What is the Optimal Trading Frequency in Financial Markets?[*][†]

Songzi Du[‡] Haoxiang Zhu[§]

December 23, 2016

Review of Economic Studies Vol. 84, No. 4 (2017): 1606-1651 DOI: 10.1093/restud/rdx006

*First version: May 2012. Earlier versions of this paper were circulated under various titles, including "Ex Post Equilibria in Double Auctions of Divisible Assets", "Dynamic Ex Post Equilibrium, Welfare, and Optimal Trading Frequency in Double Auctions", and "Welfare and Optimal Trading Frequency in Dynamic Double Auctions". For helpful comments, we are grateful to Alexis Bergès, Bruno Biais, Alessandro Bonatti, Bradyn Breon-Drish, Jeremy Bulow, Giovanni Cespa, Hui Chen, Peter DeMarzo, David Dicks, Darrell Duffie, Thierry Foucault, Willie Fuchs, Lawrence Glosten, Robin Greenwood, Lawrence Harris, Frank Hatheway, Joel Hasbrouck, Richard Haynes, Terry Hendershott, Eiichiro Kazumori, Ilan Kremer, Pete Kyle, Martin Lettau, Stefano Lovo, Andrey Malenko, Katya Malinova, Gustavo Manso, Konstantin Milbradt, Sophie Moinas, Michael Ostrovsky, Jun Pan, Andreas Park, Jonathan Parker, Parag Pathak, Michael Peters, Paul Pfleiderer, Uday Rajan, Marzena Rostek, Ioanid Rosu, Gideon Saar, Xianwen Shi, Andy Skrzypacz, Jonathan Sokobin, Chester Spatt, Sayee Srinivasan, Juuso Toikka, Lin Tong, Dimitri Vayanos, Xavier Vives, Jiang Wang, Yajun Wang, Bob Wilson, Liyan Yang, Amir Yaron, Lei Yu, and Hayong Yun, as well as seminar participants at Stanford University, Simon Fraser University, MIT, Copenhagen Business School, University of British Columbia, UNC Junior Finance Faculty Roundtable, Midwest Theory Meeting, Finance Theory Group Berkeley Meeting, Canadian Economic Theory Conference, HEC Paris, Barcelona Information Workshop, CICF, Stony Brook Game Theory Festival, SAET, Bank of Canada, Carnegie Mellon Tepper, University of Toronto, NBER Microstructure meeting, IESE Business School, Stanford Institute for Theoretical Economics, UBC Summer Workshop in Economic Theory, Toulouse Conference on Trading in Electronic Market, University of Cincinnati, the 10th Central Bank Workshop on the Microstructure of Financial Markets, FIRN Asset Pricing meeting, Imperial College High Frequency Trading Conference, UPenn Workshop on Multiunit Allocation, WFA, Econometric Society World Congress, the National University of Singapore, and the Market Microstructure Confronting Many Viewpoints Conference. We thank the CFTC for providing summary statistics on four futures contracts and NASDAQ for providing summary statistics on 117 stocks that are used to calibrate the model.

†Additional disclaimer related to CFTC: The research presented in this paper was co-authored by Haoxiang Zhu, an unpaid consultant of CFTC, who wrote this paper in his official capacity with the CFTC, and Songzi Du, an Assistant Professor of Economics at Simon Fraser University. (The majority of the paper was written prior to Haoxiang Zhu's affiliation with the CFTC.) The Office of the Chief Economist and CFTC economists and consultants produce original research on a broad range of topics relevant to the CFTCs mandate to regulate commodity futures markets, commodity options markets, and the expanded mandate to regulate the swaps markets pursuant to the Dodd-Frank Wall Street Reform and Consumer Protection Act. These papers are often presented at conferences and many of these papers are later published by peer-review and other scholarly outlets. The analyses and conclusions expressed in this paper are those of the authors and do not reflect the views of other members of the Office of Chief Economist, other Commission staff, or the Commission itself.

‡Simon Fraser University, Department of Economics, 8888 University Drive, Burnaby, B.C. Canada, V5A 1S6. songzid@sfu.ca.

§MIT Sloan School of Management and NBER, 100 Main Street E62-623, Cambridge, MA 02142. zhuh@mit.edu. Since July 2016, Haoxiang Zhu is affiliated with the CFTC as an unpaid academic economist.

What is the Optimal Trading Frequency in Financial Markets?

December 23, 2016

Abstract

This paper studies the impact of increasing trading frequency in financial markets on allocative efficiency. We build and solve a dynamic model of sequential double auctions in which traders trade strategically with demand schedules. Trading needs are generated by time-varying private information about the asset value and private values for owning the asset, as well as quadratic inventory costs. We characterize a linear equilibrium with stationary strategies and its efficiency properties in closed form. Frequent trading (more double auctions per unit of time) allows more immediate asset reallocation after new information arrives, at the cost of a lower volume of beneficial trades in each double auction. Under stated conditions, the trading frequency that maximizes allocative efficiency coincides with the information arrival frequency for scheduled information releases, but can far exceed the information arrival frequency if new information arrives stochastically. A simple calibration of the model suggests that a moderate market slowdown to the level of seconds or minutes per double auction can improve allocative efficiency for assets with relatively narrow investor participation and relatively infrequent news, such as small- and micro-cap stocks.

Keywords: trading frequency, allocative efficiency, high-frequency trading, double auction
JEL Codes: D44, D82, G14

1 Introduction

Trading in financial markets has become significantly faster over the last decade. Today, electronic transactions for equities, futures, and foreign exchange are typically conducted within milliseconds or microseconds. Electronic trading, which is faster than the manual processing of orders, is also increasingly adopted in the over-the-counter markets for debt securities and derivatives, such as corporate bonds, interest-rate swaps, and credit default swaps. Exchange traded funds, which trade at a high frequency similar to stocks, have gained significant market share over index mutual funds, which only allow buying and selling at the end of each day.

The remarkable speedup in financial markets raises important economic questions. For example, does a higher trading frequency necessarily lead to more efficient allocations of assets? What is the socially optimal frequency at which financial transactions should take place? And how does this optimal trading frequency depend on asset characteristics? Answers to these questions would provide valuable insights for the ongoing academic and policy debate on market structure, especially in the context of high-speed trading (see, for example, Securities and Exchange Commission (2010)).

In this paper, we set out to investigate the welfare consequence of speeding up transactions in financial markets. In our model, the trading process is modeled as an infinite sequence of double auctions. The shorter is the time interval between two consecutive auctions, the higher is the trading frequency of the market. As an advantage, a high-frequency market enables investors to respond quickly to new information and start the reallocation of assets. As a disadvantage, a high-frequency market is also "thinner," in the sense that strategic investors become more sensitive to price impact (in a manner described below). This tradeoff, together with the timing of information arrivals, generates the optimal trading frequency that maximizes allocative efficiency.

Model, equilibrium, and efficiency. Our model works roughly as follows. A finite number ($n \geq 3$) of strategic traders trade a divisible asset in an infinite sequence of uniform-price double auctions, held at discrete times $\{0, \Delta, 2\Delta, 3\Delta, ...\}$. At an exponentially-distributed time in the future, the asset pays a liquidating dividend, which, until that payment time, evolves according to a jump process. Over time, traders receive private informative signals about common dividend shocks, as well as idiosyncratic shocks to their private values for owning the asset. Traders' values for the assets are therefore interdependent, creating adverse selection in the trading process.[1] Traders also incur quadratic costs for holding inventories, which is equivalent to linearly decreasing marginal values. A trader's dividend signals, shocks to his private values, and his inventories are all his private information. In each double auction, traders submit demand schedules (i.e., a set of limit orders) and pay for their allocations at the market-clearing price. All traders take into account the price impact of their trades and are forward-looking about future trading opportunities.

Our model incorporates many salient features of dynamic markets in practice. For example, asymmetric and dispersed information about the common dividend creates adverse selection, whereas private-value information and convex inventory costs introduce gains from trade. These trading motives are also time-varying as news arrives over time. In this framework, the number

[1]Throughout this paper, "adverse selection" covers situations in which different traders have different pieces of information regarding the same asset. In our context "adverse selection" may also be read as "interdependent values."

of double auctions per unit of clock time, $1/\Delta$, is a simple yet realistic way to model trading frequency in dynamic markets. We emphasize that a change of trading frequency in our model does not change the fundamental properties of the asset, such as the timing and magnitude of the dividend shocks.

Our first main result is the characterization of a linear equilibrium with stationary strategies in this dynamic market, as well as its efficiency properties. In equilibrium, a trader's optimal demand in each double auction is a linear function of the price, his signal about the dividend, his most recent private value, and his private inventory. Each coefficient is solved explicitly in closed form. Naturally, the equilibrium price in each auction is a weighted sum of the average signal about the common dividend and the average private value, adjusted for the marginal holding cost of the average inventory.

Because there are a finite number of traders, demand schedules in this dynamic equilibrium are not competitive. To reduce price impact, traders engage in "demand reduction," thereby strategically understating how much they are willing to buy or sell at each price. The practical consequence of demand reduction is that a large order is split into many smaller pieces and executed slowly over time. Consequently, the equilibrium asset allocations after each auction are inefficient, although they converge gradually and exponentially over time to the efficient allocation. This convergence remains slow and gradual even in the continuous-time limit. We show that the convergence rate per double auction increases with the number of traders, the arrival intensity of the dividend, and the variance of the private-value shocks; but the convergence rate per auction decreases with the frequency of auctions and with the variance of the common-value shocks. These comparative statics are useful input for our analysis of the optimal trading frequency.

Welfare and optimal trading frequency. Characterizing the optimal trading frequency in this dynamic market is the second primary contribution of our paper.

Increasing trading frequency involves the following important tradeoff. On the one hand, a higher trading frequency allows traders to react to new information and start to adjust their asset holdings more quickly. This effect favors a faster market. On the other hand, under a higher trading frequency traders are less willing to suffer the price impact in any given period (since they anticipate more trading opportunities in future periods), which leads to less aggressive demand schedules in each period and hence a lower volume of beneficial asset reallocations. Although a higher trading frequency by definition generates more double auctions per unit of clock time, we prove that, conditional on news arrivals, the efficiency benefit of having more trading opportunities is more than offset by the inefficiency of less aggressive demand schedules in each round. This effect favors a slower market.

Given these two effects, the allocative inefficiency in this dynamic market relative to the first best can be decomposed into two components: one part due to strategic behavior and the other due to the delayed responses to new information. The optimal trading frequency should strike the best balance between maximizing beneficial asset reallocations and minimizing delays in reacting to new information.

We show that the optimal trading frequency depends critically on the nature of information arrivals. Scheduled information arrivals (e.g., earnings announcements and macroeconomic data releases) and stochastic information arrivals (e.g, mergers and geopolitical events) generate dramatically different optimal trading frequencies.

If new information about dividend and private values arrives at scheduled time intervals, we prove that the optimal trading frequency cannot be higher than the arrival frequency of new information. In the natural case that all traders are ex-ante identical, the optimal trading frequency then coincides with the information arrival frequency. Intuitively, if information arrival times are known in advance, aligning trading times with information arrival times would reap all the benefit of immediate response to new information, while maximizing the volume of beneficial asset reallocations. Under the natural condition that all traders are ex-ante identical, the optimal trading frequency coincides with the information arrival frequency.

By contrast, if new information arrives at unpredictable times, it is important to keep the market open more often to shorten delays in responding to new information. Indeed, we show that under Poisson information arrivals and if traders are ex-ante identical, the optimal trading frequency is always higher than $\left(\frac{n\alpha}{2} - \frac{1}{3}\right)\mu$, where:

- the parameter n is the number of traders;
- the (endogenous) parameter $\alpha \in (0, 1]$ is decreasing with the level of adverse selection, in a sense to be made precise in the model section; and
- the parameter μ is the expected arrival frequency of information.

That is, the optimal trading frequency $1/\Delta^*$ exceeds the information arrival frequency μ by a factor of at least $\frac{n\alpha}{2} - \frac{1}{3}$, which is a large number if the asset has broad investor participation (large n) and mild adverse selection (α close to 1). In numerical calculations, the lower bound $\left(\frac{n\alpha}{2} - \frac{1}{3}\right)\mu$ turns out to be very tight. As n or μ becomes large, the optimal trading frequency increases without an upper bound, approaching continuous trading in the limit.[2]

To concretely illustrate the application of the model and its market-design implications, we calibrate the Poisson-information-arrival version of the model to U.S. futures and equity markets. We use four liquid futures contracts (the E-mini S&P 500 futures, the 10-year Treasury futures, the Euro futures, and the crude oil futures) and a sample of 146 stocks that cover, in approximately equal proportions, large-cap, medium-cap, small-cap, and micro-cap firms. Using reasonable calibrations of the model parameters (some of which are based on proprietary summary statistics provided by the CFTC and NASDAQ), we find that all four futures contracts and the top half of the sample stocks (sorted by the calibrated optimal trading frequency) have a model-implied optimal trading frequency ranging from a few auctions per second to a few thousand auctions per second, depending on the asset and the method of calibrating news arrival rate. By contrast, stocks with relatively low investor participation or with relatively infrequent arrivals of news, such as small- and micro-cap stocks, tend to have a robustly lower optimal trading frequency, in the order of seconds or minutes per double auction. We caution that these optimal trading frequencies should be interpreted in their orders of magnitude but not the exact level.

A policy implication from our analysis is that a moderate market slowdown can improve allocative efficiency in small- and micro-cap stocks, but not necessarily for larger stocks and liquid futures contracts. More generally, sound market design in terms of trading speed should take into account the heterogeneity of assets in terms of investor participation and the arrival frequency of relevant news, among others. Our model and calibration method could also be used to evaluate the potential benefit of speed regulation in other securities and derivatives traded on electronic markets, such as ETFs, options, government securities, and foreign exchange, as

[2] To clarify, continuous trading in our model means continuous double auctions, not a continuous limit order book. The latter is effectively a discriminatory-price auction, not a uniform-price auction.

well as in other jurisdictions. For assets that are currently traded over-the-counter but are moving toward all-to-all electronic trading—such as corporate bonds, interest rate swaps, and credit default swaps—our model suggests that if exchange-like trading in these markets were deemed desirable by investors and regulators, periodic auctions could be a more natural place to start than continuous trading.

Relation to the literature. The paper closest to ours is Vayanos (1999), who studies a dynamic market in which the asset fundamental value (dividend) is public information, but agents receive periodic private inventory shocks. Traders in his model are also fully strategic. Vayanos (1999) shows that, if inventory shocks are small, then a lower trading frequency is better for welfare by encouraging traders to submit more aggressive demand schedules.[3] We make two main contributions relative to Vayanos (1999). First, our model allows interdependent values and adverse selection. Adverse selection makes trading less aggressive and reduces the optimal trading frequency. Second, our model identifies two channels of welfare losses: One channel, strategic behavior, agrees with Vayanos (1999), whereas the other, delayed responses to news, complements Vayanos (1999). The latter channel is absent in Vayanos (1999) because inventory shocks and trading times always coincide in his model. We show that the latter channel can lead to an optimal trading frequency that is much higher than the information arrival frequency if information arrivals are stochastic. Our result also generates useful predictions regarding how the optimal trading frequency varies with asset characteristics.

Rostek and Weretka (2015) study dynamic trading with multiple dividend payments. In their model, traders have symmetric information about the asset's fundamental value, and between consecutive dividend payments there is no news and no discounting. In this setting, they show that a higher trading frequency is better for welfare. Our contribution relative to Rostek and Weretka (2015) is similar to that relative to Vayanos (1999). First, our model applies to markets with adverse selection. Second, we show that the optimal trading frequency depends on the tradeoff between strategic behaviors and delayed responses to news.

Our notation of welfare and those of Vayanos (1999) and Rostek and Weretka (2015) are allocative efficiency. This is different from the welfare question in a number of high-frequency-trading studies, namely whether investments in high-speed trading technology are socially wasteful (see Biais, Foucault, and Moinas (2015), Pagnotta and Philippon (2013), Budish, Cramton, and Shim (2015), and Hoffmann (2014)).

Among recent models that study dynamic trading with adverse selection, the closest one to ours is Kyle, Obizhaeva, and Wang (2014). They study a continuous-time trading model in which agents have pure common values but "agree to disagree" on the precision of their signals. Although the disagreement component in their model and the private-value component in ours appear equivalent, they are in fact very different. As highlighted by Kyle, Obizhaeva, and Wang (2014), in a disagreement model the traders disagree not only about their values today, but also about how the values evolve over time; this behavior does not show up in a private-value model. Therefore, their model and ours answer very different economic questions: Their model generates "beauty contest" and non-martingale price dynamics, whereas our model is useful for characterizing the optimal trading frequency.

[3] Vayanos (1999) also shows that if inventory information is common knowledge, there is a continuum of equilibria. Under one of these equilibria, selected by a trembling hand refinement, welfare is increasing in trading frequency. Because our model has private information about inventories, the private-information equilibrium of Vayanos (1999) is a more appropriate benchmark for comparison.

2 Dynamic Trading in Sequential Double Auctions

This section presents the dynamic trading model and characterizes the equilibrium and its efficiency properties. Main model parameters are tabulated in Appendix A for ease of reference.

2.1 Model

Timing and the double auctions mechanism. Time is continuous, $\tau \in [0, \infty)$. There are $n \geq 3$ risk-neutral traders in the market trading a divisible asset. Trading is organized as a sequence of uniform-price divisible double auctions, held at clock times $\{0, \Delta, 2\Delta, 3\Delta, \ldots\}$, where $\Delta > 0$ is the length of clock time between consecutive auctions. The trading frequency of this market is therefore the number of double auctions per unit of clock time, i.e., $1/\Delta$. The smaller is Δ, the higher is the trading frequency. We will refer to the time interval $[t\Delta, (t+1)\Delta)$ as "period t," for $t \in \{0, 1, 2, \ldots\}$. Thus, the period-$t$ double auction occurs at the clock time $t\Delta$. The top plot of Figure 1 illustrates the timing of the double auctions.

Figure 1: Model time lines. The top plot shows times of double auctions, and the bottom plot shows the news times (dividend shocks, signals of dividend shocks, and private value shocks).

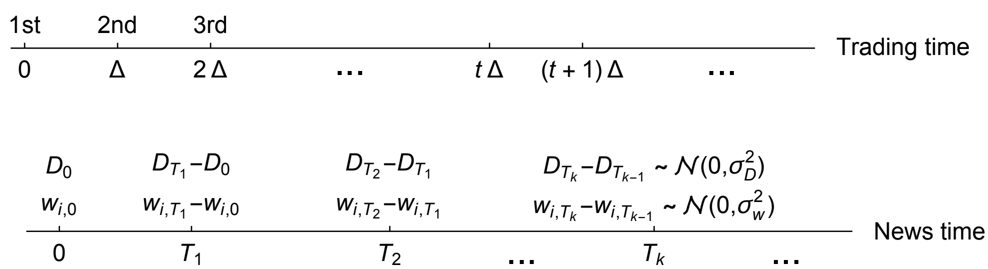

We denote by $z_{i,t\Delta}$ the inventory held by trader i immediately before the period-t double auction. The ex-ante inventories $z_{i,0}$ are given exogenously. The total ex-ante inventory, $Z \equiv \sum_i z_{i,0}$, is a commonly known constant. (In securities markets, Z can be interpreted as the total asset supply. In derivatives markets, Z is by definition zero.) As shown later, while the equilibrium characterization works for any ex-ante inventory profile $\{z_{i,0}\}_{i=1}^n$, in the analysis of trading frequency we will pay particular attention to the special case in which all traders are ex-ante identical (i.e., $z_{i,0} = Z/n$).

A double auction is essentially a demand-schedule-submission game. In period t each trader submits a demand schedule $x_{i,t\Delta}(p) : \mathbb{R} \to \mathbb{R}$. The market-clearing price in period t, $p^*_{t\Delta}$, satisfies

$$\sum_{i=1}^{n} x_{i,t\Delta}(p^*_{t\Delta}) = 0. \tag{1}$$

In the equilibrium we characterize later, the demand schedules are strictly downward-sloping in p and hence the solution $p^*_{t\Delta}$ exists and is unique. The evolution of inventory is given by

$$z_{i,(t+1)\Delta} = z_{i,t\Delta} + x_{i,t\Delta}(p^*_{t\Delta}). \tag{2}$$

After the period-t double auction, each trader i receives $x_{i,t\Delta}(p^*_{t\Delta})$ units of the assets at the price of $p^*_{t\Delta}$ per unit. Of course, a negative $x_{i,t\Delta}(p^*_{t\Delta})$ is a sale.

The asset. Each unit of the asset pays a single liquidating dividend D at a random future time $\mathcal{T}$, where $\mathcal{T}$ follows an exponential distribution with parameter $r > 0$, or mean $1/r$. The random dividend time $\mathcal{T}$ is independent of all else in the model.

Before being paid, the dividend D is unobservable and evolves as follows. At time $T_0 = 0$, $D = D_0$ is drawn from the normal distribution $\mathcal{N}(0, \sigma_D^2)$. Strictly after time 0 but conditional on the dividend time $\mathcal{T}$ having not arrived, the dividend D is shocked at each of the clock times T_1, T_2, T_3, ..., where $\{T_k\}_{k\geq 1}$ can be deterministic or stochastic. The dividend shocks at each T_k, for $k \geq 1$, are also i.i.d. normal with mean 0 and variance σ_D^2:

$$D_{T_k} - D_{T_{k-1}} \sim \mathcal{N}(0, \sigma_D^2). \tag{3}$$

We will also refer to $\{T_k\}_{k\geq 0}$ as "news times" or "information arrival times." Thus, before the dividend is paid, the unobservable dividend $\{D_\tau\}_{\tau\geq 0}$ follows a jump process:

$$D_\tau = D_{T_k}, \text{ if } T_k \leq \tau < T_{k+1}. \tag{4}$$

Therefore, at the dividend payment time $\mathcal{T}$, the realized dividend is $D_{\mathcal{T}}$.

Since the expected dividend payment time is finite $(1/r)$, for simplicity we normalize the discount rate to be zero (i.e., there is no time discounting). Allowing a positive time discounting does not change our qualitative results. Moreover, in the supplemental material to this paper, we provide an extension in which infinitely many dividends are paid sequentially and there is a time discount. The main results of this paper are robust to this extension.

Information and preference. At news time T_k, $k \in \{0, 1, 2, \ldots\}$, each trader i receives a private signal S_{i,T_k} about the dividend shock:

$$S_{i,T_k} = D_{T_k} - D_{T_{k-1}} + \epsilon_{i,T_k}, \text{ where } \epsilon_{i,T_k} \sim \mathcal{N}(0, \sigma_\epsilon^2) \text{ are i.i.d.,} \tag{5}$$

and where $D_{T_{-1}} \equiv 0$. The private signals of trader i are never disclosed to anyone else. If signals about dividend shocks were perfect, i.e. $\sigma_\epsilon^2 = 0$, the information structure of our model would be similar to that of Vayanos (1999).

In addition, beyond the common value $D_{\mathcal{T}}$, each trader i has a private value $w_{i,\mathcal{T}}$ for receiving the dividend. The private values may reflect tax or risk-management considerations. The private values are also shocked at the news times $\{T_k\}_{k\geq 0}$, and each private-value shock is i.i.d. normal random variables with mean zero and variance σ_w^2:

$$w_{i,T_k} - w_{i,T_{k-1}} \sim \mathcal{N}(0, \sigma_w^2), \tag{6}$$

where $w_{i,T_{-1}} \equiv 0$. Written in continuous time, trader i's private-value process $w_{i,\tau}$ is a jump process:

$$w_{i,\tau} = w_{i,T_k}, \text{ if } T_k \leq \tau < T_{k+1}. \tag{7}$$

The private values to trader i are observed by himself and are never disclosed to anyone else.

Therefore, if the dividend is paid at time τ, trader i receives

$$v_{i,\tau} \equiv D_\tau + w_{i,\tau} \tag{8}$$

per unit of asset held.[4]

The bottom plot of Figure 1 illustrates the news times $\{T_k\}_{k\geq 0}$, when dividend shocks, the signals of dividend shocks, and the private-value shocks all arrive. The two plots of Figure 1 make it clear that, in our model, the fundamental properties of the asset are separate from the trading frequency of the market.

Moreover, in an interval $[t\Delta, (t+1)\Delta)$ but before the dividend D is paid, trader i incurs a "flow cost" that is equal to $0.5\lambda z^2_{i,(t+1)\Delta}$ per unit of clock time, where $\lambda > 0$ is a commonly known constant. The quadratic flow cost is essentially a dynamic version of the quadratic cost used in the static models of Vives (2011) and Rostek and Weretka (2012). We can also interpret this flow cost as an inventory cost, which can come from regulatory capital requirements, collateral requirement, or risk-management considerations. (This inventory cost is not strictly risk aversion, however.) Once the dividend is paid, the flow cost no longer applies. Thus, conditional on the dividend having not been paid by time $t\Delta$, each trader suffers the flow cost for a duration of $\min(\Delta, \mathcal{T} - t\Delta)$ within period t, with the expectation

$$\mathbb{E}[\min(\Delta, \mathcal{T} - t\Delta) \mid \mathcal{T} > t\Delta] = \int_0^\infty re^{-r\tau} \min(\tau, \Delta)\, d\tau = \frac{1 - e^{-r\Delta}}{r}, \tag{9}$$

where we have used the fact that, given the memoryless property of exponential distribution, $\mathcal{T} - t\Delta$ is exponentially distributed with mean $1/r$ conditional on $\mathcal{T} > t\Delta$.

Value function and equilibrium definition. For conciseness of expressions, we let $H_{i,\tau}$ be the "history" (information set) of trader i at time τ:

$$H_{i,\tau} = \left\{ \ \{(S_{i,T_l}, w_{i,T_l})\}_{T_l \leq \tau}, \ \{z_{i,t'\Delta}\}_{t'\Delta \leq \tau}, \ \{x_{i,t'\Delta}(p)\}_{t'\Delta < \tau} \ \right\}. \tag{10}$$

That is, $H_{i,\tau}$ contains trader i's asset value-relevant information received up to time τ, trader i's path of inventories up to time τ, and trader i's demand schedules in double auctions before time τ. Notice that by the identity $z_{i,(t'+1)\Delta} - z_{i,t'\Delta} = x_{i,t'\Delta}(p^*_{t'\Delta})$, a trader can infer from $H_{i,\tau}$ the price in any past period t'. Notice also that $H_{i,t\Delta}$ does not include the outcome of the period-t double auction.

Let $V_{i,t\Delta}$ be trader i's period-t continuation value immediately before the double auction at time $t\Delta$. By the definition of $H_{i,\tau}$, trader i's information set right before the period-t double

[4] As in Wang (1994), the unconditional mean of the dividend here is zero, but one could add a positive constant to D so that the probability of $D < 0$ or $v_i < 0$ is arbitrarily small. Moreover, in the markets for many financial and commodity derivatives—including forwards, futures and swaps—cash flows can become arbitrarily negative as market conditions change over time.

auction is $H_{i,t\Delta}$. We can write $V_{i,t\Delta}$ recursively as:

$$V_{i,t\Delta} = \mathbb{E}\Bigg[- x^*_{i,t\Delta} p^*_{t\Delta} + (1 - e^{-r\Delta})(z_{i,t\Delta} + x^*_{i,t\Delta}) v_{i,t\Delta} + e^{-r\Delta} V_{i,(t+1)\Delta} \\ - \frac{1 - e^{-r\Delta}}{r} \cdot \frac{\lambda}{2} (z_{i,t\Delta} + x^*_{i,t\Delta})^2 \,\Bigg|\, H_{i,t\Delta} \Bigg], \tag{11}$$

where $x^*_{i,t\Delta}$ is a shorthand for $x_{i,t\Delta}(p^*_{t\Delta})$. The first term $-x^*_{i,t\Delta} p^*_{t\Delta}$ is trader i's net cash flow for buying $x^*_{i,t\Delta}$ units at $p^*_{t\Delta}$ each. The second term $(1 - e^{-r\Delta})(z_{i,t\Delta} + x^*_{i,t\Delta}) v_{i,t\Delta}$ says that if the dividend is paid during period t, which happens with probability $1 - e^{-r\Delta}$, then trader i receives $(z_{i,t\Delta} + x^*_{i,t\Delta}) v_{i,t\Delta}$ in expectation. (Since shocks to the common dividend and private values have mean zero, trader i's expected value is still $v_{i,t\Delta}$ even if multiple pieces of news arrive during period t.) The third term $e^{-r\Delta} V_{i,(t+1)\Delta}$ says that if the dividend is not paid during period t, which happens with probability $e^{-r\Delta}$, trader i receives the next-period continuation value $V_{i,(t+1)\Delta}$. Finally, the last term $-\frac{1-e^{-r\Delta}}{r} \cdot \frac{\lambda}{2}(z_{i,t\Delta} + x^*_{i,t\Delta})^2$ is the expected quadratic inventory cost incurred during period t for holding $z_{i,t\Delta} + x^*_{i,t\Delta}$ units of the asset (see Equation (9)).

We can expand the recursive definition of $V_{i,t\Delta}$ explicitly:

$$V_{i,t\Delta} = \mathbb{E}\Bigg[- \sum_{t'=t}^{\infty} e^{-r(t'-t)\Delta} x^*_{i,t'\Delta} p^*_{t'\Delta} + \sum_{t'=t}^{\infty} e^{-r(t'-t)\Delta} (1 - e^{-r\Delta}) v_{i,t'\Delta} (z_{i,t'\Delta} + x^*_{i,t'\Delta}) \\ - \frac{1 - e^{-r\Delta}}{r} \sum_{t'=t}^{\infty} e^{-r(t'-t)\Delta} \frac{\lambda}{2} (z_{i,t'\Delta} + x^*_{i,t'\Delta})^2 \,\Bigg|\, H_{i,t\Delta} \Bigg]. \tag{12}$$

While trader i's continuation value $V_{i,t\Delta}$ can in principle depend on everything in his information set $H_{i,t\Delta}$, in the equilibrium we characterize, $V_{i,t\Delta}$ depends on trader i's current pre-auction inventory $z_{i,t\Delta}$, his current private value $w_{i,t\Delta}$, and the sum of his private signals $\sum_{l: T_l \le t\Delta} S_{i,T_l}$ about the dividend.

Definition 1 (Perfect Bayesian Equilibrium)**.** A perfect Bayesian equilibrium is a strategy profile $\{x_{i,t\Delta}\}_{1 \le i \le n, t \ge 0}$, where each $x_{i,t\Delta}$ depends only on $H_{i,t\Delta}$, such that for every trader i and at every path of his information set $H_{i,t\Delta}$, trader i has no incentive to deviate from $\{x_{i,t'\Delta}\}_{t' \ge t}$. That is, for every alternative strategy $\{\tilde{x}_{i,t'\Delta}\}_{t' \ge t}$, we have:

$$V_{i,t\Delta}(\{x_{i,t'\Delta}\}_{t' \ge t}, \{x_{j,t'\Delta}\}_{j \ne i, t' \ge t}) \;\ge\; V_{i,t\Delta}(\{\tilde{x}_{i,t'\Delta}\}_{t' \ge t}, \{x_{j,t'\Delta}\}_{j \ne i, t' \ge t}). \tag{13}$$

2.2 The competitive benchmark equilibrium

Before solving this model with imperfect competition and strategic trading, we first solve a competitive benchmark in which all traders take prices as given. In doing so, we will also solve the traders' inference of the dividend D from equilibrium prices. The solution to this inference problem in the competitive equilibrium will be used directly in solving the strategic equilibrium later.

For clarity, we use the superscript "c" to label the strategies, allocations, and prices in the competitive equilibrium. In each period t each trader i maximizes his continuation value $V_{i,t\Delta}$, defined in Equation (12), by choosing the optimal demand schedule $x^c_{i,t\Delta}(p^c_{t\Delta})$, taking as given the period-t price and the strategies of his own and other traders in subsequent periods.

We start by conjecturing that the competitive demand schedule $x^c_{i,t\Delta}(p^c_{t\Delta})$ in period t is such that trader i's expected marginal value for holding $z^c_{i,t\Delta} + x^c_{i,t\Delta}(p^c_{t\Delta})$ units of the asset for the indefinite future is equal to the price $p^c_{t\Delta}$, for every $p^c_{t\Delta}$. That is, we conjecture that

$$\mathbb{E}\left[v_{i,t\Delta} \mid H_{i,t\Delta}, p^c_{t\Delta}\right] - \frac{\lambda}{r}(z^c_{i,t\Delta} + x^c_{i,t\Delta}(p^c_{t\Delta})) = p^c_{t\Delta}, \tag{14}$$

where the term λ/r takes into account that the marginal holding cost is incurred for an expected duration of time $1/r$. This conjecture can be rewritten as:

$$x^c_{i,t\Delta}(p^c_{t\Delta}) = -\, z^c_{i,t\Delta} + \frac{r}{\lambda}\left(\mathbb{E}\left[v_{i,t\Delta} \mid H_{i,t\Delta}, p^c_{t\Delta}\right] - p^c_{t\Delta}\right). \tag{15}$$

The bulk of the remaining derivation involves finding an explicit expression for $\mathbb{E}[v_{i,t\Delta} \mid H_{i,t\Delta}, p^c_{t\Delta}]$. After that the optimal strategy is derived and verified.

Without loss of generality, let us focus on the period-t double auction and suppose that the latest dividend shock is the k-th. Conditional on the unbiased signals $\{S_{j,T_l}\}_{0\le l\le k}$ of dividend shocks, trader j's expected value of the dividend $D_{t\Delta}$ is a multiple of $\sum_{l=0}^k S_{j,T_l}$. Moreover, his private value, w_{j,T_k}, is perfectly observable to him. We thus conjecture that each trader j uses the following symmetric linear strategy:

$$x^c_{j,t\Delta}(p) = A_1 \sum_{l=0}^{k} S_{j,T_l} + A_2 w_{j,T_k} - \frac{r}{\lambda}p - z^c_{j,t\Delta} + fZ, \tag{16}$$

where A_1, A_2, and f are constants and where we have plugged in the coefficients of p and $z_{j,t\Delta}$ from Equation (15). In particular, trader j puts a weight of A_1 on his common-value information and a weight of A_2 on his private value.

By market clearing and the fact that $\sum_j z_{j,t\Delta} = Z$ is common knowledge, each trader i is able to infer

$$\sum_{j\neq i}\left(A_1 \sum_{l=0}^{k} S_{j,T_k} + A_2 w_{j,T_k}\right) \tag{17}$$

from the equilibrium price $p^c_{t\Delta}$. Thus, each trader i infers his value $v_{i,T_k} \equiv D_{T_k} + w_{i,T_k}$ by taking the conditional expectation:

$$
\begin{aligned}
&\mathbb{E}\left[v_{i,T_k} \,\middle|\, H_{i,T_k}, \sum_{j\neq i}\left(A_1\sum_{l=0}^{k} S_{j,T_l} + A_2 w_{j,T_k}\right)\right] \\
&= w_{i,T_k} + \mathbb{E}\left[D_{T_k} \,\middle|\, \sum_{l=0}^{k} S_{i,T_l}, \sum_{j\neq i}\left(A_1\sum_{l=0}^{k} S_{j,T_l} + A_2 w_{j,T_k}\right)\right] \\
&= w_{i,T_k} + B_1\sum_{l=0}^{k} S_{i,T_l} + B_2 \underbrace{\sum_{j\neq i}\left(A_1\sum_{l=0}^{k} S_{j,T_l} + A_2 w_{j,T_k}\right)}_{\text{Inferred from } p^c_{t\Delta}},
\end{aligned}
\tag{18}
$$

where we have used the projection theorem for normal distribution and where the constants B_1 and B_2 are functions of A_1, A_2, and other primitive parameters. In particular, trader i's conditional expected value has a weight of B_1 on his common-value information $\sum_{l=0}^{k} S_{i,T_l}$ and a weight of 1 on his private value w_{i,T_k}. The third term is inferred from the price.

Because trader i's competitive strategy $x^c_{i,t\Delta}$ is linear in $\mathbb{E}[v_{i,T_k} \mid p^*_{t\Delta}, H_{i,T_k}]$, trader i's weight on his common-value information and his weight on the private value have a ratio of B_1. But by symmetric strategies, this ratio must be consistent with the conjectured strategy to start with, i.e., $B_1 = A_1/A_2$. In Appendix C.1, we explicitly calculate that this symmetry pins down the ratio to be $B_1 = A_1/A_2 \equiv \chi$, where $\chi \in (0,1)$ is the unique solution to[5]

$$
\frac{1/(\chi^2\sigma_\epsilon^2)}{1/(\chi^2\sigma_D^2) + 1/(\chi^2\sigma_\epsilon^2) + (n-1)/(\chi^2\sigma_\epsilon^2 + \sigma_w^2)} = \chi. \tag{19}
$$

On the left-hand side of Equation (19), we apply the projection theorem to Equation (18) to derive the weight B_1 as a function of $A_1/A_2 \equiv \chi$. The projection theorem weighs the precision of the noise $\chi\epsilon_{i,T_k}$ in trader i's dividend signal, against the precision of the dividend shock $\chi(D_{T_k} - D_{T_{k-1}})$ and the precision of others' dividend noise and private value $\sum_{j\neq i}(\chi\epsilon_{j,T_k} + w_{j,T_k})$.

We define the "total signal" $s_{i,t\Delta}$ by

$$
\begin{aligned}
s_{i,T_k} &\equiv \frac{\chi}{\alpha}\sum_{l=0}^{k} S_{i,T_l} + \frac{1}{\alpha} w_{i,T_k}, \\
s_{i,\tau} &= s_{i,T_k}, \ \text{ for } \tau \in [T_k, T_{k+1}),
\end{aligned}
\tag{20}
$$

where the scaling factor α is defined to be

$$
\alpha \equiv \frac{\chi^2\sigma_\epsilon^2 + \sigma_w^2}{n\chi^2\sigma_\epsilon^2 + \sigma_w^2} > \frac{1}{n}. \tag{21}
$$

Trader i's total signal incorporates the two-dimensional information $(\sum_{l=0}^{k} S_{i,T_l}, w_{i,T_k})$ in a linear

[5] The left-hand side of Equation (19) is decreasing in χ. It is $1/(1 + \sigma_\epsilon^2/\sigma_D^2) > 0$ if $\chi = 0$ and is $1/(1 + \sigma_\epsilon^2/\sigma_D^2 + (n-1)/(1+\sigma_w^2/\sigma_\epsilon^2)) < 1$ if $\chi = 1$. Hence, Equation (19) has a unique solution $\chi \in \mathbb{R}$, and such solution satisfies $\chi \in (0,1)$.

combination with weights χ/α and $1/\alpha$.

This construction of total signals leads to a very intuitive expression of the conditional expected value v_{i,T_k}. Direct calculation implies that (see details in Appendix C.1, Lemma 1)

$$\mathbb{E}\left[v_{i,T_k}\middle|H_{i,T_k},\sum\nolimits_{j\neq i}s_{j,T_k}\right]=\alpha s_{i,T_k}+\frac{1-\alpha}{n-1}\underbrace{\sum\nolimits_{j\neq i}s_{j,T_k}}_{\text{Inferred from } p^c_{t\Delta}}. \tag{22}$$

Equation (22) says that conditional on his own information and $\sum_{j\neq i}s_{j,T_k}$ (inferred from the equilibrium price), trader i's expected value of the asset is a weighted average of the total signals, with a weight of $\alpha>1/n$ on his own total signal s_{i,T_k} and a weight of $(1-\alpha)/(n-1)<1/n$ on each of the other traders' total signal s_{j,T_k}. The weights differ because other traders' total signals include both common dividend information and their private values, and others' private values are essentially "noise" to trader i (hence under-weighting).

Substituting the conditional expected value of Equation (22) into Equation (15), we have

$$x^c_{i,t\Delta}(p^c_{t\Delta})=-z^c_{i,t\Delta}+\frac{r}{\lambda}\left(\alpha s_{i,t\Delta}+\frac{1-\alpha}{n-1}\sum_{j\neq i}s_{j,t\Delta}\right)-\frac{r}{\lambda}p^c_{t\Delta}, \tag{23}$$

By market clearing, $\sum_i x^c_{i,t\Delta}(p^c_{t\Delta})=0$, we solve

$$p^c_{t\Delta}=\frac{1}{n}\sum_{j=1}^{n}s_{j,t\Delta}-\frac{\lambda}{rn}Z. \tag{24}$$

The first term of $p^c_{t\Delta}$ is the average total signal, and the second term is the marginal cost of holding the average inventory Z/n for an expected duration of time $1/r$.

Substituting $\sum_{j\neq i}s_{j,t\Delta}$ from Equation (24) back to the expression of $x^c_{i,t\Delta}(p^c_{t\Delta})$ in Equation (23), we obtain explicitly the competitive demand schedule:

$$x^c_{i,t\Delta}(p)=\frac{r(n\alpha-1)}{\lambda(n-1)}\left(s_{i,t\Delta}-p-\frac{\lambda(n-1)}{r(n\alpha-1)}z^c_{i,t\Delta}+\frac{\lambda(1-\alpha)}{r(n\alpha-1)}Z\right). \tag{25}$$

Appendix C.2 verifies that under this strategy the first-order condition of trader i's value function (12) can indeed be written in the form of Equation (15). The second-order condition is satisfied as $n\alpha>1$ by the definition of α.

The after-auction allocation in the competitive equilibrium in period t is:

$$z^c_{i,(t+1)\Delta}=z^c_{i,t\Delta}+x^c_{i,t\Delta}(p^c_{t\Delta})=\frac{r(n\alpha-1)}{\lambda(n-1)}\left(s_{i,t\Delta}-\frac{1}{n}\sum_{j=1}^{n}s_{j,t\Delta}\right)+\frac{1}{n}Z. \tag{26}$$

That is, after each double auction, each trader is allocated the average inventory plus a constant multiple of how far his total signal deviates from the average total signal. The key feature of the competitive equilibrium is that the after-auction inventory does not depend on the pre-auction inventory. This is because the pre-auction inventory enters the demand schedule with the coefficient -1 (see Equation (25)). As we will see in the next subsection, this property of the competitive equilibrium will not hold in the strategic equilibrium. We also see that the competitive inventories are martingales since total signals are martingales. We refer to this allocation as the "competitive allocation."

The following proposition summarizes the competitive equilibrium.

Proposition 1. *In the competitive equilibrium, the strategies are given by Equation (25), the price by Equation (24), and the allocations by Equation (26).*

2.3 Characterizing the strategic equilibrium

Having solved a competitive benchmark, we now turn to the equilibrium with imperfect competition and strategic behavior, i.e., traders take into account the impact of their trades on prices. The equilibrium is stated in the following proposition.

Proposition 2. *Suppose that $n\alpha > 2$, which is equivalent to*

$$\frac{1}{n/2 + \sigma_\epsilon^2/\sigma_D^2} < \sqrt{\frac{n-2}{n}}\frac{\sigma_w}{\sigma_\epsilon}. \tag{27}$$

With strategic trading, there exists a perfect Bayesian equilibrium in which every trader i submits the demand schedule

$$x_{i,t\Delta}(p; s_{i,t\Delta}, z_{i,t\Delta}) = b\left(s_{i,t\Delta} - p - \frac{\lambda(n-1)}{r(n\alpha-1)}z_{i,t\Delta} + \frac{\lambda(1-\alpha)}{r(n\alpha-1)}Z\right), \tag{28}$$

where

$$b = \frac{(n\alpha-1)r}{2(n-1)e^{-r\Delta}\lambda}\left((n\alpha-1)(1-e^{-r\Delta}) + 2e^{-r\Delta} - \sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2 + 4e^{-r\Delta}}\right) > 0. \tag{29}$$

The period-t equilibrium price is

$$p^*_{t\Delta} = \frac{1}{n}\sum_{i=1}^{n} s_{i,t\Delta} - \frac{\lambda}{rn}Z. \tag{30}$$

The derivation of the strategic equilibrium follows similar steps to that of the competitive equilibrium derived in Section 2.2. The details of equilibrium construction are delegated to Appendix C.3. Below, we discuss the key intuition of the strategic equilibrium by comparing it with the competitive one.

Let us start with common properties shared between the strategic equilibrium and the competitive one. For example, both equilibria have the same price. Because prices are equal, inference from prices is the same across both equilibria; hence, the total signals $\{s_{i,t\Delta}\}$ that consolidate traders' information about the common dividend and private values are constructed in the same way. The price aggregates the most recent total signals $\{s_{i,t\Delta}\}$. Since the total signals are martingales, the equilibrium prices over time also form a martingale.

The second term $-\lambda Z/(nr)$ in $p^*_{t\Delta}$ and $p^c_{t\Delta}$ is the expected marginal cost of holding the average inventory Z/n until the dividend is paid, i.e., for an expected duration of time $1/r$.

Although each trader learns from $p^*_{t\Delta}$ the average total signal $\sum_i s_{i,t\Delta}/n$ in period t, he does not learn the total signal or inventory of any other *individual* trader. Nor does a trader perfectly distinguish between the common-value component and the private-value component of the price. Thus, private information is not fully revealed after each round of trading. Finally, the equilibrium strategies in Equations (28) and (25) are stationary: a trader's strategy only depends on his most recent total signal $s_{i,t\Delta}$ and his current inventory $z_{i,t\Delta}$, but does not depend explicitly on t.[6]

There are two important differences between the strategic equilibrium of Proposition 2 and the competitive benchmark in Section 2.2. First, in the strategic equilibrium, rather than taking the price as given, each trader in each period effectively selects a price-quantity pair from the residual demand schedule of all other traders. To mitigate price impact, they trade less aggressively in the strategic equilibrium than in the competitive equilibrium. Formally, the endogenous coefficient b in Equation (28) is strictly smaller than $\frac{r(n\alpha-1)}{\lambda(n-1)}$ in Equation (25):

$$\frac{b}{\frac{r(n\alpha-1)}{\lambda(n-1)}} = 1 + \frac{(n\alpha-1)(1-e^{-r\Delta}) - \sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2 + 4e^{-r\Delta}}}{2e^{-r\Delta}} < 1. \tag{31}$$

This feature is the familiar "bid shading" or "demand reduction" in models of divisible auction (see Ausubel, Cramton, Pycia, Rostek, and Weretka 2014). The coefficient b captures how much additional quantity of the asset a trader is willing to buy if the price drops by one unit per period. Thus, a smaller b corresponds to a less aggressive demand schedule. As the number n of traders tends to infinity, the ratio in Equation (31) tends to 1, so the strategic equilibrium converges to the competitive equilibrium.

Intimately related to the aggressiveness of demand schedules is the extent to which a trader "liquidates" his inventory in each trading round. In the competitive equilibrium strategy $x^c_{i,t\Delta}$, the coefficient in front of $z^c_{i,t\Delta}$ is -1, meaning that each trader liquidates his inventory entirely. By contrast, under the strategy $x_{i,t\Delta}$ of Proposition 2, the coefficient in front of $z_{i,t\Delta}$ is

$$d \equiv -b\frac{\lambda(n-1)}{r(n\alpha-1)} = -1 + \frac{\sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2 + 4e^{-r\Delta}} - (n\alpha-1)(1-e^{-r\Delta})}{2e^{-r\Delta}}, \tag{32}$$

which, under the condition $n\alpha > 2$, is strictly between -1 and 0. Thus, each trader only liquidates a fraction $|d| < 1$ of his inventory, leaving a fraction $1+d \in (0,1)$. Partial liquidation of inventory implies that the strategy in period t has an impact on strategies in all future periods, and that an inefficient allocation in one period affects the inefficiency in all future periods. The

[6] Some readers may wonder why our model does not have the infinite-regress problem of beliefs about beliefs, beliefs about beliefs about beliefs, etc. The reason is that the equilibrium price reveals the average total signal in each period; thus, a trader's belief about the common dividend, as well as his potential high-order beliefs, is actually spanned by this trader's own private information and the equilibrium price. This logic was previously used by He and Wang (1995) and Foster and Viswanathan (1996) to show that the potential infinite-regress problem is resolved in their dynamic models with heterogenous information. Our assumption that the common dividend and private values evolve as random walks implies that only the current price has the most updated information and hence allows us to characterize a linear equilibrium with stationary strategies. We note that the random walk assumption is probably stronger than necessary to obtain tractability. For example, if everyone's dividend signals and private values decay to zero at a same constant rate, then the model is expected to remain tractable; we thank an anonymous referee for pointing out this possibility with decay.

next subsection investigates in details how the quantity $1 + d$ determines allocative inefficiency, which is ultimately related to the optimal trading frequency that we study in Section 3.

Relative to the competitive benchmark, the second important difference of the strategic equilibrium is that its existence requires $n\alpha > 2$. If and only if $n\alpha > 2$ is the coefficient b positive, i.e., demand is decreasing in price.[7] Intuitively, if a trader observes a higher equilibrium price, he infers that other traders have either higher private values or more favorable information about the common dividend. If the trader attributes too much of the higher price to a higher dividend, he may end up buying more conditional on a higher price, which leads to a negative b and violates the second order condition. Learning from prices does not cause such a problem in the competitive equilibrium because a higher price there also reflects traders' disregard of price impact. Thus, conditional on the same price, traders do not learn as much about the dividend in the competitive equilibrium as in the strategic one.

The condition $n\alpha > 2$ requires that adverse selection regarding the common dividend is not "too large" relative to the gains from trade over private values. We can show that it is equivalent to condition (27), written in primitive parameters. All else equal, condition (27) holds if n is sufficiently large, if signals of dividend shocks are sufficiently precise (i.e., σ_ϵ^2 is small enough), if new information on the common dividend is not too volatile (i.e., σ_D^2 is small enough), or if shocks to private values are sufficiently volatile (i.e., σ_w^2 is large enough). All these conditions reduce adverse selection.

In particular, the condition $n\alpha > 2$ is trivially satisfied if $\alpha = 1$, which applies if dividend information is public ($\sigma_\epsilon^2 = 0$ and $\sigma_w^2 > 0$) or if traders have pure private values ($\sigma_D^2 = 0$ and $\sigma_w^2 > 0$). Securities with public dividend information correspond to high-quality government bonds like those issued by the United States, Germany, Japan, or the United Kingdom. Securities that are riskier but have low degrees of adverse selection also have an α close to 1. For instance, a broad equity market index like the S&P 500 probably has a very low signal-to-noise ratio $\sigma_D^2/\sigma_\epsilon^2$, since very few asset managers can consistently beat the overall stock market. Letting $\sigma_D^2/\sigma_w^2 \to 0$ in Equation (19), we get $\chi \to 0$ and hence $\alpha \to 1$. By contrast, an individual stock's α is likely smaller because adverse selection at the single stock level is generally more severe than at the index level. These observations are consistent with the fact that trading large stock indices incurs lower bid-ask spreads than trading single stocks, especially small stocks.[8]

We close this subsection with a brief discussion of equilibrium uniqueness. Since news times and trading times are separate in our model, it could happen that no new information arrives during one or more periods. For example, if no new information arrives in the time interval $((t-1)\Delta, t\Delta]$, then the period-t double auction will have the same price as the period-$(t-1)$ double auction, i.e., the period-t double auction looks like a public-information game. Vayanos (1999) shows that public-information games admit a continuum of equilibria, and he uses a trembling-hand argument to select one of them.

[7]The existence condition for our equilibrium is analogous to Kyle, Obizhaeva, and Wang (2014)'s equilibrium existence condition that each trader believes that his signal about the asset value is roughly at least twice as precise as others traders believe it to be.

[8]For instance, SPDR reports that the average bid-ask spread for the SPY, the SPDR ETF for the S&P 500 index, is on average 1 cent, or below 1 basis point of the price level (100 basis points is 1%). See https://www.ssga.com/investment-topics/general-investing/why-spy-size-liquidity-and-low-cost-of-ownership.pdf. In contrast, Brogaard, Hendershott, and Riordan (2014) report that the relative bid-ask spreads of large, medium, and small U.S. stocks are 4.7, 14.6 and 38.1 basis points, respectively.

Our approach to equilibrium selection is to impose stationarity, i.e., the coefficients in the linear strategy are the same across all periods. Going back to the example, if no new information arrives in $((t-1)\Delta, t\Delta]$, the stationarity-selected equilibrium in the period-t double auction will be identical to one in which fresh news does arrive in $((t-1)\Delta, t\Delta]$ but the realizations of the dividend shock, the n signals of dividend shocks, and the n private-value shocks all turn out to be zero.[9] The following proposition shows that the equilibrium of Proposition 2 is unique if strategies are restricted to be linear and stationary.

Proposition 3. *The equilibrium from Proposition 2 is the unique perfect Bayesian equilibrium in the following class of strategies:*

$$x_{i,t\Delta}(p) = \sum_{T_l \leq t\Delta} a_l S_{i,T_l} + a_w w_{i,t\Delta} - bp + dz_{i,t\Delta} + f, \tag{33}$$

where $\{a_l\}_{l\geq 0}, a_w, b, d$ *and* f *are constants.*

As the proof of Proposition 2 makes clear, each trader's optimal strategy belongs to class (33) if other traders also use strategy from class (33). Therefore, Equation (33) is not a restriction on the traders' strategy space, but rather a restriction on the domain of equilibrium uniqueness. (We have not ruled out the existence of non-linear equilibria.)

2.4 Efficiency and comparative statics

We now study the allocative efficiency (or inefficiency) in the equilibrium of Proposition 2. The results of this section lay the foundation for the study of optimal trading frequency in the next section.

We denote by $\{z^*_{i,\tau}\}$ the continuous-time inventory path obtained in the strategic equilibrium of Proposition 2, and denote by $\{z^c_{i,\tau}\}$ the continuous-time inventory path obtained in the competitive equilibrium of Proposition 1. For any clock-time $\tau \in (t\Delta, (t+1)\Delta]$, they are defined by:

$$z^*_{i,\tau} = z^*_{i,(t+1)\Delta} = z^*_{i,t\Delta} + x_{i,t\Delta}(p^*_{t\Delta}; s_{i,t\Delta}, z^*_{i,t\Delta}), \tag{34}$$

$$z^c_{i,\tau} = z^c_{i,(t+1)\Delta} = z^c_{i,t\Delta} + x^c_{i,t\Delta}(p^c_{t\Delta}; s_{i,t\Delta}, z^c_{i,t\Delta}). \tag{35}$$

Figure 2 shows a sample path of $z^*_{i,\tau}$ and $z^c_{i,\tau}$. Again, $z^*_{i,(t+1)\Delta}$ and $z^c_{i,(t+1)\Delta}$ are the inventories of trader i in the time interval $(t\Delta, (t+1)\Delta)$ in the two equilibria.

[9] We note that the continuation value $V_{i,t\Delta}$ (cf. Equation (12)) is a stationary function of the total signal $s_{i,t\Delta}$ and inventory $z_{i,t\Delta}$, whether or not there is new information in $((t-1)\Delta, t\Delta]$. This stationarity is different from that in Rostek and Weretka (2015) because in our model traders get discounted flow utility as they trade, while in Rostek and Weretka (2015) traders get utility at the end of trading rounds (between consecutive dividend shocks).

Figure 2: Illustration of $z^*_{i,\tau}$ and $z^c_{i,\tau}$.

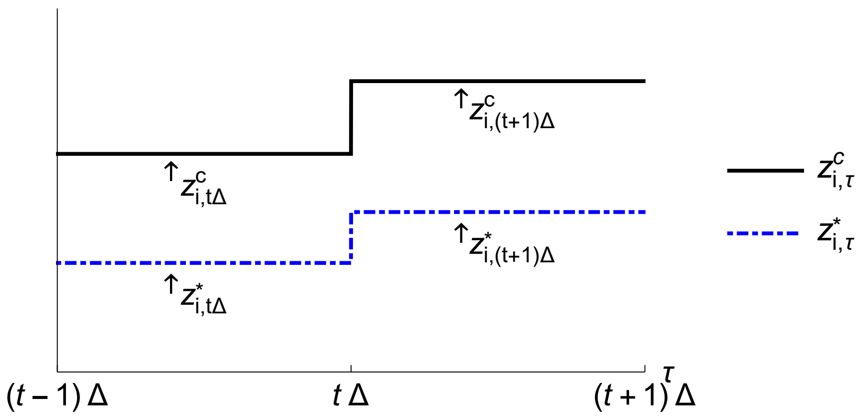

The inventories $\{z^*_{i,t\Delta}\}$ obtained by the strategic equilibrium evolve according to:

$$\begin{aligned} z^*_{i,(t+1)\Delta} &= z^*_{i,t\Delta} + x_{i,t\Delta}(p^*_{t\Delta}; s_{i,t\Delta}, z^*_{i,t\Delta}) \\ &= (1+d)z^*_{i,t\Delta} + b\left(s_{i,t\Delta} - \frac{1}{n}\sum_j s_{j,t\Delta} + \frac{\lambda(n-1)}{r(n\alpha-1)}\frac{Z}{n}\right) \\ &= b\left(s_{i,t\Delta} - \frac{1}{n}\sum_j s_{j,t\Delta}\right) + \frac{1}{n}Z + (1+d)\left(z^*_{i,t\Delta} - \frac{1}{n}Z\right) \\ &= (1+d)z^*_{i,t\Delta} - dz^c_{i,(t+1)\Delta}, \end{aligned} \tag{36}$$

where in the second line we have substituted in the equilibrium strategy $x_{i,t\Delta}$ and the equilibrium price $p^*_{t\Delta}$, the third line follows from the identity of $d/b = -\frac{\lambda(n-1)}{r(n\alpha-1)}$ in the equilibrium strategy, and in the last line we have substituted in the competitive allocation of Equation (26).

Comparing Equation (36) to Equation (26), we can see two differences. First, the after-auction allocation in the strategic equilibrium has an extra term $(1+d)(z^*_{i,t\Delta} - Z/n)$. Since $1+d \in (0,1)$, any inventory imbalance at the beginning of period t partly carries over to the next period. As discussed in the previous subsection, this is a direct consequence of demand reduction caused by strategic trading. Second, because inventories cannot be liquidated quickly due to strategic bidding, traders are more reluctant to acquire inventory. Therefore, the coefficient in front of $s_{i,t\Delta} - \sum_j s_{j,t\Delta}/n$ in the strategic allocation (36) is smaller than that in the competitive allocation (26). That is, strategic bidding makes after-auction asset allocations less sensitive to the dispersion of information (as measured by the total signals).

The above derivation directly leads to the exponential convergence to the competitive allocation over time, shown in the next proposition.

Proposition 4. *Suppose that for some* $0 \le \underline{t} \le \bar{t}$, $s_{i,t\Delta} = s_{i,\underline{t}\Delta}$ *for every* i *and every* $t \in \{\underline{t}, \underline{t}+1, \ldots, \bar{t}\}$. *Then, for every* i, *the equilibrium inventories* $z^*_{i,t\Delta}$ *satisfy:*

$$z^*_{i,(t+1)\Delta} - z^c_{i,(\underline{t}+1)\Delta} = (1+d)^{t+1-\underline{t}}(z^*_{i,\underline{t}\Delta} - z^c_{i,(\underline{t}+1)\Delta}), \quad \forall t \in \{\underline{t}, \underline{t}+1, \ldots, \bar{t}\}, \tag{37}$$

where $d \in (-1, 0)$ is given by Equation (32).

Moreover, $1+d$ is decreasing in n, r, and σ_w^2, but increasing in σ_D^2. As Δ increases, $1+d$ decreases, and the time-discounted geometric sum

$$L \equiv (1-e^{-r\Delta}) \sum_{t=0}^{\infty} e^{-r\Delta t}(1+d)^{2(t+1)} = \frac{(1-e^{-r\Delta})(1+d)^2}{1-e^{-r\Delta}(1+d)^2}, \tag{38}$$

decreases as well.

Equation (37) says that the after-auction allocation $z^*_{i,(t+1)\Delta}$, $\underline{t} \le t \le \bar{t}$, converges exponentially to $z^c_{i,(\underline{t}+1)\Delta}$, which is the competitive allocation given the set of total signals $\{s_{i,\underline{t}\Delta}\}_{i=1}^n$. (Recall that $z^c_{i,(\underline{t}+1)\Delta}$ denotes the competitive allocation *right before* the double auction at time $(\underline{t}+1)\Delta$, hence the time subscript $\underline{t}+1$ instead of $\underline{t}$.) In addition, the competitive allocation does not change from the clock time $\underline{t}\Delta$ to the clock time $\bar{t}\Delta$ because the total signals do not change during this time interval. If $\underline{t} = \bar{t} = t$, Equation (37) is equivalent to Equation (36). The case of $\underline{t} < \bar{t}$ follows by mathematical induction.

Proposition 4 reveals that the strategic equilibrium is inefficient in allocating assets, although the allocative inefficiency converges to zero exponentially over time (as long as no new information arrives). After new dividend shocks and private-value shocks, the competitive allocation changes accordingly, and the strategic allocation starts to converge toward the new competitive allocation exponentially. Exponential convergence of this kind is previously obtained in the dynamic model of Vayanos (1999) under the assumption that common-value information is public.

The comparative statics of $1+d$ with respect to n, r, σ_w^2, and σ_D^2 are all intuitive. A smaller $1+d$ means a faster convergence to efficiency. A larger n makes traders more competitive, and a larger r makes them more impatient. Both effects encourage aggressive bidding and speed up convergence. A large σ_D^2 implies a large uncertainty of a trader about the common asset value and a severe adverse selection; hence, in equilibrium the trader reduces his demand or supply relative to the fully competitive market. Therefore, a higher σ_D^2 implies less aggressive bidding and slower convergence to the competitive allocation. The effect of σ_w^2 is the opposite: a higher σ_w^2 implies larger gains from trade, and hence more aggressive bidding and faster convergence to the competitive allocation.[10] The effects of n, σ_D^2 and σ_w^2 on bidding aggressiveness are present in the earlier static models of Vives (2011) and Rostek and Weretka (2012). The effect of σ_D^2 in reducing the convergence speed to efficiency is also confirmed by Sannikov and Skrzypacz (2014) in a continuous-time trading model.

The comparative statics with respect to Δ are more novel and subtle. First, $1+d$ is smaller if Δ is larger, that is, the convergence per period is faster if the trading frequency is lower. Intuitively, if traders have more subsequent opportunities to trade, say once every second, they are less willing to suffer any price impact now. Hence, their demand schedules are less aggressive in each period.

The time-discounted sum L, as defined in Equation (38), has an intuitive interpretation. It

[10]The comparative static of $1+d$ with respect to σ_ϵ^2 is ambiguous. It can be shown that the endogenous parameters α and χ, and hence the speed of convergence, depend on the "normalized variances" $\sigma_D^2/\sigma_\epsilon^2$ and $\sigma_w^2/\sigma_\epsilon^2$. As σ_ϵ^2 increases, $\sigma_D^2/\sigma_\epsilon^2$ and $\sigma_w^2/\sigma_\epsilon^2$ both decrease. A decrease in $\sigma_D^2/\sigma_\epsilon^2$ increases the speed of convergence, while a decrease in $\sigma_w^2/\sigma_\epsilon^2$ decreases the speed of convergence. The net effect is ambiguous.

is proportional to the expected total inventory cost conditional on no change of the competitive allocation. To see the intuition, suppose for simplicity that: (i) the only piece of news is the one arriving at time 0 before the very first double auction; (ii) $Z = 0$; and (iii) $\{s_{i,0}\}$ are the same across i, so the competitive allocation given the time-0 signals is $Z/n = 0$ for all traders. In this case, trader i's expected inventory cost from time 0 to time Δ is $\frac{1-e^{-r\Delta}}{r}\frac{\lambda}{2}(1+d)^2 z_{i,0}^2$, his expected inventory cost from time Δ to time 2Δ is $e^{-r\Delta}\frac{1-e^{-r\Delta}}{r}\frac{\lambda}{2}(1+d)^4 z_{i,0}^2$, ..., his expected inventory cost from time $t\Delta$ to $(t+1)\Delta$ is $e^{-rt\Delta}\frac{1-e^{-rt\Delta}}{r}\frac{\lambda}{2}(1+d)^{2(t+1)} z_{i,0}^2$, ..., and so on. Here, we have used the exponential convergence of inventory shown in Proposition 4, and the time-discount $e^{-rt\Delta}$ is the probability that the liquidating dividend is not yet paid by time $t\Delta$ and hence holding costs are incurred afterwards. Summing up these terms, we see that each trader i's total expected inventory hold cost is proportional to L:

$$\frac{\lambda}{2r} z_{i,0}^2 \left[(1-e^{-r\Delta}) \sum_{t=0}^{\infty} e^{-rt\Delta}(1+d)^{2(t+1)} \right] = \frac{\lambda}{2r} z_{i,0}^2 \cdot L. \tag{39}$$

Proposition 4 states that L is smaller if Δ is larger, i.e., if the trading frequency is lower. Intuitively, a lower trading frequency reduces $1+d$, hence "front-loading" the asset reallocation among traders toward earlier periods while reducing the expected asset reallocation in later periods.[11] But because of the time discounting $e^{-rt\Delta}$, early-period allocative efficiency is more important than later-period allocative efficiency. Thus, L is decreasing in Δ. Indeed, Equation (32) implies that

$$\frac{(1-e^{-r\Delta})(1+d)}{1-e^{-r\Delta}(1+d)^2} = \frac{1}{n\alpha - 1}, \tag{40}$$

and hence

$$L = \frac{1+d}{n\alpha - 1}, \tag{41}$$

which is decreasing in Δ because $1+d$ is decreasing in Δ.

We emphasize that the relation between L and the expected inventory cost holds only if the competitive allocation does not change between trading periods (this is why each term involves the same $z_{i,0}^2$ in the derivation above). Whether the competitive allocation changes between auctions depends critically on whether or not the times of new information arrivals are predictable. For this reason, scheduled versus stochastic news arrivals have drastically different implications on how trading frequency affects welfare, as we show in the next section.

[11] To see this, note that, in this example, trader i's trading volume in double auction $t \in \{0, 1, 2, ...\}$, conditional on the liquidating dividend not being realized up to time $t\Delta$, is equal to $|z_{i,0}|\left((1+d)^t - (1+d)^{t+1}\right) = |z_{i,0}|(1+d)^t(-d)$. For any $\Delta \in [0, \infty)$, these trading volumes sum up or integrate to $|z_{i,0}|$ over t. It is easy to verify that for $t < \frac{1+d}{-d}$, a smaller $1+d$, equivalently a larger $-d > 0$, leads to a larger $(1+d)^t(-d)$. Thus, a smaller $1+d$ pushes trading volumes toward earlier periods and away from later periods, even conditional on the liquidating dividend having not realized. The possible realization of the liquidating dividend further reduces the expected later-period trading volume (i.e., asset reallocation) because of the discounting $e^{-rt\Delta}$.

3 Welfare and Optimal Trading Frequency

In this section, we use the model framework developed in Section 2 to analyze the welfare implications of trading frequency. Throughout this section we conduct the analysis based on the perfect Bayesian equilibrium of Proposition 2, which requires the parameter condition $n\alpha > 2$.

3.1 Welfare definition and notations

We define the equilibrium welfare as the sum of the ex-ante expected utilities over all traders:

$$W(\Delta) = \mathbb{E}\left[\sum_{i=1}^{n}(1-e^{-r\Delta})\sum_{t=0}^{\infty}e^{-rt\Delta}\left(v_{i,t\Delta}z^*_{i,(t+1)\Delta} - \frac{\lambda}{2r}(z^*_{i,(t+1)\Delta})^2\right)\right], \tag{42}$$

where $\{z^*_{i,t\Delta}\}$ is the inventory path in the equilibrium of Proposition 2, defined by Equation (36). As usual, the price terms are canceled out as they are transfers. We denote the Δ that maximizes $W(\Delta)$ as Δ^*.

Analogously, we can define the welfare in the competitive equilibrium of Section 2.2 as:

$$W^c(\Delta) = \mathbb{E}\left[\sum_{i=1}^{n}(1-e^{-r\Delta})\sum_{t=0}^{\infty}e^{-rt\Delta}\left(v_{i,t\Delta}z^c_{i,(t+1)\Delta} - \frac{\lambda}{2r}(z^c_{i,(t+1)\Delta})^2\right)\right]. \tag{43}$$

Although the competitive equilibrium is more efficient than the strategic one, it is still not fully efficient because new information may arrive between two double auctions. To explicitly take into account the possible misalignment between trading times and information arrival times, we use the following allocation as a benchmark:

$$z^e_{i,\tau} \equiv \frac{r(n\alpha-1)}{\lambda(n-1)}\left(s_{i,\tau} - \frac{1}{n}\sum_{j=1}^{n}s_{j,\tau}\right) + \frac{1}{n}Z, \text{ for every } \tau \geq 0. \tag{44}$$

The allocation $z^e_{i,\tau}$ is obtained in an idealized world in which a competitive double auction is held immediately after each news arrival. For this reason, we will refer to $z^e_{i,\tau}$ as the "competitive allocation without delay" or "zero-delay competitive allocation." If there were no adverse selection (asymmetric information about the common dividend), then $z^e_{i,\tau}$ would be the fully efficient allocation and would be a good benchmark. With adverse selection, however, $z^e_{i,\tau}$ is not the fully efficient allocation because the dispersion in $\{z^e_{i,\tau}\}$ across traders reflects their different private signals about the common dividend. If those signals about the common value were public, the dispersion in traders' fully efficient allocations would only reflect their different private values. Of course, this source of inefficiency in $z^e_{i,\tau}$ is only a consequence of the information structure but not the trading process, so in our analysis of the optimal trading frequency, $z^e_{i,\tau}$ is still a reasonable benchmark. The superscript "e" in $z^e_{i,\tau}$ indicates that the allocation $\{z^e_{i,\tau}\}$ is more efficient than $\{z^*_{i,t\Delta}\}$ and $\{z^c_{i,t\Delta}\}$ (see Lemma 2 in Section C.6), both of which involve potentially delayed response to new information.

Figure 3 illustrates a possible sample path of $z^*_{i,\tau}$, $z^c_{i,\tau}$, and $z^e_{i,\tau}$. In this example, a piece of new information arrives strictly between period-$(t-1)$ and period-t double auctions. We

Figure 3: Illustration of $z^*_{i,\tau}$, $z^c_{i,\tau}$, and $z^e_{i,\tau}$.

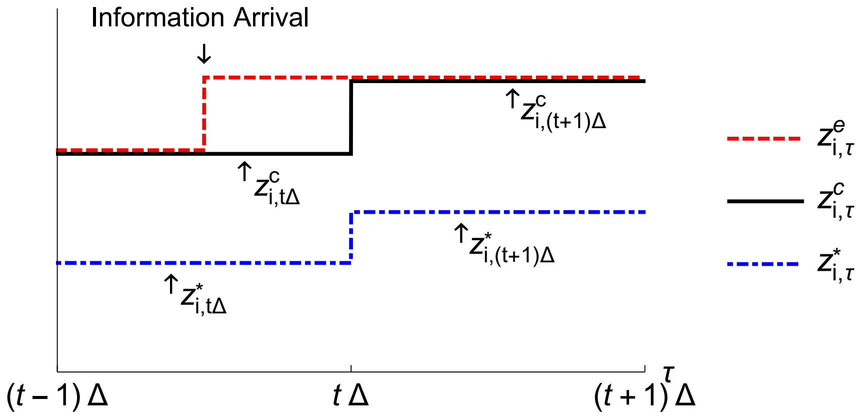

see that $z^e_{i,\tau}$ responds immediately to the new information, but $z^*_{i,\tau}$ and $z^c_{i,\tau}$ only change at the next trading opportunity. The gap between $z^*_{i,\tau}$ and $z^c_{i,\tau}$ represents the inefficiency caused by strategic behavior, and the gap between $z^c_{i,\tau}$ and $z^e_{i,\tau}$ represents the inefficiency caused by the misalignment between trading times and news times.

The ex-ante welfare under the competitive allocation without delay is:

$$W^e = \mathbb{E}\left[\sum_{i=1}^{n} \int_{\tau=0}^{\infty} re^{-r\tau}\left(v_{i,\tau} z^e_{i,\tau} - \frac{\lambda}{2r}(z^e_{i,\tau})^2\right) d\tau\right], \tag{45}$$

which is independent of Δ. Again, W^e would be the maximum possible welfare without adverse selection.

Because W^e is invariant to Δ, it is without loss of generality to use W^e as a benchmark in assessing the impact of trading frequency on welfare. We can thus define the following metric of the allocative inefficiency in the strategic equilibrium of Proposition 2:

$$X(\Delta) \equiv W^e - W(\Delta) = \underbrace{[W^c(\Delta) - W(\Delta)]}_{X_1(\Delta),\text{ welfare cost of strategic behavior}} + \underbrace{[W^e - W^c(\Delta)]}_{X_2(\Delta),\text{ welfare cost of trading delay}}. \tag{46}$$

The first part of the above decomposition, call it $X_1(\Delta)$, is due to strategic behavior and demand reduction. The second part, call it $X_2(\Delta)$, is due to the potential misalignment between trading times and news times. This decomposition highlights the important tradeoff in increasing trading frequency:

- A smaller Δ allows investors to react quickly to new information, reducing $X_2(\Delta)$.
- A smaller Δ also reduces the aggressiveness of demand schedules in each double auction, increasing $X_1(\Delta)$. This channel is a consequence of the fact that L decreases in Δ (see Proposition 4).

This tradeoff is a fundamental determinant of the optimal trading frequency, as we explain in the rest of this section.

Finally, we define:

$$\sigma_z^2 \equiv \sum_{i=1}^{n} \mathbb{E}[(z^e_{i,T_k} - z^e_{i,T_{k-1}})^2] = \left(\frac{r(n\alpha - 1)}{\lambda(n-1)}\right)^2 \frac{(n-1)(\chi^2(\sigma_D^2 + \sigma_\epsilon^2) + \sigma_w^2)}{\alpha^2} > 0, \tag{47}$$

$$\sigma_0^2 \equiv \sum_{i=1}^{n} \mathbb{E}[(z_{i,0} - z^e_{i,0})^2] = \sum_{i=1}^{n} \mathbb{E}[(z_{i,0} - z^c_{i,\Delta})^2]. \tag{48}$$

The first variance σ_z^2 describes the extent to which each arrival of new information changes the competitive allocation without delay. The second variance σ_0^2 describes the distance between the ex-ante inventory and the competitive zero-delay allocation given the new information that arrives at time 0. If $z_{i,0} = Z/n$ for every trader i (all traders are ex-ante identical), then $\sigma_0^2 = \sigma_z^2$.[12] One may naturally view "time 0" as a reduced-form representation of a steady state, in which case σ_0^2 and σ_z^2 should be equal. In the results below we will keep σ_0^2 as a generic parameter but highlight results for the most natural case of $\sigma_0^2 = \sigma_z^2$.

3.2 Scheduled arrivals of new information

We first consider scheduled information arrivals. In particular, we suppose that shocks to the common dividend and shocks to private values occur at regularly spaced clock times $T_k = k\gamma$ for a positive constant γ, where $k \geq 0$ is an integer. Examples of scheduled information arrivals include macroeconomic data releases and corporate earnings announcements.

Proposition 5. *Suppose $T_k = k\gamma$ for a positive constant γ. Then $W(\Delta) < W(\gamma)$ for any $\Delta < \gamma$. That is, $\Delta^* \geq \gamma$.*

Proposition 5 shows that if new information repeatedly arrives at scheduled times, then the optimal trading frequency cannot be higher than the frequency of information arrivals.

Now we provide a heuristic argument and describe the intuition behind Proposition 5. First, we show in Appendix C.6 (Lemma 3) that for any information arrival process,

$$X_1(\Delta) = \frac{1+d}{n\alpha - 1}\left(\frac{\lambda}{2r}\sigma_0^2 + \frac{\lambda}{2r}\sum_{t=1}^{\infty} e^{-rt\Delta} \sum_{i=1}^{n} \mathbb{E}[(z^c_{i,(t+1)\Delta} - z^c_{i,t\Delta})^2]\right). \tag{49}$$

In the expression of $X_1(\Delta)$, the term $\frac{\lambda}{2r}\sigma_0^2$ is the time-0 allocative inefficiency, whereas the term $\frac{\lambda}{2r}\sum_{t=1}^{\infty} e^{-rt\Delta}\sum_{i=1}^{n}\mathbb{E}[(z^c_{i,(t+1)\Delta} - z^c_{i,t\Delta})^2]$ measures the time-discounted cumulative change in the competitive allocation caused by arrivals of new information. Combined, the terms in the big parenthesis represent the time-discounted total amount of allocative inefficiency, measured at trading times, that can be potentially eliminated by holding double auctions. The leading multiplier, $\frac{1+d}{n\alpha-1} = L$, is the fraction of the total inefficiency that remains in the market due to the strategic behavior of traders (see Proposition 4 and the discussion of the multiplier L following the proposition).

[12] To see this, note that Z/n is the efficient allocation to each trader if the initial total signals of all traders are zero. For each trader i, the total signal $s_{i,0}$ received at time 0 has the same distribution as the innovation $s_{i,T_k} - s_{i,T_{k-1}}$. Thus, $z^e_{i,0} - Z/n$ has the same distribution as $z^e_{i,T_k} - z^e_{i,T_{k-1}}$ for any $k \geq 1$. Thus, $\sigma_0^2 = \sigma_z^2$ if $z_{i,0} = Z/n$ for all i.

Next, we set $\Delta = \gamma$. In this special case, auction times and information arrival times are perfectly aligned, and (49) reduces to

$$X_1(\gamma) = \frac{1+d}{n\alpha - 1}\left(\frac{\lambda}{2r}\sigma_0^2 + \frac{\lambda}{2r}\sum_{t=1}^{\infty} e^{-rt\gamma}\sigma_z^2\right), \tag{50}$$

where we have substituted in σ_z^2 from Equation (47) and used the fact that $z^c_{i,T_k+\Delta} = z^e_{i,T_k}$. Since there is no delay between the arrival of new information and trading, $X_2(\gamma) = 0$.

Finally, we double the trading frequency by letting $\Delta = \gamma/2$ while holding γ fixed. Auctions at times 0, $\gamma = 2\Delta$, $2\gamma = 4\Delta$, ..., still happen right after new information arrives, so $X_2(\Delta)$ remains zero. But auctions at time $0.5\gamma = \Delta$, $1.5\gamma = 3\Delta$, $2.5\gamma = 5\Delta$, ..., happen strictly between information arrivals. Obviously, $z^c_{i,(t+1)\gamma} - z^c_{i,(t+0.5)\gamma} = 0$ and $z^c_{i,t\gamma+0.5\gamma} - z^c_{i,t\gamma} \neq 0$. Therefore, after doubling the trading frequency, the terms in the big parenthesis of Equation (50) remain the same, but the leading multiplier, $\frac{1+d}{n\alpha-1}$, goes up because $1+d$ is larger if Δ is smaller (see Proposition 4). Thus, doubling the trading frequency increases the inefficiency X_1. By the same argument, $X_1(\gamma/k) < X_1(\gamma)$ for any positive integer $k \geq 2$.

In the appendix, we prove that choosing any $\Delta < \gamma$ is suboptimal even if γ/Δ is not an integer, that is, even if trading times do not entirely cover information arrival times. This general case is not trivial because it involves tradeoffs between $X_1(\Delta)$ and $X_2(\Delta)$.

Proposition 5 establishes that the optimal trading frequency cannot be strictly higher than information frequency. Next, we ask if the optimal trading frequency can be strictly lower than information frequency. As Δ increases beyond γ, $X_2(\Delta)$ is generally positive. Thus, traders face the basic tradeoff we discussed at the beginning of this section: a large $\Delta > \gamma$ induces more aggressive trading per period, but incurs the cost that traders cannot react quickly to new information.

The welfare $W(\Delta)$ is hard to analyze if Δ/γ is not an integer. For analytical tractability but at no cost of economic intuition, for the case of $\Delta > \gamma$ we restrict attention to $\Delta = l\gamma$ for a positive integer l. Let $l^* \in \operatorname{argmax}_{l \in \mathbb{Z}_+} W(l\gamma)$.

Proposition 6. *Suppose that $T_k = k\gamma$ for a positive constant γ. The following results hold.*

1. *If $z_{i,0} = Z/n$ for every trader i (i.e., $\sigma_0^2 = \sigma_z^2$), then $l^* = 1$.*
2. *If σ_0^2/σ_z^2 remains bounded as $n \to \infty$, then $l^* = 1$ as $n \to \infty$.*

Part 1 of Proposition 6 states a sharp result: for the steady-state specification $\sigma_0^2 = \sigma_z^2$, the optimal trading frequency is equal to the information frequency. To see the intuition, consider, for instance, slowing down trading from $\Delta = \gamma$ to $\Delta = 2\gamma$. Reducing the trading frequency by a half will make demand schedules more aggressive at times $0, 2\gamma, 4\gamma, \ldots$, at the cost of entirely disabling reaction to new information at times $\gamma, 3\gamma, 5\gamma, \ldots$. But because new information at each arrival time is equally informative and in expectation shocks the zero-delay competitive allocation by the same magnitude, there is no reason to let traders trade very aggressively over half of the news but shut down trading for the other half. Instead, it is better to allow equal opportunities respond to all information arrivals. That is, $l^* = 1$. This intuition applies to any $\Delta = l\gamma$ for an integer $l > 1$. The proof of Proposition 6 makes this intuitive argument formal.

Part 2 of Proposition 6 allows the time-0 information to be different from information that arrives later. A sufficient condition for σ_0^2/σ_z^2 remaining bounded as $n \to \infty$ is $\mathbb{E}[\sum_{i=1}^{n}(z_{i,0} -$

$Z/n)^2] = O(n)$. If $\sigma_0^2 > \sigma_z^2$, for instance, it is possible that the optimal $l^* > 1$ so that eliminating the time-0 allocative inefficiency is more important than allowing immediate reaction to less important news later. That said, in a large market it is still asymptotically optimal to align trading times with information arrival times. The intuition is that as n increases sufficiently, the market becomes almost competitive, and the inefficiency associated with strategic demand reduction diminishes. In the limit as $n \to \infty$, $X_1(\Delta) \to 0$, and the allocation efficiency is entirely determined by how fast traders can react to new information. Thus, the optimal $l^* = 1$.

Discussion of Vayanos (1999). We close this subsection with a discussion on the similarity and difference between our welfare results under scheduled information arrivals and those of Vayanos (1999). In Vayanos's model, trading times and information times (inventory shocks) are perfectly aligned, and there is no asymmetric information about the common dividend. Thus, his model corresponds to the case of $\Delta = \gamma$ and $\alpha = 1$ in our model. Moreover, Vayanos's measure of welfare loss as inventory shocks go to zero is equivalent to the leading multiplier $\frac{1+d}{n\alpha - 1} = L$ in Equation (50) in our model.[13] Thus, his result that faster trading reduces welfare is equivalent to the comparative statics of $\frac{\partial L}{\partial \Delta} < 0$, which is also how we prove that $X_1(\gamma/k) > X_1(\gamma)$. In addition, his result on the convergence rate to efficiency as the market gets large can be reproduced in our model by setting $\alpha = 1$.[14]

Nonetheless, our results on scheduled information arrival complement Vayanos (1999)'s in two ways. First, we separately model trading times and information arrival times. This general setting allows us to *prove* that the trading process should never outpace the information arrival process and that trading times should be aligned to news times under certain conditions (Proposition 5 and Proposition 6). These results are far from obvious ex ante. Second, our model incorporates adverse selection about the common dividend ($\alpha < 1$). Appendix B.2 shows that the presence of adverse selection, however small, slows down the speed of convergence to efficiency as the market becomes large. Specifically, as $n \to \infty$ and with adverse selection, X_1/n converges to zero at the rate of $n^{-4/3}$ for any fixed $\Delta > 0$, but the convergence rate of $(\lim_{\Delta \to 0} X_1)/n$ is $n^{-2/3}$; the corresponding rate without adverse selection is n^{-2} and n^{-1}, respectively.

Our most novel results about trading frequency, relative to Vayanos (1999), come from stochastic arrivals of information. When information arrival times are no longer predictable, it is in general not possible to align trading times with information arrival times. It is in those situations that the decoupling of trading times and information arrival times provides the most interesting economics and practical relevance. We turn to stochastic information arrivals next.

[13] Vayanos's measure of welfare loss as inventory shocks go to zero has the same functional form as L defined in Proposition 4 (see his Proposition 6.1), but $L = \frac{1+d}{n\alpha - 1}$ by Equation (41), which corresponds to his Equation (E.2).

[14] In our model, as $\Delta \to 0$, $1+d \to 1$, so the per-trader inefficiency measure X_1/n (cf. Equation (50)) becomes of order $1/(n\alpha - 1)$. On the other hand, as $\Delta \to \infty$, $1+d \to 1/(n\alpha - 1)$, so X_1/n becomes of order $1/(n\alpha - 1)^2$. (Using X_1/n, instead of X_1, is necessary here because σ_0^2 and σ_z^2 both measure inefficiency across all n traders and hence are of order $O(n)$.) In the special case of $\alpha = 1$, these orders of magnitude reduce to $O(1/n)$ and $O(1/n^2)$, which are the same as those in Vayanos (1999) (see his Corollary 6.2).

3.3 Stochastic arrivals of new information

We now consider stochastic arrivals of information. Examples of stochastic news include unexpected corporate announcements (e.g., mergers and acquisitions), regulatory actions, and geopolitical events. There are many possible specifications for stochastic information arrivals, and it is technically hard to calculate the optimal trading frequency for all of them. Instead, we analyze the simple yet natural case of a Poisson process for news arrivals. We expect the economic intuition of the results to apply to more general signal structures.

Suppose that the timing of the news shocks $\{T_k\}_{k\geq 1}$ follows a homogeneous Poisson process with intensity $\mu > 0$. (The first shock still arrives at time $T_0 = 0$.) Since the time interval between two consecutive news shocks has the expectation $1/\mu$, μ is analogous to $1/\gamma$ from Section 3.2. There are in expectation $\Delta\mu$ arrivals of new information during an interval of length Δ, and each arrival of information shocks the squared difference in the zero-delay competitive allocation by σ_z^2 (see Equation (47)). Thus,

$$\sum_{i=1}^{n} \mathbb{E}[(z^c_{i,(t+1)\Delta} - z^c_{i,t\Delta})^2] = \Delta\mu\sigma_z^2, \tag{51}$$

$$\sum_{i=1}^{n} \mathbb{E}[(z^e_{i,\tau} - z^c_{i,\tau})^2] = (\tau - t\Delta)\mu\sigma_z^2, \quad \tau \in (t\Delta, (t+1)\Delta). \tag{52}$$

To gain further intuition, we now focus on the natural case that all traders are ex-ante identical (i.e., $\sigma_0^2 = \sigma_z^2$) and explicitly spell out $X_1(\Delta)$ and $X_2(\Delta)$ from the decomposition (46):

$$X_1(\Delta) = \frac{\lambda(1+d)}{2r(n\alpha - 1)} \left(1 + \frac{\Delta e^{-r\Delta}}{1 - e^{-r\Delta}}\mu\right)\sigma_z^2, \tag{53}$$

$$X_2(\Delta) = \frac{\lambda}{2r}\mathbb{E}\left[\int_{\tau=0}^{\infty} re^{-r\tau}(z^c_{i,\tau} - z^e_{i,\tau})^2\, d\tau\right] = \frac{\lambda}{2r}\left(\frac{1}{r} - \frac{\Delta e^{-r\Delta}}{1 - e^{-r\Delta}}\right)\mu\sigma_z^2. \tag{54}$$

Here, the expression $X_1(\Delta)$ is obtained by substituting the squared difference in Equation (51) into Equation (49) of $X_1(\Delta)$ (Equation (49) holds for any information arrival process). We have also used Lemma 2 (Appendix C.6) to write $X_2(\Delta)$ as the expected squared difference between $z^c_{i,\tau}$ and $z^e_{i,\tau}$, for which Equation (52) applies. In fact, the final expression of $X_2(\Delta)$ can be rewritten as the more intuitive form:

$$X_2(\Delta) = \frac{\lambda}{2} \cdot \underbrace{\left(\int_0^{\infty} e^{-r\tau} d\tau - \sum_{t=1}^{\infty} e^{-rt\Delta}\Delta\right)}_{\text{Misalignment of information arrival times and trading times}} \cdot \frac{\mu\sigma_z^2}{r}.$$

The first term $\lambda/2$ is the multiplier of quadratic holding cost. The middle term in the bracket represents the misalignment of information arrival times and trading times, for it is the difference between an integral and its Δ-discrete counterpart, a summation. The third term is the expected variance of the change in the zero-delay competitive allocation per unit of time, $\mu\sigma_z^2$, multiplied by the expected waiting time until the dividend is paid, $1/r$. Note that the misalignment term only involves the Poisson information arrivals after time 0, since the first information arrival time coincides with the first trading time, time 0. The misalignment term and hence $X_2(\Delta)$ vanish as $\Delta \to 0$, i.e., there is zero welfare cost from trading delay if trading is continuous.

The total inefficiency of the strategic equilibrium is

$$X(\Delta) \equiv X_1(\Delta) + X_2(\Delta) = \frac{\lambda}{2r}\sigma_z^2 \left[\frac{\mu}{r} + 1 - \left(1 - \frac{1+d}{n\alpha - 1}\right)\left(1 + \frac{\Delta e^{-r\Delta}}{1 - e^{-r\Delta}}\mu\right)\right]. \tag{55}$$

We see that under Poisson news arrivals both terms $X_1(\Delta)$ and $X_2(\Delta)$ are generally positive. We thus expect an interior optimal Δ^*. Clearly, the optimal Δ^* is independent of λ.

The following proposition, the main result of this subsection, characterizes the optimal trading frequency in the case of ex-ante identical traders.

Proposition 7. *Suppose that $\{T_k\}_{k\geq 1}$ is a Poisson process with intensity μ, and $z_{i,0} = Z/n$ for every trader i (i.e., $\sigma_0^2 = \sigma_z^2$). Then the following holds.*

1. *The optimal trading frequency has the lower bound*

$$\frac{1}{\Delta^*} \geq \left(\frac{n\alpha}{2} - \frac{1}{3}\right)\mu. \tag{56}$$

   In particular, we always have $\frac{1}{\Delta^*} > \frac{2}{3}\mu$ since $n\alpha > 2$ by assumption.

2. *Δ^* strictly decreases in μ, n and σ_w^2, and strictly increases in σ_D^2.*

Part 1 of Proposition 7 shows a lower bound of the optimal trading frequency that is only a function of n, μ, and α but is independent of r. The multiplier in front of μ can be quite large. For examples, we expect α to be much closer to 1 than to $2/n$ for "liquid" assets, namely those with wide investor participation (large n), low information asymmetry (low σ_D^2 or low σ_ϵ^2), and high liquidity-driven trading motives (high σ_w^2). In particular, as we discussed in Section 2.2, $\alpha = 1$ if $\sigma_D^2 = 0$ (pure private value) or if $\sigma_\epsilon^2 = 0$ (public information about common value). These liquid assets include major equity indices, government securities, and foreign currencies, as well as the corresponding futures contracts.

A comparison between Part 1 of Proposition 7 and Proposition 5 reveals the major difference between scheduled and stochastic information arrivals. Take the case of ex-ante identical traders. Under scheduled information arrivals, the optimal trading frequency is equal to the information arrival frequency. Under stochastic information arrival, the optimal trading frequency can be much higher than the information arrival frequency.

The comparative statics in Proposition 7 can be proven by directly calculating the mixed partial derivative:

$$\frac{\partial^2 X(\Delta)}{\partial\Delta\partial\mu} = -\frac{\partial}{\partial\Delta}\left[\left(1 - \frac{1+d}{n\alpha - 1}\right)\frac{\Delta e^{-r\Delta}}{1 - e^{-r\Delta}}\right] > 0. \tag{57}$$

Thus, Δ^* is strictly decreasing in μ. Intuitively, if information arrives more frequently, the optimal trading frequency should also increase in order to allow traders faster response to new information.

By the same method, we can show that the optimal Δ^* is strictly decreasing in $n\alpha$. Since α measures the *lack of adverse selection*, α is strictly increasing in the variance of private-value shocks, σ_w^2, and is strictly decreasing in the variance of common-value shocks, σ_D^2. Hence, Δ^* is also strictly decreasing in σ_w^2 and is strictly increasing in σ_D^2. Moreover, we can show that $n\alpha$ is strictly increasing in n (even though α itself decreases with n), so Δ^* is strictly decreasing in n. The intuition is that if trading is motivated less by private information about the common dividend and more by idiosyncratic private values, or if there are more traders, then traders will submit more aggressive demand schedules. In those situations a higher-frequency market is better because reducing delays in responding to new information becomes more pressing than counteracting demand reduction.

We close this section with the following proposition on comparative statics when σ_0^2 and σ_z^2 are potentially different.

Proposition 8. *Suppose that $\{T_k\}_{k\geq 1}$ is a Poisson process with intensity μ. The following comparative statics holds:*

1. *If $\sigma_0^2 > 0$, then Δ^* strictly decreases in μ from ∞ (as $\mu \to 0$) to 0 (as $\mu \to \infty$).*
2. *If σ_0^2/σ_z^2 remains bounded as $n \to \infty$, then $\Delta^* \to 0$ as $n \to \infty$.*

The intuition for Proposition 8 is similar to that of Proposition 7. In particular, continuous trading becomes optimal in the limit as the market becomes large or as the arrival rate of new information increases without bound.

3.4 Calibration of optimal trading frequency under stochastic news arrivals

In this subsection we conduct a simple calibration exercise for the optimal trading frequency, using U.S. futures and equities as examples. This exercise illustrates the practical use of the model and its market-design implications. The resulting numbers, however, should be interpreted in their orders of magnitude instead of their exact levels. We will use actual data to calibrate n and μ, whereas α and r will be set directly to reasonable numbers. Note that λ does not affect the optimal trading frequency, so it is not a part of the input.

Futures market. We select four liquid contracts in U.S. futures markets for calibration: the E-mini S&P 500 futures, the 10-year Treasury futures, the Euro futures, and the crude oil futures. The sample period is from January 2013 to August 2016. All model parameters are calibrated to the daily frequency as follows:

- n. For each of the four contracts and on each day, we set n to be the average number of clearing accounts[15] that trade the futures contract in question. This proprietary statistic is provided by the CFTC.
- μ is calibrated in two ways.

[15] A firm could have multiple clearing accounts, and each clearing account may cover multiple trading desks or funds. Since all margin and collateral are managed at the clearing account level, this level of granularity seems to be suitable for our model because the inventory cost is partly motivated by margin and collateral constraints.

Our first way of calibrating μ is to set it equal to the average number of transaction price changes per day. In the model, the price changes if and only if news arrives. Hence, for each futures contract, the number of news arrivals is taken to be

$$\mu = \frac{1}{T}\sum_{t=1}^{T}\sum_{k} 1(p_{t,k} \neq p_{t,k+1}), \tag{58}$$

where $p_{t,k}$ is the price of the k-th transaction on day t of the futures contract. This number is directly provided by the CFTC.[16]

Our second way of calibrating μ is to set it to the daily average number of news articles that are related to the relevant contract. Using the Reuters News Analytics database from 2003 to 2005, Hendershott, Livdan, and Schürhoff (2015) enumerate the number of news articles that are related to various topics, ranging from macroeconomic conditions (e.g., unemployment) to firm-specific news (e.g., earnings). Since the four futures contracts are affected by the overall economic conditions, we count news articles in the following set of topics: Major Breaking News, Macro News, Business Activities, Regulatory Issues, Legislation, and Labor/(un)employment. Moreover, for the E-mini, 10-year Treasury, and Euro futures, we respectively include news articles in the categories of Stock Markets, Debt Markets, and Forex Markets. Importantly, news articles counted this way ignore idiosyncratic liquidity needs (that do not make to the news) and hence should be treated as a lower bound of the actual news arrival rate.

- α. In our model $1-\alpha$ measures adverse selection. Since each of the four futures contracts is about a broad market, we expect low adverse selection. So we set $\alpha = 1$. Calibrating α directly is much simpler and probably more robust than calibrating σ_D, σ_w, and σ_ϵ separately.

- r. In our model r represents the arrival rate of the liquidating dividend. For futures contracts, the liquidating dividend could be interpreted as the mark-to-market value of the contract at expiry. Since a futures contract tends to be most liquid for one month (when it is the front-month contract), we set $r = 1/30$. In fact, the calibrated optimal frequency is very insensitive to r, so this input is not critical.

Table 1 shows the calibrated optimal trading frequencies of the four futures contracts, together with the model inputs. In column 2 we observe that these four contracts attract wide market participation in the thousands of clearing accounts per day, with the E-mini having almost 10 thousand participants per day. Given the lower bound of $1/\Delta^*$ in Proposition 7, the large n already implies that the optimal trading frequency should be far larger than the news arrival frequency by a factor between one thousand and five thousand, approximately.

Columns 3–5 show, respectively, the average number of price changes per day, the numerically calculated optimal trading frequency per second, and the lower bound of the optimal frequency per second implied by Proposition 7. We convert all daily frequencies to per-second frequency using 23 trading hours per day. We observe that these four futures contracts have a

[16] Note that this number can also be obtained by purchasing trade-by-trade data from the CME or other vendors.

Table 1: Calibration of the optimal trading frequency for four futures contracts. We take $r = 1/30$ and $\alpha = 1$.

Contract	n	$\mu = \#$ transaction price changes/day			$\mu = \#$ market-wide news/day		
		μ (per day)	Optimal frequency (per second)	Lower bound (per second)	μ (per day)	Optimal frequency (per second)	Lower bound (per second)
E-mini S&P 500	9840	32716	1943.8	1943.8	162.7	9.7	9.7
10-year Treasury	2647	4581	73.2	73.2	163.6	2.6	2.6
Euro	1914	12222	141.2	141.2	128.6	1.5	1.5
Crude oil	3964	31332	749.9	749.9	126.3	3.0	3.0

large μ ranging from 4.6 thousand to 32.7 thousand per day, implying that the optimal trading frequency ranges from 73 to 1944 auctions per second, i.e., millisecond-level trading. The lower bound in Proposition 7, $(n\alpha/2 - 1/3)\mu$, turns out to be very tight.

The last three columns of Table 1 show the optimal trading frequency if news arrival rates are calculated from the number of news articles. Depending on the contract, there are between 126 and 164 relevant news articles per day on average, implying that the optimal trading frequency ranges from 1.5 to 9.7 double auctions per second. Unsurprisingly, due to a much smaller μ, the frequency in column 7 is much lower than that in column 4. Again, the lower bound $(n\alpha/2 - 1/3)\mu$ is very tight.

Overall, the futures market calibration suggests an optimal trading frequency ranging from a few auctions per second to two thousand auctions per second.

Equity markets. Our sample period is October 2010, since this is the sample period for the proprietary summary statistics provided by NASDAQ (described below). We use 146 stocks for the calibration. Among those, 117 stocks are the same as the ones used in Menkveld, Yueshen, and Zhu (2016). These 117 of stocks contain large-cap, medium-cap and small-cap firms in approximately equal proportions. The remaining 29 stocks are a stratified sample of very illiquid micro-cap stocks that have fewer than 100 trades per day during the sample period.[17] For each stock, the model parameters are calibrated to the daily frequency as follows:

- n. For each of the 117 stocks used in Menkveld, Yueshen, and Zhu (2016), NASDAQ provides the daily average number of NASDAQ member firms that trade this stock. The highest count is 219 (Apple), and the lowest is 22 (Delek). Since the NASDAQ counts only include broker-dealers but not investors, they are likely to vastly understate the total number of active participants in that stock. We adjust for this limitation by using the

[17] Specifically, we obtain the list of stocks that are in the basket for the iShares micro-cap ETF, as of September 2016. Then, we download all trades of these stocks in October 2010. The added 29 stocks are a stratified sample of the subset of stocks that have fewer than 100 trader per day on average.

following formula:

$$n = \min \Big(50 \times \text{The daily average number of NASDAQ members trading the stock}, \quad 2 \times \text{Average daily number of trades} \Big). \tag{59}$$

In this formula, the factor 50 is chosen so that the most liquid stocks like Apple and Amazon have roughly the same n as the E-mini S&P 500 futures contract. But for less liquid stocks, multiplying by 50 is likely to overstate the number of market participants. For instance, Delek has 22 NASDAQ member count, and multiplying by 50 gives 1100. However, Delek stock only has 208 trades per day on average. If each side of each of the 208 trades is a distinct investor, there would be only $2 \times 208 = 416$ investors actively trading Delek on a typical day, rather than 1100. This observation motivates the second part of the formula for n.

For each of the 29 illiquid micro-cap stocks, we simply set

$$n = 2 \times \text{Average daily number of trades}, \tag{60}$$

since the NASDAQ summary statistics do not cover them.

- μ. As before, we calibrate μ in two ways. The first is to set μ to be the average number of exchange transaction price changes per day, calculated from TAQ data.[18] Our second way is to set μ to be the average number of news articles per day that we use to calibrate the E-mini futures contract, that is, 162.7 (see Table 1). The second method significantly underestimates the actual μ because it only accounts for systematic information in news articles, but not information idiosyncratic to the particular stock or liquidity needs of particular traders. This μ should be interpreted as a (not too tight) lower bound.

- α. Individual stocks tend to have some degree of adverse selection. Since σ_ϵ^2 and σ_w^2 are unobservable, a direct calibration of α is difficult. We thus conduct separate calibrations for two values of α, 0.9 and 0.1, that are near the boundaries of its range, $[0, 1]$.

- r. Different from futures, stocks do not have an "expiration date," so the date of the "liquidating dividend" in the model should be interpreted as the date on which substantial uncertainty regarding the stock's fundamental value is resolved (even if temporarily). Since the quarterly earnings announcements provide the most important fundamental information about stocks, we set $r = 1/90$. Again, the calibrated optimal frequency is insensitive to r, so this input is not critical.

Table 2 shows the calibrated optimal trading frequencies together with the model inputs for stocks located at (or nearest to) key percentiles of the sample, sorted by the optimal trading frequency in column 5 (see the discussion below). In column 3, we observe remarkable heterogeneity in the (estimated) number of active participants across stocks, ranging from about 6 at the bottom to more than 10 thousand at the top.

[18]Different from futures contracts, some of stock trades are done off-exchange, or "in the dark," with limited or no price discovery. We thus eliminate those trades. See Zhu (2014) and Menkveld, Yueshen, and Zhu (2016) for more discussions of dark trading in U.S. equity markets.

Table 2: Calibrated optimal trading frequency for stocks at (or nearest to) various percentiles of the sample, sorted by the calibrated optimal trading frequency in column 5. We set $r = 1/90$.

	Ticker	n	$\mu = \#$ of transaction price changes/day			$\mu = \#$ stock market-wide news/day (162.7)	
			μ (per day)	Optimal frequency (per second), $\alpha = 0.9$	Optimal frequency (per second), $\alpha = 0.1$	Optimal frequency (per second), $\alpha = 0.9$	Optimal frequency (per second), $\alpha = 0.1$
Max	AAPL	10950	41286.1	8693.3	965.4	34.3	3.8
Pct 95	MOS	6250	7947.5	955.1	106.0	19.6	2.2
Pct 90	PFE	6150	4967.0	587.4	65.2	19.2	2.1
Pct 75	GLW	5000	3037.6	292.0	32.4	15.6	1.7
Median	SF	2731	668.7	35.1	3.9	8.5	0.9
Pct 25	KNOL	823	188.9	2.9881	0.3296	2.5749	0.2840
Pct 10	RVSB	101	11.4	0.0220	0.0023	0.3136	0.0328
Pct 5	STS	42	6.7	0.0053	0.0005	0.1288	0.0123
Min	KINS	5.6	1.1	0.0001	na	0.0153	na

Column 4 shows the model input μ, calibrated to the average number of transaction price changes per day. Columns 5 and 6 show the corresponding optimal frequency per second[19] for $\alpha = 0.9$ and $\alpha = 0.1$, respectively. We observe that the top half of the stocks have an optimal trading frequency ranging from several double auctions to several thousand double auctions per second. The bottom quartile of stocks, however, tend to have much lower optimal frequencies, generally slower than one auction per second and sometimes slower than one auction per minute (one auction per minute corresponds to an optimal frequency of $1/60 = 0.167$). The bottom stock has $n = 5.6$, so under $\alpha = 0.1$ the linear equilibrium does not exist since $n\alpha < 2$.

The last two columns show the sample stocks' optimal trading frequencies per second if μ is set to the number of news articles per day that we use to calibrate the E-mini, that is, 162.7 (see Table 1). Unsurprisingly, under this much lower μ, the optimal frequency for the top half of the stocks drops significantly, ranging from one to 34 double auctions per second. For the bottom quartile of the stocks, however, the optimal frequency is similar in order of magnitude to that in columns 5 and 6, i.e., in seconds or minutes per double auction.

The robust pattern from the stocks calibration is that, as in futures contracts, stocks that have broad market participation are optimally traded with sub-second delays. But stocks that have narrow participation can benefit from a moderate market slowdown to seconds or even minutes per double auction.

Summary. A policy implication coming out of this calibration exercise is that, if regulators or marketplaces were to implement mechanisms that slow down the market, it makes the most sense to start with stocks that have the lowest investor participation, which tend to be small- and micro-cap stocks. By contrast, stocks and futures contracts with broad investor participation have a wide range of calibrated optimal trading frequency, from a few auctions to a few thousand auctions per second, so a market slowdown is not necessarily warranted. Again, the calibration illustrates and reinforces the model implication that the optimal trading frequency depends on properties of the specific asset, such as the level of investor participation and the arrival

[19]Since trading hours in U.S. equity markets last from 9:30am to 4:00pm ET, the per-second frequency is converted from the daily frequency by dividing by 390×60 seconds per day.

frequency of relevant news. Hence, policies that aim to adjust trading speed should take such heterogeneity into account.

Given the relevant data, the calibration exercise shown here can be conducted on electronic markets for other assets, such as ETFs, options, government securities, and foreign exchange, as well as in other jurisdictions.

4 Concluding Remarks

In this paper, we present and solve a dynamic model in which a finite number of traders receive private information over time and trade strategically with demand schedules in a sequence of double auctions. We characterize a linear equilibrium with stationary strategies in closed form. The equilibrium price aggregates a weighted sum of signals about the common value and the private values, but the two components cannot be separated from the price. Due to imperfect competition, the equilibrium allocation is not fully efficient, but it converges to the efficient allocation exponentially over time. The presence of adverse selection—asymmetric information regarding the common-value component of the asset—slows down this convergence speed.

We use this modeling framework to study the optimal trading frequency that maximizes allocative efficiency. Trading frequency is measured as the number of double auctions per unit of clock time. A higher trading frequency reduces the aggressiveness of demand schedules, but allows more immediate reactions to new information. The allocative inefficiency in this dynamic market can be decomposed into two parts: one part due to strategic behavior and the other due to delayed responses to new information. The optimal tradeoff between these two effects depends on the nature of information. If new information arrives at scheduled time intervals, the optimal trading frequency is never higher than the information frequency, and these two frequencies coincide if traders are ex-ante identical. By contrast, if new information arrives as a Poisson process, the optimal trading frequency can be much higher than the information arrival frequency, and we explicitly characterize a lower bound for the optimal trading frequency.

We illustrate the application of the model by calibrating the optimal trading frequency of four liquid futures contracts and a sample of 146 stocks in U.S. markets. Using reasonable and data-derived proxies for the model parameters, we find that the model-implied optimal trading frequencies for the futures contracts and the top half of sample stocks (in terms of investor participation) are optimally traded with between-auction delays less than a second and sometimes as low as milliseconds. By contrast, stocks with lower investor participation, such as small- and micro-cap stocks, have an optimal frequency in the order of seconds or minutes per auction. This calibration suggests that a moderate market slowdown can improve allocative efficiency for small- and micro-cap stocks, but not necessarily for large stocks, futures contracts, or other liquid securities with broad investor participation. More broadly, our analysis suggests that market design in terms of speed should take into account asset heterogeneity, such as heterogeneity in the level of investor participation and the arrival frequency of news.

Our results are useful not only for markets that are already centralized and electronic, but also for over-the-counter (OTC) markets that are moving toward all-to-all electronic trading, such as those for corporate bonds, interest rate swaps, and credit default swaps. If exchange-like trading for these OTC instruments were deemed desirable by investors and regulators, periodic auctions could be a more natural place to start than continuous trading.

While our model captures two important effects of increasing market speed, it does not capture all effects. For instance, our model ignores the investments spent on speeding up the market. Biais, Foucault, and Moinas (2015) suggest that too much such investment is made and prescribe a tax. Budish, Cramton, and Shim (2015) argue that continuous limit order books should be replaced by frequent double auctions. Our result that assets with narrow investor participation can benefit from a moderate market slowdown is consistent with their conclusions. We go beyond their results in this direction by providing explicit formulas that can be used to calibrate the optimal trading frequency.

A discussion of heterogeneous trading speeds. Heterogeneous trading speeds would be an interesting direction for future research. That certain traders are faster than others has been a persistent feature of financial markets, and this issue has generated renewed and passionate debate in the context of high-frequency trading. We conclude this article with a discussion of heterogeneous speeds that is mainly based on calculations that appear in an early version of the paper. Since those calculations are conducted under somewhat restrictive assumptions, we will focus on intuition that we believe is robust in more general settings.

In a model with sequential double auctions, a trader's speed could be defined by how frequently he accesses the market. For example, a fast trader can participate in all double auctions no matter how frequently the auctions are held, but a slower trader can only participate in auctions at, say, one-second time intervals.[20] The latter implies that during any time interval of one second, each slow trader joins the market exactly once.[21] Moreover, this way of modeling speed heterogeneity does not require heterogeneous information about asset fundamentals; in particular, it does not endow the fast traders with superior information of any kind.[22] Instead, trading is generated by shocks to private values or inventories.

Modeled this way, speed heterogeneity creates a discrepancy between fast and slow traders' preferred market designs. For intuition, let us start with double auctions that are held once per second, at times 1, 2, 3, ... (in unit of seconds). For the reason mentioned above, all slow traders can participate in each double auction together with all fast traders. Now, let us speed up the market ten times so that one auction is held every 100 milliseconds, at times 0.1, 0.2, 0.3, ... (again in seconds). Unless all slow traders happen to come to the market within a 100-millisecond time interval, they are effectively "partitioned" by the more frequent auctions.

[20] In practical terms, the fast traders could be high-frequency traders who install their computer servers next to the stock exchange's server, whereas the slow traders could be other investors or broker-dealers whose computer system has a one-second delay in accessing the exchange's matching engine.

[21] Within each one-second time interval, whether slow trades show up at random or deterministic times is a modeling choice, but this technical choice is unlikely to affect the economic intuition discussed below.

[22] The existing literature on high-frequency traders has so far focused on the *information* advantage of HFTs. For instance, in Biais, Foucault, and Moinas (2015) and Jovanovic and Menkveld (2012), HFTs have outright fundamental information. In Budish, Cramton, and Shim (2015); Hoffmann (2014), HFTs obtain public news faster than others so they can quickly use market orders to make profit or cancel stale limit orders to avoid a loss. In Yang and Zhu (2016), HFTs observe signals of executed orders of informed investors and hence learn fundamental information. To the best of our knowledge, two papers in the HFT literature model the feature that faster traders have more opportunities to trade. In Biais, Foucault, and Moinas (2015), fast institutions find trading opportunities for sure, whereas slow institutions can only trade with some probability. In Cespa and Vives (2016), high-frequency traders trade in two periods, whereas other dealers trade only in one period. In both models, market participants form continuums and hence are competitive. Our model is strategic.

For instance, we only expect a fraction of slow traders to show up in the market in the time interval $[0, 0.1]$, and they can trade in the first double auction or delay until a later double auction. Trading in the first auction saves a slow trader inventory cost, but also incurs the trader a higher price impact cost because only a fraction of the participants are in the market by the time of the first auction. In particular, these early-arriving slow traders cannot commit to never trading in the first double auction, because a slow trader receiving a sufficiently large inventory shock would prefer immediately liquidating part of the inventory to holding everything until a later auction (recall inventory cost is quadratic). But given the one-second time delay, once this early-arriving slow trader participates in the first double auction at time 0.1, he can only return to the market at time 1.1, thus skipping the nine auctions at times 0.2, 0.3, ..., 1 and inadvertently making price impact higher in those auctions by his absence. By similar reasoning, no slow trader who arrives in the interval $(0.1, 0.2]$ can commit to never trading in the second double auction; if he does trade in the second auction, he would skip auctions at times 0.3, 0.4, 0.5, $\dots$, 1.1 and inadvertently make price impact higher in those auctions by his absence. This argument applies to all slow traders.

Summarizing, since no individual slow trader can commit to waiting when early trading with fast traders is an option, a more frequent market effectively partitions the slow traders into more non-overlapping groups in expectation. As a consequence, each double auction has a higher price impact in expectation. We can show that, at least in a simplistic and restrictive setting, the slow traders collectively prefer to bunch together and synchronize their trading times to obtain a lower price impact, but the lack of commitment to delaying defeats this purpose for reasons mentioned above. Fast traders, on the other hand, benefit from a higher market frequency because partitioning slow traders into multiple groups increases price impact and hence their profits of providing liquidity. In fact, in a simplistic setting, we can show that the fast traders wish to hold double auctions so frequently that they only interact with exactly one slow trader at a time. By contrast, slow traders prefer a strictly lower market frequency, which enables them to synchronize their trading and reducing price impact. The broader takeaway of this analysis, to be made more rigorous in subsequent work, is that a market speed-up has heterogeneous impacts on various groups of participants, with the fastest market participants benefiting the most.

Future work on speed heterogeneity may require a careful (re)selection of model components that approximate institutional reality reasonably well and simultaneously retain tractability. The "winning" model that best tackles heterogeneous speeds may turn out to be quite different from the current model of this paper that analyzes the optimal trading frequency under homogeneous speed. For instance, should trading times be modeled as deterministic or random? How do traders keep track of or infer each others' inventories over time, especially because not all traders are present at all times? Is it possible to avoid the forecasting-the-forecast-of-others problem and an explosion in the number of state variables, and if not, are there reasonable heuristics that one can use to obtain tractability? Finally, can we embed information asymmetry about the fundamental value into a model of speed heterogeneity? Building such a model is far beyond the scope of this paper and is left for future research.

A List of Model Variables

Variable	Explanation
	Sections 2–3, Exogenous Variables
t	Discrete trading period , $t \in \{0, 1, 2, 3, \dots\}$
τ	Continuous clock time, $\tau \in [0, \infty)$
Δ	Length of each trading period
$\mathcal{T}, r$	The clock time $\mathcal{T}$ of dividend payment has an exponential distribution with intensity $r > 0$.
$\{T_k\}_{k\in\{0,1,2,\dots\}}$	Times of shocks to the common dividend and private values
D_{T_k}	The common dividend value immediately after the k-th shock
σ_D^2	Each dividend shock $D_{T_k} - D_{T_{k-1}}$ has the distribution $\mathcal{N}(0, \sigma_D^2)$.
S_{i,T_k}	Trader i's signal about the k-th dividend shock
σ_ϵ^2	The noise in trader i's dividend signal regarding the k-th dividend shock, $S_{i,T_k} - (D_{T_k} - D_{T_{k-1}})$, has the distribution $\mathcal{N}(0, \sigma_\epsilon^2)$.
w_{i,T_k}	Trader i's private value for the asset immediately after the k-th shock
σ_w^2	Shocks to each trader i's private value, $w_{i,T_k} - w_{i,T_{k-1}}$, has the distribution $\mathcal{N}(0, \sigma_w^2)$.
$v_{i,\tau}$	$D_{T_k} + w_{i,T_k}$ if T_k is the last shock before τ
λ	Before the dividend is paid, the flow cost for holding asset position q is $0.5\lambda q^2$ per unit of clock time for each trader.
Z	The total inventory held by all traders, $Z \equiv \sum_{1\le j\le n} z_{j,0}$
γ	Time interval of scheduled information arrivals
μ	Intensity of stochastic information arrivals
	Sections 2–3, Endogenous Variables
$z_{i,t\Delta}$	Trader i's inventory level right before the period-t double auction
$x_{i,t\Delta}(p)$	Trader i's demand schedule in the period-t double auction
$p^*_{t\Delta}$	The equilibrium price in period-t double auction
$H_{i,t\Delta}$	Trader i's history (information set) up to time $t\Delta$ but before the period-t double auction, defined in Equation (10)
s_{i,T_k}	Trader i's total signal right after the k-th shock, defined in Equation (20)
$V_{i,t\Delta}$	The expected utility of trader i in period t, conditional on $H_{i,t\Delta}$
χ, α	Constants defined in Section 2.2
$z^c_{i,t\Delta}$	The competitive allocation immediately before trading in period-t auction
$z^e_{i,\tau}$	Zero-delay competitive allocation
σ_z^2, σ_0^2	Constants defined in Equations (47) and (48)
$W(\Delta)$	Welfare under homogeneous speed and trading interval Δ

B Additional Results

B.1 The continuous-time limit of Proposition 2

In this appendix we examine the limit of the equilibrium in Proposition 2 as $\Delta \to 0$ (i.e., as trading becomes continuous in clock time) and its efficiency properties.

Proposition 9. *As $\Delta \to 0$, the equilibrium of Proposition 2 converges to the following perfect Bayesian equilibrium:*

1. *Trader i's equilibrium strategy is represented by a process $\{x^{\infty}_{i,\tau}\}_{\tau \geq 0}$. At the clock time τ, $x^{\infty}_{i,\tau}$ specifies trader i's* rate *of order submission and is defined by*

$$x^{\infty}_{i,\tau}(p; s_{i,\tau}, z_{i,\tau}) = b^{\infty} \left(s_{i,\tau} - p - \frac{\lambda(n-1)}{r(n\alpha - 1)} z_{i,\tau} + \frac{\lambda(1-\alpha)}{r(n\alpha-1)} Z \right), \tag{61}$$

where

$$b^{\infty} = \frac{r^2(n\alpha-1)(n\alpha-2)}{2\lambda(n-1)}. \tag{62}$$

Given a clock time $\tau > 0$, in equilibrium the total amount of trading by trader i in the clock-time interval $[0, \tau]$ is

$$z^*_{i,\tau} - z_{i,0} = \int_{\tau'=0}^{\tau} x^{\infty}_{i,\tau'}(p^*_{\tau'}; s_{i,\tau'}, z^*_{i,\tau'})\, d\tau'. \tag{63}$$

2. *The equilibrium price at any clock time τ is*

$$p^*_{\tau} = \frac{1}{n}\sum_{i=1}^{n} s_{i,\tau} - \frac{\lambda}{nr} Z. \tag{64}$$

3. *Given any $0 \leq \underline{\tau} < \overline{\tau}$, if $s_{i,\tau} = s_{i,\underline{\tau}}$ for all i and all $\tau \in [\underline{\tau}, \overline{\tau}]$, then the equilibrium inventories $z^*_{i,\tau}$ in this interval satisfy:*

$$z^*_{i,\tau} - z^e_{i,\underline{\tau}} = e^{-\frac{1}{2}r(n\alpha-2)(\tau-\underline{\tau})} \left(z^*_{i,\underline{\tau}} - z^e_{i,\underline{\tau}} \right), \tag{65}$$

where

$$z^e_{i,\underline{\tau}} \equiv \frac{r(n\alpha-1)}{\lambda(n-1)} \left(s_{i,\underline{\tau}} - \frac{1}{n}\sum_{j=1}^{n} s_{j,\underline{\tau}} \right) + \frac{1}{n} Z \tag{66}$$

is the competitive allocation at clock time $\underline{\tau}$ (cf. Equation (26)).

Proof. The proof follows by directly calculating the limit of Proposition 2 as $\Delta \to 0$ using L'Hopitâl's rule. □

Proposition 9 reveals that even if trading occurs continuously, in equilibrium the competitive allocation is not reached instantaneously. The delay comes from traders' price impact and the associated demand reduction. This feature is also obtained by Vayanos (1999). Although submitting aggressive orders allows a trader to achieve his desired allocation sooner, aggressive bidding also moves the price against the trader and increases his trading cost. Facing this tradeoff, each trader uses a finite rate of order submission in the limit. As in Proposition 4, the rate of convergence to the competitive allocation in Proposition 9, $r(n\alpha - 2)/2$, is increasing in n, r, and σ^2_w but decreasing in σ^2_D. (The proof of Proposition 4 shows that $\partial(n\alpha)/\partial\sigma^2_w > 0$, $\partial(n\alpha)/\partial\sigma^2_D < 0$, and $\partial(n\alpha)/\partial n > 0$.)

B.2 Convergence rate to efficiency in large markets

To further explore the effect of adverse selection for allocative efficiency, and to compare with the literature (in particular with Vayanos (1999)), we consider the rate at which inefficiency vanishes as the number of traders becomes large, with and without adverse selection. Adverse selection exists if $\sigma_D^2 > 0$ and $\sigma_\epsilon^2 > 0$. For fixed $\sigma_\epsilon^2 > 0$ and $\sigma_w^2 > 0$, we compare the convergence rate in the case of a fixed $\sigma_D^2 > 0$ to that in the case of $\sigma_D^2 = 0$.

We consider the inefficiency caused by strategic behavior, that is, the difference between the total ex-ante utility in the strategic equilibrium of Proposition 2 and the total ex-ante utility in the competitive equilibrium:

$$X_1(\Delta) \equiv \mathbb{E}\Bigg[\sum_{t=0}^{\infty}(e^{-rt\Delta} - e^{-r(t+1)\Delta})\sum_{i=1}^{n}\Bigg(\Big(v_{i,t\Delta}z^*_{i,(t+1)\Delta} - \frac{\lambda}{2r}(z^*_{i,(t+1)\Delta})^2\Big) - \Big(v_{i,t\Delta}z^c_{i,(t+1)\Delta} - \frac{\lambda}{2r}(z^c_{i,(t+1)\Delta})^2\Big)\Bigg)\Bigg], \tag{67}$$

where $\{z^*_{i,(t+1)\Delta}\}$ is strategic allocation given by Equation (36), and $z^c_{i,(t+1)\Delta}$ is the competitive allocation given by Equation (26). This $X_1(\Delta)$ is the same as that defined in Section 3. As usual, prices do not enter the welfare criterion as they are transfers.

Proposition 10. *Suppose that the news times $\{T_k\}_{k\geq 1}$ either satisfy $T_k = k\gamma$ for a constant $\gamma > 0$ or are given by a homogeneous Poisson process. Suppose also that $\sigma_\epsilon^2 > 0$, $\sigma_w^2 > 0$, and $\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[(z_{i,0} - z^e_{i,0})^2]$ is bounded as n becomes large. Then, the following convergence results hold:*

1. *If $\sigma_D^2 > 0$, then as $n \to \infty$:*

$$\frac{X_1(\Delta)}{n} \text{ converges to zero at the rate } n^{-4/3} \text{ for any } \Delta > 0,$$
$$\lim_{\Delta\to 0}\frac{X_1(\Delta)}{n} \text{ converges to zero at the rate } n^{-2/3}.$$

2. *If $\sigma_D^2 = 0$, then as $n \to \infty$:*

$$\frac{X_1(\Delta)}{n} \text{ converges to zero at the rate } n^{-2} \text{ for any } \Delta > 0,$$
$$\lim_{\Delta\to 0}\frac{X_1(\Delta)}{n} \text{ converges to zero at the rate } n^{-1}.$$

The convergence rates under $\sigma_D^2 = 0$ (i.e., pure private values) are also obtained in the model of Vayanos (1999), who is the first to show that convergence rates differ between discrete-time trading and continuous-time trading. Relative to the results of Vayanos (1999), Proposition 10 reveals that the rate of convergence is slower if traders are subject to adverse selection. For any fixed $\Delta > 0$ and as $n \to \infty$, the inefficiency $X_1(\Delta)/n$ vanishes at the rate of $n^{-4/3}$ if $\sigma_D^2 > 0$, but the corresponding rate is n^{-2} if $\sigma_D^2 = 0$. If one first takes the limit of $\Delta \to 0$, then the convergence rates as n becomes large are $n^{-2/3}$ and n^{-1} with and without adverse selection, respectively. (The limiting behavior of the strategic equilibrium as $\Delta \to 0$ is stated in Appendix B.1.) Interestingly, the asymptotic rates do not depend on the size of σ_D^2 but only depend on whether σ_D^2 is positive or not.

C Proofs

C.1 Construction of total signals

In this appendix we show details of the construction of the total signals in Equation (20). The total signals are subsequently used in the strategies of the competitive benchmark and the strategic equilibrium.

Lemma 1. *For any constant x, we have:*

$$\begin{aligned}
&\mathbb{E}\left[v_{i,T_k}\middle|H_{i,T_k}\cup\left\{\sum\nolimits_{j\neq i}\left(x\sum\nolimits_{l=0}^{k}S_{j,T_l}+w_{j,T_k}\right)\right\}\right] \\
&=w_{i,T_k}+\frac{1/(x^2\sigma_\epsilon^2)}{1/(x^2\sigma_D^2)+1/(x^2\sigma_\epsilon^2)+(n-1)/(x^2\sigma_\epsilon^2+\sigma_w^2)}\sum_{l=0}^{k}S_{i,T_l} \\
&\quad+\frac{1/(x^2\sigma_\epsilon^2+\sigma_w^2)}{1/(x^2\sigma_D^2)+1/(x^2\sigma_\epsilon^2)+(n-1)/(x^2\sigma_\epsilon^2+\sigma_w^2)}\cdot\frac{1}{x}\left(\sum_{j\neq i}\left(x\sum_{l=0}^{k}S_{j,T_l}+w_{j,T_k}\right)\right).
\end{aligned}\tag{68}$$

Proof. Define

$$\tilde{S}_{i,T_l}\equiv xS_{i,T_l}+w_{i,T_l}-w_{i,T_{l-1}}.\tag{69}$$

By the projection theorem for multivariate normal distribution:

$$\begin{aligned}
&\mathbb{E}\left[D_{T_l}-D_{T_{l-1}}\mid S_{i,T_l},\sum\nolimits_{j\neq i}\tilde{S}_{j,T_l}\right] \\
&=(x\sigma_D^2,(n-1)x\sigma_D^2)\cdot\begin{pmatrix}x^2(\sigma_D^2+\sigma_\epsilon^2) & (n-1)x^2\sigma_D^2\\ (n-1)x^2\sigma_D^2 & (n-1)(x^2(\sigma_D^2+\sigma_\epsilon^2)+\sigma_w^2)+(n-1)(n-2)x^2\sigma_D^2\end{pmatrix}^{-1} \\
&\quad\cdot\left(xS_{i,T_l},\sum\nolimits_{j\neq i}\tilde{S}_{j,T_l}\right)'.
\end{aligned}\tag{70}$$

We compute:

$$\begin{aligned}
&\begin{pmatrix}x^2(\sigma_D^2+\sigma_\epsilon^2) & (n-1)x^2\sigma_D^2\\ (n-1)x^2\sigma_D^2 & (n-1)(x^2(\sigma_D^2+\sigma_\epsilon^2)+\sigma_w^2)+(n-1)(n-2)x^2\sigma_D^2\end{pmatrix}^{-1} \\
&=\begin{pmatrix}(n-1)(x^2\sigma_\epsilon^2+\sigma_w^2)+(n-1)^2x^2\sigma_D^2 & -(n-1)x^2\sigma_D^2\\ -(n-1)x^2\sigma_D^2 & x^2(\sigma_D^2+\sigma_\epsilon^2)\end{pmatrix} \\
&\quad\cdot\frac{1}{(n-1)x^2(x^2\sigma_\epsilon^2+\sigma_w^2)(\sigma_D^2+\sigma_\epsilon^2)+(n-1)^2x^4\sigma_D^2\sigma_\epsilon^2},
\end{aligned}$$

and

$$\begin{aligned}
&\mathbb{E}\left[D_{T_l}-D_{T_{l-1}}\mid S_{i,T_l},\sum\nolimits_{j\neq i}\tilde{S}_{j,T_l}\right] \\
&=\frac{(n-1)x^2\sigma_D^2(x^2\sigma_\epsilon^2+\sigma_w^2)S_{i,T_l}+(n-1)x^3\sigma_D^2\sigma_\epsilon^2\sum_{j\neq i}\tilde{S}_{j,T_l}}{(n-1)x^2(x^2\sigma_\epsilon^2+\sigma_w^2)(\sigma_D^2+\sigma_\epsilon^2)+(n-1)^2x^4\sigma_D^2\sigma_\epsilon^2} \\
&=\frac{(1/x^2\sigma_\epsilon^2)S_{i,T_l}+(1/(x^2\sigma_\epsilon^2+\sigma_w^2))\frac{1}{x}\sum_{j\neq i}\tilde{S}_{j,T_l}}{1/(x^2\sigma_\epsilon^2)+1/(x^2\sigma_D^2)+(n-1)/(x^2\sigma_\epsilon^2+\sigma_w^2)}.
\end{aligned}$$

Summing the above equation across $l\in\{0,1,\dots,k\}$ and adding w_{i,T_k} gives Equation (68). □

By Equation (68), we have

$$
\begin{aligned}
&\mathbb{E}\left[v_{i,T_k}\middle| H_{i,T_k}\cup\left\{\sum\nolimits_{j\neq i}\left(\chi\sum\nolimits_{l=0}^{k}S_{j,T_l}+w_{j,T_k}\right)\right\}\right]\\
=&\, w_{i,T_k}+\chi\sum_{l=0}^{k}S_{i,T_l}+\frac{1/(\chi^2\sigma_\epsilon^2+\sigma_w^2)}{1/(\chi^2\sigma_D^2)+1/(\chi^2\sigma_\epsilon^2)+(n-1)/(\chi^2\sigma_\epsilon^2+\sigma_w^2)}\cdot\frac{1}{\chi}\left(\sum_{j\neq i}\left(\chi\sum_{l=0}^{k}S_{j,T_l}+w_{j,T_k}\right)\right)\\
=&\,\alpha s_{i,T_k}+\frac{1-\alpha}{n-1}\sum_{j\neq i}s_{j,T_k},
\end{aligned}
$$

where in the second line we used the definition of χ in Equation (19), and in the third line we used the definition of s_{i,T_k} in Equation (20), and the definition of α:

$$
\alpha\equiv\frac{1}{1+\frac{(n-1)/(\chi^2\sigma_\epsilon^2+\sigma_w^2)}{1/(\chi^2\sigma_D^2)+1/(\chi^2\sigma_\epsilon^2)+(n-1)/(\chi^2\sigma_\epsilon^2+\sigma_w^2)}\cdot\frac{1}{\chi}}=\frac{\chi^2\sigma_\epsilon^2+\sigma_w^2}{n\chi^2\sigma_\epsilon^2+\sigma_w^2}. \tag{71}
$$

C.2 Verification of the competitive equilibrium strategy

The value function of trader i, rewritten from Equation (12), is:

$$
\begin{aligned}
\max_{\{x_{i,t'\Delta}\}_{t'\geq t}}\sum_{t'=t}^{\infty}e^{-r(t'-t)\Delta}\mathbb{E}\Bigg[&(1-e^{-r\Delta})\left(v_{i,t'\Delta}(z^c_{i,t'\Delta}+x_{i,t'\Delta}(p^c_{t'\Delta}))-\frac{\lambda}{2r}(z^c_{i,t'\Delta}+x_{i,t'\Delta}(p^c_{t'\Delta}))^2\right)\\
&-p^c_{t'\Delta}\cdot x_{i,t'\Delta}(p^c_{t'\Delta})\;\Bigg|\;H_{i,t'\Delta},p^c_{t'\Delta}\Bigg].
\end{aligned}\tag{72}
$$

The first-order condition of (72) with respect to $x_{i,t\Delta}$ at the competitive equilibrium $\{x^c_{i,t'\Delta}\}_{t'\geq t}$ is

$$
\begin{aligned}
\mathbb{E}\Bigg[\sum_{t'=t}^{\infty}e^{-r(t'-t)\Delta}\Bigg(&(1-e^{-r\Delta})\left(v_{i,t'\Delta}-\frac{\lambda}{r}(z^c_{i,t'\Delta}+x^c_{i,t'\Delta}(p^c_{t'\Delta}))\right)\frac{\partial(z^c_{i,t'\Delta}+x^c_{i,t'\Delta}(p^c_{t'\Delta}))}{\partial x^c_{i,t\Delta}(p^c_{t\Delta})}\\
&-p^c_{t'\Delta}\frac{\partial x^c_{i,t'\Delta}(p^c_{t'\Delta})}{\partial x^c_{i,t\Delta}(p^c_{t\Delta})}\Bigg)\;\Bigg|\;H_{i,t\Delta},p^c_{t\Delta}\Bigg]=0.
\end{aligned}\tag{73}
$$

Under the derived strategy $x^c_{i,t\Delta}$ in Equation (25),

$$
\frac{\partial x^c_{i,t'\Delta}(p^c_{t'\Delta})}{\partial x^c_{i,t\Delta}(p^c_{t\Delta})}=\begin{cases}-1, & \text{if } t'=t+1\\ 0, & \text{if } t'>t+1\end{cases}, \tag{74}
$$

and

$$
\frac{\partial(z^c_{i,t'\Delta}+x^c_{i,t'\Delta}(p^c_{t'\Delta}))}{\partial x^c_{i,t\Delta}(p^c_{t\Delta})}=0,\quad t'>t. \tag{75}
$$

So the first-order condition reduces to

$$
\mathbb{E}\left[(1-e^{-r\Delta})\left(v_{i,t\Delta}-\frac{\lambda}{r}(z^c_{i,t\Delta}+x^c_{i,t\Delta}(p^c_{t\Delta}))\right)-p^c_{t\Delta}+e^{-r\Delta}p^c_{(t+1)\Delta}\;\middle|\;H_{i,t\Delta},p^c_{t\Delta}\right]=0. \tag{76}
$$

Because the price is a martingale, i.e. $\mathbb{E}\left[p^c_{(t+1)\Delta} \mid H_{i,t\Delta}, p^c_{t\Delta}\right] = p^c_{t\Delta}$, the above equation reduces to the conjecture (15).

C.3 Proof of Proposition 2

We conjecture that traders use the following linear, symmetric and stationary strategy:

$$x_{j,t\Delta}(p; s_{j,t\Delta}, z_{j,t\Delta}) = as_{j,t\Delta} - bp + dz_{j,t\Delta} + fZ. \tag{77}$$

This conjecture implies the market-clearing prices of

$$p^*_{t\Delta} = \frac{a}{nb}\sum_{j=1}^{n} s_{j,t\Delta} + \frac{d+nf}{nb}Z. \tag{78}$$

Fix a history $H_{i,t\Delta}$ and a realization of $\sum_{j\neq i} s_{j,t\Delta}$. We use the single-deviation principle to construct an equilibrium strategy (77): under the conjecture that other traders $j \neq i$ use strategy (77) in every period $t' \geq t$, and that trader i returns to strategy (77) in period $t' \geq t+1$, we verify that trader i has no incentive to deviate from strategy (77) in period t.[23]

If trader i uses an alternative demand schedule in period t, he faces the residual demand $-\sum_{j\neq i} x_{j,t\Delta}(p_{t\Delta})$ and is effectively choosing a price $p_{t\Delta}$ and getting $x_{i,t\Delta}(p_{t\Delta}) = -\sum_{j\neq i} x_{j,t\Delta}(p_{t\Delta})$. Therefore, by differentiating trader i's expected utility in period t with respect to $p_{t\Delta}$ and evaluating it at $p_{t\Delta} = p^*_{t\Delta}$ in Equation (78), we obtain the following first order condition in period t of trader i:

$$\begin{aligned}\mathbb{E}\Bigg[&(n-1)b\cdot\left((1-e^{-r\Delta})\sum_{k=0}^{\infty} e^{-rk\Delta}\frac{\partial(z_{i,(t+k)\Delta} + x^*_{i,(t+k)\Delta})}{\partial x^*_{i,t\Delta}}\left(v_{i,(t+k)\Delta} - \frac{\lambda}{r}(z_{i,(t+k)\Delta} + x^*_{i,(t+k)\Delta})\right)\right.\\ &\left.-\sum_{k=0}^{\infty} e^{-rk\Delta}\frac{\partial x^*_{i,(t+k)\Delta}}{\partial x^*_{i,t\Delta}} p^*_{(t+k)\Delta}\right) - \sum_{k=0}^{\infty} e^{-rk\Delta} x^*_{i,(t+k)\Delta}\frac{\partial p^*_{(t+k)\Delta}}{\partial p_{t\Delta}}\,\Bigg|\, H_{i,t\Delta}\cup\left\{\sum\nolimits_{j\neq i} s_{j,t\Delta}\right\}\Bigg] = 0,\end{aligned} \tag{79}$$

where we write $x^*_{i,(t+k)\Delta} = x_{i,(t+k)\Delta}(p^*_{(t+k)\Delta}; s_{i,(t+k)\Delta}, z_{i,(t+k)\Delta})$ for the strategy $x_{i,(t+k)\Delta}(\,\cdot\,)$ defined in Equation (77), and by definition $z_{i,(t+k+1)\Delta} = z_{i,(t+k)\Delta} + x^*_{i,(t+k)\Delta}$.

Since all traders follow the conjectured strategy in Equation (77) from period $t+1$ and onwards, we have the following evolution of inventories: for any $k \geq 1$,

$$\begin{aligned} z_{i,(t+k)\Delta} + x^*_{i,(t+k)\Delta} =& as_{i,(t+k)\Delta} - bp^*_{(t+k)\Delta} + fZ + (1+d)z_{i,(t+k)\Delta} \\ =&(as_{i,(t+k)\Delta} - bp^*_{(t+k)\Delta} + fZ) + (1+d)(as_{i,(t+k-1)\Delta} - bp^*_{(t+k-1)\Delta} + fZ) \\ &+\cdots + (1+d)^{k-1}(as_{i,(t+1)\Delta} - bp^*_{(t+1)\Delta} + fZ) + (1+d)^k(x_{i,t\Delta} + z_{i,t\Delta}). \end{aligned} \tag{80}$$

The evolution of prices and inventories, given by Equations (78) and (80), reveals that by changing the demand or price in period t, trader i has the following effects on inventories and prices in period

[23] For a description of the single-deviation principle, also called "one-stage deviation principle", see Theorem 4.1 and 4.2 of Fudenberg and Tirole (1991). We can apply their Theorem 4.2 because the payoff function in our model, which takes the form of a "discounted" sum of period-by-period payoffs, satisfies the required "continuity at infinity" condition.

$t+k$, $k \geq 1$:

$$\frac{\partial(z_{i,(t+k)\Delta} + x^*_{i,(t+k)\Delta})}{\partial x^*_{i,t\Delta}} = (1+d)^k, \tag{81}$$

$$\frac{\partial x^*_{i,(t+k)\Delta}}{\partial x^*_{i,t\Delta}} = (1+d)^{k-1} d, \tag{82}$$

$$\frac{\partial p^*_{(t+k)\Delta}}{\partial p_{t\Delta}} = \frac{\partial p^*_{(t+k)\Delta}}{\partial x^*_{i,t\Delta}} = 0. \tag{83}$$

As we verify later, the equilibrium value of d satisfies $-1 < d < 0$, so the partial derivatives (81) and (82) converge.

The first order condition (79) simplifies to:

$$\mathbb{E}\Bigg[(n-1)b\Bigg((1-e^{-r\Delta})\sum_{k=0}^{\infty} e^{-rk\Delta}(1+d)^k \left(v_{i,(t+k)\Delta} - \frac{\lambda}{r}(z_{i,(t+k)\Delta} + x^*_{i,(t+k)\Delta})\right) \\ - p^*_{t\Delta} - \sum_{k=1}^{\infty} e^{-rk\Delta}(1+d)^{k-1} d\, p^*_{(t+k)\Delta}\Bigg) - x^*_{i,t\Delta} \,\Bigg|\, H_{i,t\Delta} \cup \left\{\sum\nolimits_{j\neq i} s_{j,t\Delta}\right\}\Bigg] = 0, \tag{84}$$

where we have (cf. Lemma 1, Equations (78) and (80)):

$$\mathbb{E}\left[p^*_{i,(t+k)\Delta} \mid H_{i,t\Delta} \cup \left\{\sum\nolimits_{j\neq i} s_{j,t\Delta}\right\}\right] = p^*_{t\Delta}, \tag{85}$$

$$\mathbb{E}\left[v_{i,(t+k)\Delta} \mid H_{i,t\Delta} \cup \left\{\sum\nolimits_{j\neq i} s_{j,t\Delta}\right\}\right] = \mathbb{E}\left[v_{i,t\Delta} \mid H_{i,t\Delta} \cup \left\{\sum\nolimits_{j\neq i} s_{j,t\Delta}\right\}\right] \\ = \alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1}\sum_{j\neq i} s_{j,t\Delta}, \tag{86}$$

$$\mathbb{E}\left[z_{i,(t+k)\Delta} + x^*_{i,(t+k)\Delta} \mid H_{i,t\Delta} \cup \left\{\sum\nolimits_{j\neq i} s_{j,t\Delta}\right\}\right] \\ = (a s_{i,t\Delta} - b p^*_{t\Delta} + fZ)\left(\frac{1}{-d} - \frac{(1+d)^k}{-d}\right) + (1+d)^k (x^*_{i,t\Delta} + z_{i,t\Delta}). \tag{87}$$

Substituting Equations (78), (85), (86) and (87) into the first-order condition (84) and using the notation $\bar{s}_{t\Delta} = \sum_{1\leq j\leq n} s_{j,t\Delta}/n$, we get:

$$(n-1)b(1-e^{-r\Delta})\Bigg[\frac{1}{1-e^{-r\Delta}(1+d)}\left(\alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1}\sum_{j\neq i} s_{j,t\Delta} - \left(\frac{a}{b}\bar{s}_{t\Delta} + \frac{d+nf}{nb}Z\right)\right) \\ - \sum_{k=0}^{\infty} \frac{\lambda}{r} e^{-rk\Delta}(1+d)^k \left(\frac{1}{-d} - \frac{(1+d)^k}{-d}\right)\left(a s_{i,t\Delta} - b\left(\frac{a}{b}\bar{s}_{t\Delta} + \frac{d+nf}{nb}Z\right) + fZ\right) \\ - \frac{\lambda}{(1-e^{-r\Delta}(1+d)^2)r}(x^*_{i,t\Delta} + z_{i,t\Delta})\Bigg] - x^*_{i,t\Delta} = 0. \tag{88}$$

Rearranging the terms gives:

$$\begin{aligned}
&\left(1+\frac{(n-1)b(1-e^{-r\Delta})\lambda}{(1-e^{-r\Delta}(1+d)^2)r}\right)x^*_{i,t\Delta} \qquad (89)\\
&=(n-1)b(1-e^{-r\Delta})\Bigg[\frac{1}{1-e^{-r\Delta}(1+d)}\left(\frac{n\alpha-1}{n-1}s_{i,t\Delta}+\frac{n-n\alpha}{n-1}\bar{s}_{t\Delta}-\frac{a}{b}\bar{s}_{t\Delta}\right)\\
&\qquad-\frac{\lambda e^{-r\Delta}(1+d)}{r(1-(1+d)e^{-r\Delta})(1-(1+d)^2e^{-r\Delta})}a(s_{i,t\Delta}-\bar{s}_{t\Delta})\\
&\qquad-\frac{\lambda}{(1-e^{-r\Delta}(1+d)^2)r}z_{i,t\Delta}\\
&\qquad-\left(\frac{1}{1-e^{-r\Delta}(1+d)}\left(\frac{d+nf}{nb}+\frac{\lambda}{rn}\right)-\frac{\lambda}{(1-(1+d)^2e^{-r\Delta})nr}\right)Z\Bigg].
\end{aligned}$$

On the other hand, substituting Equation (78) into the conjectured strategy (77) gives:

$$x^*_{i,t\Delta}=a(s_{i,t\Delta}-\bar{s}_{t\Delta})+dz_{i,t\Delta}-\frac{d}{n}Z. \qquad (90)$$

We match the coefficients in Equation (90) with those in Equation (89). First of all, we clearly have

$$a=b. \qquad (91)$$

We also obtain two equations for b and d:

$$\begin{aligned}
&\left(1+\frac{(n-1)b(1-e^{-r\Delta})\lambda}{(1-e^{-r\Delta}(1+d)^2)r}\right)=\frac{(1-e^{-r\Delta})(n\alpha-1)}{1-e^{-r\Delta}(1+d)}-\frac{(n-1)b(1-e^{-r\Delta})e^{-r\Delta}(1+d)\lambda}{(1-(1+d)e^{-r\Delta})(1-(1+d)^2e^{-r\Delta})r},\\
&\left(1+\frac{(n-1)b(1-e^{-r\Delta})\lambda}{(1-e^{-r\Delta}(1+d)^2)r}\right)d=-\frac{(n-1)b(1-e^{-r\Delta})\lambda}{(1-e^{-r\Delta}(1+d)^2)r}. \qquad (92)
\end{aligned}$$

There are two solutions to the above system of equations. One of them leads to unbounded inventories, so we drop it.[24] The other solution leads to converging inventories and is given by

$$b=\frac{(n\alpha-1)r}{2(n-1)e^{-r\Delta}\lambda}\left((n\alpha-1)(1-e^{-r\Delta})+2e^{-r\Delta}-\sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2+4e^{-r\Delta}}\right), \qquad (93)$$

$$d=-\frac{1}{2e^{-r\Delta}}\left((n\alpha-1)(1-e^{-r\Delta})+2e^{-r\Delta}-\sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2+4e^{-r\Delta}}\right). \qquad (94)$$

Lastly, matching the coefficient of Z gives:

$$f=-\frac{d}{n}-\frac{b\lambda}{nr}. \qquad (95)$$

Under the condition $n\alpha>2$, we can show that $b>0$ and $-1<d<0$, that is, the demand schedule is downward-sloping in price and the inventories evolutions (81)–(82) converge.

[24] This dropped solution to Equation (92) has the property of $(1+d)e^{-r\Delta}<-1$, which leads to an unbounded path of inventories (cf. Equation (80)) and utilities.

By Equation (19), the condition $n\alpha > 2$ is equivalent to the condition

$$\chi^2 < \frac{(n-2)\sigma_w^2}{n\sigma_\epsilon^2} \tag{96}$$

which is equivalent to the following condition on the fundamentals:

$$\frac{1}{n/2 + \sigma_\epsilon^2/\sigma_D^2} < \sqrt{\frac{n-2}{n}} \frac{\sigma_w}{\sigma_\epsilon}. \tag{97}$$

Finally, we verify the second-order condition. Under the linear strategy in Equation (77) with $b > 0$, differentiating the first-order condition (79) with respect to p_0 gives

$$(n-1)b(1-e^{-r\Delta})\left(-\frac{\lambda}{r}(n-1)b\sum_{k=0}^{\infty} e^{-rk\Delta}(1+d)^{2k} - 1\right) - (n-1)b < 0. \tag{98}$$

This completes the construction of a perfect Bayesian equilibrium.

C.4 Proof of Proposition 3

Suppose that every trader i use the strategy:

$$x_{i,t\Delta}(p) = \sum_{T_l \le t\Delta} a_l S_{i,T_l} + a_w w_{i,t\Delta} - bp + dz_{i,t\Delta} + f, \tag{99}$$

where $\{a_l\}_{l\ge 0}, a_w, b, d$ and f are constants. We show that for everyone using strategy (99) to be a perfect Bayesian equilibrium (PBE), the constants must be the ones given by Proposition 2. We divide our arguments into two steps.

Step 1. Define $x_l \equiv a_l/a_w$. [25] As a first step, we show that if strategy (99) is a symmetric PBE, then we must have $x_l = \chi$ for every l, where χ is defined in Equation (19).

Suppose that $(t-1)\Delta \in [T_{k'}, T_{k'+1})$ and $t\Delta \in [T_k, T_{k+1})$, so there are $k - k' \ge 1$ dividend shocks between time $(t-1)\Delta$ and time $t\Delta$. [26] Without loss of generality, assume $k' = 0$. Since all other traders $j \ne i$ are using strategy (99), by computing the difference $p^*_{t\Delta} - p^*_{(t-1)\Delta}$, trader i can infer from the period-t price the value of

$$\sum_{j\ne i}\sum_{l=1}^{k} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}.$$

By the projection theorem for normal distribution, we have

$$\begin{aligned} &\mathbb{E}\left[D_{T_k} - D_{T_0} \mid H_{i,t\Delta} \cup \left\{\sum\nolimits_{j\ne i}\sum\nolimits_{T_l\le t\Delta} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}\right\}\right] \\ =\,&\mathbb{E}\left[D_{T_k} - D_{T_0} \mid \{S_{i,T_l}\}_{l=1}^{k} \cup \left\{\sum\nolimits_{j\ne i}\sum\nolimits_{l=1}^{k} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}\right\}\right] \\ =\,&\mathbf{u}\,\mathbf{\Sigma}^{-1}\cdot\left(S_{i,T_1},\ldots,S_{i,T_k},\sum\nolimits_{j\ne i}\sum\nolimits_{l=1}^{k} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}\right)', \end{aligned} \tag{100}$$

[25] Clearly, we cannot have $a_w = 0$, since players use their private values in any equilibrium.
[26] In period $t = 0$, we take $D_{T_{-1}} = w_{i,T_{-1}} = 0$.

where $\mathbf{\Sigma}$ is the covariance matrix of $\left(S_{i,T_1}, \ldots, S_{i,T_k}, \sum_{j\neq i}\sum_{l=1}^{k} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}\right)$: for $1 \le l \le k+1$ and $1 \le m \le k+1$,

$$\mathbf{\Sigma}_{l,m} = \begin{cases} \sigma_D^2 + \sigma_\epsilon^2 & 1 \le l = m \le k \\ 0 & 1 \le l \neq m \le k \\ (n-1)(\sum_{l'=1}^{k} x_{l'}^2)(\sigma_D^2 + \sigma_\epsilon^2) + (n-1)k\sigma_w^2 & l = m = k+1 \\ \quad +(n-1)(n-2)\sum_{l'=1}^{k} x_{l'}^2 \sigma_D^2 & \\ (n-1)x_l\sigma_D^2 & 1 \le l \le k, m = k+1 \end{cases}, \tag{101}$$

and $\mathbf{\Sigma}_{k+1,l} = \mathbf{\Sigma}_{l,k+1}$. And $\mathbf{u}$ is a row vector of covariances between $\left(S_{i,T_1}, \ldots, S_{i,T_k}, \sum_{j\neq i}\sum_{l=1}^{k} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}\right)$ and $D_{T_k} - D_{T_0}$:

$$\mathbf{u} = (\sigma_D^2, \ldots, \sigma_D^2, (n-1)\sum_{l=1}^{k} x_l \sigma_D^2). \tag{102}$$

Therefore, we have

$$\begin{aligned} &\mathbb{E}\left[v_{i,t\Delta} \mid H_{i,t\Delta} \cup \left\{\sum_{j\neq i}\sum_{T_l \le t\Delta} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}\right\}\right] \\ =\,& w_{i,T_k} + \mathbb{E}[D_{T_0} \mid \{S_{i,T_0}\} \cup \left\{\sum_{j\neq i} x_0 S_{j,T_0} + w_{j,0}\right\}] \\ &+ \mathbf{u}\,\mathbf{\Sigma}^{-1} \cdot \left(S_{i,T_1}, \ldots, S_{i,T_k}, \sum_{j\neq i}\sum_{l=1}^{k} x_l S_{j,T_l} + w_{j,T_l} - w_{j,T_{l-1}}\right)'. \end{aligned} \tag{103}$$

Since we look for a symmetric equilibrium in which everyone plays strategy (99), trader i's conditional value in Equation (103) must place a weight of x_l on S_{i,T_l}, $1 \le l \le k$, which implies that

$$\mathbf{u}\,\mathbf{\Sigma}^{-1} = \mathbf{x}, \tag{104}$$

where $\mathbf{x} = (x_1, \ldots, x_k, y)$ and y is an arbitrary number. Clearly, Equation (104) is equivalent to

$$\mathbf{u} = \mathbf{x}\,\mathbf{\Sigma},$$

which implies (from the first k entries of the row vector)

$$\sigma_D^2 = x_l(\sigma_D^2 + \sigma_\epsilon^2) + y(n-1)x_l\sigma_D^2, \quad 1 \le l \le k,$$

i.e.,

$$x_1 = \cdots = x_k = \frac{\sigma_D^2}{\sigma_D^2 + \sigma_\epsilon^2 + y(n-1)\sigma_D^2}.$$

Now define $x \equiv x_1 = \cdots = x_k$. Applying Lemma 1 to the conditional value in Equation (103) implies that for the conditional value in Equation (103) to place a weight of x on S_{i,T_l}, $1 \le l \le k$, we must have $x = \chi$.

Step 2. Given Step 1, we can rewrite the strategy (99) as

$$x_{i,t\Delta}(p) = a_w \cdot \alpha\, s_{i,t\Delta} - bp + dz_{i,t\Delta} + f, \tag{105}$$

where $s_{i,t\Delta}$ is the total signal defined in Equation (20) and α is defined in Equation (21). The equilibrium construction in Appendix C.3 then uniquely determines the values of a_w, b, d and f. This concludes the proof of Proposition 3.

C.5 Proof of Proposition 4

The exponential convergence to efficient allocation follows directly from Equation (36).

Now we prove the comparative statics. We write

$$\eta \equiv n\alpha - 1. \tag{106}$$

and recall that

$$\begin{aligned} 1+d &= \frac{1}{2e^{-r\Delta}}\left(\sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2+4e^{-r\Delta}}-(n\alpha-1)(1-e^{-r\Delta})\right), \\ &= \frac{2}{\sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2+4e^{-r\Delta}}+(n\alpha-1)(1-e^{-r\Delta})}. \end{aligned} \tag{107}$$

We first note that $\frac{\partial(1+d)}{\partial\eta} < 0$.

1. The comparative statics with respect to r follow by straightforward calculations showing that $\frac{\partial(1+d)}{\partial r} < 0$.
2. As σ_D^2 increases, the left-hand side of Equation (19) increases, and hence the solution χ to Equation (19) increases, which means that $n\alpha$ decreases because according to Equation (21) $n\alpha$ is a decreasing function of χ^2. Thus, $\frac{\partial\eta}{\partial\sigma_D^2} < 0$, and $\frac{\partial(1+d)}{\partial\sigma_D^2} > 0$.
3. As σ_w^2 increases, the left-hand side of Equation (19) increases, and hence the solution χ to Equation (19) increases; by Equation (19) this means that σ_w^2/χ^2 must increase as well. Thus, $n\alpha$ increases because according to Equation (21) $n\alpha$ is an increasing function of σ_w^2/χ^2. Hence, $\frac{\partial\eta}{\partial\sigma_w^2} > 0$ and $\frac{\partial(1+d)}{\partial\sigma_w^2} < 0$.
4. We can rewrite Equation (19) as
$$\frac{1}{\frac{1}{\alpha}+\frac{\sigma_\epsilon^2}{\sigma_D^2}} = \chi, \tag{108}$$
and Equation (21) as
$$\chi = \sqrt{\frac{1-\alpha}{n\alpha-1}}\frac{\sigma_w}{\sigma_\epsilon}, \tag{109}$$
and hence
$$\frac{1}{\frac{n}{\eta+1}+\frac{\sigma_\epsilon^2}{\sigma_D^2}} = \sqrt{\frac{n-\eta-1}{n\eta}}\frac{\sigma_w}{\sigma_\epsilon}. \tag{110}$$
From Equation (110) is it straightforward to show that η must increase with n. Thus, $1+d$ decreases in n.
5. For the comparative statics with respect to Δ, direction calculation shows that $\partial(1+d)/\partial\Delta < 0$.

C.6 Proofs of Propositions 5, 6, 7, 8 and 10

We first establish some general properties of the equilibrium welfare, before specializing to the optimal trading frequency given scheduled (Appendices C.6.1 and C.6.2) and stochastic (Appendix C.6.3) arrivals of new information, as well as to the rate that inefficiency vanishes as $n \to \infty$ (Appendix C.6.4).

Lemma 2. *For any profile of inventories* $(z_1, z_2, \ldots, z_n)$ *satisfying* $\sum_{i=1}^n z_i = Z$ *and any profile of total signals* $(s_{1,t\Delta}, s_{2,t\Delta}, \ldots, s_{n,t\Delta})$*, we have:*

$$\begin{aligned}
&\sum_{i=1}^n \left(\left(\alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1} \sum_{j \neq i} s_{j,t\Delta} \right) z^c_{i,(t+1)\Delta} - \frac{\lambda}{2r} (z^c_{i,(t+1)\Delta})^2 \right) \\
&- \sum_{i=1}^n \left(\left(\alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1} \sum_{j \neq i} s_{j,t\Delta} \right) z_i - \frac{\lambda}{2r} (z_i)^2 \right) \\
= &\frac{\lambda}{2r} \sum_{i=1}^n (z^c_{i,(t+1)\Delta} - z_i)^2.
\end{aligned} \tag{111}$$

Remark. Recall that $z^c_{i,(t+1)\Delta} = z^e_{i,t\Delta}$ is the competitive allocation given total signals $\{s_{i,t\Delta}\}_{i=1}^n$.

Proof of Lemma 2. Since $(z_i)^2 = (z^c_{i,(t+1)\Delta})^2 + 2z^c_{i,(t+1)\Delta}(z_i - z^c_{i,(t+1)\Delta}) + (z_i - z^c_{i,(t+1)\Delta})^2$, we have:

$$\begin{aligned}
&\sum_{i=1}^n \left(\left(\alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1} \sum_{j \neq i} s_{j,t\Delta} \right) z_i - \frac{\lambda}{2r} (z_i)^2 \right) \\
= &\sum_{i=1}^n \left(\left(\alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1} \sum_{j \neq i} s_{j,t\Delta} \right) z^c_{i,(t+1)\Delta} - \frac{\lambda}{2r} (z^c_{i,(t+1)\Delta})^2 \right) \\
&+ \sum_{i=1}^n \left(\left(\alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1} \sum_{j \neq i} s_{j,t\Delta} \right) - \frac{\lambda}{r} z^c_{i,(t+1)\Delta} \right) (z_i - z^c_{i,(t+1)\Delta}) \\
&- \frac{\lambda}{2r} \sum_{i=1}^n (z_i - z^c_{i,(t+1)\Delta})^2.
\end{aligned} \tag{112}$$

The middle term in Equation (112) is zero because

$$\left(\alpha s_{i,t\Delta} + \frac{1-\alpha}{n-1} \sum_{j \neq i} s_{j,t\Delta} \right) - \frac{\lambda}{r} z^c_{i,(t+1)\Delta} = p^c_{t\Delta} \tag{113}$$

for the competitive equilibrium price $p^c_{t\Delta}$ (cf. Equations (15) and (23)), and $\sum_{i=1}^n p^c_{t\Delta}(z_i - z^c_{i,(t+1)}) = p^c_{t\Delta}(Z - Z) = 0$. □

Lemma 3.

$$X_1(\Delta) = \frac{\lambda(1+d)}{2r(n\alpha - 1)} \left(\sum_{i=1}^n \mathbb{E}[(z_{i,0} - z^c_{i,\Delta})^2] + \sum_{t=1}^\infty e^{-rt\Delta} \sum_{i=1}^n \mathbb{E}[(z^c_{i,(t+1)\Delta} - z^c_{i,t\Delta})^2] \right). \tag{114}$$

Proof of Lemma 3. We first simplify the squared difference between the strategic and competitive equilibrium allocation, for $t \geq 1$:

$$\begin{aligned}\mathbb{E}[(z^*_{i,(t+1)\Delta} - z^c_{i,(t+1)\Delta})^2] &= (1+d)^2\mathbb{E}[(z^*_{i,t\Delta} - z^c_{i,(t+1)\Delta})^2] \\ &= (1+d)^2\mathbb{E}[(z^*_{i,t\Delta} - z^c_{i,t\Delta})^2] + (1+d)^2\mathbb{E}[(z^c_{i,(t+1)\Delta} - z^c_{i,t\Delta})^2],\end{aligned} \tag{115}$$

where the first equality follows from Proposition 4, and the second equality follows from the fact that $z^*_{i,t\Delta}$ and $z^c_{i,t\Delta}$ are measurable with respect to the information in period $t-1$, and that $\{z^c_{i,t\Delta}\}_{t\geq 0}$ is a martingale, so $\mathbb{E}[(z^*_{i,t\Delta} - z^c_{i,t\Delta})(z^c_{i,(t+1)\Delta} - z^c_{i,t\Delta})] = 0$ by the law of iterated expectations. Then by induction, we have:

$$\mathbb{E}[(z^*_{i,(t+1)\Delta} - z^c_{i,(t+1)\Delta})^2] = (1+d)^{2(t+1)}\mathbb{E}[(z_{i,0} - z^c_{i,\Delta})^2] + \sum_{t'=1}^{t}(1+d)^{2(t-t'+1)}\mathbb{E}[(z^c_{i,(t'+1)\Delta} - z^c_{i,t'\Delta})^2]. \tag{116}$$

The above equation says that after auction t, allocative inefficiency is a linear combination of the inefficiency in initial allocations and the time variations in the competitive allocation up to time $t\Delta$.

Applying Lemma 2, we express X_1 as the weighted quadratic difference between the strategic and competitive allocations:

$$X_1(\Delta) = (1 - e^{-r\Delta}) \cdot \frac{\lambda}{2r} \sum_{i=1}^{n}\sum_{t=0}^{\infty} e^{-rt\Delta}\mathbb{E}[(z^*_{i,(t+1)\Delta} - z^c_{i,(t+1)\Delta})^2]. \tag{117}$$

Substituting Equation (116) into the expression of X_1 in Equation (117), we get:

$$\begin{aligned}
&X_1(\Delta) \\
=& \frac{\lambda(1-e^{-r\Delta})}{2r}\sum_{i=1}^{n}\sum_{t=0}^{\infty} e^{-rt\Delta}\left((1+d)^{2(t+1)}\mathbb{E}[(z_{i,0} - z^c_{i,\Delta})^2] + \sum_{t'=1}^{t}(1+d)^{2(t-t'+1)}\mathbb{E}[(z^c_{i,(t'+1)\Delta} - z^c_{i,t'\Delta})^2]\right) \\
=& \frac{\lambda}{2r}\frac{(1-e^{-r\Delta})(1+d)^2}{1-(1+d)^2e^{-r\Delta}}\sum_{i=1}^{n}\mathbb{E}[(z_{i,0} - z^c_{i,\Delta})^2] \\
&+ \frac{1-e^{-r\Delta}}{r}\frac{\lambda}{2}\sum_{i=1}^{n}\sum_{t'=1}^{\infty}\mathbb{E}[(z^c_{i,(t'+1)\Delta} - z^c_{i,t'\Delta})^2]\sum_{t=t'}^{\infty}e^{-rt\Delta}(1+d)^{2(t-t'+1)} \\
=& \frac{\lambda}{2r}\frac{(1-e^{-r\Delta})(1+d)^2}{1-(1+d)^2e^{-r\Delta}}\sum_{i=1}^{n}\mathbb{E}[(z_{i,0} - z^c_{i,\Delta})^2] \\
&+ \frac{\lambda}{2r}\frac{(1-e^{-r\Delta})(1+d)^2}{1-(1+d)^2e^{-r\Delta}}\sum_{i=1}^{n}\sum_{t'=1}^{\infty}\mathbb{E}[(z^c_{i,(t'+1)\Delta} - z^c_{i,t'\Delta})^2]e^{-rt'\Delta}.
\end{aligned} \tag{118}$$

We can simplify the constant in the above equations by a direct calculation:

$$\begin{aligned}
&e^{-r\Delta}(1+d)^2 \\
=& \frac{2(n\alpha-1)^2(1-e^{-r\Delta})^2 + 4e^{-r\Delta} - 2(n\alpha-1)(1-e^{-r\Delta})\sqrt{(n\alpha-1)^2(1-e^{-r\Delta})^2 + 4e^{-r\Delta}}}{4e^{-r\Delta}} \\
=& 1 - (n\alpha-1)(1-e^{-r\Delta})(1+d),
\end{aligned} \tag{119}$$

which implies:

$$\frac{(1-e^{-r\Delta})(1+d)^2}{1-(1+d)^2e^{-r\Delta}} = \frac{1+d}{n\alpha-1}. \tag{120}$$

$\square$

C.6.1 Proof of Proposition 5

For any $\tau > 0$, we let $\bar{t}(\tau) = \min\{t \geq 0 : t \in \mathbb{Z}, t\Delta \geq \tau\}$. That is, if new signals arrive at the clock time τ, then $\bar{t}(\tau)\Delta$ is the clock time of the next trading period (including time τ).

For any $\Delta \leq \gamma$, by the assumption of Proposition 5 there is at most one new signal profile arrival in each interval $[t\Delta, (t+1)\Delta)$. Thus, we only need to count the changes in competitive allocation between period $\bar{t}((k-1)\gamma)$ and $\bar{t}(k\gamma)$, for $k \in \mathbb{Z}_+$. Using this fact, we can rewrite $X_1(\Delta)$ and $X_2(\Delta)$ as:

$$\begin{aligned} X_1(\Delta) &= \frac{\lambda(1+d)}{2r(n\alpha-1)}\left(\sum_{i=1}^{n}\mathbb{E}[(z_{i,0}-z^e_{i,0})^2] + \sum_{i=1}^{n}\sum_{k=1}^{\infty} e^{-r\bar{t}(k\gamma)\Delta}\mathbb{E}[(z^e_{i,k\gamma}-z^e_{i,(k-1)\gamma})^2]\right) \\ &= \frac{\lambda(1+d)}{2r(n\alpha-1)}\left(\sum_{i=1}^{n}\mathbb{E}[(z_{i,0}-z^e_{i,0})^2] + \sum_{i=1}^{n}\sum_{k=1}^{\infty} e^{-rk\gamma}\mathbb{E}[(z^e_{i,k\gamma}-z^e_{i,(k-1)\gamma})^2]\right) \\ &\quad - \frac{\lambda(1+d)}{2r(n\alpha-1)}\sum_{i=1}^{n}\sum_{k=1}^{\infty}(e^{-rk\gamma}-e^{-r\bar{t}(k\gamma)\Delta})\mathbb{E}[(z^e_{i,k\gamma}-z^e_{i,(k-1)\gamma})^2]. \end{aligned} \tag{121}$$

and

$$\begin{aligned} X_2(\Delta) &= \frac{\lambda}{2r}\sum_{i=1}^{n}\sum_{t=0}^{\infty}\int_{\tau=t\Delta}^{(t+1)\Delta} re^{-r\tau}\mathbb{E}[(z^e_{i,t\Delta}-z^e_{i,\tau})^2]\,d\tau \\ &= \frac{\lambda}{2r}\sum_{i=1}^{n}\sum_{k=1}^{\infty}(e^{-rk\gamma}-e^{-r\bar{t}(k\gamma)\Delta})\mathbb{E}[(z^e_{i,k\gamma}-z^e_{i,(k-1)\gamma})^2]. \end{aligned} \tag{122}$$

Note that all the expectations in the expressions of $X_1(\Delta)$ and $X_2(\Delta)$ do not depend on Δ. To make clear the dependence of d on Δ, we now write $d = d(\Delta)$. Since $(1+d(\Delta))/(n\alpha-1) < 1$, we have for any $\Delta < \gamma$:

$$\begin{aligned} X(\Delta) &= X_1(\Delta) + X_2(\Delta) \\ &> \frac{\lambda(1+d(\Delta))}{2r(n\alpha-1)}\left(\sum_{i=1}^{n}\mathbb{E}[(z_{i,0}-z^e_{i,0})^2] + \sum_{i=1}^{n}\sum_{k=1}^{\infty} e^{-rk\gamma}\mathbb{E}[(z^e_{i,k\gamma}-z^e_{i,(k-1)\gamma})^2]\right) \\ &> \frac{\lambda(1+d(\gamma))}{2r(n\alpha-1)}\left(\sum_{i=1}^{n}\mathbb{E}[(z_{i,0}-z^e_{i,0})^2] + \sum_{i=1}^{n}\sum_{k=1}^{\infty} e^{-rk\gamma}\mathbb{E}[(z^e_{i,k\gamma}-z^e_{i,(k-1)\gamma})^2]\right) \\ &= X(\gamma), \end{aligned} \tag{123}$$

where the last inequality holds because $d(\Delta)$ decreases with Δ (which can be verified by taking derivative the $d'(\Delta)$) and where the last equality holds because $\bar{t}(k\gamma)\Delta = k\gamma$ if $\gamma = \Delta$. Therefore, we have $W(\Delta) < W(\gamma)$ for any $\Delta < \gamma$. This proves Proposition 5.

Notice that for this lower bound of $\Delta^* \geq \gamma$ we make no use of the assumption that $\mathbb{E}[(z^e_{i,k\gamma} - z^e_{i,(k-1)\gamma})^2]$ is a constant independent of k. Thus $\Delta^* \geq \gamma$ also holds if traders have no common value but have private value shocks $w_{i,k\gamma} - w_{i,(k-1)\gamma}$ that are non-stationary in k.

C.6.2 Proof of Proposition 6

If $\Delta = l\gamma$, where $l \geq 1$ is an integer, we have:

$$X_1(l\gamma) = \frac{\lambda(1+d(l\gamma))}{2r(n\alpha-1)}\left(\sigma_0^2 + \sum_{t=0}^{\infty} e^{-r(t+1)l\gamma} l\sigma_z^2\right) = \frac{\lambda(1+d(l\gamma))}{2r(n\alpha-1)}\left(\sigma_0^2 + \frac{e^{-rl\gamma}}{1-e^{-rl\gamma}} l\sigma_z^2\right), \tag{124}$$

$$\begin{aligned} X_2(l\gamma) &= \frac{\lambda}{2r}\frac{1}{1-e^{-rl\gamma}}\left((e^{-\gamma r} - e^{-2\gamma r}) + 2(e^{-2\gamma r} - e^{-3\gamma r}) + \cdots + (l-1)(e^{-(l-1)\gamma r} - e^{-l\gamma r})\right)\sigma_z^2 \\ &= \frac{\lambda}{2r}\frac{1}{1-e^{-rl\gamma}}\left(e^{-\gamma r} + e^{-2\gamma r} + e^{-3\gamma r} + \cdots + e^{-(l-1)\gamma r} - (l-1)e^{-l\gamma r}\right)\sigma_z^2 \\ &= \frac{\lambda}{2r}\frac{1}{1-e^{-rl\gamma}}\left(\frac{1-e^{-rl\gamma}}{1-e^{-\gamma r}} - 1 - (l-1)e^{-rl\gamma}\right)\sigma_z^2 \\ &= \frac{\lambda}{2r}\left(\frac{1}{1-e^{-\gamma r}} - 1 - l\frac{e^{-rl\gamma}}{1-e^{-rl\gamma}}\right)\sigma_z^2. \end{aligned} \tag{125}$$

Hence, if $\Delta = l\gamma$, $l \in \mathbb{Z}_+$, we have:

$$X(l\gamma) = \frac{\lambda(1+d(l\gamma))}{2r(n\alpha-1)}\sigma_0^2 - \frac{\lambda}{2r}\left(1 - \frac{1+d(l\gamma)}{n\alpha-1}\right)\frac{le^{-rl\gamma}}{1-e^{-rl\gamma}}\sigma_z^2 + \frac{\lambda e^{-\gamma r}}{2r(1-e^{-\gamma r})}\sigma_z^2. \tag{126}$$

By taking derivative, we can show that the function (involved in the first term in Equation (126))

$$\frac{1+d(\Delta)}{n\alpha-1} = \frac{1}{2e^{-r\Delta}}\left(\sqrt{(1-e^{-r\Delta})^2 + \frac{4e^{-r\Delta}}{(n\alpha-1)^2}} - (1-e^{-r\Delta})\right)$$

is strictly decreasing in Δ, while

$$\left(1 - \frac{1+d(\Delta)}{n\alpha-1}\right)\frac{\Delta e^{-r\Delta}}{1-e^{-r\Delta}}$$

involved in the second term in Equation (126) is also strictly decreasing in Δ.

We first prove part 2 of Proposition 6. As n tends to infinity, the proof of Proposition 10 implies that $n\alpha$ tends to infinity as well. As $n\alpha \to \infty$, $(1+d(l\gamma))/(n\alpha-1) \to 0$ for every $l \in \mathbb{Z}_+$, and by assumption σ_0^2/σ_z^2 remains bounded, so the second term in Equation (126) dominates, and hence $X(l\gamma)$ is minimized at $l^* = 1$.

For part 1 of Proposition 6, suppose $z_{i,0} = Z/n$ for every i, so we have $\sigma_0^2 = \sigma_z^2$. Minimizing $X(l\gamma)$ over positive integers l is equivalent to maximizing $\tilde{W}(l\gamma)$ over l:

$$\tilde{W}(l\gamma) \equiv \log\left(1 - \frac{1+d(l\gamma)}{n\alpha-1}\right) + \log\left(1 + \frac{le^{-rl\gamma}}{1-e^{-rl\gamma}}\right). \tag{127}$$

We have:

$$\tilde{W} = \log\left(1 + \delta - \sqrt{(1-\delta)^2 + 4\delta y}\right) - \log(2\delta) + \log\left(1 - \frac{\log(\delta)\delta}{r\gamma(1-\delta)}\right), \tag{128}$$

where

$$\delta \equiv e^{-rl\gamma}, \quad y \equiv \frac{1}{(n\alpha-1)^2}. \tag{129}$$

We calculate:

$$\frac{d\tilde{W}}{d\delta} = \frac{1-\delta-\sqrt{1+\delta(-2+4y+\delta)}}{2\delta\sqrt{1+\delta(-2+4y+\delta)}} + \frac{\delta-1-\log(\delta)}{(1-\delta)(r\gamma(1-\delta)-\delta\log(\delta))}. \tag{130}$$

Clearly, $\frac{1-\delta-\sqrt{1+\delta(-2+4y+\delta)}}{2\delta\sqrt{1+\delta(-2+4y+\delta)}}$ is decreasing in y, and $y \in (0,1)$, so

$$\frac{1-\delta-\sqrt{1+\delta(-2+4y+\delta)}}{2\delta\sqrt{1+\delta(-2+4y+\delta)}} > -\frac{1}{1+\delta}, \tag{131}$$

where the right-hand side is obtained from substituting $y=1$ to the left-hand side.

Thus,

$$\frac{d\tilde{W}}{d\delta} > -\frac{1}{1+\delta} + \frac{\delta-1-\log(\delta)}{(1-\delta)(r\gamma(1-\delta)-\delta\log(\delta))} > 0, \tag{132}$$

if and only if

$$-\frac{(1+\delta^2)\log(\delta)}{(1-\delta)^2} - \frac{1-\delta^2}{(1-\delta)^2} > r\gamma, \tag{133}$$

which is satisfied whenever $0 \le \delta \le e^{-1.5r\gamma}$: the left-hand side is decreasing in δ, and is equal to $\frac{1-e^{3r\gamma}+1.5(1+e^{3r\gamma})r\gamma}{(e^{1.5r\gamma}-1)^2} > r\gamma$ when $\delta = e^{-1.5r\gamma}$. This proves that the l^* that maximizes $\tilde{W}$ (and hence minimizes X) satisfies $l^* \le 2$.

To show that $l^* = 1$, we calculate that:

$$\int_{e^{-2r\gamma}}^{e^{-r\gamma}} \frac{d\tilde{W}}{d\delta}\, d\delta > \int_{e^{-2r\gamma}}^{e^{-r\gamma}} -\frac{1}{1+\delta} + \frac{\delta-1-\log(\delta)}{(1-\delta)(r\gamma(1-\delta)-\delta\log(\delta))}\, d\delta = \log\left(\frac{r\gamma(\delta-1)+\delta\log(\delta)}{1-\delta^2}\right)\Bigg|_{e^{-2r\gamma}}^{e^{-r\gamma}} = 0. \tag{134}$$

C.6.3 Proofs of Proposition 7 and Proposition 8

We have:

$$\begin{aligned} X_1(\Delta) &= \frac{\lambda(1+d)}{2r(n\alpha-1)}\left(\sigma_0^2 + \sum_{t=0}^{\infty} e^{-r(t+1)\Delta}\Delta\mu\sigma_z^2\right) \\ &= \frac{\lambda(1+d)}{2r(n\alpha-1)}\left(\sigma_0^2 + \frac{\Delta e^{-r\Delta}}{1-e^{-r\Delta}}\mu\sigma_z^2\right) \end{aligned} \tag{135}$$

and

$$X_2(\Delta) = \frac{\lambda}{2r}\sum_{t=0}^{\infty} e^{-rt\Delta}\int_{\tau=0}^{\Delta} re^{-r\tau}\tau\mu\sigma_z^2\, d\tau = -\frac{\lambda}{2r}\frac{\Delta e^{-r\Delta}}{1-e^{-r\Delta}}\mu\sigma_z^2 + \frac{\lambda}{2r^2}\mu\sigma_z^2. \tag{136}$$

Therefore,

$$X(\Delta) = \frac{\lambda(1+d)}{2r(n\alpha-1)}\sigma_0^2 - \frac{\lambda}{2r}\left(1-\frac{1+d}{n\alpha-1}\right)\frac{\Delta e^{-r\Delta}}{1-e^{-r\Delta}}\mu\sigma_z^2 + \frac{\lambda}{2r^2}\mu\sigma_z^2. \tag{137}$$

We note that the above is the same expression as Equation (126) in the proof of Proposition 6, replacing μ with $1/\gamma$ and Δ with $l\gamma$, and ignoring the last term which is independent of Δ. The result (Part 2 of Proposition 8) for $n \to \infty$ has the same proof as that in Proposition 6. For Part 1 of Proposition 8, we note that as established in the proof of Proposition 6, $\frac{1+d}{n\alpha-1}$ and $\left(1-\frac{1+d}{n\alpha-1}\right)\frac{\Delta e^{-r\Delta}}{1-e^{-r\Delta}}$ are both decreasing in Δ, so $\frac{\partial^2 X}{\partial\Delta\partial\mu} > 0$ and as μ becomes larger the second term becomes more important than the first term.

For the proof of Proposition 7, suppose that $z_{i,0} = Z/n$ for every trader i. Then we have $\sigma_0^2 = \sigma_z^2$. Minimizing $X(\Delta)$ over Δ is equivalent to maximizing $\tilde{W}(\Delta)$ over Δ, where:

$$\tilde{W}(\Delta) \equiv \log\left(1 - \frac{1+d(\Delta)}{n\alpha - 1}\right) + \log\left(1 + \frac{\mu\Delta e^{-r\Delta}}{1-e^{-r\Delta}}\right). \tag{138}$$

We have:

$$\tilde{W} = \log\left(1+\delta-\sqrt{(1-\delta)^2+4\delta y}\right) - \log(2\delta) + \log\left(1 - \frac{\mu\log(\delta)\delta}{r(1-\delta)}\right), \tag{139}$$

where

$$\delta \equiv e^{-r\Delta}, \quad y \equiv \frac{1}{(n\alpha-1)^2}. \tag{140}$$

We calculate:

$$\frac{d\tilde{W}}{d\delta} = \frac{1-\delta-\sqrt{1+\delta(-2+4y+\delta)}}{2\delta\sqrt{1+\delta(-2+4y+\delta)}} + \frac{\delta-1-\log(\delta)}{(1-\delta)((1-\delta)r/\mu - \delta\log(\delta))}, \tag{141}$$

We note that the first term in the righthand side of Equation (141) is negative, while the second term is positive. Moreover,

$$-\frac{1-\delta-\sqrt{1+\delta(-2+4y+\delta)}}{2\delta\sqrt{1+\delta(-2+4y+\delta)}} \Bigg/ \frac{\delta-1-\log(\delta)}{(1-\delta)((1-\delta)r/\mu - \delta\log(\delta))}$$

is increasing in δ, tends to 0 as $\delta \to 0$, and tends to $1 + r/\mu$ as $\delta \to 1$. Therefore, there exists a unique δ^* at which $\frac{d\tilde{W}}{d\delta} = 0$, and such δ^* maximizes $\tilde{W}$.

Equation (141) implies that

$$\frac{d^2\tilde{W}}{d\delta\, dy} = -\frac{(1-\delta)}{(1+\delta(4y+\delta-2))^{3/2}} < 0 \tag{142}$$

for every $\delta \in (0,1)$ and $y \in (0,1)$. Therefore, the optimal δ^* that maximizes $\tilde{W}$ is strictly decreasing with y, i.e., the optimal Δ^* that maximizes $\tilde{W}$ is strictly decreasing with $n\alpha$. We have previously established that α is increasing with σ_w^2 and is decreasing with σ_D^2, and $n\alpha$ is increasing with n. This concludes the proof of Part 2 for Proposition 7.

From Equation (141), we have $\frac{d\tilde{W}}{d\delta} > 0$ if

$$y < \frac{(1-\delta)^2}{4\delta}\left(\left(\frac{(1-\delta)((1-\delta)r/\mu - \delta\log(\delta))}{(1-\delta)^2 r/\mu + 2\delta(1-\delta) + (\delta^2+\delta)\log(\delta)}\right)^2 - 1\right). \tag{143}$$

After substituting $\delta = e^{-r/(l\mu)}$, the righthand side of the above equation is increasing in r/μ if $l \geq 2/3$.

As $r/\mu \to 0$, the righthand side tends to $9/(6l-1)^2$. This shows that for any $l \geq 2/3$, we have $\Delta^* < 1/(l\mu)$ if $n\alpha > 2l + 2/3$. Thus we have proved Part 1 of Proposition 7.

C.6.4 Proofs of Proposition 10

Suppose that $T_0 = 0$ and $\{T_k\}_{k\geq 1}$ is a homogeneous Poisson process with intensity $\mu > 0$. (The proof for scheduled information arrivals $T_k = k\gamma$ is analogous and omitted.)

Lemma 3 then implies that

$$\frac{X_1(\Delta)}{n} = \frac{\lambda(1+d(\Delta))}{2r(n\alpha-1)} \cdot \left(\frac{\sum_{i=1}^n \mathbb{E}[(z_{i,0} - z_{i,0}^e)^2]}{n} + \frac{e^{-r\Delta}\mu\Delta}{1-e^{-r\Delta}} \frac{\sum_{i=1}^n \mathbb{E}[(z_{i,T_k}^e - z_{i,T_{k-1}}^e)^2]}{n} \right), \tag{144}$$

where for any $k \geq 1$,

$$\frac{\sum_{i=1}^n \mathbb{E}[(z_{i,T_k}^e - z_{i,T_{k-1}}^e)^2]}{n} = \left(\frac{r(n\alpha-1)}{\lambda(n-1)}\right)^2 \frac{(n-1)(\chi^2(\sigma_D^2+\sigma_\epsilon^2)+\sigma_w^2)}{n\alpha^2}, \tag{145}$$

by Equation (47).

Equation (145) tends to a positive constant as $n \to \infty$ (since $\chi \to 0$ as $n \to \infty$), and $\lim_{\Delta\to 0} \frac{e^{-r\Delta}\mu\Delta}{1-e^{-r\Delta}} = \frac{\mu}{r}$. By assumption, $\sum_{i=1}^n \mathbb{E}[(z_{i,0}-z_{i,0}^e)^2]/n$ is bounded as $n \to \infty$. Thus, for $\lim_{n\to\infty} X_1(\Delta)/n$ it suffices to analyze

$$\frac{1+d(\Delta)}{n\alpha-1} = \frac{1}{2e^{-r\Delta}} \left(\sqrt{(1-e^{-r\Delta})^2 + \frac{4e^{-r\Delta}}{(n\alpha-1)^2}} - (1-e^{-r\Delta}) \right). \tag{146}$$

Suppose $\sigma_D^2 > 0$. Equation (110) (where $\eta \equiv n\alpha - 1$) implies that $n\alpha$ is of order $n^{2/3}$ as $n \to \infty$. To see this, first note that $\eta \to \infty$ and $\eta/n \to 0$ as $n \to \infty$, for otherwise the left-hand side and right-hand side of Equation (110) cannot match. Suppose that as n becomes large, η is of order n^y for some $y < 1$. The left-hand side of Equation (110) is of order n^{y-1}, and the right-hand side is of order $n^{-y/2}$. Thus, $y = 2/3$.

For any fixed $\Delta > 0$, it is straightforward to use Taylor expansion to calculate that, as n becomes large,

$$\frac{1+d(\Delta)}{n\alpha-1} = \frac{1}{1-e^{-r\Delta}}(n\alpha-1)^{-2} + O((n\alpha-1)^{-4}).$$

Therefore, $(1+d(\Delta))/(n\alpha-1)$ and hence $X_1(\Delta)/n$ are of order $n^{-4/3}$.

But if we first take the limit $\Delta \to 0$, we clearly have

$$\lim_{\Delta\to 0} \frac{1+d(\Delta)}{n\alpha-1} = \frac{1}{n\alpha-1},$$

so $\lim_{\Delta\to 0}(1+d(\Delta))/(n\alpha-1)$ and hence $\lim_{\Delta\to 0} X_1(\Delta)/n$ are of order $n^{-2/3}$.

If $\sigma_D^2 = 0$, then $n\alpha = n$. The same calculation as above shows that $X_1(\Delta)/n$ is of order n^{-2} for a fixed $\Delta > 0$ but is of order n^{-1} if we first take the limit $\Delta \to 0$.

References

Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek, and M. Weretka (2014): "Demand Reduction, Inefficiency and Revenues in Multi-Unit Auctions," *Review of Economic Studies*, 81, 1366–1400.

Biais, B., T. Foucault, and S. Moinas (2015): "Equilibrium Fast Trading," *Journal of Financial Economics*, 116, 292–313.

Brogaard, J., T. Hendershott, and R. Riordan (2014): "High-Frequency Trading and Price Discovery," *Review of Financial Studies*, 27, 2267–2306.

Budish, E., P. Cramton, and J. Shim (2015): "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 130, 1547–1621.

Cespa, G. and X. Vives (2016): "High Frequency Trading and Fragility," Working paper.

Foster, F. D. and S. Viswanathan (1996): "Strategic Trading When Agents Forecast the Forecasts of Others," *Journal of Finance*, 51, 1437–1478.

Fudenberg, D. and J. Tirole (1991): *Game Theory*, MIT Press.

He, H. and J. Wang (1995): "Differential informational and dynamic behavior of stock trading volume," *Review of Financial Studies*, 8, 919–972.

Hendershott, T., D. Livdan, and N. Schürhoff (2015): "Are Institutions Informed about News?" *Journal of Financial Economics*, 117, 249–287.

Hoffmann, P. (2014): "A dynamic limit order market with fast and slow traders," *Journal of Financial Economics*, 113, 156–169.

Jovanovic, B. and A. J. Menkveld (2012): "Middlemen in Limit-Order Markets," Working paper.

Kyle, A. S., A. A. Obizhaeva, and Y. Wang (2014): "Smooth Trading with Overconfidence and Market Power," Working paper.

Menkveld, A. J., B. Z. Yueshen, and H. Zhu (2016): "Shades of Darkness: A Pecking Order of Trading Venues," *Forthcoming, Journal of Financial Economics.*

Pagnotta, E. and T. Philippon (2013): "Competing on Speed," Working paper.

Rostek, M. and M. Weretka (2012): "Price Inference in Small Markets," *Econometrica*, 80, 687–711.

——— (2015): "Dynamic Thin Markets," *Review of Financial Studies*, 28, 2946–2992.

Sannikov, Y. and A. Skrzypacz (2014): "Dynamic Trading: Price Inertia, Front-Running and Relationship Banking," Working paper.

Securities and Exchange Commission (2010): "Concept Release on Equity Market Structure; Proposed Rule," *Federal Register*, 75, 3593–3614.

Vayanos, D. (1999): "Strategic Trading and Welfare in a Dynamic Market," *The Review of Economic Studies*, 66, 219–254.

Vives, X. (2011): "Strategic Supply Function Competition with Private Information," *Econometrica*, 79, 1919–1966.

Wang, J. (1994): "A Model of Competitive Stock Trading Volume," *Journal of Political Economy*, 102, 127–168.

Yang, L. and H. Zhu (2016): "Back-Running: Seeking and Hiding Fundamental Information in Order Flows," Working paper.

Zhu, H. (2014): "Do Dark Pools Harm Price Discovery?" *Review of Financial Studies*, 27, 747–789.