT is excessive in equilibrium just as in this paper. While both papers share this result, the nature of the underlying externality is very different. Different from Biais et al. (2012), FTs are not a *source* of adverse selection in our model but instead help to *reduce* adverse selection by being able to revise limit orders after news arrivals. While this dampens an existing inefficiency, it constitutes a loss in bargaining power for slow market participants and leads them to submit limit orders with a lower execution probability.

Also closely related, Jovanovic and Menkveld (2012) study competitive middlemen that intermediate between early limit order traders and late market order traders. As in this study, HFTs' speed advantage may reduce adverse selection by updating quotes quickly and therefore increase trade. On the other hand, HFTs' ability to process (hard) information quickly can introduce an additional adverse selection problem that reduces trade. A calibration exercise reveals a slight increase in welfare.

A number of other papers study HFT from a theoretical perspective. Cartea and Penalva (2011) propose a model where their increased speed allows HFT to impose a haircut on liquidity traders, which increases trading volume and price volatility, but lowers the welfare of liquidity traders. Foucault et al. (2013) study the trading strategy of an informed trader that is able to react faster than others to news. They conclude that this speed advantage makes the informed trader's order flow more volatile and increases his relative share in trading volume. Rosu and Martinez (2012) study HFTs as strategic informed traders who instantaneously react to new information and ensure that it is reflected in prices immediately. Finally, Pagnotta and Philippon (2012) provide a model where competing exchanges invest into speed and compete for investors. They provide conditions under which competition, fragmentation and speed can improve or reduce welfare.

Several studies empirically examine the impact of algorithmic and high-frequency trading on market quality. In summary, this stream of the literature concludes that automated trading strategies improve liquidity (Hendershott et al. (2011), Oeck and Saar (2011)), are highly profitable (e.g.

Menkveld (2012)), and contribute to price discovery (e.g. Hendershott and Riordan (2012), Hendershott et al. (2012)). In line with our findings on the effects of speed in a limit order market, Hendershott and Riordan (2012) find that algorithmic traders "supply liquidity when it is expensive and consume liquidity when it is cheap". Hagströmer and Norden (2013) and Malinova et al. (2012) use order level data and present evidence that is consistent with HFTs primarily engaging in market making activities. Moallemi and Sağlam (2011) provide some empirical estimates of the "cost of latency" and find a dramatic increase between 1995 and 2005. Chaboud et al. (2013) study computer- and human-generated order flow in the FX market and conclude that the trading strategies of automated traders are more correlated with each other than those of human market participants. Kirilenko et al. (2011) examine the recent "flash crash" in U.S. equity markets and find that HFT may have exacerbated volatility during this brief liquidity crisis, although they are not to blame for the crash itself. While most empirical studies are only restricted to a single market, Boehmer et al. (2012) examine a sample of 39 exchanges around the world. In sum, they confirm the view that algorithmic trading has a positive effect on liquidity and price efficiency, but also find that it increases volatility (this effect is not due to improved price efficiency).

This paper is organized as follows. Section 2 provides an outline of the model, whose solution is presented subsequently in Section 3. Section 4 details the implications for trading profits and the impact of HFT on trading volume as well as other statistics of the order flow. The proportion of fast traders is endogenized in Section 5, followed by a discussion and conclusion. All proofs are relegated to the appendix.

2 The model

2.1 The limit order market

We consider an infinite-horizon⁵ version of Foucault's (1999) dynamic limit order market. There is a single risky asset whose fundamental value follows a random walk, i.e.

$$v_t = v_{t-1} + \varepsilon_t$$

where the innovations are i.i.d. and can take values of $+\sigma$ and $-\sigma$ with equal probability. At each time point $t = 1, 2, \dots$ a new trader arrives to the marketplace in order to buy or sell the asset. In this model, trading arises due to differences in traders' private values for the asset. Specifically, we

⁵This assumption is merely for convenience as it simplifies the algebra. Foucault (1999) assumes that the terminal date is stochastic, as the trading process stops after each period with constant probability $1 - \rho > 0$. An infinite horizon may be interpreted as the limiting case where $\rho \rightarrow 1$.

assume that at time t' , a trader arriving at time $t \leq t'$ values the asset at

$$R_{t'} = v_{t'} + y_t$$

which is the sum of the asset's current fundamental value and the time-invariant private valuation y_t . We assume that this private valuation can take two values $y_h = +L$ and $y_l = -L$ with equal probability, where $L > 0$. Moreover, the private valuations are i.i.d. across traders and independent of the asset value innovations. We call traders with a high (low) private valuation buyers (sellers).

Trading takes place in a limit order market and we assume for tractability that the order size is fixed at one unit and limit orders expire after one period. This implies that at each point in time the limit order book either a) contains an ask quote, b) a bid quote or c) is empty. Let A_t^m and B_t^m denote the best quotes at time t , where we write $A_t^m = \infty$ ($B_t^m = -\infty$) if there is no ask (bid) quote in the respective limit order book.

Unlike Foucault (1999), we assume that there are two types of agents. Agents may be either fast traders (FTs) or slow traders (STs), and the proportion of FTs is denoted by α . It is well-known that limit orders face the risk of being "picked off" because the arrival of new information may render them stale and thereby grant a free trading option to others (see e.g. Copeland and Galai (1983)). This is also sometimes termed the winner's curse because limit orders are more likely to be executed after unfavourable price movements than after favourable ones. In this context, being fast can be extremely valuable because it may allow for a timely cancellation or revision of limit orders and therefore significantly reduces the risk of being adversely selected. In order to model this advantage in the most parsimonious way we assume that FTs can revise (or update) their limit orders after the arrival of new information (i.e. the realization of ε_{t+1}) yet before the arrival of the next agent, but only provided he is a ST. If the next trader is a FT as well, the order cannot be revised. STs can never cancel their orders (as in Foucault (1999)). Notice that this implies that being fast is a purely *relative* advantage, it is only valuable if there is someone slow in the market. Moreover, it implies that the model has the same solution for all agents being slow ($\alpha = 0$) and all agents being fast ($\alpha = 1$). While we take α as exogenous for most of our analysis, Section 5 endogenizes investment into trading speed, similar in spirit to Biais et al. (2012).

It is important to note that the assumption that a FT can revise his limit order if and only if the next agent is a ST does not require that he actually knows the next agent's type. To see this, suppose that there is an outstanding limit order and ε is revealed. The news arrival triggers a race between the limit order trader and the next arriving agent, who both rush to be the first mover (the former to revise the limit order and the latter to execute it via a market order). In this context, the assumption made in Foucault (1999) is that the next arriving agent always wins this race, and the natural extension in our setting is that the next arriving agent wins *unless* the limit order trader is faster than him.

2.2 Payoffs and strategies

For ease of exposition, we will henceforth assume that the limit order trader is always a buyer (with private valuation y_h) and the market order trader is always a seller (with private valuation y_l). This is without loss of generality because the symmetric arguments apply to limit order sellers and market order buyers. Additionally, and for this subsection only, we will assume that sellers (buyers) do never buy (sell) the asset. While it turns out that this is always the case in equilibrium⁶, assuming it for the moment greatly simplifies the presentation here.

Now consider a seller who enters the market at time t and let $V_{t,k}^{LO}$ be the expected profit he would obtain when choosing to post a limit order, where $k \in \{ST, FT\}$ refers to his type. Clearly, he will opt for a market sell order if the best available bid price B_t^m is such that

$$B_t^m - (v_t - L) \geq V_{t,k}^{LO} \quad k \in \{ST, FT\} \quad (1)$$

In other words, the expected profits obtained from posting limit orders constitute an endogenous outside option when deciding upon whether or not to submit a market order. Thus the seller's order choice is entirely determined by whether or not the best available bid is above his sell cutoff price, $\hat{B}_k(v_t)$, which is the bid price at which he is indifferent between submitting a limit order or a market order

$$\hat{B}_k(v_t) = (v_t - L) + V_{t,k}^{LO} \quad k \in \{ST, FT\} \quad (2)$$

When placing limit orders, traders generally face the tradeoff that more aggressive limit orders imply a higher probability of execution but a lower profit conditional on execution. Because a limit order placed at time t can only be executed at $t + 1$, its execution does not only depend on the type of trader that arrives in the next period but also on the realization of the asset value innovation ε_{t+1} . Therefore, the execution probability is pinned down by the ordering of cutoff prices, which in principle need not be fixed across different parametrizations. For example, suppose that the cutoff prices of sellers arriving at time $t + 1$ are ordered as $\hat{B}_{ST}(v_t - \sigma) \leq \hat{B}_{FT}(v_t - \sigma) \leq \hat{B}_{ST}(v_t + \sigma) \leq \hat{B}_{FT}(v_t + \sigma)$. Then, the unconditional execution probability $p(B)$ of a ST's buy limit order with bid price B is given by

$$p(B) = \begin{cases} 0 & \text{if } B < \hat{B}_{ST}(v_t - \sigma) \\ \frac{(1-\alpha)}{4} & \text{if } \hat{B}_{ST}(v_t - \sigma) \leq B < \hat{B}_{FT}(v_t - \sigma) \\ \frac{1}{4} & \text{if } \hat{B}_{FT}(v_t - \sigma) \leq B < \hat{B}_{ST}(v_t + \sigma) \\ \frac{2-\alpha}{4} & \text{if } \hat{B}_{ST}(v_t + \sigma) \leq B < \hat{B}_{FT}(v_t + \sigma) \\ \frac{1}{2} & \text{if } \hat{B}_{FT}(v_t + \sigma) \leq B \end{cases} \quad (3)$$

⁶More precisely, a buyer that decides to post a limit order at time t will always choose the bid price such that the probability of execution conditional on a buyer arriving in period $t + 1$ is equal to zero.

Because $p(B)$ is an increasing step function, there are multiple bid prices with the same execution probability. Then, optimality implies that the bid price is chosen among the set of cutoff prices $\mathcal{B} = \{\hat{B}_{ST}(v_t - \sigma), \hat{B}_{FT}(v_t - \sigma), \hat{B}_{ST}(v_t + \sigma), \hat{B}_{FT}(v_t + \sigma)\}$. Thus the objective function of a ST that decides to submit a limit order can be written as

$$V_{t,ST}^{LO} = \max_{B_{t,ST} \in \mathcal{B}} \{p(B_{t,ST})(v_t + E_{Ex}[\varepsilon_{t+1}] + L - B_{t,ST})\} \quad (4)$$

where $E_{Ex}[\cdot]$ denotes expectation conditional on execution. The decision problem of a FT opting for limit orders is slightly more complex because he may revise his limit order upon the realization of ε_{t+1} conditional on the next trader being a ST. Hence they effectively choose a tuple of three bid prices $(B_{t,FT}, B_{t,FT}^{+\sigma}, B_{t,FT}^{-\sigma})$. Now let $q_{|k, \varepsilon_{t+1}}(B)$ denote the execution probability of a FT's limit order with bid price B conditional on the next trader being of type $k \in \{ST, FT\}$ and the asset value innovation being $\varepsilon_{t+1} \in \{+\sigma, -\sigma\}$. Then a FT's objective function can be written as

$$V_{t,FT}^{LO} = \max_{B_{t,FT}, B_{t,FT}^{+\sigma}, B_{t,FT}^{-\sigma} \in \mathcal{B}} \left\{ \begin{array}{l} \alpha q_{|FT}(B_{t,FT})(v_t + E_{Ex}[\varepsilon_{t+1}] + L - B_{t,FT}) \\ + (1 - \alpha) q_{|ST, +\sigma}(B_{t,FT}^{+\sigma})(v_t + \sigma + L - B_{t,FT}^{+\sigma}) \\ + (1 - \alpha) q_{|ST, -\sigma}(B_{t,FT}^{-\sigma})(v_t - \sigma + L - B_{t,FT}^{-\sigma}) \end{array} \right\} \quad (5)$$

It is easy to see that our assumption on FTs' ability to revise their limit orders implies that $B_{t,FT}^{+\sigma}$ and $B_{t,FT}^{-\sigma}$ are set in perfect knowledge about both the next trader's type and the realization of ε_{t+1} . Hence they must be optimally chosen to be equal to a ST's cutoff price at $t + 1$.

Lemma 1 *In equilibrium, FTs' revised bid quotes are given by*

$$B_{t,FT}^{-\sigma*} = \hat{B}_{ST}^*(v_t - \sigma) \text{ and } B_{t,FT}^{+\sigma*} = \hat{B}_{ST}^*(v_t + \sigma).$$

Proof. See the Appendix. ■

Moreover, the initial quote $B_{t,FT}$ is de-facto only accessible to FTs which implies that we must have $B_{t,FT} \in \mathcal{B}' = \{\hat{B}_{FT}(v_t - \sigma), \hat{B}_{FT}(v_t + \sigma)\}$. Hence we can rewrite a FT's objective function as

$$V_{t,FT}^{LO} = \max_{B_{t,FT} \in \mathcal{B}'} \{\alpha q_{|FT}(B_{t,FT})(v_t + E_{Ex}[\varepsilon_{t+1}] + L - B_{t,FT})\} \quad (6)$$

where

$$q_{|FT}(B) = \begin{cases} 0 & \text{if } B < \hat{B}_{FT}(v_t - \sigma) \\ \frac{1}{4} & \text{if } \hat{B}_{FT}(v_t - \sigma) \leq B < \hat{B}_{FT}(v_t + \sigma) \\ \frac{1}{2} & \text{if } \hat{B}_{FT}(v_t + \sigma) \leq B \end{cases} \quad (7)$$

As in Foucault (1999) and Colliard and Foucault (2012), we focus on stationary Markov-perfect equilibria, which is natural because traders' profits do not depend on the history of the game but

only on the state of the market upon their arrival. The equilibrium is found by simultaneously solving equations (2), (4) and (6) under the condition that the ordering of cutoff prices used for obtaining the execution probability $p(B)$ is satisfied. The following Lemma states that there is a unique ordering of sellers' cutoff prices in equilibrium, which significantly facilitates its computation.

Lemma 2 *In equilibrium, $\hat{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t - \sigma) \leq \hat{B}_{ST}^*(v_t + \sigma) \leq \hat{B}_{FT}^*(v_t + \sigma)$*

Proof. See the Appendix. ■

Lemma 2 implies (via equation (2)) that the endogenous outside option of posting limit orders is more valuable for FTs than for STs, which is natural given their ability to revise their limit orders conditional on the arrival of a ST in the next period. Additionally, the advantage of being fast is directly related to the severity of the picking off risk, σ , and therefore ensures a unique ordering. Importantly, in the limit when $\sigma \rightarrow 0$, being fast is not beneficial anymore because limit orders are no further exposed to the risk of being picked off.

3 Equilibrium

As in Colliard and Foucault (2012), it is straightforward to categorize agents' strategies by their respective limit order execution probabilities. Let $p_{|k, \varepsilon_{t+1}}(B)$ denote the execution probability of a ST's limit order with bid price B conditional the type of the next trader and the asset value innovation. We say that a ST uses a high fill-rate strategy if his quote has a strictly positive execution probability for each possible realization of ε_{t+1} , that is $p_{|\varepsilon_{t+1}}(B_{ST}^*) > 0$ for all ε_{t+1} . Otherwise the strategy is said to have a low fill-rate. FTs' initial quotes B_{FT}^* are categorized into a low fill-rate and a high fill-rate strategy in exactly the same fashion. Moreover, we say that a ST uses a specialized strategy if $p_{|FT, \varepsilon_{t+1}}(B_{ST}^*) \neq p_{|ST, \varepsilon_{t+1}}(B_{ST}^*)$ for some $\varepsilon_{t+1} \in \{+\sigma, -\sigma\}$, that is the limit order is not attracting an execution from both types of agents for at least one possible realization of ε_{t+1} . Notice that because FTs de facto face a binary choice, there is no need to categorize their strategies further. Equipped with this typology of strategies, we can now state the following.

Proposition 1 *For fixed parameters (α, σ, L) , there exists a unique Markov-perfect equilibrium in the limit order market. In equilibrium*

- a) *STs use a high fill-rate strategy for $\sigma < \sigma_{ST}^*(\alpha)$ and a low fill-rate strategy otherwise.*
 - b) *STs use a specialized strategy for $\alpha < \alpha_S^*(\sigma)$ and an unspecialized strategy otherwise.*
 - c) *FTs use a high fill-rate strategy for $\sigma < \sigma_{FT}^*(\alpha)$ and a low fill-rate strategy otherwise.*
- The definitions of $\sigma_{FT}^*(\alpha)$, $\sigma_{ST}^*(\alpha)$ and $\alpha_S^*(\sigma)$ are provided in the Appendix.*

Proof. See the Appendix. ■

It is straightforward to interpret Proposition 1 in the light of a trade-off between execution probability and expected profit conditional on execution. In line with Foucault (1999), parts a) and c) state that a high (low) level of volatility induces agents to post limit orders with a low (high) execution probability. Intuitively, larger innovations imply a more severe adverse selection risk for a limit order trader, and the natural reaction is then to protect himself from unfavourable price movements by posting less aggressive limit orders. Part b) is similarly intuitive: For a low level of α , it is not very attractive for STs to target FTs with their limit orders because this leads only to a small increase in execution probabilities but requires considerably more aggressive quotes (due to a more valuable outside option). Hence they only use an unspecialized strategy when α is sufficiently large and otherwise partially shade their orders.

Figure 1 in the Appendix depicts the functions $\sigma_{FT}^*(\alpha)$, $\sigma_{ST}^*(\alpha)$ and $\alpha_S^*(\sigma)$ in the (α, σ) -space, where we have set $L = 1$ (this is without loss of generality as only the ratio of σ and L is relevant). One can see that we have $\sigma_{FT}^*(\alpha) = \sigma_{ST}^*(\alpha)$ for $\alpha \geq \alpha_S^*(\sigma)$ and $\sigma_{FT}^*(\alpha) < \sigma_{ST}^*(\alpha)$ otherwise, which implies that there are 5 distinct types of equilibria in total. To see why this is the case, first consider the situation where $\alpha > \alpha_S^*(\sigma)$ such that it is optimal for STs to use an unspecialized strategy. It is easy to see that in this case both STs and FTs face the same choices, which is to pick a bid quote from the restricted set $\mathcal{B}' = \{\hat{B}_{FT}(v_t - \sigma), \hat{B}_{FT}(v_t + \sigma)\}$. Hence both types of traders use the same strategy, and accordingly there are two types of equilibria (one with a high fill-rate and another one with a low fill-rate). On the other hand, STs optimally use a specialized strategy for $\alpha < \alpha_S^*(\sigma)$, which implies that both types of traders face different choices (FTs choose from \mathcal{B}' , STs choose from $\mathcal{B}'' = \{\hat{B}_{ST}(v_t - \sigma), \hat{B}_{ST}(v_t + \sigma)\}$) and consequently the level of σ for which they are indifferent between a high fill-rate and a low fill-rate strategy does not coincide such that we have three different types of equilibria for $\alpha < \alpha_S^*(\sigma)$. One may show that $\sigma_{FT}^*(\alpha) < \sigma_{ST}^*(\alpha)$ directly follows from the fact that $\hat{B}_{FT}^*(v_t) > \hat{B}_{ST}^*(v_t)$.

The main goal of our analysis is to examine how trading speed affects the order flow, trading profits, and welfare by comparing equilibrium outcomes arising from parameters (α, σ, L) , where $\alpha > 0$, with those arising in the absence of FTs, i.e the equilibrium outcomes for $(0, \sigma, L)$. Unfortunately, the fact that ε is assumed to be both discrete and bounded (which is necessary to obtain a closed-form solution) complicates this endeavour slightly because it restricts traders to choose between a low and a high fill-rate and thus gives rise to bang-bang solutions. If agents are close to being indifferent between those two strategies, a slight change in α can lead to a sharp change in the equilibrium execution probability of limit orders and thereby induce a move from one extreme outcome (low trade) to another (high trade). In order to avoid complications stemming from these discontinuities, we will frequently assume that σ is not "too close" to the level where STs are indifferent between

a high fill-rate and a low fill-rate strategy in the absence of FTs, such that a change in α cannot trigger a switch from an equilibrium with little trade to an equilibrium with a lot of trade or vice versa. In order to formalize this assumption, we define $\underline{\sigma} \equiv \min_{\alpha} \sigma_{FT}^*(\alpha)$ and $\bar{\sigma} \equiv \max_{\alpha} \sigma_{ST}^*(\alpha)$ and make the following assumption throughout the remainder of the paper

Assumption 1 (technical) $\sigma \in \Sigma \equiv \underline{\Sigma} \cup \bar{\Sigma}$ where $\underline{\Sigma} = [0, \underline{\sigma})$ and $\bar{\Sigma} = [\bar{\sigma}, \infty)$.

Figure 1 also depicts the excluded interval $[\underline{\sigma}, \bar{\sigma})$ in the (α, σ) -space. As can be seen, this assumption effectively rules out the equilibrium where STs use a specialized high fill-rate strategy and FTs use a low fill-rate strategy, such that we are left with 4 distinct types of equilibria. While this may seem restrictive at first glance, it turns out that this equilibrium is very similar to the equilibrium where STs follow the same strategy but FTs choose a high fill-rate strategy instead; as there are very few FTs in either case ($\alpha < \alpha_S^*(\sigma)$), the difference between both equilibria is of second order. In the following, we refer to the different types of equilibria by the strategy chosen by STs for parsimony. We abbreviate the unspecialized low (high) fill-rate equilibrium as ULFR (UHFR) and the specialized low (high) fill-rate equilibrium as SLFR (SHFR). While only a subset of our results relies on Assumption 1, we retain it throughout the text as it additionally helps to simplify the exposition.⁷

4 Trading profits and the order flow

4.1 The outside option

The limit order market analyzed here can be seen as a sequential bargaining game over a total surplus of $2L$, where one agent posts a take-it-or-leave-it offer (a limit order) to the next agent, who then can either accept this offer (via a market order) or instead make another offer (via limit order) to the agent arriving one period later. As usual in these situations, each agent's bargaining power is then determined by the size of his outside option. Now because FTs face a reduced risk of being picked off due to their ability to cancel some of their limit orders, the alternative of posting a limit order is more valuable to them than to STs (see Lemma 2). This directly implies that they are able to extract a higher share of the total surplus at the expense of STs, whose outside option value, V_{ST}^{LO*} , accordingly deteriorates compared to the situation without fast agents (denoted by V_0^{LO*} , with associated limit order execution probability p_0^*)⁸.

⁷In order to verify that Assumption 1 does not distort the major economic tradeoffs in the model we alternatively solve the model numerically under the assumption of normally distributed asset value innovations. The resulting conclusions are qualitatively in line with those derived here in closed form.

⁸We henceforth adopt the convention that the subscript 0 denotes equilibrium quantities for $\alpha = 0$.

Corollary 1 *FTs' (STs') limit order profits are always above (below) those obtained by agents in a market with only STs, that is we have $V_{FT}^{LO*} > V_0^{LO*} > V_{ST}^{LO*}$ for all $\alpha \in (0, 1)$. Moreover, STs submit limit orders with a lower execution probability than in the absence of FTs, i.e. $p^* \leq p_0^*$ for all $\alpha \in (0, 1)$.*

Proof. See the Appendix. ■

Effectively, the presence of FTs creates a dilemma for STs, which is best seen by starting from the situation where all agents are slow (i.e. $\alpha = 0$). Now an increase in α presents STs with two possible choices: Either they keep the execution probability of their limit orders constant at $p^* = p_0^*$ by targeting both STs and FTs (i.e. use an unspecialized strategy) at the expense of a higher bid price, or alternatively they accept a lower execution probability $p^* < p_0^*$ by posting limit orders that are at least in some cases only accepted by STs (a specialized strategy). It turns out that in either case, they are worse off compared to the situation without FTs. The above Corollary offers an interesting empirical implication: The proliferation of HFT leads slower market participants to submit limit orders with a lower execution probability. To our knowledge, this prediction has not been tested in the literature so far.

In order to gain some more intuition, it is instructive to examine how FTs' advantage at posting limit orders, $\delta^* \equiv V_{FT}^{LO*} - V_{ST}^{LO*}$, varies with α and σ . Clearly, in this model the value of being fast directly derives from avoiding the winner's curse, and therefore it comes at no surprise that δ^* is increasing in σ for all α . However, the relationship between δ^* and α is slightly more complex. First, it is easy to see that an increase in α raises the exposure of FTs' limit orders to the risk of being picked off as they can only revise their quotes when interacting with STs. Hence we have $\partial V_{FT}^{LO*} / \partial \alpha < 0$ for all α . Now consider what happens to STs' outside option as the proportion of FTs in the market increases. On the one hand, a larger α increases the likelihood to interact with a market participant that has a higher bargaining power, which lowers V_{ST}^{LO*} as a higher share of the surplus $2L$ is ceded to the trading counterparty on average. On the other hand, an increase in α simultaneously hurts the profitability of FTs' limit orders and therefore in turn also benefits STs. Because the first (second) effect dominates if α is sufficiently small (large), V_{ST}^{LO*} is first decreasing and then increasing in α . As a result, δ^* may actually be increasing in α provided there are enough STs in the market. A sufficient condition for this is that V_{FT}^{LO*} does not decrease too quickly in α , which is the case when the risk of being picked off is not too severe, i.e. if σ is sufficiently small.

Corollary 2 *δ^* is increasing in σ for all $\alpha \in (0, 1)$. Moreover, there exists a threshold $\tilde{\sigma}$ such that δ^* is decreasing in α for all $\alpha \in (0, 1)$ and $\sigma > \tilde{\sigma}$.*

Proof. See the Appendix. ■

4.2 The order flow

We next turn to the analysis of the order flow composition, i.e. the equilibrium mix between limit and market orders. Consider the following four mutually exclusive events that may occur on the equilibrium path. The arriving agent can be 1) a ST submitting a limit order, 2) a ST submitting a market order, 3) a FT submitting a limit order, or 4) a FT submitting a market order. Now let $\varphi^* = (\varphi_{ST}^{LO*}, \varphi_{ST}^{MO*}, \varphi_{FT}^{LO*}, \varphi_{FT}^{MO*})$ denote the stationary probability distribution of the above market events in equilibrium, where we naturally have $\varphi_{ST}^{LO*} + \varphi_{ST}^{MO*} = 1 - \alpha$ and $\varphi_{FT}^{LO*} + \varphi_{FT}^{MO*} = \alpha$. Then one can show the following.

Proposition 2 *The stationary probability distribution of equilibrium events is given by*

$$\begin{aligned}\varphi_{ST}^{LO*} &= \frac{1-\alpha}{\chi} [1 - \alpha(\frac{1}{2} - q_{|FT}^*)] \\ \varphi_{ST}^{MO*} &= \frac{1-\alpha}{\chi} [\chi - 1 + \alpha(\frac{1}{2} - q_{|FT}^*)] \\ \varphi_{FT}^{LO*} &= \frac{\alpha}{\chi} [1 + (1-\alpha)(p_{|ST}^* - p_{|FT}^*)] \\ \varphi_{FT}^{MO*} &= \frac{\alpha}{\chi} [\chi - 1 - (1-\alpha)(p_{|ST}^* - p_{|FT}^*)]\end{aligned}$$

where $\chi = (1 + \alpha q_{|FT}^*)(1 + (1-\alpha)p_{|ST}^*) - \frac{\alpha(1-\alpha)}{2} p_{|FT}^*$

Proof. See the Appendix. ■

Given this probability distribution, we may define the trading rate (or per-period expected trading volume) as the unconditional probability of observing a trade in a given period, that is

$$TR^* \equiv \varphi_{ST}^{MO*} + \varphi_{FT}^{MO*} \quad (8)$$

Similarly one can define the trader type-specific trading rates as $TR_{ST}^* \equiv \varphi_{ST}^{MO*}/(1-\alpha)$ and $TR_{FT}^* \equiv \varphi_{FT}^{MO*}/\alpha$. In comparison to a market that is solely populated by STs, the existence of FTs affects the trading rate in two distinct ways.

First, FTs' ability to revise their limit orders when trading with STs can help to reduce the inefficiency that is rooted in the adverse selection problem faced by limit orders. To see this, begin by considering the case where all agents are slow and suppose that $\sigma \in \bar{\Sigma}$. In this situation it is optimal for agents to submit limit orders with a low fill-rate (execution probability equal to 1/4) because the risk of being picked off is too severe. This implies that some potential gains between buyers and sellers are not realized because limit order traders post cautious quotes in an effort to protect themselves against adverse price movements. Now compare this with the case where some,

but not all, traders are fast (i.e. $0 < \alpha < 1$). Because FTs may revise their limit orders after the materialization of new information conditional on a ST arriving in the next period, the winner's curse is eliminated in these situations and more gains from trade can be reaped. Importantly, this efficiency gain diminishes as α approaches unity because being fast allows to avoid being picked off only when facing slow traders.

Second, the presence of FTs negatively affects STs' outside option (see Corollary 1). As discussed in Section 3.1, the latter may react to this loss in bargaining power by either *i*) stepping up the aggressiveness of their quotes (thus maintaining a constant probability of execution) or *ii*) accepting a lower execution probability. While their expected profits decline in either case, the second choice (which is optimal for $\alpha < \alpha_S^*(\sigma)$) introduces an inefficiency by reducing the likelihood that the total surplus of $2L$ is shared. In contrast, the first option only affects the way the surplus is shared.

Now the total impact of having FTs in the market on the trading rate depends on which of the two effects dominates. It is easy to see that the first effect is absent for $\sigma \in \underline{\Sigma}$ because in this case σ is sufficiently low to induce traders to submit limit orders with a high execution probability. Hence there is no inefficiency to begin with and as a consequence the presence of FTs implies a drop in the trading rate if α is sufficiently low to induce STs to use a specialized strategy. Conversely, a high level of volatility, σ , leads the first effect to dominate and thus allows more gains from trade to be realized.

Corollary 3 *The presence of FTs increases trading volume except in a specialized high fill-rate (SHFR) equilibrium, i.e. $TR^* < TR_0^*$ for $\sigma \in \underline{\Sigma}$ and $\alpha < \alpha_S^*(\sigma)$ and $TR^* \geq TR_0^*$ otherwise.*

Proof. See the Appendix. ■

There are several other papers that make predictions about the impact of high-frequency trading on trading volume. Biais et al. (2012) develop a model where FTs are more likely to locate trading opportunities (which raises trading volume) but at the same time possess private information and therefore create adverse selection (which reduces trading volume). Focussing on the Pareto-dominant equilibrium, they show that a sufficiently *large* proportion of FTs can reduce trade in comparison to a market with only STs. In the present model, on the other side, only a sufficiently *small* level of α leads to a decrease in trading volume as it induces STs to submit limit orders with a lower execution probability. For α sufficiently large, trading volume always increases compared to the case where $\alpha = 0$. Similarly, Jovanovich and Menkveld (2012) present a model where the introduction of competitive HFT middlemen can have either a positive or a negative effect on trading volume in a limit order market. While greater speed allows intermediaries to avoid the winner's curse (similar in spirit to this paper), HFTs ability to process (hard) information quickly (e.g. real-time datafeeds on index futures) introduces an additional adverse selection problem that did not affect trade in their

absence. The associated empirical evidence based on HFT entry into the market for Dutch equities is mixed as it led to an increase in trade frequency, but a reduction in trading volume. Cartea and Penalva (2011) suggest that HFT increases trading volume mechanically because more and more trades get intermediated. While the accelerated increase in trading volume over the past decades (see e.g. Chordia et al. (2011)) has been accompanied by the advent of HFT, more empirical work based on natural experiments is needed to establish causality.

Besides the trading rate, the stationary probability distribution of market events, φ^* , allows us to analyze how the presence of FTs affects the equilibrium mix of limit and market orders. This is particularly relevant because up to date there is little theoretical work that connects high-frequency trading to specific properties of the order flow, but at the same time the availability of detailed data has already allowed empirical researchers to examine this relationship to some extent, as discussed further below.

To this end, we define the limit-to-market order ratio, LtM^* , as the average number of limit orders per period divided by the average number of market orders. Taking care of the fact that FTs revise their limit orders in case they are followed by a ST (i.e. with probability $(1 - \alpha)$), we have

$$\begin{aligned} LtM^* &= \frac{\varphi_{ST}^{LO*} + \varphi_{FT}^{LO*}(2 - \alpha)}{\varphi_{ST}^{MO*} + \varphi_{FT}^{MO*}} \\ &= \frac{\varphi_{ST}^{MO*}}{\varphi_{ST}^{MO*} + \varphi_{FT}^{MO*}} LtM_{ST}^* + \frac{\varphi_{FT}^{MO*}}{\varphi_{ST}^{MO*} + \varphi_{FT}^{MO*}} LtM_{FT}^* \end{aligned}$$

where $LtM_{ST}^* = \varphi_{ST}^{LO*}/\varphi_{ST}^{MO*}$ and $LtM_{FT}^* = (2 - \alpha)\varphi_{FT}^{LO*}/\varphi_{FT}^{MO*}$ denote the trader type-specific ratios.

Corollary 4 *We have $LtM_{FT}^* > LtM_{ST}^*$ and $LtM^* > LtM_0^*$ for all $\alpha \in (0, 1)$. Moreover, LtM^* and LtM_{FT}^* are increasing in σ for all $\alpha \in (0, 1)$. If $\alpha > \max_{\sigma} \alpha_S^*(\sigma)$, then LtM_{ST}^* is increasing in σ for $\alpha \leq (\sqrt{41} - 5)/2$ and decreasing in σ otherwise.*

Proof. See the Appendix. ■

The fact that FTs display a higher limit-to-market order ratio than STs comes at no surprise. Because of their higher outside option, FTs will reject some quotes that STs find acceptable and are hence more likely to submit limit orders. Moreover, their ability to revise some of their outstanding limit orders further increases LtM_{FT}^* . While there are relatively few empirical studies that are able to trace the order submission strategies of individual traders, the existing evidence is consistent with the above prediction. Hagströmer and Norden (2013) examine trading in Swedish stocks and document that that HFTs display a order-to-trade ratio that is roughly 5 times as large as that of nonHFT. Similarly in spirit, Malinova et al. (2012) find that HFTs generate around 80% of all

messages but are only involved in 25% of all trades in a dataset containing detailed message traffic from the Toronto Stock Exchange. We also find that the aggregate limit-to-market order ratio is higher with FTs than without them. Even though FTs may enable more trades to occur, this is due to their ability to revise limit orders such that in sum $LtM^* > LtM_0^*$. This is consistent with the enormous growth in message traffic in electronic markets as documented e.g. in Hendershott et al. (2011).

Because an increase in volatility implies a higher risk of adverse selection, agents protect themselves against the winner's curse by posting limit orders with lower execution probabilities, which in turn lowers (increases) the chance of observing a market (limit) order. While this intuition holds for LtM^* and LtM_{FT}^* , we find that LtM_{ST}^* may actually be decreasing in σ if there are sufficiently many FTs in the market. In order to see the intuition behind this result, suppose that α is very close to unity. This implies that most limit orders stem from FTs, and therefore a slow seller will accept virtually every buy limit order he encounters when arriving at the market. Hence any variation in LtM_{ST}^* is mainly driven by the likelihood of finding a non-empty limit order book. It is easy to see that the chance of encountering a trading opportunity must be higher in a low fill-rate equilibrium because FTs are more likely to submit limit orders and therefore LtM_{ST}^* is decreasing in σ if α is sufficiently large.

While the LtM ratios are concerned with pure order submissions, many limit orders do ultimately not result in trades. Therefore it is interesting to restrict the analysis to actual trades in order to examine whether a particular trader type is more likely to act as a maker (i.e. liquidity provider) or as a taker (i.e. liquidity consumer). Hence we define the make-take ratio as the probability of trading via a limit order divided by the probability of trading via a market order (for a given trader type), that is

$$MT_k^* = \frac{\omega_{k,ST}^* + \omega_{k,FT}^*}{\omega_{ST,k}^* + \omega_{FT,k}^*}$$

where $\omega_{k,j}^*$ denotes the equilibrium probability of observing a trade that is due to a market order from a type- j trader executing against the limit order from a type- k trader, with $j, k \in \{ST, FT\}$. Because transactions between two traders of the same type are neutral for this breakdown, it is easy to see that we have $MT_{FT}^* \geq 1 \geq MT_{ST}^*$ if and only if $\omega_{FT,ST}^* \geq \omega_{ST,FT}^*$, i.e. a FT (ST) will only display a make-take ratio of more (less) than 1 if he is more (less) likely to supply liquidity to STs (FTs) than to consume liquidity from them.

It turns out that this condition is always satisfied, which is the combined result of two effects. First, FTs enjoy the maximal limit order execution probability of 1/2 conditional on the arrival of a ST as their ability to revise quotes in the light of new information has eliminated the winner's curse. Second, FTs are less likely than STs to submit market orders when arriving to the market because some quotes that a ST finds worth accepting may be rejected by a FT (due to his higher outside option). This is easily verified by looking at the stationary probability distribution in Proposition

2.

Corollary 5 *FTs are always more likely to trade via limit order than STs, that is we have $MT_{FT}^* \geq 1 \geq MT_{ST}^*$ for all $\alpha \in (0, 1)$. Moreover, MT_{ST} is decreasing in σ and MT_{FT} is increasing in σ for all $\alpha \in (0, 1)$.*

Proof. See the Appendix. ■

The prediction that FTs are more likely to act as takers than as makers mirrors the fact that most of the well-known HFT firms (e.g. Getco, Knight, Citadel, Optiver, etc.) are engaged in substantial market making activities. Moreover it is roughly consistent with the empirical evidence available so far. Menkveld (2012) studies a large HFT that takes part in almost 15% of all transactions in Dutch stocks traded on Euronext and Chi-X and trades via limit order roughly 80% of the time. Similarly, Hagströmer and Norden (2013) and Malinova et al. (2012) provide evidence that is consistent with HFT mainly acting as makers in Swedish and Canadian equities, respectively. Chaboud et al. (2013) study trading in three different FX pairs and find that in the two most liquid pairs (EUR/USD and JPY/USD) humans are more likely to consume liquidity from computers than vice versa. This result does not obtain for the third pair (JPY/EUR), which is likely due to the fact that most of these trades are aimed at exploiting triangular arbitrage opportunities. Contrary to all these papers, Hendershott et al. (2012) document that HFTs are more likely to be takers in a sample of Nasdaq stocks.

The intuition behind MT_k^* 's dependence on σ is straightforward. While an increase in volatility induces traders to submit less aggressive limit orders, FTs' effectively only change their strategy towards other FTs while continuing to provide liquidity to STs for both possible realizations of ε (see Lemma 1). Hence a higher level of σ increases the relative odds of FTs being makers and conversely reduces the chances of STs trading passively. To our knowledge, this is a novel empirical implication of our model that has not yet been tested in the literature.

4.3 Market orders

We now turn to the analysis of the equilibrium profits obtained from posting market orders, which we denote V_k^{MO*} for $k \in \{ST, FT\}$. Recalling the LHS of inequality (1), we may write

$$V_k^{MO*} = L - E(\tau_k^*) \quad (9)$$

where we follow Foucault (1999) in defining a seller's trading cost as the discount he cedes to the limit order buyer with respect to the asset's fair value, that is

$$\tau_k = v_t - B_t^m$$

It is important to note that the trading cost crucially depends on both the type of the market order trader and the limit order trader whose quote is executed as well as on the most recent realization of the asset value innovation. Hence the expectation $E(\tau_k^*)$ above is taken jointly over all possible states for both v_t and B_t^m . Now let $\tau_{j,k}^\varepsilon$ denote the trading cost incurred by a type- k seller that executes the bid quote of a type- j buyer conditional on the most recent realization of ε , where $j, k \in \{ST, FT\}$, and let $\omega_{j,k}^{\varepsilon*}$ be the associated probability of this event in equilibrium conditional on ε . Then, the equilibrium expected trading cost for a type- k trader is given by

$$E(\tau_k^*) = \frac{\omega_{ST,k}^{-\sigma*} \tau_{ST,k}^{-\sigma*} + \omega_{ST,k}^{+\sigma*} \tau_{ST,k}^{+\sigma*} + \omega_{FT,k}^{-\sigma*} \tau_{FT,k}^{-\sigma*} + \omega_{FT,k}^{+\sigma*} \tau_{FT,k}^{+\sigma*}}{\omega_{ST,k}^{+\sigma*} + \omega_{ST,k}^{-\sigma*} + \omega_{FT,k}^{+\sigma*} + \omega_{FT,k}^{-\sigma*}} \quad (10)$$

Finally, we may also consider the expected trading cost of the average transaction, which can be computed as a linear combination of $E(\tau_{ST}^*)$ and $E(\tau_{FT}^*)$, where the weights are given by the respective share in the overall trading rate

$$E(\tau^*) = \frac{\varphi_{ST}^{MO*}}{\varphi_{ST}^{MO*} + \varphi_{FT}^{MO*}} E(\tau_{ST}^*) + \frac{\varphi_{FT}^{MO*}}{\varphi_{ST}^{MO*} + \varphi_{FT}^{MO*}} E(\tau_{FT}^*) \quad (11)$$

At this point it is useful to recall (see inequality (1)) that market orders are only submitted in case they offer a payoff which at least matches the expected profit that can be obtained by posting a limit order instead, i.e. we have $V_k^{MO*} \geq V_k^{LO*}$. Particularly profitable trading opportunities arise in the case where the asset value moves "against" the limit order trader, which allows the market order trader to pick off an extra windfall profit of 2σ on top of his outside option value. Notice that these situations can only arise when limit orders are sufficiently aggressive such that they may be executed for either realization of ε .

With this in mind, it is easy to see how the presence of FTs can lead to an increase in $E(\tau_{ST}^*)$ relative to the case where all agents are slow. Because FTs may revise their orders when being followed by STs, the latter now face a decreased likelihood of enjoying additional "picking off"-profits. In fact, Lemma 1 states that STs just obtain their outside option value V_{ST}^{LO*} when trading against the limit orders of FTs, which implies (Corollary 1) that the costs incurred for these transactions are higher than those prevailing in a market with only STs (henceforth denoted $E(\tau_0^*)$). Importantly, this is true even in the case where agents use low fill-rate strategies and market orders do not enjoy the above benefits arising from the winner's curse. The reason behind this is that FTs' speed advantage additionally enables them to discriminate between trader types.⁹ To see this, consider a FT submitting a buy limit order at time t and suppose that σ is sufficiently high such that he

⁹Notice that it is not profitable to discriminate between different types of traders for $\sigma = 0$ because in this case all agents have the same outside option.

initially posts $B_{FT}^* = \hat{B}_{FT}(v_t - \sigma)$. Now consider what happens if the asset value decreases prior to the arrival of a seller in the next period, i.e. $\varepsilon_{t+1} = -\sigma$. If the following trader is fast, he will submit a market order and thus obtains a profit of $\hat{B}_{FT}(v_t - \sigma) - (v_t - \sigma - L) = V_{FT}^{LO*}$, which is just his outside option. Now if the following trader is slow, the fast limit order trader still has an incentive to revise his quote downwards despite having ruled out the winner's curse ex-ante, because STs have a lower outside option and are therefore willing to accept a profit of $\hat{B}_{ST}(v_t - \sigma) - (v_t - \sigma - L) = V_{ST}^{LO*} < V_{FT}^{LO*}$ instead.

Unlike STs, FTs do not lose profitable trading opportunities and therefore it is easy to see that we must have $E(\tau_{FT}^*) < E(\tau_{ST}^*)$. Moreover, because their outside option value V_{FT}^{LO*} exceeds V_0^{LO*} , their trading costs are in fact below those prevailing in a market with only STs. The following Corollary summarizes.

Corollary 6 *$E(\tau_{FT}^*) < E(\tau_{ST}^*)$ and $E(\tau_{FT}^*) < E(\tau_0^*)$ for all $\alpha \in (0, 1)$. Moreover, if $\sigma \in [8/15, \underline{\sigma})$ then $E(\tau_{ST}^*) > E(\tau^*) > E(\tau_0^*)$ for all $\alpha \in (0, 1)$.*

Proof. See the Appendix. ■

The fact that we do not observe an increase in STs expected trading costs across the entire parameter space is once again due to the specific distributional assumption concerning the asset value innovations. In fact, because ε is discrete, STs increase the execution probability of their limit order conditional on the arrival of a FT without increasing the likelihood that the order is executed by a ST when switching from a specialized to an unspecialized strategy. This implies that they offer a higher bid price to slow sellers without actually improving the chances of having their limit order filled by them, and the resulting discrete drop in the cost of trading with other STs can cause $E(\tau_{ST}^*)$ (and eventually also $E(\tau^*)$) to drop below $E(\tau_0^*)$ for intermediate values of α .¹⁰

There is overwhelming empirical evidence that is consistent with the view that being slow (fast) leads to trading at less (more) favourable prices. Oeck & Saar (2009) document that the "lifetimes" of limit orders have decreased considerably over the past years, which suggests that an increasingly large fraction of limit orders is in fact not accessible for market participants that are not able to react sufficiently fast to quote updates. In fact, several exchanges (e.g. Direct Edge, Bats, and Nasdaq) have at least temporarily facilitated price discrimination via trading speed by introducing so-called "flash orders" (see Skjeltorp et al. (2011) for details), a practice that has been banned by the SEC in the meantime. Garvey and Wu (2010) provide more direct evidence by showing that geographical distance to the market center is negatively related to execution speed and positively related to transactions costs. Hendershott and Riordan (2012) study transactions data from the German Stock Exchange and find that "algorithmic traders consume liquidity when it is cheap" in the sense

¹⁰ Solving the model numerically for normally distributed innovations yields $E(\tau_{ST}^*) > E(\tau^*) > E(\tau_0^*)$ for all parameters.

that they pay lower effective spreads than slower human traders. In line with this, Moallemi and Sağlam (2011) develop a model of the "cost of latency" and estimate that it has increased threefold in the period 1995-2005. Finally, McNish and Upson (2012) study the Benchmark Quote exception to the Order Protection Rule and provide evidence that is consistent with FTs exploiting their speed advantage at the expense of STs.

Several empirical papers (see e.g. Hendershott et al. (2011), Boehmer et al. (2012), and Hasbrouck and Saar (2011)) have documented a negative relationship between algorithmic trading and traditional measures of market liquidity (e.g. quoted and effective spreads). While some of these studies also provide evidence for a causal effect, it is difficult to disentangle the effect of trading speed (i.e. HFT) from other features of algorithmic trading (e.g. improved search and trade execution), the removal of competitive barriers in terms of liquidity provision, decimalization, and inter-market competition. Therefore these findings are not necessarily at odds with the implication of Corollary 6 that high-frequency trading leads to an increase in trading costs. Additionally, it is important to stress that our measure of trading costs merely reflects the distribution of bargaining power between limit and market order traders, which is not necessarily in line with the more traditional notion of the "cost of immediacy" payable to an intermediary. In fact, trading costs can be negative in our model (as e.g. in Goettler et al. (2009)) because we may have $B_t^m > v_t$, while most empirical measures assume that the midquote is equal to the true value and hence $B_t^m < v_t$ by assumption.

5 Welfare

So far, we have taken the proportion of FTs as exogenous. While this helps to learn about the effect of HFT in a limit order market in general, a rigorous evaluation of the associated welfare effects demands that α is determined endogenously. In order to do so, we assume (as is customary in the literature on information acquisition, see e.g. Grossman and Stiglitz (1980)) that all agents are born slow but prior to entering the market have the opportunity to become fast at an exogenously given cost $c > 0$. Now let $W_{ST}^*(\alpha)$ and $W_{FT}^*(\alpha)$ denote the equilibrium expected trading profits for STs and FTs, respectively, as a function of α . Then, an interior equilibrium requires that the profits of STs are equal to those of FTs net of the cost for becoming fast, that is

$$\Delta W^*(\alpha^*) \equiv W_{FT}^*(\alpha^*) - W_{ST}^*(\alpha^*) = c \quad (12)$$

Obviously, corner equilibria may arise either because the cost c is prohibitively high, i.e.

$$\max_{\alpha} \Delta W^*(\alpha) < c$$

in which case all agents choose to remain slow, $\alpha^* = 0$, or because the incremental benefit of being fast is high enough to justify the cost for any level of α (hence $\alpha^* = 1$), which is the case iff

$$\min_{\alpha} \Delta W^*(\alpha) > c$$

Using the results from the previous sections, the expected equilibrium trading profits for each trader type can easily be calculated as a weighted average of the profits obtained via each order type

$$W_{ST}^* = \frac{\varphi_{ST}^{LO*}}{\varphi_{ST}^{LO*} + \varphi_{ST}^{MO*}} V_{ST}^{LO*} + \frac{\varphi_{ST}^{MO*}}{\varphi_{ST}^{LO*} + \varphi_{ST}^{MO*}} V_{ST}^{MO*} \quad (13)$$

$$W_{FT}^* = \frac{\varphi_{FT}^{LO*}}{\varphi_{FT}^{LO*} + \varphi_{FT}^{MO*}} V_{FT}^{LO*} + \frac{\varphi_{FT}^{MO*}}{\varphi_{FT}^{LO*} + \varphi_{FT}^{MO*}} V_{FT}^{MO*} \quad (14)$$

Then, emphasizing the dependence of all equilibrium outcomes on α , equilibrium welfare is given by

$$\begin{aligned} W^*(\alpha^*) &= (1 - \alpha^*) W_{ST}^*(\alpha^*) + \alpha^* (W_{FT}^*(\alpha^*) - c) \\ &= TR^*(\alpha^*) \times (2L) - \alpha^* c \end{aligned} \quad (15)$$

Note that an equilibrium may fail to exist because $\Delta W^*(\alpha)$ is not continuous in α (due to the discrete change when moving from one type of equilibrium to another). Additionally, an interior equilibrium need not be unique because there may be several values of α that satisfy equation $\Delta W^*(\alpha) = c$ for fixed (σ, L, c) . Straightforward computations reveal that STs are always worse off than in the absence of FTs. Hence equation (12) directly implies the following.

Corollary 7 *Suppose that the parameters (σ, L, c) are such that an equilibrium with endogenous investment in HFT exists. Then, any positive equilibrium level of investment $\alpha^* > 0$ exceeds the socially optimal level of investment α^+ and moreover yields a social welfare loss, that is we have $W^*(\alpha^*) < W^*(0)$.*

Proof. See the Appendix. ■

As in Biais et al. (2012), the equilibrium level of investment into HFT is excessive from a social point of view because FTs exert a negative externality on STs. Yet, the nature of these externalities are quite different from each other. In Biais et al. (2012), FTs are a *source* of adverse selection because their speed allows them to act swiftly on new information. As is customary, this induces dealers to widen the bid-ask spread payable by everyone and consequently implies a welfare loss to STs. Here, FTs are not a *source* of adverse selection but instead have the potential to *avoid* adverse selection. Yet, because agents trade directly with each other in a dynamic setting, the ability to

avoid the winner's curse increases their bargaining power relative to STs and consequently leaves the latter with a smaller share of the surplus $2L$. Additionally, market orders become less profitable for STs because they can no longer pick off FTs.

It is important to stress that Corollary 7 does not imply that $\alpha = 0$ is socially optimal. In fact, Corollary 3 shows that the presence of some FTs allows more gains from trade to be realized for $\sigma \in \bar{\Sigma}$, which implies that some investment in HFT can be optimal for c sufficiently small. For the purpose of illustration, consider the following numerical example. Suppose that $\sigma = L = 1$ and $c = 0.1$. Then straightforward calculations reveal that the socially optimal proportion of FTs is equal to $\alpha^+ = 0.347$, which is strictly positive. Compared to the case without FTs, welfare increases by more than 10% from $W^*(0) = 0.4$ to $W^*(\alpha^+) = 0.446$. In line with Corollary 7, the equilibrium level of HFT is considerably higher at $\alpha^* = 0.759$ and implies a welfare loss of close to 5% ($W^*(\alpha^*) = 0.383$). Finally, it is worth stressing that it can never be optimal to have a market with only FTs, because in this case the equilibrium that obtains is exactly identical to the one with only slow traders with the difference that all agents have spent c to become fast. In contrast, $\alpha = 1$ may be optimal in Biais et al. (2012) if STs are sufficiently inefficient in finding trading opportunities.

Because FTs are the source of a negative externality, there is scope for regulatory intervention. Among the numerous suggestions that have been tabled in the public debate are minimum resting times for limit orders and restrictions concerning the number of electronic messages individual market participants may send. But it is precisely FTs' ability to revise outstanding limit orders instantaneously that is the source of an important efficiency gain, and therefore it is socially desirable to implement a regulation that improves upon the market outcome but at the same time attempts to reap some of the benefits that HFT may bring along. Any direct form of taxation on HFT activity would in principle be suited to fulfill this objective.

One salient feature of our model which is at odds with reality is that limit orders are always submitted with the aim of being executed. In reality, this is clearly not the case as some market participants try to slow down others by submitting a large number of limit orders deep in the book. This practice, called "quote stuffing"¹¹, and other related strategies may be responsible for a large share of the total growth in message traffic over the past 5-10 years that has become a burden on both market regulators and exchanges alike. While several market centers (e.g. Nasdaq, Euronext, Deutsche Börse, etc.) have started to fine users for excessive system usage, there is little attempt to discriminate those that want to trade from those that aim to disrupt markets and benefit at the expense of others. One notable exception is the policy adopted in early 2011 by the Intercontinental Exchange (ICE), which only penalizes message traffic beyond the best quote.¹² Such a weighting may be a sensible way to curb the excessive submission of limit orders without loosing the potential benefits of HFT.

¹¹see Financial Times, "High-frequency trading: Up against a bandsaw", September 2, 2010

¹²see <http://ir.theice.com/releasedetail.cfm?releaseid=652686>

Another question is whether there are other ways of reaching the efficiency gain from avoiding the winner's curse without suffering from the negative externalities of HFT. One potential solution could be the use of "pegged" limit orders, i.e. orders whose limit price is indexed to the fundamental value of the underlying asset. While welfare would be maximal in this case, it would obviously be difficult to observe v_t in practice.

6 Conclusion

This paper contributes to the ongoing controversy on the pros and cons of HFT by presenting a stylized model of trading in a limit order market where investors differ in their trading speed. We show that HFT may allow more gains from trade to be realized because being fast reduces the risk of being "picked off" and thus eliminates the need for posting cautious limit orders. Yet, at the same time, slow market participants suffer from a worse bargaining position and this may lead them to submit limit orders with a lower execution probability and therefore reduce trade. This negative externality ensures that there is too much investment in HFT in equilibrium. We additionally show how the presence of FTs affects the order flow and develop several testable implications, some novel and others that are consistent with the existing literature.

While our model confirms the view held by policy makers that HFT should be regulated, our analysis suggests that some of the proposed measures such as minimum order resting times or undifferentiated limits to message traffic would deprive society of some potential gains.

References

- [1] Biais, B., Foucault, T., and Moinas, S. (2012), "Equilibrium High Frequency Trading", Working Paper
- [2] Biais, B., and Woolley, P. (2011), "High Frequency Trading", Working Paper
- [3] Boehmer, E., Fong, K. Y. L., and Wu, J. (2012) "International Evidence on Algorithmic Trading", Working Paper
- [4] Cartea, A., and Penalva, J. (2011) "Where is the value in high-frequency trading", Quarterly Journal of Finance, forthcoming
- [5] Chaboud, A., Chiquoine, B., Hjalmarsson, E., and Vega, C. (2013), "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market", Working Paper

- [6] Colliard, J.-E., and Foucault, T. (2012) "Trading Fees and Efficiency in limit order markets", Review of Financial Studies, forthcoming
- [7] Copeland, T. E., and Galai, D. (1983), "Information effects on the bid-ask spread", Journal of Finance, Vol. 38, No. 5: 1457-69
- [8] Foucault, T. (1999) "Order Flow Composition and Trading Costs in a dynamic limit order market", Journal of Financial Markets, Vol. 2, No. 2: 99-134.
- [9] Foucault, T. (2012), "Algorithmic Trading: Issues and Preliminary Evidence", forthcoming in "Market Microstructure: Confronting Many Viewpoints", John Wiley & Sons
- [10] Foucault, T., Hombert, J., and Rosu, I. (2013) "News trading and speed", Working Paper
- [11] Garvey, R. and Wu, F. (2010) "Speed, distance, and electronic trading: New evidence on why location matters", Journal of Financial Markets, Vol. 13, No 4: 367–396.
- [12] Glosten, L. and Milgrom, P. (1985) "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders", Journal of Financial Economics, Vol. 14, No. 1: 71–100
- [13] Goettler, R., Parlour, C. A., and Rajan, U. (2009) "Informed Traders and Limit Order Markets", Journal of Financial Economics, Vol. 93, No. 1: 67–87
- [14] Hagströmer, B. and Norden, L. (2013) "The diversity of high-frequency traders", Working Paper
- [15] Hasbrouck, J., and Saar, G. (2009) "Trading and technology: The blurring of traditional definitions", Journal of Financial Markets, Vol. 12, No. 2: 143-172
- [16] Hasbrouck, J., and Saar, G. (2011) "Low-latency trading", Working Paper
- [17] Hendershott, T., Jones, C., and Menkveld, A. (2011) "Does algorithmic trading improve liquidity?", Journal of Finance, Vol. 66, No. 1: 1-33.
- [18] Hendershott, T., and Riordan, R. (2012) "Algorithmic Trading and the Market for Liquidity ", Journal of Financial and Quantitative Analysis, forthcoming
- [19] Hendershott, T., and Riordan, R., and Brogaard, J. (2012) "High-frequency trading and price discovery", Working Paper
- [20] Jovanovic, B., and Menkveld, A. (2012), "Middlemen in Securities Markets", Working Paper
- [21] Kirilenko, A., Kyle, A., Samadi, M., and Tuzun, T. (2011) "The Flash Crash: The Impact of High-Frequency Trading on an Electronic Market", Working Paper

- [22] Malinova, K, Parks, A., and Riordan, R. (2012) "Do retail traders suffer from high-frequency traders?", Working Paper
- [23] McInish, T, and Upson, J. (2012) "Strategic Liquidity Supply in a Market with Fast and Slow Traders", Working Paper
- [24] Menkveld, A. (2012) "High-frequency trading and the new-market makers", Working Paper
- [25] Moallemi, C. C., and Sağlam, M. (2011) "The cost of latency", Working Paper
- [26] Pagnotta, E., and Philippon, T. (2012) "Competing on Speed", Working Paper
- [27] Rosu, I., and Martinez, V. (2013) "High-Frequency Traders, News and Volatility", Working Paper
- [28] Skjeltorp, J. A., Sojli, E., and Tham, W. W. (2011) "Sunshine Trading: Flashes of Trading Intent at the NASDAQ", Working Paper

7 Appendix

7.1 Proof of Lemma 1

Consider a FT that has placed a bid at time t . If he observes the innovation ε_{t+1} and can still modify his order, he knows that the next agent (provided he is a seller) is a ST with sell cutoff price $\hat{B}_{ST}^*(v_t + \varepsilon_{t+1})$. Clearly, this is the optimal bid price since a lower (higher) bid has a zero (the same) execution probability. ■

7.2 Proof of Lemma 2

Suppose that, in equilibrium, a ST posts a buy limit order that executes only if the asset value decreases. Then, a FT can always do better by posting the same bid price and revise his order according to Lemma 1 in case the asset value increases and the next trader is a ST. Similarly, if a ST posts a limit order that executes irrespectively of the asset value innovation, a FT can do better by revising his order according to Lemma 1 and obtaining higher profits by incorporating the latest innovation into his limit price whenever possible (i.e. when a ST follows). Hence we conclude that $V_{FT}^{LO*} > V_{ST}^{LO*}$, and equation (2) then implies that $\hat{B}_{FT}^*(v_t) > \hat{B}_{ST}^*(v_t)$.

It remains to show that $\hat{B}_{FT}^*(v_t - \sigma) \leq \hat{B}_{ST}^*(v_t + \sigma)$. First, notice that L is the *expected* total gains from trade per period (if two agents with different private valuations trade they share a surplus of $2L$, but this occurs at most with probability $1/2$) we must have $L \geq V_k^{LO*} \geq 0$ for $k \in \{ST, FT\}$. Now suppose that $\sigma \geq L/2$. Using equation (2), we have $v_t + \sigma \geq \hat{B}_{ST}^*(v_t + \sigma) \geq v_t + \sigma - L$ and $v_t - \sigma \geq \hat{B}_{FT}^*(v_t - \sigma) \geq v_t - \sigma - L$, which directly implies $\hat{B}_{ST}^*(v_t + \sigma) \geq \hat{B}_{FT}^*(v_t - \sigma)$. Now assume that $\sigma < L/2$ and consider a fast buyer submitting a buy limit order. It is easy to see that in this case we have $\frac{\alpha}{4}[v - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)] + \frac{\alpha}{4}[v + \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)] \geq \frac{\alpha}{4}[v - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma)]$ such that his optimal choice is $B_{FT}^* = \hat{B}_{FT}^*(v_t + \sigma)$. Notice that a buyer arriving at $t+1$ will never execute this order because $v - \sigma + L > \hat{B}_{FT}^*(v_t + \sigma)$, that is the bid price B_{FT}^* is below his lowest possible valuation. Now consider a slow buyer and suppose he posts a buy limit order with $B_{ST}^* = \hat{B}_{FT}^*(v_t + \sigma)$. As this is not necessarily his equilibrium strategy we have that $V_{ST}^{LO*} \geq \frac{1}{2}[v + L - \hat{B}_{FT}^*(v_t + \sigma)]$. But we just concluded that $V_{FT}^{LO*} = \frac{\alpha}{2}[v + L - \hat{B}_{FT}^*(v_t + \sigma)] + \frac{1-\alpha}{4}[v - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma)] + \frac{1-\alpha}{4}[v + \sigma + L - \hat{B}_{ST}^*(v_t + \sigma)]$, and therefore $V_{FT}^{LO*} - V_{ST}^{LO*} \leq \frac{1-\alpha}{4}[\hat{B}_{FT}^*(v_t + \sigma) - \hat{B}_{ST}^*(v_t - \sigma)] + \frac{1-\alpha}{4}[\hat{B}_{FT}^*(v_t + \sigma) - \hat{B}_{ST}^*(v_t + \sigma)] = \frac{1-\alpha}{2}[\hat{B}_{FT}^*(v_t) - \hat{B}_{ST}^*(v_t)] + \frac{1-\alpha}{2}\sigma$. Then, using equation (2), we obtain $\hat{B}_{FT}^*(v_t) - \hat{B}_{ST}^*(v_t) \leq \frac{1-\alpha}{1+\alpha}\sigma$, which lets us conclude $\hat{B}_{ST}^*(v_t) - \hat{B}_{FT}^*(v_t) + 2\sigma > 0$ as desired. ■

7.3 Proof of Proposition 1

Although Lemma 2 implies that the equilibrium sell cutoff prices for sellers always satisfy $\hat{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t - \sigma) \leq \hat{B}_{ST}^*(v_t + \sigma) \leq \hat{B}_{FT}^*(v_t + \sigma)$, there is no one-to-one mapping between quotation strategies and limit order execution probabilities because one additionally needs to consider the possibility that buyers submit sell market orders. While we have briefly ruled this out in Subsection X.X in order to simplify the exposition of agents' strategies, it needs to be considered in finding an equilibrium. Because Lemma 2 is essentially silent on the relative positioning of the equilibrium sell cutoff prices of *buyers*, which we henceforth denote by $\check{B}_k^*(v_{t+1})$, there are four distinct possibilities for the joint ordering of buyers' and sellers' equilibrium sell cutoff prices:

Ordering 1: $\hat{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t - \sigma) \leq \check{B}_{ST}^*(v_t - \sigma) \leq \check{B}_{FT}^*(v_t - \sigma) \leq \hat{B}_{ST}^*(v_t + \sigma) \leq \hat{B}_{FT}^*(v_t + \sigma) \leq \check{B}_{ST}^*(v_t + \sigma) \leq \check{B}_{FT}^*(v_t + \sigma)$

Ordering 2: $\hat{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t - \sigma) \leq \check{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{ST}^*(v_t + \sigma) \leq \check{B}_{FT}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t + \sigma) \leq \check{B}_{ST}^*(v_t + \sigma) \leq \check{B}_{FT}^*(v_t + \sigma)$

Ordering 3: $\hat{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t - \sigma) \leq \hat{B}_{ST}^*(v_t + \sigma) \leq \check{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t + \sigma) \leq \check{B}_{FT}^*(v_t - \sigma) \leq \check{B}_{ST}^*(v_t + \sigma) \leq \check{B}_{FT}^*(v_t + \sigma)$

Ordering 4: $\hat{B}_{ST}^*(v_t - \sigma) \leq \hat{B}_{FT}^*(v_t - \sigma) \leq \hat{B}_{ST}^*(v_t + \sigma) \leq \hat{B}_{FT}^*(v_t + \sigma) \leq \check{B}_{ST}^*(v_t - \sigma) \leq \check{B}_{FT}^*(v_t - \sigma) \leq \check{B}_{ST}^*(v_t + \sigma) \leq \check{B}_{FT}^*(v_t + \sigma)$

In order to find an equilibrium for fixed parameters, we follow the following steps:

- 1) Conjecture an ordering of cutoff prices.
- 2) Conjecture equilibrium strategies and solve for the equilibrium cutoff prices.
- 3) Verify that
 - a) the assumed strategies are best replies (i.e. deviations are not profitable) and
 - b) the cutoff prices satisfy the assumed ordering.

Type 1 equilibrium:

Case A:

Step 1: Assume Ordering 1.

Step 2: Assume $B_{ST}^* = \hat{B}_{ST}^*(v_t - \sigma)$ and $B_{FT}^* = \hat{B}_{FT}^*(v_t - \sigma)$ which implies that $p^* = \frac{1-\alpha}{4}$ and $q_{|FT}^* = \frac{1}{4}$. Using the the optimally revised FT quotes from Lemma 1 with conditional execution probabilities $q_{|ST,+\sigma} = q_{|ST,+\sigma} = \frac{1}{4}$, the equilibrium expected profits from posting limit orders are given by

$$\begin{aligned}
V_{ST}^{LO*} &= \frac{1-\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma) \right] \\
V_{FT}^{LO*} &= \frac{1-\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma) \right] \\
&\quad + \frac{1-\alpha}{4} \left[v_t + \sigma + L - \hat{B}_{ST}^*(v_t + \sigma) \right] \\
&\quad + \frac{\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma) \right]
\end{aligned}$$

Substituting $V_{t,k}^{LO} = \hat{B}_k(v_t) - (v_t - L)$ (equation (2)) one obtains a system of 2 equations in 2 unknowns that can be solved for the equilibrium cutoff prices $\hat{B}_{ST}^*(v_t) = v_t - L + (2L)\frac{1-\alpha}{5-\alpha}$ and $\hat{B}_{FT}^*(v_t) = v_t - L + (2L)\frac{8-\alpha(3+\alpha)}{(5-\alpha)(4+\alpha)}$.

Step 3: Notice that it can never be optimal to choose a bid price equal to the sell cutoff price of a buyer, as one can always fare strictly better by choosing the next lowest sell cutoff price of a seller (thus avoiding the loss when trading with the buyer). Hence we must have $\hat{B}_{ST}^* \in \mathcal{B}$ and $\hat{B}_{FT}^* \in \mathcal{B}'$. Thus, the strategy conjectured for STs is a best reply iff

$$\begin{aligned}
\frac{1-\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma) \right] &\geq \frac{1}{4} \left[v_t - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma) \right] \\
\frac{1-\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma) \right] &\geq \frac{1}{2} \left[v_t - \sigma + L - \hat{B}_{ST}^*(v_t + \sigma) \right] \\
&\quad + \frac{1-\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{ST}^*(v_t + \sigma) \right] \\
\frac{1-\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma) \right] &\geq \frac{1}{2} \left[v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma) \right] \\
&\quad + \frac{1}{4} \left[v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma) \right]
\end{aligned}$$

Similarly, fast buyers have no incentives to deviate iff:

$$\begin{aligned}
\frac{\alpha}{4} \left[v_t - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma) \right] &\geq \frac{\alpha}{2} \left[v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma) \right] \\
&\quad + \frac{\alpha}{4} \left[v_t + \sigma + L - \hat{B}_{FT}^*(v_t + \sigma) \right]
\end{aligned}$$

Brute-force algebra reveals that these inequalities and the assumed ordering of cutoff prices are satisfied if and only $\alpha \leq \sqrt{5} - 2$ and $\sigma \geq \frac{24+\alpha(1-\alpha)}{(5-\alpha)(4+\alpha)}L$.

Case B:

Step 1: Assume Ordering 2.

Step 2: Conjecture the same equilibrium strategies as in Case A, which implies identical cutoff prices.

Step 3: The proposed strategies are best replies (for buyers) iff

$$\begin{aligned}
\frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma)] &\geq \frac{1}{4} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma)] \\
\frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma)] &\geq \frac{2-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t + \sigma)] \\
&\quad + \frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t + \sigma)] \\
\frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma)] &\geq \frac{1}{2} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)] \\
&\quad + \frac{1}{4} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)] \\
\frac{\alpha}{4} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma)] &\geq \frac{\alpha}{2} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)] \\
&\quad + \frac{\alpha}{4} [v_t + \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)]
\end{aligned}$$

These inequalities together with the conjectured ordering of cutoff prices are satisfied iff $\alpha \leq \sqrt{5} - 2$ and $\frac{24+\alpha(1-\alpha)}{(5-\alpha)(4+\alpha)}L > \sigma \geq L$.

Case C:

Step 1: Assume Ordering 3.

Step 2: Conjecture the same equilibrium strategies as in Case A, which implies identical cutoff prices.

Step 3: The proposed strategies are best replies (for buyers) iff

$$\begin{aligned}
\frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma)] &\geq \frac{1}{4} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma)] \\
\frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma)] &\geq \frac{1}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t + \sigma)] \\
&\quad + \frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t + \sigma)] \\
\frac{1-\alpha}{4} [v_t - \sigma + L - \hat{B}_{ST}^*(v_t - \sigma)] &\geq \frac{2-\alpha}{4} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)] \\
&\quad + \frac{1}{4} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t + \sigma)] \\
\frac{\alpha}{4} [v_t - \sigma + L - \hat{B}_{FT}^*(v_t - \sigma)] &\geq \frac{\alpha}{2} [v_t + L - \hat{B}_{FT}^*(v_t + \sigma)]
\end{aligned}$$

These inequalities together with the conjectured ordering of cutoff prices are satisfied iff $\alpha \leq \sqrt{5} - 2$ and $L > \sigma \geq L_{\frac{4}{5-\alpha}}$. Combining Cases A, B and C, we conclude $B_{ST}^* = \hat{B}_{ST}^*(v_t - \sigma)$ and $B_{FT}^* = \hat{B}_{FT}^*(v_t - \sigma)$ and the associated order choice strategy constitute an equilibrium iff $\alpha \leq \sqrt{5} - 2$ and $\sigma \geq L_{\frac{4}{5-\alpha}}$.

The proof for the remaining equilibrium types follows exactly the same logic, such that we omit them for brevity. The following tables indicates the respective orderings that give rise to the

individual equilibria together with the equilibrium strategies, cutoff prices, execution probabilities and parameter conditions, where we use the following definitions: $\alpha_1^* \equiv \sqrt{5} - 2$, $\alpha_2^* \equiv \frac{\sqrt{33}-5}{2}$, $\sigma_1^*(\alpha) \equiv L \frac{4}{5-\alpha}$, $\sigma_2^*(\alpha) \equiv L \frac{2\alpha(1+\alpha)}{3-4\alpha}$, $\sigma_3^*(\alpha) \equiv L \frac{4(4+\alpha)}{26-\alpha^2}$, $\sigma_4^*(\alpha) \equiv L \frac{2(1-\alpha)(4+\alpha)}{7+3\alpha}$, and $\sigma_5^*(\alpha) \equiv L \frac{4(1+\alpha)}{7+3\alpha}$.

Equilibrium	Ordering	B_{ST}^*	B_{FT}^*	p^*	q_{FT}^*	Condition α	Condition σ
1 (SLFR)	1-3	$\hat{B}_{ST}^*(v_t - \sigma)$	$\hat{B}_{FT}^*(v_t - \sigma)$	$\frac{1-\alpha}{4}$	$\frac{1}{4}$	$\alpha \leq \alpha_1^*$	$\sigma \geq \sigma_1^*$
2 (ULFR)	1-4	$\hat{B}_{FT}^*(v_t - \sigma)$	$\hat{B}_{FT}^*(v_t - \sigma)$	$\frac{1}{4}$	$\frac{1}{4}$	$\alpha_1^* < \alpha \leq \alpha_2^*$	$\sigma \geq \sigma_4^*$
						$\alpha_2^* < \alpha$	$\sigma \geq \sigma_5^*$
3	3 & 4	$\hat{B}_{ST}^*(v_t + \sigma)$	$\hat{B}_{FT}^*(v_t - \sigma)$	$\frac{2-\alpha}{4}$	$\frac{1}{4}$	$\alpha \leq \alpha_1^*$	$\sigma_1^* > \sigma \geq \sigma_3^*$
						$\alpha_1^* < \alpha \leq \alpha_2^*$	$\sigma_4^* > \sigma \geq \sigma_3^*$
4 (SHFR)	4	$\hat{B}_{ST}^*(v_t + \sigma)$	$\hat{B}_{FT}^*(v_t + \sigma)$	$\frac{2-\alpha}{4}$	$\frac{1}{2}$	$\alpha \leq \alpha_2^*$	$\sigma_3^* > \sigma \geq \sigma_2^*$
5 (UHFR)	4	$\hat{B}_{FT}^*(v_t + \sigma)$	$\hat{B}_{FT}^*(v_t + \sigma)$	$\frac{1}{2}$	$\frac{1}{2}$	$\alpha \leq \alpha_2^*$	$\sigma_2^*(\alpha) > \sigma$
						$\alpha > \alpha_2^*$	$\sigma_5^*(\alpha) > \sigma$

Equilibrium	$\hat{B}_{ST}^*(v_t)$	$\hat{B}_{FT}^*(v_t)$
1 (SLFR)	$v_t - L + (2L) \frac{1-\alpha}{5-\alpha}$	$v_t - L + (2L) \frac{8-\alpha(3+\alpha)}{(5-\alpha)(4+\alpha)}$
2 (ULFR)	$v_t - L + (2L) \frac{1+\alpha}{7+3\alpha}$	$v_t - L + (2L) \frac{3-\alpha}{7+3\alpha}$
3	$v_t - L + (2L) \frac{2-\alpha}{6-\alpha} - \frac{2}{6-\alpha} \sigma$	$v_t - L + (2L) \frac{8-\alpha(2+\alpha)}{(6-\alpha)(4+\alpha)} + \frac{4(1-\alpha)}{(6-\alpha)(4+\alpha)} \sigma$
4 (SHFR)	$v_t - L + (2L) \frac{2-\alpha}{6-\alpha} - \frac{2}{6-\alpha} \sigma$	$v_t - L + (2L) \frac{4+\alpha(2-\alpha)}{(6-\alpha)(2+\alpha)} + \frac{2-\alpha(8-\alpha)}{(6-\alpha)(2+\alpha)} \sigma$
5 (UHFR)	$v_t - L + (2L) \frac{1}{3} - \frac{2}{3(1+\alpha)} \sigma$	$v_t - L + (2L) \frac{1}{3} + \frac{1-3\alpha}{3(1+\alpha)} \sigma$

Moreover, the functions $\sigma_{ST}^*(\alpha)$, $\sigma_{FT}^*(\alpha)$ and $\alpha_S^*(\sigma)$ are defined as

$$\begin{aligned}
\sigma_{ST}^*(\alpha) &\equiv \begin{cases} \sigma_1^*(\alpha) & \text{if } \alpha \leq \alpha_1^* \\ \sigma_4^*(\alpha) & \text{if } \alpha_1^* < \alpha \leq \alpha_2^* \\ \sigma_5^*(\alpha) & \text{if } \alpha_2^* < \alpha \end{cases} \\
\sigma_{FT}^*(\alpha) &\equiv \begin{cases} \sigma_3^*(\alpha) & \text{if } \alpha \leq \alpha_2^* \\ \sigma_5^*(\alpha) & \text{if } \alpha_2^* < \alpha \end{cases} \\
\alpha_S^*(\sigma) &\equiv \begin{cases} \alpha_1^* & \text{if } \sigma \geq \sigma_1^*(\alpha_1^*) \\ \sigma_4^{*-1}(\sigma) & \text{if } \sigma_1^*(\alpha_1^*) > \sigma \geq \sigma_3^*(\alpha_2^*) \\ \sigma_2^{*-1}(\sigma) & \text{if } \sigma_3^*(\alpha_2^*) > \sigma \end{cases}
\end{aligned}$$

Finally, it is straightforward, albeit tedious to show that no other equilibria exist, which establishes uniqueness. ■

7.4 Proof of Corollary 1

The result follows immediately from the cutoff prices and equilibrium execution probabilities computed in the proof of Proposition 1, where $V_0^{LO*} = \lim_{\alpha \rightarrow 0} V_{ST}^{LO*}$ and $p_0^* = \lim_{\alpha \rightarrow 0} p^*$. ■

7.5 Proof of Corollary 2

The table below collects δ^* for each type of equilibrium, which is calculated directly by using the cutoff prices obtained in the proof of Proposition 1.

Equilibrium	δ^*
SLFR	$(2L) \frac{4}{(5-\alpha)(4+\alpha)}$
ULFR	$(2L) \frac{2(1-\alpha)}{7+3\alpha}$
SHFR	$(2L) \frac{2\alpha}{(6-\alpha)(2+\alpha)} + \frac{6-\alpha(6-\alpha)}{(6-\alpha)(2+\alpha)} \sigma$
UHFR	$\frac{1-\alpha}{1+\alpha} \sigma$

Now because δ^* is continuous in both α and σ it suffices to perform comparative statics within a fixed type of equilibrium. From the table above it can easily be deduced that we have $\partial \delta^* / \partial \sigma \leq 0$ for all $\alpha \in (0, 1)$. Moreover, it is straightforward to verify that $\partial \delta^* / \partial \alpha < 0$ may only be violated in a specialized high fill-rate equilibrium, where this condition boils down to $\sigma > \frac{12+\alpha^2}{48-18\alpha+\alpha^2} (2L)$. Hence for δ^* to be decreasing in α for all $\alpha \in (0, 1)$ we require $\sigma > \tilde{\sigma}$, where $\tilde{\sigma} \simeq 0.5788$ is the solution to the fixed point problem $\sigma = \frac{12+(\sigma_2^{*-1}(\sigma))^2}{48-18\sigma_2^{*-1}(\sigma)+(\sigma_2^{*-1}(\sigma))^2} (2L)$. ■

7.6 Proof of Proposition 2

On the equilibrium path, the transitions from one state to another follow a Markov chain with transition matrix

$$P^* = \begin{bmatrix} (1-\alpha)(1-p_{|HT}^*) & (1-\alpha)p_{|HT}^* & \alpha(1-p_{|AT}^*) & \alpha p_{|AT}^* \\ 1-\alpha & 0 & \alpha & 0 \\ (1-\alpha)\frac{1}{2} & (1-\alpha)\frac{1}{2} & \alpha(1-q_{|FT}^*) & \alpha q_{|FT}^* \\ 1-\alpha & 0 & \alpha & 0 \end{bmatrix}$$

As in Colliard and Foucault (2012), the stationary probability distribution $\varphi^* = (\varphi_{HT}^{LO*}, \varphi_{HT}^{MO*}, \varphi_{AT}^{LO*}, \varphi_{AT}^{MO*})$ is then simply given by the left eigenvalue of P^* . ■

7.7 Proof of Corollary 3

It is straightforward to compute the trading rates for each type of equilibrium by substituting the respective execution probabilities from the proof of Proposition 1 into the formula for the stationary probability distribution derived in Proposition 2. The following table collects the trading rates for all types of equilibria.

Equilibrium	TR^*
SLFR	$\frac{4+\alpha(1-\alpha)}{(4+\alpha)(5-\alpha)}$
ULFR	$\frac{4+3\alpha(1-\alpha)}{20-\alpha+\alpha^2}$
SHFR	$\frac{4-\alpha(1-\alpha)}{12+\alpha-\alpha^2}$
UHFR	$\frac{1}{3}$

Notice that $TR_0^* = 1/3$ for $\sigma \in \underline{\Sigma}$ and $TR_0^* = 1/5$ for $\sigma \in \bar{\Sigma}$. The result follows immediately. ■

7.8 Proof of Corollary 4

It directly follows from Proposition 2 that we have $LtM_{FT}^* > LtM_{ST}^*$ for all $\alpha \in (0, 1)$. As the limit-to-market order ratios do not *explicitly* depend on σ , the remaining results require to compare the relevant ratio across different types of equilibria. Under Assumption 1, increases in σ given a fixed level of (α, L) may lead to the following transitions: UHFR→SHFR, SHFR→SLFR, SHFR→ULFR, and UHFR→ULFR. The LtM ratios for the different equilibria are easily calculated by substituting the respective equilibrium execution probabilities from the proof of Proposition 1 into the stationary probability distribution of Proposition 2 and contained in the following table.

Equilibrium	LtM_{ST}^*	LtM_{FT}^*	LtM^*
SLFR	$\frac{4(4-\alpha)}{4+5\alpha-\alpha^2}$	$\frac{4(2-\alpha)}{\alpha}$	$\frac{16+4\alpha(5-\alpha)(1-\alpha)}{4-\alpha(1-\alpha)}$
ULFR	$\frac{4(4-\alpha)}{4+3\alpha+\alpha^2}$	$\frac{16(2-\alpha)}{4-\alpha(1-\alpha)}$	4
SHFR	$\frac{8}{4+\alpha-\alpha^2}$	$\frac{2(5-\alpha)(2-\alpha)}{2+\alpha(3-\alpha)}$	$\frac{8+2\alpha(6-\alpha)(1-\alpha)}{4-\alpha(1-\alpha)}$
UHFR	2	$2(2-\alpha)$	$2(1+\alpha(1-\alpha))$

It is easy to see that any of the above transitions leads to an increase in the LtM_{FT}^* and LtM^* ratios. Clearly, for $\alpha > \max_{\sigma} \alpha_S^*(\sigma)$, increases in σ may only lead to a transition from a ULFR to a UHFR equilibrium. Therefore it directly follows that LtM_{ST}^* is increasing in σ for $\alpha < (\sqrt{41}-5)/2$ and decreasing otherwise. ■

7.9 Proof of Corollary 5

First, we need to show that $\omega_{FT,ST}^* = \varphi_{FT}^{LO*}(1-a)q_{|ST}^* \geq \varphi_{ST}^{LO*}\alpha p_{|FT}^* = \omega_{ST,FT}^*$. Lemma 1 implies that $q_{|ST}^* = 1/2$ and using the definition of the trader type-specific trading rates we can rewrite this as $(1 - TR_{FT}^*) \geq (1 - TR_{ST}^*)2p_{|FT}^*$. Because $2p_{|FT}^* \leq 1$, Proposition 2 then implies that this inequality is always satisfied.

The following table collects MT_{ST}^* and MT_{FT}^* for all equilibria.

Equilibrium	MT_{ST}^*	MT_{FT}^*
SLFR	$\frac{4-5\alpha+\alpha^2}{4+5\alpha-\alpha^2}$	$\frac{2-\alpha}{\alpha}$
ULFR	$\frac{4-\alpha}{4+3\alpha+\alpha^2}$	$\frac{8-4\alpha}{4-\alpha+\alpha^2}$
SHFR	$\frac{4-2\alpha}{4+\alpha-\alpha^2}$	$\frac{20-4\alpha}{8+12\alpha-4\alpha^2}$
UHFR	1	1

In order to show that MT_{FT}^* (MT_{ST}^*) is increasing (decreasing) in σ , we again simply require that an equilibrium transition due to an increase in σ yields an increase (decrease) in the ratio (see the Proof of Corollary 5). It is easily verified that this is always the case such that the result follows. ■

7.10 Proof of Corollary 6

First, notice that for the assumed range of σ two different types of equilibria may emerge (SHFR and UHFR) and that we have $E(\tau^*) = \frac{1}{3}L - \frac{2}{3}\sigma$ (see Foucault (1999)). Now consider the UHFR equilibrium, which can only arise if $\alpha > 1/3 = \alpha_S^*(8/15)$. In this case we have $\varphi^* = (\frac{2(1-\alpha)}{3}, \frac{(1-\alpha)}{3}, \frac{2\alpha}{3}, \frac{\alpha}{3})$, which implies that $\frac{\omega_{ST,k}^*}{\omega_{ST,k}^* + \omega_{FT,k}^*} = \frac{\varphi_{ST}^{MO*}}{\varphi_{ST}^{MO*} + \varphi_{FT}^{MO*}} = (1-\alpha)$. Using equations (9) and (??), it is straightforward to show that $E(\tau_{ST}^*) = (1-\alpha)(L - V_{FT}^{LO*} - \sigma) + \alpha(L - V_{ST}^{LO*}) = \frac{1}{3}L - \frac{4-6\alpha}{3(1+\alpha)}\sigma$ and $E(\tau_{FT}^*) = L - V_{FT}^{LO*} - \sigma = \frac{1}{3}L - \frac{4}{3(1+\alpha)}\sigma$, from which we can easily deduce that $E(\tau^*) = \frac{1}{3}L - \frac{4-6\alpha(1-a)}{3(1+\alpha)}\sigma$. Now because $\alpha > 1/3$ we obtain $E(\tau_{FT}^*) > E(\tau^*) > E(\tau_{ST}^*) > E(\tau_0^*)$ as required. The calculations for the SHFR equilibrium involve somewhat more tedious algebra such that we omit them for brevity. ■

7.11 Proof of Corollary 7

Clearly we can never have $\alpha^+ = 1$ because $W^*(1) = W^*(0) - c$. Moreover, if $\sigma \in \underline{\Sigma}$ the socially optimal level of investment is $\alpha^+ = 0$ because the trading rate is at its maximum level of $1/3$ even without FTs, rendering investment inefficient. Hence assume that $\sigma \in \bar{\Sigma}$. Notice that in this

case we have $\Delta W^*(1) = 0$ which implies that we must have $\alpha^* < 1$. Thus consider an interior equilibrium $\alpha^* \in (0, 1)$ and assume that $\alpha^+ > 0$ (otherwise $\alpha^* > \alpha^+$ is trivial). By definition $\frac{\partial W^*(\alpha)}{\partial \alpha} \Big|_{\alpha=\alpha^+} = 0$ and $\Delta W^*(\alpha^*) = c$. Now notice that the chain rule implies that $\frac{\partial W^*(\alpha)}{\partial \alpha} = \Delta W^*(\alpha) + (1-\alpha) \frac{\partial W_{ST}^*(\alpha)}{\partial \alpha} + \alpha \frac{\partial W_{ST}^*(\alpha)}{\partial \alpha} - c$ and it is easy (albeit very cumbersome) to verify that $(1-\alpha) \frac{\partial W_{ST}^*(\alpha)}{\partial \alpha} + \alpha \frac{\partial W_{ST}^*(\alpha)}{\partial \alpha} < 0$ for all α , such that we conclude that $\Delta W^*(\alpha^+) > c$. Finally, because for $\sigma \in \bar{\Sigma}$ we have $\frac{\partial \Delta W^*(\alpha)}{\partial \alpha} < 0$ for all α , this implies $\alpha^* > \alpha^+$.

To show that $W^*(\alpha^*) < W^*(0)$ for $\alpha^* > 0$ it suffices to consider the case where $\sigma \in \bar{\Sigma}$ and $\alpha^* < 1$. While again involving some lengthy computations, it is straightforward to verify that $W^*(0) = 2L/5 > W_{ST}^* = \frac{\varphi_1^*}{\varphi_1^* + \varphi_2^*} V_{ST}^{LO*} + \frac{\varphi_2^*}{\varphi_1^* + \varphi_2^*} (L - E(\tau_{ST}^*))$. ■

8 Appendix B: Figures

8.1 Figure 1: Equilibrium Map

This graph depicts the functions $\sigma_{FT}^*(\alpha)$ (blue), $\sigma_{ST}^*(\alpha)$ (red) and $\alpha_S^*(\sigma)$ (black) in the (α, σ) -space, where we have set $L = 1$. The green lines indicate the interval $[\underline{\sigma}, \bar{\sigma}]$.

