

Belief-free Price Formation*

Johannes Hörner[†], Stefano Lovo[‡], Tristan Tomala[§]

September 21, 2016

Abstract

We analyze security price formation in a dynamic setting in which long-lived dealers repeatedly compete for the opportunity to trade with short-lived retail traders. We characterize equilibria in which dealers' pricing strategies are optimal irrespective of the private information that each dealer may possess. Thus, our model's predictions are robust to different specifications of the dealers' information structure. These equilibria reconcile, in a unified and parsimonious framework, price dynamics that are reminiscent of well-known stylized facts: excess price volatility, price/trading-flow correlation, stochastic volatility and inventory-related trading.

Keywords: Financial Market Microstructure, Informed Dealers, Price Volatility, Belief-free Equilibria.

JEL Codes: G1, G12, C72, C73.

*We thank Vincent Fardeau, Thierry Foucault and Ioanid Rosu for their useful comments, as well as the seminar audiences at HEC Paris, EDHEC, Warwick Business School, the SEAT 2012 meeting in Brisbane and the EFA 2012 Meeting in Copenhagen, the Roy-Adres Seminar. S. Lovo and T. Tomala gratefully acknowledge financial support from the HEC Foundation. J. Hörner gratefully acknowledges financial support from NSF Grant SES 092098. S. Lovo acknowledges financial support from ANR/Investissements d'Avenir under grant ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047.

[†]Yale University, 30 Hillhouse Ave., New Haven, CT06520, USA. johannes.horner@yale.edu.

[‡]HEC, Paris and GREGHEC, 78351 Jouy-en-Josas, France. lovo@hec.fr.

[§]HEC, Paris and GREGHEC, 78351 Jouy-en-Josas, France. tomala@hec.fr.

Introduction

In this paper, we consider a class of market microstructure models in which some long-lived market participants (“dealers”) repeatedly trade a risky asset with short-lived market participants (“traders”). We characterize price formation equilibria that are robust to all possible specifications of information asymmetries across dealers, irrespective of how simple or complex these asymmetries might be. We show that these equilibria can explain, in a single parsimonious model, a number of well-known stylized facts concerning stock price dynamics.

Since Glosten and Milgrom (1985) and Kyle (1985), one of the most common assumptions in the financial market microstructure literature is that the equilibrium price of an asset results from Bertrand competition among dealers, who share identical information about the asset fundamentals and the motivation that leads a trader to supply or demand the asset. While the assumption of equally uninformed dealers leads to tractable models, little is known about the equilibrium price dynamics if this assumption is dropped. This is an important question for two reasons. First, because actual dealers have access to different sources of information, they are unlikely to possess identical information (Ellis, Michaely, and O’Hara (2002)). Second, because information asymmetries are not directly observable, it is impossible to assess the extent to which a model’s informational assumptions capture real-life situations.

In this paper, we take an agnostic approach concerning the dealers’ information and knowledge. Instead of specifying an ad hoc information structure and then solving for an equilibrium, we characterize dealers’ strategies that form an equilibrium *irrespective of* each dealer’s information and beliefs about the fundamentals of the economy. That is, we make virtually no assumption about whether and how dealers are asymmetrically informed. We will refer to these equilibria as ‘belief-free equilibria’.¹

¹See Hörner and Lovo (2009), Fudenberg and Yamamoto (2010) and Hörner, Lovo and Tomala (2010) for the general definition and analysis of belief-free equilibria in repeated games of incomplete information.

Our approach is particularly suited to certain markets, such as the bond market. At any point in time, each dealer might have some private information about the motivations of his clients. This translates into private information concerning the bond's fundamental value. Thus, when setting his public quotes, each dealer takes his own private information into account. Because different dealers have access to different pieces of information, observing his competitors' quotes further affects a dealer's beliefs about the bond's value. In a dynamic setting, in which the same dealers repeatedly meet over time, a dealer's quotes affect not only the dealer's current profit but also his competitors' beliefs. Thus, these quotes alter his competitors' future behavior and hence, the dealer's future profits. Assuming that dealers are uninformed does not fit such markets. The standard approach requires specifying dealers' information structure and solving for a Bayesian equilibrium given this structure. Unfortunately, except for very simple information structures, this approach leads to intractable models.² Moreover, nothing guarantees that the equilibrium strategy computed for a hypothetical information structure remains an equilibrium if the actual information structure is different. By contrast, in a belief-free equilibrium, each dealer's strategy is *ex post* optimal, *i.e.*, independent of the actual fundamental value of the bond. *Ex post* optimality guarantees that, irrespective of the dealers' information structure and their beliefs, the strategy profile remains optimal and forms a subgame perfect Nash equilibrium of the underlying complete information game. Thus, a belief-free equilibrium (henceforth, BFE) remains an equilibrium even in those markets in which private information differs across dealers and varies over time. In addition, a BFE also remains a subgame perfect Nash equilibrium under the informational assumption that dealers themselves are agnostic with respect to other dealers' beliefs about fundamentals. That is, as we illustrate below, in a BFE, to find the optimal quoting strategy, dealers merely need to follow a simple Markovian mapping from the public history of quotes and trades. Because this mapping is independent of the true asset value, in such equilibria,

²The most common ways to recover tractability are to either assume that dealers have no private information or that whatever private information a dealer has at a given time becomes obsolete at all later times.

knowing what other dealers believe is neither necessary nor useful to maximize a dealer's profit.

In the first part of the paper, we illustrate the mechanism behind BFE within the framework of a Glosten and Milgrom economy, in which a finite number of long-lived risk-neutral dealers make a market for a population of short-lived mean-variance traders. The fundamental value of the asset is $\tilde{W} = \tilde{v} + \tilde{e}$, where $\tilde{v} \in \{v_1, v_2\}$, and \tilde{e} has zero mean and positive variance. Traders know the realization of \tilde{v} but not the realization of \tilde{e} . Compared to the standard Glosten and Milgrom economy, we remain agnostic with respect to each dealer's private information about \tilde{W} . The purpose is to focus on BFE that are as tractable as the standard zero-profit equilibrium. First, we characterize four necessary conditions that a price formation strategy must satisfy to form a BFE.

1) Because a BFE must be an equilibrium even when dealers are uninformed, dealers' strategies must be measurable with respect to the public information, namely, the information that results from the observation of the traders' order flow (Lemma 3).

2) Each dealer can always guarantee zero profit by abstaining from trade. Thus, in every BFE and irrespective of the asset true value, each dealer makes strictly positive profit over time. That is, dealers might lose money in the short run, but their average long-run profit must be strictly positive, independently of the asset value \tilde{W} (Lemma 4).

3) Because no large inventory comes at zero cost, a large inventory necessarily translates into a negative profit, under certain beliefs regarding asset fundamentals. Thus, to guarantee that, irrespective of a dealer's belief about \tilde{W} , the dealer expects to make a positive profit; in every BFE, dealers maintain balanced inventories even if they are risk-neutral. In other words, dealers' inventories are mean-reverting. The smaller the average trading volume and the larger the residual uncertainty concerning the asset value, the tighter are the bounds on dealers' inventories (Lemma 5).

4) In a BFE, and even if market participants are truly Bayesian, it cannot be the case that equilibrium quotes reflect Bayesian beliefs about fundamentals. Namely, price sensitivity to

trading volume does not fade as public information accumulates. Thus, long-term price volatility remains large even without exogenous shocks to fundamentals (Lemma 7).

5) Although price volatility does not decline over time, most trading occurs at quotes that are close to \tilde{v} , the expected asset value after aggregating traders' private information (Lemma 6). This, together with necessary condition (4), implies that quotes recurrently diverge from the fundamental value of the asset.

Second, we show that if dealers are patient, the following Markov strategy for dealers forms a BFE. In essence, the strategy comprises two components, which correspond to exploring and exploiting phases. Exploring phases are periods in which dealers' quotes are set to probe traders' demand to learn \tilde{v} , that is, what traders know about the true asset value. Exploiting phases are periods of time in which dealers exploit the information from the exploring phases and make profits through the intermediation of traders' demand and supply. The Markov variable governing the dealers' quotes and the phase transition process is what we call the market measure (a probability measure over the possible values of \tilde{v}). The market measure is updated based on the observation of how traders react to dealers' quotes. We say that a market measure points at a given $v \in \{v_1, v_2\}$ when it attaches a sufficiently large probability to v . Whenever the market measure points at a given v , an exploiting phase starts, and dealers' bid and ask quotes straddle v , such that, if actually $\tilde{v} = v$, then they profitably intermediate traders' demand while maintaining balanced inventories. When the market measure is not concentrated on any single $v \in \{v_1, v_2\}$, dealers set quotes to prompt informative order flows from traders. When this occurs, we say that dealers are in an exploring phase. Crucially, how the market measure is updated does not reflect Bayesian updating. Instead, the updating rule is such that, irrespective of the past trading history, the market measure indicates the true \tilde{v} relatively quickly and with high probability.

Why is this quoting strategy a BFE if dealers are sufficiently patient? First, irrespective of a dealer's belief about the true \tilde{W} and irrespective of the past history, how the market measure is

updated leads each dealer to expect equilibrium quotes to move relatively rapidly toward those that prevail during the exploiting phase, corresponding to the true state. Therefore, each dealer expects future profits to be strictly positive irrespective of his current beliefs and the past history. To prevent dealers from deviating, we specify continuation strategies after a deviation such that the deviating dealer is punished while other dealers are rewarded.³

In the second part of the paper, we generalize the trading protocol and the nature of uncertainty. Our general model embeds different trading mechanisms, including quote-driven markets and limit-order markets. Furthermore, it embeds situations in which dealers face uncertainty over both the fundamental value of the asset and the composition of the population of traders. We show that similar necessary and sufficient conditions apply for a strategy profile to form a BFE in this more general model.

Applications and empirical implications

Under what market structures are these equilibria plausible? Very little is required for a market microstructure economy to admit at least one BFE. First, there must be room for trade across traders to avoid the no-trade theorem, a condition present in all market microstructure models. Second, the discount rate between two trading rounds must be small, which is a condition naturally satisfied by market microstructure economies in which, depending on the type of asset, the trading frequency is between days and milliseconds. Third, there must be a finite number of long-lived dealers.

Thus, from a theoretical perspective, virtually all market microstructure economies that have long-lived and short-lived market participants admit a BFE. However, in these equilibria, dealers share the surplus coming from intermediation of traders' demand. Hence, in practice, these equilibria are more plausible in those markets in which a small number of dealers intermediate the

³Because dealers make positive profits and deviations are deterred by punishment phases, one can regard a BFE as a collusive equilibrium.

largest fraction of trades. Whereas this was formerly the case in the NASDAQ (Ellis, Michaely, and O’Hara (2002)), which indeed displayed anti-competitive practices (Christie, Harris, and Schultz (1994), Christie and Schultz (1994), Weston (2000)), the structure of this market has changed over time as dealers faced more direct competition from traders’ limit orders. By contrast, in markets for corporate bonds, municipal bonds, off-the-run Treasuries, and credit default swaps, a small number of dealers make 90% of the market.⁴ ⁵ Atkeson, Eisfeldt and Weill (2013) theoretically demonstrate how in OTC markets few banks endogenously emerge as dealers and trade mainly to provide intermediation services.

What are the empirical implications of belief-free equilibria? Our model is parsimonious (especially the general model of Section 5) yet supports a variety of stylized facts. Although some of these facts can be explained by existing models, we are not aware of a model that would deliver them all at once.

First, as in the standard zero-profit equilibrium, public news and the trading volume are the main drivers of price changes. The relationship between trading volume and prices has been extensively documented in several markets. See, for instance, Chordia, Roll and Subrahmanyam (2002) and Boehmer and Wu (2008) for the stock market, Pasquariello and Vega (2005) for the bond market, Evans and Lyons (2002) for the currency market, and Fleming, Kirby and Ostdiek (2006) for weather-sensitive commodity markets.

Second, whereas canonical market microstructure models predict that, absent exogenous shocks, equilibrium quotes converge to fundamentals and quote volatility vanishes, quote volatility cannot vanish in a BFE. Even after an arbitrarily long trading history, quotes will recurrently drift away from the asset value, although they will be close to it in expectation. This generates an endogenous pattern of alternating regimes. The economy recurrently switches back and

⁴For example, according to International Swaps and Derivatives Association (ISDA) surveys, in 2010, the fourteen largest global derivatives dealers accounted for 82% of the total combined notional amount outstanding of interest rate, credit, and equity derivatives. In 2013, the fifteen largest dealers accounted for approximately 97%.

⁵We thank an anonymous referee for highlighting this.

forth between a high-volatility, high-mispricing regime (exploring phases) and a low-volatility, low-mispricing regime (exploiting phases). This prediction is consistent with the phenomenon of excess and stochastic regime-shift in price volatility. The first papers providing empirical evidence that stock price volatility does not coincide with commensurate volatility in corporations' fundamentals are Shiller (1981) and LeRoy and Porter (1981). In the asset pricing literature, results reveal that introducing stochastic price volatility in the underlying price process (see for example Heston (1993)) or regime shifts in volatilities (Calvet and Fisher (2004)) helps explain actual derivative prices.

In addition, the BFE provides a new testable implication relating the switch in volatility regime to the trading volume. A switch from the low to the high price volatility regime is preceded by periods of increasing order imbalance. The switch from high to low volatility should be preceded by a stabilization of dealers' inventories and small order imbalance.

Third, in contrast to the canonical market microstructure models that imply zero profits, BFE pricing implies that dealers make strictly positive profits. This is consistent with the evidence provided for the NASDAQ (Christie, Harris, and Schultz (1994), Christie and Schultz (1994), Huang and Stoll (1996), Kandel and Marx (1997), Weston (2000), Ellis, Michaely, and O'Hara (2002)), the corporate bond market (Bessembinder and Maxwell (2008), Goldstein and Hotchkiss (2007)) and the municipal bond market (Harris and Piwowar (2006), Green, Hollifield and Schüroff (2007), Green, Li and Schüroff (2010), Li and Schüroff (2014)).

Fourth, in a BFE, dealers' inventories are mean-reverting. This is true even even when dealers are risk-neutral and not subject to institutional constraints on inventory size. Evidence of mean-reverting inventories is provided by Madhavan and Smidt (1993), Hasbrouck and Sofianos (1993), Hansch, Naik, and Viswanathan (1999), Reiss and Werner (1998), Naik and Yadav (2003), and more recently, by Hendershott and Menkveld (2014).

Fifth, a BFE implies a positive relationship between a dealer's average inventory and trading volume. That is, the larger the trading volume that a dealer is expected to intermediate, the

larger the inventory he can maintain. This is consistent with the findings of Li and Schüroff (2014) for the municipal bond market, demonstrating that peripheral dealers intermediate few deals and hold small inventories, whereas central dealers account for a larger intermediation volume and hold larger inventories. A BFE also implies a positive relationship between the average profitability of a deal and the dealer's maximum inventory size. This is in line with the findings of Bessembinder and Maxwell (2008) and Goldstein and Hotchkiss (2007), who show that dealers' average inventory decreased following the reduction in dealers' profits caused by an increase in transparency in the corporate bond market.

Related Literature

The applications of repeated games to market microstructure that are closest to our work are Dutta and Madhavan (1997), Benveniste et al. (1992), Desgranges and Foucault (2005) and Carlin, Lobo and Viswanathan (2007). These papers assume either no information asymmetry or short-lived information asymmetries. They construct trigger-strategy equilibria in which a monopolistic optimum is sustained by the threat of reverting to the static Nash equilibrium. Episodically, the short-run stakes can be so large that the threat is not effective, and thus, a punishment phase emerges on the equilibrium path. These equilibria are not belief-free because both the monopolistic optimum and the static Nash equilibrium depend on market participants' beliefs. Clearly, these would not be equilibria in the presence of long-lived asymmetric information. In this paper, instead, the state of nature is perfectly persistent, and thus, asymmetric information is long-lived.⁶ Nevertheless, the idea of sustaining positive profits through the threat of punishment phases is common to these papers and our BFE. Compared to the standard trigger-strategy equilibrium, a BFE differs through the on-path alternation of exploring and exploiting phases. One of the implications of our model is that, if one recognizes that non-trivial information

⁶The same results hold if the frequency of trading is high compared to the frequency with which the state of nature changes.

asymmetries across dealers matter, then equilibria in which dealers' profits are positive are the only equilibria that are independent of the dealers' information. In these situations, zero-profit competitive equilibria should be regarded as the exception rather than the norm.

Few theoretical papers analyze the effect of asymmetric information among dealers. Even fewer do so within a dynamic framework. Some static models in which dealers or, more generally, liquidity providers are asymmetrically informed are Roëll (1988), Bloomfield and O'Hara (2000), de Frutos and Manzano (2005) and Boulatov and George (2013). Within a dynamic framework, Moussa Saley and De Meyer (2003) and Calcagno and Lovo (2006) study the case of one better-informed price maker. De Meyer (2010) considers the case of two-sided incomplete information. However, their constructions are not belief-free, to the extent that they rely on the specific assumptions made on the dealers' information structure. The results obtained by Du and Zhu (2012) are closer in spirit to our work. Within the framework of a double auction, they show that for a specific additive functional form of bidders' values, the static auction has an *ex post* equilibrium and that this property extends to the repeated auction, giving rise to an "almost" belief-free equilibrium.

This paper is organized as follows. Section 1 introduces the baseline model. Section 2 describes the benchmark zero-profit equilibrium for the baseline model. Section 3 presents the BFE that is defined in 3.1. Section 3.2 presents the necessary condition that a BFE must satisfy, and Section 3.3 presents a simple "Markov" BFE. Section 4 presents simulations of the BFE and its empirical implications. Section 5 presents the general model and studies its BFE. Section 6 discusses extensions to imperfect monitoring of dealers' actions, non-stationary states of nature and dealers' strategies based on private information. Section 7 concludes. Appendix A contains the proofs for the baseline model, and Appendix B contains the proofs for the general model.

1 A Model of Price Formation

In this section, we illustrate the definition, the logic and the main features of a BFE in a simple, well-known financial market microstructure framework à la Glosten and Milgrom (1985). In Section 5, we generalize the model and illustrate how the equilibrium construction must be amended to account for some institutional features of real markets that are not captured by this simple model.

A risky asset is exchanged for money over an infinite number of periods $t = 1, 2, \dots$. At time 0 and once and for all, Nature randomly chooses the state ω from some finite set Ω . We denote by $W(\omega)$ the fundamental value of the asset in state ω . Assume that

$$W(\omega) = v(\omega) + e(\omega).$$

That is, the asset's fundamental value is a random variable $\tilde{W} = \tilde{v} + \tilde{e}$, where \tilde{v}, \tilde{e} are two independent random variables. We assume that $\tilde{v} \in \{v_1, v_2\}$ and that \tilde{e} takes a value in the finite set $\mathcal{E} \subset [-\bar{e}, \bar{e}]$, *i.e.*, $\bar{e} > 0$ and $-\bar{e}$ are the maximum and the minimum possible values that \tilde{e} can take. The random variable \tilde{e} has zero mean and strictly positive variance σ^2 (given mean-variance preferences, nothing more is needed). As in Back and Baruch (2004), a public release of information occurs at a random time τ , and conditional on this not yet having occurred, the probability that it occurs in the next period is constant over time. After the public announcement, all positions are liquidated at price $W(\omega)$.

There are two groups of market participants. The first consists of professional financial intermediaries, such as dealers, market makers and brokers. There are finitely many of them, and they consistently monitor and participate in the market. We call this group of market participants *dealers* and model them as n infinitely-lived risk-neutral agents. The second group consists of households, fund managers and institutional investors that occasionally participate

in the market to rebalance their portfolios and/or to exploit some private information that they possess about asset fundamentals. We call this group *traders* and model them as an infinite sequence of short-lived, risk-averse agents.

In every period, trade unfolds as follows. First, dealers simultaneously set their bid and ask quotes, at which each dealer stands ready to buy and sell one unit of the asset. Then, a randomly selected trader comes to the market, observes the dealers' quotes and chooses whether to buy one unit, sell one unit, or not to trade. The trader's market order is executed against the best quotes from dealers, and then the trader leaves the market. We denote by α_i^t and β_i^t the ask and bid quotes set by dealer i in period t . Quotes belong to an interval $G := [0, M]$, where M is finite and can be arbitrarily large relative to $v_2 + \bar{e}$. Moreover, we denote by $a^t := \{\alpha_i^t, \beta_i^t\}_{i=1, \dots, n}$ the set of dealers' quotes at time t . Thus, $A := \times_{i=1, \dots, n} G^2$ is the compact set of all possible profiles of bid and ask quotes, and ΔA denotes the set of probability distributions over the set A . For any given a^t in A , we denote by $\alpha(a^t) := \min_i \alpha_i^t$ and $\beta(a^t) := \max_i \beta_i^t$ the best ask and bid quotes in period t . Let $s \in S := \{-1, 0, 1\}$ be a trader's market order. We adopt the convention that $s = 1$ corresponds to the trader selling.

Traders are short-lived mean-variance investors with utility $E[\tilde{x}] - \frac{\gamma}{2} \text{Var}[\tilde{x}]$, where \tilde{x} is a trader's post-trade wealth. Traders only differ in their initial inventory y . We denote by $Z(x)$ the probability that the time- t trader's inventory is less than x . We assume that y is uniformly distributed on the interval $[-\frac{\phi}{2}, \frac{\phi}{2}]$, with $\phi > 1$.

Suppose time- t trader's inventory is y . Then, his wealth resulting from a market order $s \in S$, given dealers' quotes $a \in A$, is $\tilde{x}(a, s, y) = (y - s)\tilde{W} + s(\beta(a)1_{\{s=1\}} + \alpha(a)1_{\{s=-1\}})$, where $1_{\{\cdot\}}$ is the indicator function. Thus, the trader chooses the order s that solves

$$\max_{s \in S} E^t[\tilde{x}(a, s, y)] - \frac{\gamma}{2} \text{Var}^t[\tilde{x}(a, s, y)], \quad (1)$$

where E^t and Var^t stand for the expectation and variance, computed using the trader's beliefs

about \tilde{W} .

Because dealers are long lived, we distinguish between the profit or loss made in a single trading round, which we call the “reward,” and the discounted sum of the sequence of current and future rewards, which we call the “payoff.” Suppose that at time t , the dealers’ action is $a \in A$ and the trader’s order is $s \in S$. Then dealer i buys (sells) the asset if his bid is the highest (resp., ask is the lowest) and the trader sells (buys). In state ω , this translates into a reward for dealer i equal to

$$u_i(\omega, a, s) = (v(\omega) + e(\omega) - \beta_i)1_{\{\beta_i = \beta(a), s=1\}}\eta_\beta(a) + (\alpha_i - v(\omega) - e(\omega))1_{\{\alpha_i = \alpha(a), s=-1\}}\eta_\alpha(a),$$

where $\eta_\beta > 0$ and $\eta_\alpha > 0$ are exogenous tie-breaking rules applied in the event that more than one dealer sets the best bid or ask, respectively. The rewards of the dealers are discounted at the common factor $\delta < 1$, and the payoff is the average discounted sum of rewards. The discount factor δ accounts for both the dealer’s impatience and the possibility that public information is released in the current period.⁷ Thus, in state ω , an infinite sequence of quotes and orders $h = \{a^t, s^t\}_{t=0}^\infty$ translates into a payoff for dealer i equal to

$$V_i(\omega, h) = \sum_{t=0}^{\infty} (1 - \delta)\delta^t u_i(\omega, a^t, s^t). \quad (2)$$

1.1 Information structure

Traders’ information We assume that traders know the realization of the \tilde{v} component of the asset value \tilde{W} but not the realization of the \tilde{e} component. They commonly share some belief about \tilde{e} . There is no correlation between the distribution of traders’ inventory and \tilde{e} . This implies that information about a trader’s inventory is useless for learning \tilde{e} . Thus, \tilde{e} can be

⁷Allowing for a stochastic discount factor complicates the exposition but does not affect results as long as the expected discount factor is close enough to one.

interpreted as the residual uncertainty over the fundamental asset value after aggregating all possible information from the population of traders, whereas \tilde{v} is what traders collectively know about the fundamentals. More formally, traders' initial information partition over the possible states Ω is $\hat{\Omega} := \{\hat{\omega}_{v_1}, \hat{\omega}_{v_2}\}$, where $\hat{\omega}_v := \{\omega \in \Omega \mid v(\omega) = v\}$.

Dealers' information The main contribution of this paper consists in providing predictions that do not rely on specific assumptions regarding dealers' information about ω . Here, we depart from the canonical approach in the microstructure literature that consists in, first, assuming that dealers are equally uninformed and, second, focusing on the zero-profit Bayesian equilibrium consistent with this assumption.

By contrast, our purpose is to find a price formation strategy that is a Bayesian equilibrium *for all possible specifications of information asymmetries across dealers*, irrespective of whether these asymmetries are present and irrespective of how simple or complex these asymmetries might be. That is, we focus on a strategy profile for dealers that forms a Bayesian equilibrium regardless of what each individual dealer knows about the fundamentals and whether other dealers are informed. In particular, each dealer might have specific private information about the state of nature ω ; he might have information about \tilde{e} , \tilde{v} , or both. Information about one or the other component makes no difference in a BFE.⁸ At any time, he might be uncertain about whether other dealers have received some private information. Nevertheless, each dealer will find it optimal to use the same pricing strategy. This type of equilibria is known as belief-free equilibria, formally defined below. In a BFE, each dealer's strategy must be a best-response, regardless of the true state ω . Hence, it is optimal irrespective of each dealer's belief about ω .

⁸This is not necessarily the case in a "standard" Bayesian equilibrium, in which best-replies depend on beliefs and hence on each dealer's specific information about \tilde{v} and \tilde{e} .

2 The canonical equilibrium

As a benchmark, it is useful to first describe the “textbook” canonical equilibrium that is obtained by, first, assuming that dealers start from a common prior belief about the fundamental value of the asset and that no dealer ever receives any private information and, second, focusing on a Bayesian equilibrium where, in each period, each dealer’s expected reward is nil because of Bertrand competition.

Let us first study traders’ behavior. Recall that traders know the \tilde{v} component of \tilde{W} but not the \tilde{e} component. Thus, if dealers have no private information, then no market participant has any information about \tilde{e} . Hence, if the state is ω , a time- t trader’s inventory is y , and dealers’ quote profile is a , then the trader’s expected utility from order s is equal to

$$E^t[\tilde{x}(a, s, y)] - \frac{\gamma}{2} \text{Var}^t[\tilde{x}(a, s, y)] = (y - s)v(\omega) + s(\beta(a)1_{\{s=1\}} + \alpha(a)1_{\{s=-1\}}) - \frac{\gamma}{2}(y - s)^2\sigma^2.$$

Note that this expression does not depend on time or on the value of $e(\omega)$. Denote by $F(\omega, a, s)$ the probability that a time- t trader’s order is s , conditional on the state being ω and time- t dealers’ quotes being $a \in A$. Then, we have

Lemma 1 *If dealers have no private information and their quote profile is a , then in state of nature ω , at any time t , the distribution of a time- t trader’s order is*

$$F(\omega, a, 1) = \begin{cases} 0 & \text{for } \beta(a) < v(\omega) - \rho \\ \frac{\rho + \beta(a) - v(\omega)}{\phi\gamma\sigma^2} & \text{for } v(\omega) - \rho \leq \beta(a) \leq v(\omega) + \rho\frac{\phi+1}{\phi-1} \\ 1 & \text{for } \beta(a) > v(\omega) + \rho\frac{\phi+1}{\phi-1}, \end{cases}$$

and

$$F(\omega, a, -1) = \begin{cases} 1 & \text{for } \alpha(a) < v(\omega) - \rho \frac{\phi+1}{\phi-1} \\ \frac{\rho + v(\omega) - \alpha(a)}{\phi \gamma \sigma^2} & \text{for } v(\omega) - \rho \frac{\phi+1}{\phi-1} \leq \alpha(a) \leq v(\omega) + \rho \\ 0 & \text{for } \alpha(a) > v(\omega) + \rho, \end{cases}$$

as well as

$$F(\omega, a, 0) = \max\{1 - F(\omega, a, 1) - F(\omega, a, -1), 0\},$$

where $\rho := (\phi - 1)\gamma\sigma^2/2 > 0$.

That is, the probability of observing a buy (sell) order decreases in the ask price (resp., increases in the bid price) and increases (resp., decreases) in $v(\omega)$. The probability of a buy order is strictly positive, as long as the best ask price does not exceed $v(\omega)$ by more than ρ . Similarly, a sell order occurs with positive probability if the best bid is larger than $v(\omega) - \rho$. In other words, the greater the traders' risk aversion γ , their dispersion of their inventories ϕ , and the uncertainty over \tilde{e} , σ^2 , the greater the potential for trade between traders and dealers.

Because $F(\omega, \cdot, \cdot)$ depends on the realization of \tilde{v} but not on the realized \tilde{e} , we can focus on the function $F(v, a, s)$ that is equal to $F(\omega, a, s)$ when ω is such that $v(\omega) = v$. That is, at any time t , traders' behavior only depends on $v(\omega)$ and dealers' current quotes. Namely, if the state is ω and time- t quotes are a , then, the probability that dealer i buys in period t is

$$Q_i^+(v(\omega), a) := F(v(\omega), a, 1)1_{\{\beta_i = \beta(a)\}}\eta_\beta(a); \quad (3)$$

the probability that dealer i sells in period t is

$$Q_i^-(v(\omega), a) := F(v(\omega), a, -1)1_{\{\alpha_i = \alpha(a)\}}\eta_\alpha(a); \text{ and} \quad (4)$$

the expected net cash flow for dealer i in t is

$$P_i(v(\omega), a) := -\beta_i Q_i^+(v(\omega), a) + \alpha_i Q_i^-(v(\omega), a) \leq v(\omega) + \rho, \quad (5)$$

where the inequality follows from Lemma 1, namely the fact that in state ω , no trader ever pays more than $v(\omega) + \rho$ for the asset. Thus, we can re-write dealer i 's reward in period t by taking expectations with respect to the traders' order. Dealer i 's expected reward in state ω if the current dealers' quote is a is equal to

$$\begin{aligned} u_i(\omega, a) &:= \sum_{s \in S} u_i(\omega, a, s) F(v(\omega), a, s) \\ &= W(\omega)(Q_i^+(v(\omega), a) - Q_i^-(v(\omega), a)) + P_i(v(\omega), a). \end{aligned} \quad (6)$$

That is, dealer i 's time- t reward in state ω from quotes a is equal to the fundamental value of the asset, times the expected change in the dealer's inventory at time t , plus the expected net cash flow at time t .

The canonical equilibrium relies on the assumption that all dealers are uninformed about \tilde{W} . Let h^t denote the public history of past trades and quotes up to time t . Let H^t denote the set of all public histories of length t . Given a public history $h^t \in H^t$, let p^t denote the dealers' common belief that $\tilde{v} = v_2$ given h^t . Then, p^t evolves according to Bayes' rule:

$$p^{t+1} = \psi_B(p^t, a^t, s^t) := \frac{p^t F(v_2, a^t, s^t)}{p^t F(v_2, a^t, s^t) + (1 - p^t) F(v_1, a^t, s^t)}. \quad (7)$$

Given traders' beliefs at time t , in the canonical equilibrium, each dealer's expected profit at time t is nil. That is, at any time t , bid and ask quotes are $\beta^t = E[\tilde{W} \mid h^t, s^t = 1]$ and $\alpha^t = E[\tilde{W} \mid h^t, s^t = -1]$, respectively. To avoid the trivial case in which all traders' information is disclosed at the first order, we assume that $\rho > v_2 - v_1$.⁹ Then, we have

⁹If $\rho \leq v_2 - v_1$, then one would have $\beta^1 = v_1$ and $\alpha^1 = v_2$, and $\beta^t = \alpha^t = v(\omega)$ in all periods following the

Lemma 2 Canonical Equilibrium: Assume that $\rho > v_2 - v_1$, and that dealers are equally uninformed. Then, there is a Bayesian equilibrium in which in any period t , the following hold:

1. Bid and ask quotes are

$$\alpha^t = \alpha(p^t) := E[\tilde{v} \mid h^t] + \frac{\rho}{2} - \sqrt{\frac{\rho^2}{4} - \text{Var}[\tilde{v} \mid h^t]}, \quad (8)$$

$$\beta^t = \beta(p^t) := E[\tilde{v} \mid h^t] - \frac{\rho}{2} + \sqrt{\frac{\rho^2}{4} - \text{Var}[\tilde{v} \mid h^t]}, \quad (9)$$

where $E[\tilde{v} \mid h^t] = p^t v_2 + (1 - p^t) v_1$ and $\text{Var}[\tilde{v} \mid h^t] = p^t(1 - p^t)(v_2 - v_1)^2$.

2. Dealers' common beliefs about \tilde{v} evolve according to $p^{t+1} = \psi_B(p^t, a^t, s^t)$.

3. Each dealer's expected reward computed under belief p^t is nil.

This equilibrium has a simple “Markovian” structure: first, in every period t , the best bid and ask quotes depend solely on the dealers' common belief p^t ; second, the dealers' common posterior belief p^{t+1} depends only on the common time- t prior p^t and on (a^t, s^t) , the dealers' quotes and the trader's order at time t . However, for almost any other dealer information, these strategies would not form an equilibrium. Specifically, suppose that dealer i has some private information $h_i^t \neq h^t$ such that $E[\tilde{W} \mid h_i^t] \neq E[\tilde{W} \mid h^t]$, then this dealer can earn a strictly positive profit by either setting an ask strictly lower than $\alpha(p^t)$ or a bid strictly larger than $\beta(p^t)$, a profitable deviation from the canonical equilibrium. By contrast, strategies that form a BFE would remain an equilibrium regardless of h_i^t .

first trade.

3 Belief-free equilibrium

3.1 Equilibrium Concept

Let us drop the assumption that dealers are equally uninformed. We are interested in the belief-free equilibria of the repeated game. These are Bayesian equilibria that are subgame perfect Nash equilibria for any possible underlying state. Thus, belief-free equilibria do not rely on the specification of each dealer's information about the true state ω .

Let $H := \cup_t H^t$ denote the set of all public histories of any length. A public strategy profile (strategy, henceforth) is a mapping $\sigma : H \rightarrow \times_i \Delta A_i$ that maps each public history h^t (history, henceforth) onto the (possibly mixed) action profile that dealers play at time t . Traders' behavior can be represented as a mapping $F : \hat{\Omega} \times H \times A \rightarrow \Delta S$ specifying the probability of each market order given \tilde{v} , the history and the current dealers' quote profile, where F is such that equation (1) is satisfied at all t . For any given state ω and any history h^t , a strategy σ induces a probability distribution over future histories in the standard fashion and, hence, an occupation measure over action profiles, which we denote $\nu_{(\omega, \sigma, h^t)} \in \Delta(A \times S)$. Formally, the occupation measure $\nu_{(\omega, \sigma, h^t)}(a, s)$ is the discounted expected frequency with which the action profile (a, s) is played after history h^t , if the dealers' strategy is σ and the state is ω .¹⁰

$$\nu_{(\omega, \sigma, h^t)}(a, s) := E_{\sigma} \left[\sum_{\tau \geq t} (1 - \delta) \delta^{t-\tau} 1_{\{(a^{\tau}, s^{\tau}) = (a, s)\}} \middle| \omega, h^t \right].$$

Let $V_i(\omega, \sigma \mid h^t)$ denote dealer i 's expected continuation payoff after history h^t , given ω and σ .

We have

$$V_i(\omega, \sigma \mid h^t) = \sum_{(a, s) \in A \times S} \nu_{(\omega, \sigma, h^t)}(a, s) u_i(\omega, a, s).$$

¹⁰It is worth stressing that the occupation measure is not just the simple “long-run” frequency with which actions are played but the discounted frequency. That is, *ceteris paribus*, actions that are played earlier have more weight in the occupation measure than actions played later.

We look for a dealers' strategy profile σ such that, at any time t , and after any history h^t , for any dealer i , it is optimal to choose an action according to $\sigma_i(h^t)$, regardless of what the dealer believes about the true state ω . Formally, we have the following:

Definition 1 *A belief-free equilibrium is a strategy profile σ^* such that, for every state ω , σ^* is a subgame perfect Nash equilibrium of the repeated game with rewards $u(\omega, \cdot)$, that is, of the repeated game with complete information in which the state ω is common knowledge among dealers.*

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i} V_i(\omega, \sigma_i, \sigma_{-i}^* \mid h^t), \quad (10)$$

for all players i , all states $\omega \in \Omega$, all periods t and all histories $h^t \in H^t$.

Some remarks are in order. First, a BFE is a perfect Bayesian equilibrium given any initial prior distribution of dealers' belief about ω and any additional private information a dealer might possess.¹¹ Thus, a BFE is a subgame perfect Nash equilibrium irrespective of the specific dealers' information structure. Second, a BFE is an equilibrium regardless of whether dealers are Bayesian.

3.2 Necessary conditions

We now identify the features of dealers' strategies that are necessary for these strategies to form a BFE. This is useful to distinguish the predictions of the market microstructure theories that do not depend on informational uncertainty from those that rely on specific assumptions regarding the dealers' information structure. In other words, a strategy that forms a perfect

¹¹To see this, note that in a perfect Bayesian equilibrium, dealers' strategies satisfy

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i} \mathbb{E} [V_i(\omega, \sigma_i, \sigma_{-i}^* \mid h^t) \mid I_i],$$

where expectations are taken with respect to both the possible states ω and the possible realizations of traders' orders $\{s^t\}_{t=1}^\infty$, and I_i is dealer i 's private information. Hence, a BFE is a perfect Bayesian equilibrium, but a perfect Bayesian equilibrium need not be belief-free.

Bayesian equilibrium for some specification of a dealer's beliefs but does not satisfy at least one of the properties described below would not form an equilibrium for a different dealers' information structure.

Consider a dealers' strategy profile that is a BFE. Then, the same strategy profile must remain a subgame perfect Nash equilibrium for all possible configurations of dealers' information. This includes the following two polar cases: first, the situation in which dealers have no private information about ω and, second, the situation in which there is at least one dealer who is perfectly informed about ω , that is, he knows both \tilde{v} and \tilde{e} .

Consider the first case. If dealers are uninformed, then their quotes disclose no additional information about \tilde{W} . Thus, traders' behavior F must be as described in Lemma 1. That is, as in the canonical equilibrium, in a BFE, the history of traders' orders and the evolution of dealers' quotes depend on the value of \tilde{v} but not on the value of \tilde{e} . More formally,

Lemma 3 (Measurability with respect to traders' behavior) *If σ forms a BFE, then the following hold:*

1. *Traders' behavior F is as described in Lemma 1.*
2. *The equilibrium occupation measure ν is such that, for any history h^t and any pair of states $\omega, \omega' \in \Omega$, it holds that $\nu_{(\omega, \sigma, h^t)} \neq \nu_{(\omega', \sigma, h^t)}$ only if $v(\omega) \neq v(\omega')$. In particular, after any history h^t , the equilibrium occupation measure of future dealers' quotes depends only on h^t and $v(\omega)$.*

In words, Lemma 3 can be interpreted as follows. In a BFE, the following statements hold: 1. Traders behave as if dealers have no private information; 2. because in two states ω and ω' with $v(\omega) = v(\omega')$, traders behave in the same way, then dealers' behavior in state ω is the same as dealers' behavior in state ω' . In other words, the behaviors of both traders and dealers do not depend on the realization of \tilde{e} . We denote by $\tilde{a}(v, h^t) \in \Delta A$ the equilibrium occupation measure

of future dealers' quotes after history h^t in all states in which $v(\omega) = v$. We can then decompose the continuation payoff $V_i(\omega, \sigma \mid h^t)$ in the same way that we decomposed a dealer's reward in equation (6). After history h^t , dealer i 's continuation payoff in state ω from strategy σ can be decomposed as the fundamental value of the asset, times the expected (discounted) change in inventory after history h^t , plus the expected net (discounted) change in cash after history h^t . Formally,

$$\begin{aligned} V_i(\omega, \sigma \mid h^t) &= \sum_{a \in A} \tilde{a}(v(\omega), h^t)(a) u_i(\omega, a) \\ &= W(\omega)(Q_i^+(v(\omega), h^t) - Q_i^-(v(\omega), h^t)) + P_i(v(\omega), h^t), \end{aligned} \quad (11)$$

where we used result 1 in Lemma 3 and equations (3)–(5) to define

$$\begin{aligned} Q_i^+(v(\omega), h^t) &:= \sum_{a \in A} \tilde{a}(v(\omega), h^t)(a) Q_i^+(v(\omega), a) \in (0, 1), \\ Q_i^-(v(\omega), h^t) &:= \sum_{a \in A} \tilde{a}(v(\omega), h^t)(a) Q_i^-(v(\omega), a) \in (0, 1), \\ P_i(v(\omega), h^t) &:= \sum_{a \in A} \tilde{a}(v(\omega), h^t)(a) P_i(v(\omega), a). \end{aligned}$$

The quantity $Q_i^+(v(\omega), h^t) - Q_i^-(v(\omega), h^t)$ is the expected (discounted) change in dealer i 's inventory in state ω after history h^t . $P_i(v(\omega), h^t)$ is the expected (discounted) change in dealer i 's cash holdings in state ω after history h^t . Observe that given Lemma 3, the functions $Q_i^+(\cdot, h^t)$, $Q_i^-(\cdot, h^t)$ and $P_i(\cdot, h^t)$ differ in two states only if \tilde{v} differs for those two states. However, because $W(\omega) = v(\omega) + e(\omega)$, a dealer's continuation payoff $V_i(\omega, \cdot)$ is affected by both $v(\omega)$ and $e(\omega)$.

Now, let us turn to the other polar case, that is, the scenario in which at least one dealer is fully informed of ω and, hence, knows both \tilde{v} and \tilde{e} . Observe that if a strategy profile forms a BFE, then the same strategy profile must remain a subgame perfect Nash equilibrium even in this scenario. Namely, take any state ω and any dealer i ; we cannot exclude the possibility that

dealer i knows that the state is ω . Note that, regardless of the state, dealer i can guarantee a payoff of 0 by setting a large enough bid-ask spread and encounter no volume from traders.¹² This implies that in a BFE, for each state $\omega \in \Omega$, each dealer's payoff cannot be strictly negative. Otherwise, there would be a state ω and a dealer i who, if informed, would deviate to the no-trade action. Thus, in a BFE, each dealer's continuation payoff is non-negative regardless of the fundamental value of the asset and the past history. Recall from Definition 1 that $V_i(\omega, \sigma \mid h^t)$ is the continuation payoff after history h^t . Formally,

Lemma 4 (Positive dealers' payoffs) *If σ forms a BFE, then $V_i(\omega, \sigma \mid h^t) \geq 0$ for all dealers i , all states ω and all histories h^t . Moreover, if the equilibrium leads to a change in inventory for dealer i , then his continuation payoff is nil in at most two states and strictly positive in all other states.*

Note that Lemma 4 is an individual rationality requirement that is stronger than the usual Bayesian requirement. It must hold irrespective of what belief each dealer holds and not just in expectation using some specification of dealers' beliefs. That is, each dealer must make positive profit *in each state of nature and after every history*.

In summary, Lemma 3 implies that for a given value of $\tilde{v} \in \{v_1, v_2\}$, the equilibrium occupation measure of dealers' quotes and traders' orders does not depend on the value of \tilde{e} . Lemma 4 implies that this single occupation measure must lead to non-negative payoffs for all possible values of \tilde{e} . What are the occupation measures guaranteeing that for a given value of \tilde{v} , each dealer's payoff is strictly positive, regardless of the value of \tilde{e} ? Intuitively, because \tilde{e} affects the fundamental value of the asset, but not the traders' behavior, one way for a dealer to make profits irrespective of \tilde{e} is to set bid and ask quotes across \tilde{v} as to intermediate traders' demand and supply, gain the bid-ask spread and take no position in the asset. Having no net position in the asset, a dealer's profit does not depend on \tilde{e} . We formalize and generalize this idea below.

¹²From Results 1 and 2 in Lemma 1 and from Lemma 3, if dealer i sets $\beta_i < v_1 - \rho$ and $\alpha_i > v_2 + \rho$, then he is certain not to trade irrespective of ω .

For any occupation measure of dealers' quote $\tilde{a} \in \Delta A$, let $u_i(\omega, \tilde{a}) := \sum_{a \in A} \tilde{a}(a) u_i(\omega, a)$ be dealer i 's expected continuation payoff resulting from \tilde{a} in state ω . Define $A^*(v)$ and $A_+^*(v)$ as the sets of $\tilde{a} \in \Delta A$ such that, for each dealer i , and all ω such that $v(\omega) = v$, one has $u_i(\omega, \tilde{a}) \geq 0$ and $u_i(\omega, \tilde{a}) > 0$, respectively. That is, $A^*(v)$ and $A_+^*(v)$ are the sets of occupation measures over dealers' quotes that guarantee that, in all states that satisfy $v(\omega) = v$, each dealer's payoff is non-negative or strictly positive. Then, Lemma 3, together with Lemma 4, implies that in state ω , in a BFE, after any history h^t , the equilibrium occupation measure $\tilde{a}(v(\omega), h^t)$ must belong to $A^*(v(\omega))$. Now we can state the implication that BFE has on dealers' inventories.

Lemma 5 (Bounded dealer inventories) *If σ forms a BFE, then each dealer's expected trading volume is relatively balanced and maintains bounded inventories for dealers. Formally,*

$$|Q_i^+(v(\omega), h^t) - Q_i^-(v(\omega), h^t)| \leq (Q_i^+(v(\omega), h^t) + Q_i^-(v(\omega), h^t)) \frac{\rho}{\bar{e}}, \quad (12)$$

for all dealers i , all states ω , and all histories h^t .

The left-hand side of inequality (12) can be interpreted as the inventory, in absolute terms, that dealer i expects to accumulate after history h^t in state ω . The quantity $Q_i^+(v(\omega), \sigma, h^t) + Q_i^-(v(\omega), \sigma, h^t)$ can be interpreted as the volume of traders' orders that dealer i expects to execute. In particular, inequality (12) shows that the larger \bar{e} is, *i.e.*, the greater the potential uncertainty resulting from the \tilde{e} component, and the smaller the trading volume between traders and dealers is, the tighter the bound imposed on dealers' average inventory. The intuition is simple. Each dealer's payoff can be decomposed into two components. The first is what the dealer gains from intermediating traders' demand and supply without taking a net position in the asset. The second is what dealers gain or lose from loading on net positions in the asset at prices that differ from $W(\omega)$. The first component is bounded from above by $(Q_i^+(v(\omega), h^t) + Q_i^-(v(\omega), h^t)) \rho$, that is, the trading volume times the maximum bid-ask spread that is consistent with traders

buying and selling. The second component depends on \tilde{W} . In particular, we cannot exclude the possibility that a dealer believes that the value of \tilde{e} is extreme. Recall that \bar{e} measures the maximum and minimum possible values of \tilde{e} , *i.e.*, $\tilde{e} \in [-\bar{e}, \bar{e}]$. Thus, for such beliefs, the second component can result in a loss as large as $|Q_i^+(v(\omega), h^t) - Q_i^-(v(\omega), h^t)|\bar{e}$. For dealers' continuation payoff to be positive in all states, a dealer's average inventory must be bounded. Because the maximum potential loss on a net position is proportional to \bar{e} , the larger \bar{e} is, the closer to 0 each dealer's inventory is, on average. Moreover, an increase in ρ means that it is easier for dealers to buy from sellers at low prices and to sell at high prices. Thus, a larger ρ is associated with potentially larger profits for dealers. The inequality of Lemma 5 implies that an increase in intermediation profits softens the constraint on bounded inventories.

Another important implication of Lemma 5 is that dealers' profits come from intermediating a substantial, but balanced, flow of orders from traders. From Lemmas 1 and 3, we know that the traders' order flow is affected by \tilde{v} . Hence, to induce an order flow that is both substantial and balanced, the level of dealers' quotes must be appropriately tuned to the value of \tilde{v} . Namely, if the bid and ask quotes straddle \tilde{v} , but the spread is too large, then the traders' order flow is nil. If the bid-ask spread is small but quotes are consistently larger (or smaller) than \tilde{v} , then there is trading volume, but dealers accumulate an inventory that is large in absolute terms, in excess of the bounds imposed by Lemma 5. Thus, on average, the bid and ask quotes should straddle $v(\omega)$ and the inside spread should not be too large. More formally, let $m(\omega, h^t) := \sum_{a \in A} \tilde{a}(v(\omega), h^t)(a) \frac{\alpha(a) + \beta(a)}{2}$ denote the average discounted mid-quote computed using the equilibrium occupation measure following history h^t in state ω . Then, we have

Lemma 6 (*Average information efficiency*) *If σ forms a BFE, then, for all ω and all histories h^t ,*

$$|m(\omega, h^t) - v(\omega)| < \frac{2\rho^2\phi}{(\phi - 1)\bar{e}}.$$

Lemma 6 states that, regardless of the past history and the state ω , the discounted average

of future mid-quotes remains close enough to $v(\omega)$. This can be interpreted as a form of market informational efficiency: the mid-quote does not differ too much, or for too long, from $v(\omega)$, that is, from the fair asset value, conditional on traders' information. How close is the average mid-quote to $v(\omega)$ in a BFE? This depends on the parameters of the model. In particular, \bar{e} plays a crucial role. The larger \bar{e} is, the closer to $v(\omega)$ the bid and ask quotes must be. In fact, the larger \bar{e} is, the tighter the constraint on the magnitude of the inventory imposed by Lemma 5, requiring the mid-quotes to be close enough to $v(\omega)$ to minimize inventory imbalance.¹³

What differentiates this result from the prediction of the canonical equilibrium is that the condition must hold *after all histories and for all states*. To appreciate this difference, consider first the canonical equilibrium and suppose that, after a finite history h^t , the dealers' common belief is $p^t = \Pr(\tilde{v} = v_2 \mid h^t) = 1 - \epsilon$, with $\epsilon > 0$. For a sufficiently small ϵ , bid and ask quotes are close to v_2 . However, if the true state is such that $\tilde{v} = v_1$, the flow of trades eventually moves the price toward v_1 . Under the canonical equilibrium, how much time is required for the quotes to adjust and be close to v_1 ? The time required for this adjustment can be arbitrarily long if ϵ is very small. This implies that, in the canonical equilibrium, and after some histories, the mid-quote takes a very long time to be close to $v(\omega)$, and hence, its discounted average can be far from $v(\omega)$. This cannot occur in a BFE. If after some history h^t , quotes took a long time to adjust, then for a long period, the order flow would be unbalanced. We cannot exclude the possibility that a dealer knows the true state and hence expects to make a loss due to the large inventory accumulated at the wrong prices during the long adjustment period. However, this contradicts Lemma 4: each dealer makes positive profit in every state and after every history. Lemma 6 states that in a BFE, irrespective of the past evidence from the order flow, if the true \tilde{v} is not what past evidence suggests, it will not take long for the mid-quote to adjust and be relatively close to the true \tilde{v} . Thus, in a BFE, regardless of the past history of trade, quotes'

¹³Clearly, for $\bar{e} = \infty$, a BFE would not exist because Lemma 6 would require the mid-quote to be equal to the true \tilde{v} from the very first trading round onward, which cannot hold simultaneously for both $\tilde{v} = v_1$ and $\tilde{v} = v_2$. For this reason if, for example, \tilde{e} is assumed to be normally distributed, the economy has no BFE.

sensitivity to new trades does not vanish. As a consequence, dealers' quotes do not reflect a Bayesian expectation of \tilde{v} and, compared to the canonical equilibrium quotes, display excess volatility. To further formalize this idea, let $\hat{e} := \frac{\phi(\phi-1)\gamma^2\sigma^4}{v_2-v_1}$. When $\bar{e} > \hat{e}$, the mid-quote cannot be sufficiently close to both v_2 and v_1 , and thus, $A^*(v_1) \cap A^*(v_2) = \emptyset$. In this case, quotes that would guarantee all dealers positive profits when $\tilde{v} = v_1$ would generate negative profits for at least one dealer if $\tilde{v} = v_2$. This does not prevent the existence of a BFE. However, it requires dealers to make use of the information coming from the trading flow to tune the quotes to the true \tilde{v} . More formally, denote by $\tilde{\tau}(\omega, h^t)$ the additional time required by the mid-quote to be $\frac{2\rho^2\phi}{(\phi-1)\bar{e}}$ -close to $v(\omega)$ after history h^t , conditional on the state being ω . Then we have the following:

Lemma 7 (*Excess volatility*) *If σ forms a BFE, and $\bar{e} > \hat{e}$, then there exists $\kappa \in (0, 1)$ such that for all ω , after any history h^t , for any T ,*

$$\Pr(\tilde{\tau}(\omega, h^t) < T) \geq \frac{\kappa - \delta^T}{1 - \delta^T}.$$

Again, let us compare the prediction of this lemma with that of the canonical equilibrium. As explained above, if quotes follow Bayesian beliefs, the time $\tilde{\tau}(\omega, h^t)$ is arbitrarily long for some h^t and some ω . Namely, for any finite time T and any positive ϵ , one can find a finite history h^t such that $\Pr(\tilde{\tau}(\omega, h^t) < T) < \epsilon$. Lemma 7 states that in a BFE the, time $\tilde{\tau}(\omega, h^t)$ cannot be arbitrary long, regardless of the past history h^t and the state of nature.

3.3 Sufficient Conditions for BFE Pricing

In this section, we establish the existence of a BFE by constructing one. We first introduce the components of our candidate strategy profile. We then show how to combine these components to obtain a BFE.

Ingredients: We are interested in constructing a BFE that maintains the simple Markovian structure of the canonical equilibrium. To this end, we define a *market measure* π over the partition $\hat{\Omega}$ and a *market measure updating rule* ψ mapping time- t market measure, quotes and traders' orders into a time- $t + 1$ market measure. On the equilibrium path and in each period t , dealers' quotes depend on π^t only.

Specifically, recall that $\hat{\Omega} = \{\hat{\omega}_1, \hat{\omega}_2\}$, where $\hat{\omega}_k$ is the set of states $\omega \in \Omega$ such that $v(\omega) = v_k$, $k = 1, 2$. Fix some small $\varepsilon > 0$, and let $\Pi := [\varepsilon/4, 1 - \varepsilon/4]$. Let $\pi^t \in \Pi$ denote the probability that the market measure assigns to $\hat{\omega}_2$ at time t . Fix an arbitrary $\pi^0 \in [\varepsilon, 1 - \varepsilon]$ as the initial market measure. Thereafter, the market measure evolves according to the following updating rule $\psi : \Pi \times A \times S \rightarrow \Pi$:

$$\pi^{t+1} = \psi(\pi^t, a^t, s^t) := \arg \min_{\pi \in \Pi} \|\pi - \psi_B(\pi^t, a^t, s^t)\|, \quad (13)$$

where $\psi_B(\pi^t, a^t, s^t)$ is the Bayesian posterior as in expression (7) and p_t is replaced by π^t . In words, $\psi(\pi^t, a^t, s^t)$ maps a probability $\pi^t \in \Pi$ and a quote-order profile (a^t, s^t) onto the probability $\pi^{t+1} \in \Pi$ that is closest to the Bayesian posterior computed using a prior equal to π^t and the information provided by a trader's order s^t , given dealers' quotes a^t . We say that π^t *points to* state $\hat{\omega}_2$ if $\pi^t \geq 1 - \varepsilon$ and to state $\hat{\omega}_1$ if $\pi^t < \varepsilon$.

From the definition of a BFE, note that the market measure need not reflect any of the dealers' beliefs. In particular, π^0 can be chosen arbitrarily in the interval $[\varepsilon, 1 - \varepsilon]$, and thereafter, π^t does not follow Bayes' rule. Nevertheless, the level of the market measure affects the equilibrium quotes set by rational Bayesian dealers.

We can now define $\sigma : \Pi \rightarrow \Delta A$ as the partial strategy that maps, on the equilibrium path, the market measure onto dealers' quotes.¹⁴ Fix $d \in (0, \rho)$ and n strictly positive weights $\{\theta_1, \theta_2, \dots, \theta_n\}$, with $\sum_i \theta_i = 1$. When π^t *points to* state $\hat{\omega}$, we say that the market is in an

¹⁴The complete strategy for the dealers also describes dealers' behavior following a deviation and is detailed in the proof of Proposition 1.

exploiting phase and a^t is such that the best ask and bid quotes satisfy

$$\alpha(a^t) = v(\hat{\omega}) + d, \quad (14)$$

$$\beta(b^t) = v(\hat{\omega}) - d. \quad (15)$$

If $\tilde{v} = v(\hat{\omega})$, then the dealers' aggregate expected reward during one period of the exploiting phase is $\frac{2d(\rho-d)}{\phi\gamma\sigma^2}$, which is independent of $e(\omega)$ and strictly positive because $d \in (0, \rho)$.

When $\pi^t \in (\varepsilon, 1 - \varepsilon)$, the market is in an *exploring phase* and a^t is such that the best ask and bid quotes satisfy

$$\alpha^t = \alpha(\pi^t) + d, \quad (16)$$

$$\beta^t = \beta(\pi^t) - d, \quad (17)$$

where $\alpha(\cdot)$ and $\beta(\cdot)$ are the functions defined in equations (8) and (9), respectively. The best bid and ask quotes are set by dealer i a fraction θ_i of the time. Thus, each dealer i obtains the strictly positive fraction θ_i of dealers' aggregate profit and inventory.

The (on-path) equilibrium play can then be decomposed into an alternation of two types of phases: exploring phases and exploiting phases. Whenever $\pi^t \in [\varepsilon, 1 - \varepsilon]$, the market is in an exploring phase: dealers' quotes induce an informative flow of trades and are sensitive to the order. Thus, as time passes, the market measure attaches a growing weight to the true state $\hat{\omega}$. An exploiting phase begins once the market measure attaches enough weight to a specific state. In an exploiting phase, the bid and ask straddle the most likely value of \tilde{v} according to the market measure. If the true \tilde{v} is the one that the market measure points to, then dealers' quotes induce a trader's order flow that is balanced on average over time. This allows dealers to make profits through the non-competitive bid-ask spread $2d$ and keep their inventories balanced. If the true \tilde{v} is not that to which the market measure points, then traders' order flow is not balanced,

which moves the market measure back to the exploring phase. By definition, the market measure never attaches a probability greater than $1 - \frac{3}{4}\varepsilon$ to any state in $\hat{\Omega}$, regardless of the past history. Hence, the time required to move from the “wrong” exploiting phase to the “right” exploiting phase cannot grow without bound. However, for the same reason, the probability of temporarily passing from the right exploiting phase to an exploring phase is bounded away from 0. In particular, because the market measure is never too concentrated, even in an exploiting phase, an unbalanced order flow of finite length is sufficient to trigger exploring. For example, suppose that the market measure assigns the maximum weight of $1 - \varepsilon/4$ to state v_1 , which happens to be the true state. A finite sequence of buy orders suffices to move the market measure into the range $(\varepsilon, 1 - \varepsilon)$.¹⁵ That is, irrespective of the past history, the probability of moving from any exploiting phase to an exploring phase is bounded away from 0 and does not dampen over time. As a result, exploring phases recurrently emerge and the economy randomly switches between a regime of low volatility with quotes close to the true \tilde{v} and a regime of volatile quotes that temporarily drift away from \tilde{v} . Note that the price volatility does not dampen in subsequent exploring phases. In each exploring phase, quotes only depend on the level of the market measure (that is in the interval $(\varepsilon, 1 - \varepsilon)$) and the arrival of the traders’ orders, not on the total length of the trading history. This is illustrated in the simulation depicted in Figure 3, Section 4.

It remains to demonstrate that the simple Markov strategy profile described above forms a BFE. For this purpose, it is useful to define two crucial properties of the pair (ψ, σ) . We say that a strategy is ε -learning if, over many periods, the market measure points to the true $\hat{\omega}$ with a frequency that is at least $1 - \varepsilon$. In other words, the market measure is rarely far away from the truth in terms of long-run frequency. Formally the following hold:

Definition 2 *The pair (ψ, σ) is ε -learning, for $\varepsilon > 0$, if for any $\hat{\omega} \in \hat{\Omega}$, any $\omega \in \hat{\omega}$ and any $\pi^0 \in \Pi$,*

¹⁵This might be because, by chance, a sequence of traders who need to buy for hedging come to the market.

$$\Pr_{\omega, \sigma} \left[\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T 1_{\{\pi^t(\hat{\omega}) > 1-\varepsilon\}} < 1 - \varepsilon \right] < \varepsilon, \quad (18)$$

We say that a pair (ψ, σ) is ε -exploiting if, whenever the market measure points at some $\hat{\omega}$, the play is such that each dealer's payoff is strictly positive in all states ω included in $\hat{\omega}$. Formally, we have the following:

Definition 3 *The pair (ψ, σ) is ε -exploiting, for $\varepsilon > 0$, if, for all $\hat{\omega} \in \hat{\Omega}$ and all h^t such that $\pi^t(\hat{\omega}) \geq 1 - \varepsilon$, we have $\Pr_{\sigma} [a^t \in A_+^*(v(\hat{\omega})) \mid h^t] > 1 - \varepsilon$.*

The following proposition shows that a pair (ψ, σ) that is both ε -exploring and ε -exploiting forms a BFE if dealers are patient enough.

Proposition 1 *The pair (ψ, σ) defined in equations (13)–(17) is ε -learning and ε -exploiting. Furthermore, there exists $\underline{\delta} < 1$ such that the outcome induced by σ is a belief-free equilibrium outcome for all $\delta \in (\underline{\delta}, 1)$.*

That is, there exists a BFE σ^* that specifies the same action profile as the partial strategy σ after any history after which no player has deviated. Observe that a pair (ψ, σ) that is both ε -exploring and ε -exploiting forms a strategy profile satisfying the necessary conditions for a BFE, as described in Section 3.2. Namely, first, how dealers set their actions is clearly measurable with respect to traders' behavior (Lemma 3) because dealers' actions at t only depend on π^t , which is itself a function of the public history and the order flow. Second, this strategy leads to positive profits (Lemma 4) because the market measure frequently points to the right $\hat{\omega}$ (ε -exploring), and when this happens, the dealers' payoff is positive (ε -exploiting). Third, the fact that dealers' inventories are bounded (Lemma 5) and average mid quotes close to \tilde{v} (Lemma 6) is a consequence of the fact that dealers' payoffs remain positive for all values of $W(\omega)$. Fourth,

the strategy generates a sensitivity of quotes to the trading flow that never vanishes (Lemma 7). Because the market measure is never too concentrated on a state, it remains sensitive to the trading flow regardless of the past history. In the proof of Proposition 1, we show that dealers have no incentive to deviate. For this purpose, we show that there are strategies that are played after a deviation and that punish the deviating dealer while rewarding the other dealers. For this threat of punishment to be an effective deterrent, dealers should care enough about their future payoffs, *i.e.*, be patient enough.

4 Simulations and empirical implications

To illustrate some of the empirical implications of belief-free equilibria and their salient differences from the canonical equilibrium (hereafter, CE), we simulate price behavior resulting from these two equilibria.¹⁶

Dealers' profits: One of the necessary conditions for an equilibrium to be belief-free is that dealers' long-term profits are strictly positive in all states. In a BFE, this is achieved by maintaining a spread that is larger than that predicted in CE. Note, for example, that in the BFE, the spread remains bounded away from 0 even when the market measure is relatively concentrated. As a result, while in CE, the average dealers' aggregate per-period profit quickly converges to 0, in a BFE, it is of the same magnitude as d (see Figure 1). Note that a dealer's *ex post* profit also depends on the value of $e(\omega) \in \{-\bar{e}, \bar{e}\}$. Figure 2 represents the *ex post* cumulative profit for $e(\omega) = -\bar{e}$ and $e(\omega) = \bar{e}$. In CE (panel (a) of Figure 2), the dealers' cumulative profit remains negative for at least one realization of $e(\omega)$. In a BFE, the dealers' cumulative profit eventually becomes positive regardless of the realized $e(\omega)$ (panel (b) of Figure 2).

¹⁶For this simulation, the parameters of the fundamentals are $v_1 = \$36$, $v_2 = \$40$, $\bar{e} = 3$, $\text{Var}[\bar{e}] = 4$, $\gamma = 5$, $\phi = 3$ (thus, $\rho = 15$), whereas the BFE parameters are $d = \$0.02$, $\varepsilon = 0.02$, and $\pi^0 = p^0 = 0.5$. The figures report the time series for 10,000 trades.

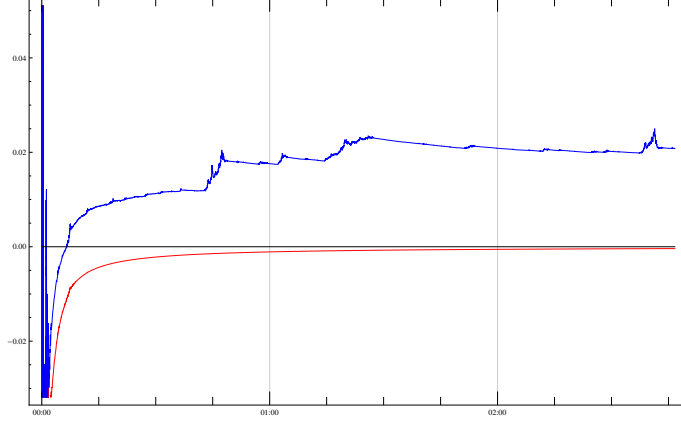
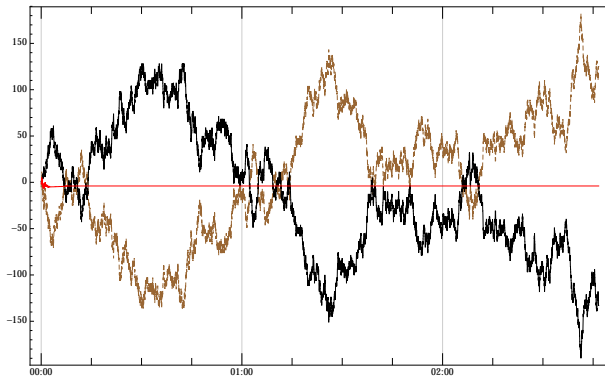
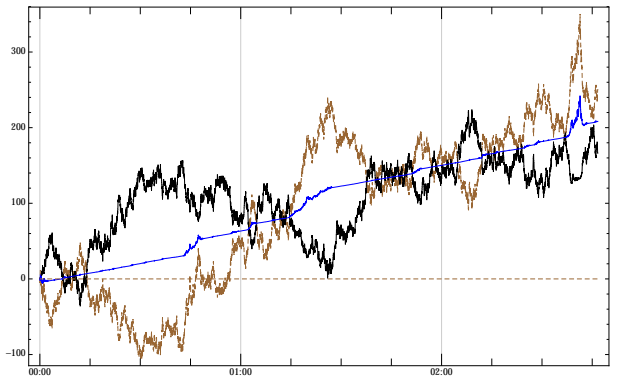


Figure 1: Evolution of the average per-period profit taking $e(\omega) = 0$ in CE (red line) and in BFE (blue line).



(a) Dealers' cumulative profits in CE



(b) Dealers' cumulative profits in BFE

Figure 2: Evolution of dealers' cumulative profits in CE and BFE for $e(\omega) = 0$ (red and blue lines), $e(\omega) = \bar{e}$ (black dashed) and $e(\omega) = -\bar{e}$ (brown dashed line).

Excess price volatility: In Figure 3, we report a simulation of the two equilibria for $v(\omega) = v_1$ in CE (left panel) and BFE (right panel). The sequence of traders' inventories in the two simulations is identical. However, the resulting sequences of equilibrium quotes differ substantially. In both BFE and CE, quote changes are correlated with the trading flow. However, whereas in CE, quotes' sensitivity to trades vanishes over time, this is not the case for BFE. As a result, in a BFE, quotes are intrinsically more volatile than they are in CE. This is because in CE, dealers' quotes reflect the common Bayesian belief, which eventually assigns a probability arbitrarily close to 1 to the true value of v . This leads to a vanishing volatility and bid-ask spread, with quotes that remain arbitrarily close to v . This cannot happen for the BFE market measure, which is never too concentrated on a given state and hence remains unstable. Thus, the market measure's and quotes' sensitivity to trading volume never vanish, and the economy exhibits stochastic regime shifting between exploring and exploiting phases. As a result, whereas the bid-ask spread generally straddles the true $v(\omega)$, quotes recurrently diverge from it.

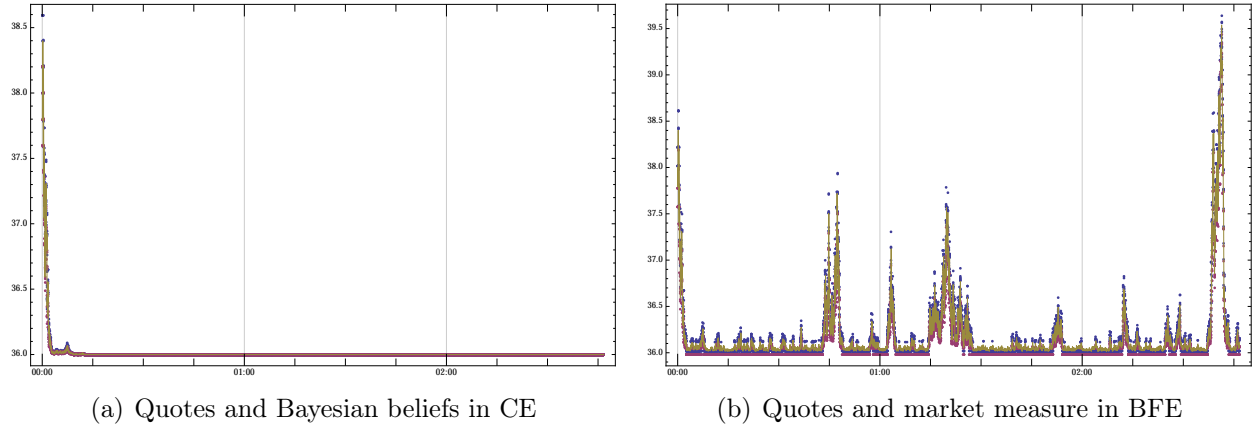
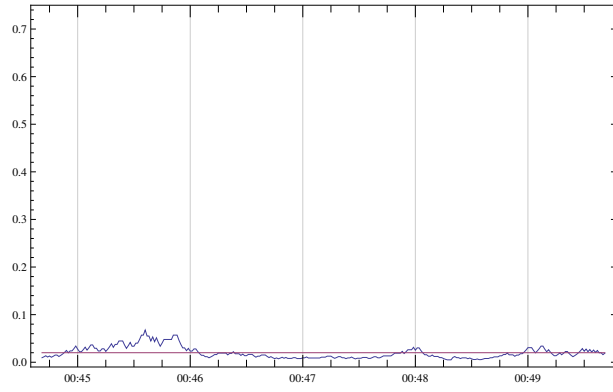
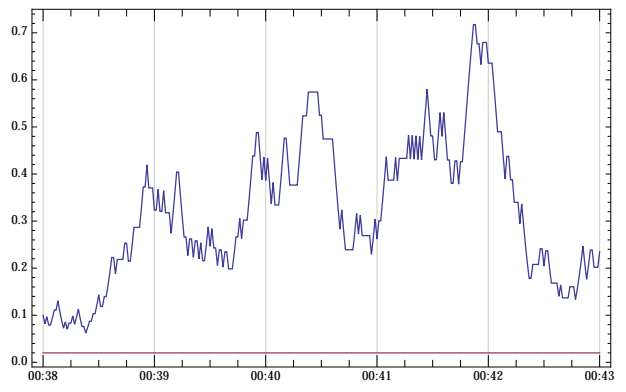


Figure 3: Evolution of bid quotes (red dots) and ask quotes (blue dots) in CE (panel (a)) and BFE (panel (b)). Yellow lines report the Bayesian expected value of \tilde{v} in CE (panel (a)) and the expected value of \tilde{v} using the market measure in BFE (panel b).

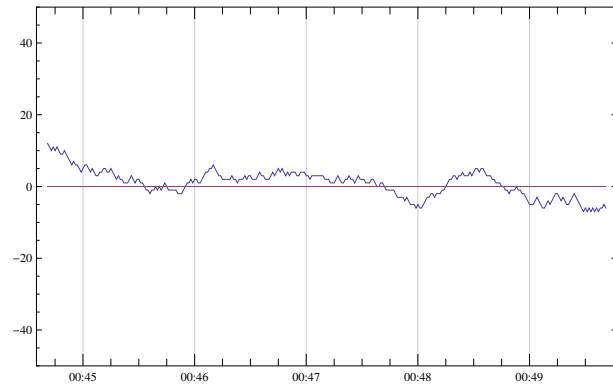
Volatility clustering: The recurrence of exploring and exploiting phases gives rise to price volatility clusters. In exploring phases, dealers attract an informative and unbalanced order flow from traders, whereas in exploiting phases, dealers make profits from intermediating a relatively



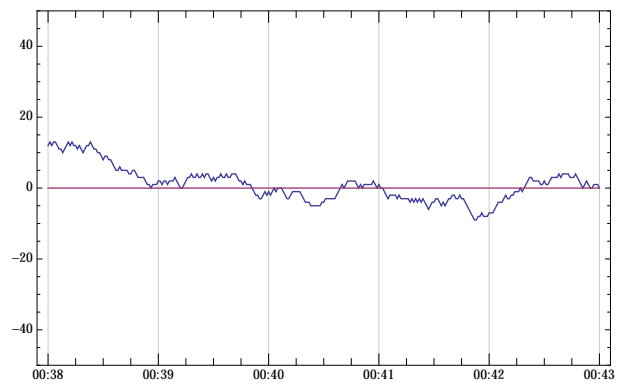
(a) Market measure in an exploiting phase



(b) Market measure in an exploring phase



(c) Evolution of inventory in an exploiting phase



(d) Evolution of inventory in an exploring phase

Figure 4: Market measure and dealers' inventory in an exploiting phase (panels (a) and (c)) and an exploring phase (panels (b) and (d)).

balanced order flow. In exploring phases, quotes react more sharply to the trading volume. Thus, quotes' volatility is higher in exploring than in exploiting phases. The volatility clustering effect is apparent in the right panel of Figure 3. The alternation of these phases endogenously generates price volatility regime shifts, a pattern that has been extensively documented in the asset-pricing literature. Price sensitivity to the order flow in exploiting and exploring phases is illustrated in Figure 4, which shows how the market measure reacts to the trading volume in an exploiting phase (panel (a)) and in an exploring phase (panel (b)). Interestingly, volatility regime shifts are anticipated by precise patterns in the order flow and evolution of dealers' inventory. A shift from low to high price volatility tends to be preceded by consistent imbalances in the traders' order flow and significant changes in dealers' inventory. The transition from high to low volatility phases follows the fading of traders' order imbalance and a stabilization of dealers' inventory. Note that in a BFE, the volatility regime shift is completely endogenous and occurs in the absence of news. Furthermore, note that the total length of the trading history affects neither the quotes' volatility during an exploiting phase nor the frequency of these phases. This is very different from the canonical equilibrium, which predicts that in the absence of news, price volatility is bound to fade away (See panel (a) of Figure 3).

Quote volatility vs. trading flow, bid-ask spread and profits: Observe that belief-free equilibria are associated with bid-ask spreads that are larger than those that would emerge in a CE, in which beliefs correspond to the market measure. However, this spread need not be as large as the spread that would emerge in the presence of a monopolistic dealer. For example, take our numerical example for the case in which \tilde{v} is perfectly known to all dealers. In a CE, the spread would be nil, and in the presence of a monopolistic dealer, it would be $\rho/2 = 10$; in our numerical example, it is $2d = 0.04$. That is, belief-free equilibria predict spreads that are larger than the competitive spread and do not require spreads to be as large as a monopolistic dealer model would predict. Such a "Markovian" BFE has some interesting implications regarding the correlation among price volatility, liquidity (measured by the bid-

ask spread), dealers' aggregate inventory and profits. In an exploring phase, orders are more informative; compared with exploiting phases, exploring phases are associated with larger bid-ask spreads, price volatility and aggregate inventory, as well as lower profits. This is consistent with the empirical regularities observed by Comerton-Forde et al. (2010): liquidity is negatively correlated with dealers' profits and inventories, as well as with price volatility.

News and volatility: Although this benchmark model does not explicitly allow for exogenous information shocks, it is straightforward to extend it to allow for the exogenous arrival of public news concerning fundamentals. Our BFE can accommodate this by having the market measure depend on all public information, *i.e.*, on public news and on the order flow. Unexpected news arriving when the market is in an exploiting phase moves the market measure and may trigger an exploring phase. As a result, upon receiving news, price volatility increases. This may generate price overshooting, and/or undershooting, with respect to the level of quotes that is reached once a new exploiting phase begins.

Dealers' inventories: A strictly positive profit in all states can only be achieved when the aggregate inventory does not grow unboundedly. For this reason, in BFE, dealers' inventory must remain bounded. This is not required in CE. For example, in the simulation, $v(\omega) = v_1$, and thus traders tend to sell rather than buy the asset. The right panel of Figure 5 reports the evolution of aggregate inventory in CE (red curve) and BFE (blue curve).

5 A general model of price formation

In this section, we generalize the Glosten and Milgrom model studied above along two dimensions. First, we have a richer set of states of nature that can affect both the asset's fundamental value and the distribution of the traders' characteristics. Second, we consider a more general trading protocol to encompass a wider range of trading systems. We then show that, at least qualitatively, the properties of the BFE derived in the simple model also hold in the general

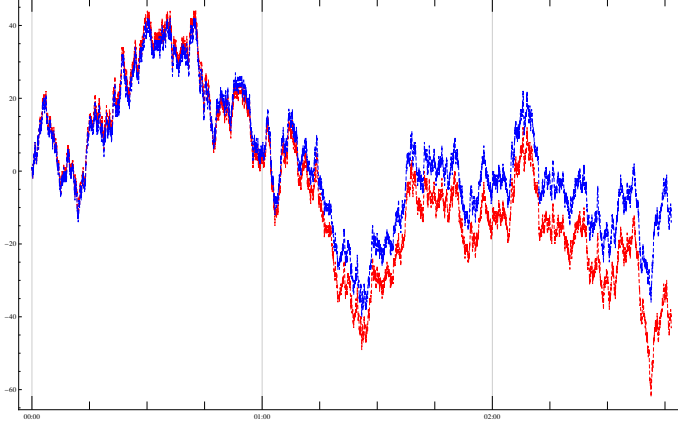


Figure 5: Evolution of dealers' aggregate inventory in CE (red line) and in BFE (blue line).

model.

5.1 Set-up

State of Nature At time 0, Nature chooses a state ω from an arbitrary set Ω . The state affects the fundamental value of the asset (fundamental uncertainty) and/or the composition of the population of traders (non-fundamental uncertainty). Formally, the fundamental value of the asset is written as $W(\omega) = v(\omega) + e(\omega)$ and is assumed to take bounded values, *i.e.*, $v(\omega) \in [\underline{v}, \bar{v}]$, $e(\omega) \in [-\bar{e}, \bar{e}]$, with $\underline{v} - \bar{e} \geq 0$ and $\bar{v} + \bar{e}$ being finite. The composition of the population of traders is denoted $Z(\omega) \in \Delta\Theta$, where Θ is a compact set, and an element $\theta \in \Theta$ describes a trader's utility function and initial holdings of cash and the asset. We assume that Nature chooses ω in Ω according to a probability distribution such that \tilde{e} is orthogonal to \tilde{v} and \tilde{Z} (*i.e.*, $E[\tilde{e}\tilde{v}] = E[\tilde{e}\tilde{Z}] = 0$) and is such that $E[\tilde{e}] = 0$, $\text{Var}[\tilde{e}] = \sigma^2 > 0$.

Trading protocol In each period $t = 1, 2, \dots, n$ risk-neutral long-lived dealers first choose an action profile from the finite set A . Second, a short-lived trader reacts by choosing an action from the finite set S . Dealers' actions and traders' reactions are publicly observable. For a given action-reaction profile $(a, s) \in A \times S$, let $q_j(a, s)$ and $p_j(a, s)$ be the resulting amounts of the

risky asset and money that other market participants transfer to agent j . That is, if the state of nature is ω , an action-reaction profile (a, s) leads to a change in agent j 's wealth given by

$$W(\omega)q_j(a, s) + p_j(a, s). \quad (19)$$

We assume that the trading protocol is such that no agent is forced to trade.

This trading protocol generalizes that of Section 1 and encompasses most trading protocols analyzed in the literature. We provide two examples below.

Quote-driven market (Biais, Martimort and Rochet (2000)) The set of actions available to each dealer is a trade schedule $T(\cdot)$, which specifies his willingness to trade q shares of the asset against the transfer of cash $T(q)$. Quantities and prices belong to some finite sets G_q and G . Thus, the set of actions is $A_i = [G]^{M_q}$, where M_q denotes the cardinality of G_q . The trader observes the dealers' schedules and chooses how many shares to trade with each dealer; thus, $S = [G_q]^n$.

Limit order markets There is no agreed-upon way to model limit order markets (see, for example, Foucault (1999), Goettler et al. (2005) and (2009), Foucault et al. (2005), Rosu (2009)). We present one possible specification that captures the functioning of a limit order market. In each period t , each dealer first submits a limit order that enters the book at the specified price. These limit orders form the initial book. Second, the trader chooses whether to submit a market order that trades against the initial book and/or a limit order that enters the book. This transforms the initial book into the interim book. Third, each dealer submits a market order that trades against the interim book. Fourth, all limit orders that are not executed are cancelled.¹⁷

¹⁷This fourth step is without loss of generality because the short-lived trader would cancel his order, whereas dealers can immediately post just-cancelled limit orders at the beginning of the following trading round.

Traders Traders know the realization of \tilde{v} but have no information about \tilde{e} . Since a trader does not learn anything about \tilde{e} by knowing his type θ , if the state is ω and no public information about \tilde{e} is available, then a type- θ trader's reaction to dealers' action a is given by

$$D(\omega, a, \theta) := \operatorname{argmax}_{s \in S} E [g_\theta((v(\omega) + \tilde{e})(q_{Tr}(a, s) + y_\theta) + p_{Tr}(a, s) + c_\theta)],$$

where we set $j = Tr$ for the trader, and $(g_\theta, c_\theta, y_\theta)$ is the triple describing the trader's utility function and initial cash and asset holdings. The function $D : \Omega \times A \times \Theta \rightarrow \Delta S$ and the distribution of types $Z(\omega) \in \Delta \Theta$ induce the probability with which the time- t trader chooses the reaction s to the dealers' action profile a , if the state of nature is ω and no information about \tilde{e} is available. This probability is

$$F(\omega, a, s) := \Pr [D(\omega, a, \tilde{\theta}) = s]. \quad (20)$$

The exact shape of the distribution $F : \Omega \times A \rightarrow \Delta S$ depends on the specification of Z , Θ and the distribution of \tilde{e} .

It is easy to see that, as long as the set Θ of traders' types includes risk-averse traders with positive inventories and risk-averse traders with negative inventories, the function F satisfies the following properties:

Room for trade (RFT) *There are $\bar{\rho} \geq \underline{\rho} > 0$ such that for all $\omega \in \Omega$,*

1. *If (a, s) translates into the trader buying at price $x > v(\omega) + \bar{\rho}$, then $F(\omega, a, s) = 0$.*
2. *If (a, s) translates into the trader selling at price $x < v(\omega) - \bar{\rho}$, then $F(\omega, a, s) = 0$.*
3. *If (a, s) translates into the trader buying at price $x < v(\omega) + \underline{\rho}$, then $F(\omega, a, s) > 0$.*
4. *If (a, s) translates into the trader selling at price $x > v(\omega) - \underline{\rho}$, then $F(\omega, a, s) > 0$.*

Conditions 1 and 2 state that traders never buy the asset at a price that is too high nor sell at a price that is too low relative to $v(\omega)$. Conditions 3 and 4 state that there is always a positive probability that the time- t trader is willing to buy or sell as long as the trading price is not too far from $v(\omega)$. In other words, if traders were to meet simultaneously in the market, there would be room for trade in all states ω . Note that the RFT property is satisfied in the market microstructure literature and guarantees that the no-trade theorem (Milgrom and Stokey (1982)) does not apply.

Dealers Dealers are risk-neutral long-lived agents. In state ω , dealer i 's reward given an action-reaction profile (a, s) is $u_i(\omega, a, s) = W(\omega)q_i(a, s) + p_i(a, s)$. With abuse of notation, we denote

$$\begin{aligned} q_i(\omega, a) &:= \sum_{s \in S} F(\omega, a, s)q_i(a, s), \\ p_i(\omega, a) &:= \sum_{s \in S} F(\omega, a, s)p_i(a, s), \end{aligned}$$

and let

$$u_i(\omega, a) := W(\omega)q_i(\omega, a) + p_i(\omega, a). \quad (21)$$

Thus, $u_i(\omega, a)$ is dealer i 's reward in state ω , given the dealers' action profile a and provided that the traders' behavior is described by F . Rewards are discounted at the common discount factor $\delta < 1$.¹⁸

Information We make no assumption regarding what each individual dealer knows about the true state ω .

As in the previous section, we assume that traders know the \tilde{W} component of \tilde{W} but not \tilde{e} . We

¹⁸As in the example from Section 1, δ can be interpreted as a measure of both time preference and the probability that no public announcement disclosing ω is made within the period.

are interested in the partition of Ω that corresponds to what dealers can statistically learn from a long-run observation of how traders react to their quotes. For any pair of states $\omega, \omega' \in \Omega$, we denote by $A(\omega, \omega') \subseteq \Delta A$ the set of action profiles a satisfying $F(\omega, a) \neq F(\omega', a)$. A state ω can be statistically distinguished from ω' only if $A(\omega, \omega') \neq \emptyset$. In fact, by consistently choosing their actions a^t in $A(\omega, \omega')$ and observing the distribution of traders' reactions, dealers can distinguish between the two states. Let $\hat{\Omega}$ be the partition of Ω induced in this way. That is, $\omega, \omega' \in \hat{\omega}$ if and only if $A(\omega, \omega') = \emptyset$. We assume that $\hat{\Omega}$ is finite with cardinality M and denote by $\hat{\omega}(\omega)$ the element of $\hat{\Omega}$ containing $\omega \in \Omega$.

5.2 Properties of the one-shot trading round

In a situation in which dealers have no private information, the traders' behavior is given by the function F defined in (20), and each dealer's reward function in state ω given by action profile a is $u_i(\omega, a)$, as defined in (21). We are interested in an economy in which $u_i(\omega, a)$ is "well behaved."

Definition 4 *The dealer's reward function $u_i(\omega, a) : \Omega \rightarrow \Delta A$ is regular if, for any given $\hat{\omega} \in \hat{\Omega}$, the following hold:*

1. *Positive maximum payoffs: There exists a non-empty set $A_+^*(\hat{\omega}) \subseteq \Delta A$ such that $u_i(\omega, a) > 0$, for all $a \in A_+^*(\hat{\omega})$, $\omega \in \hat{\omega}$ and dealer i .*
2. *Negative minimum payoffs: There exists an action profile $\underline{a}(\hat{\omega}) \in \Delta A$ such that $u_i(\omega, \underline{a}(\hat{\omega})) < 0$, for all $\omega \in \hat{\omega}$ and dealer i .*
3. *Non-positive expected payoffs: For any dealer i and probability measure $\mu_{\hat{\omega}} \in \Delta \hat{\omega}$, there exists $\underline{a}_{-i}^i(\mu_{\hat{\omega}}) \in \times_{j \neq i} \Delta A_j$ such that*

$$\max_{a_i} \sum_{\omega \in \hat{\omega}} \mu_{\hat{\omega}}(\omega) u_i(\omega, a_i, \underline{a}_{-i}^i(\mu_{\hat{\omega}})) \leq 0.$$

4. *Non-equivalent payoffs:* There exist n action profiles $\{a^1(\hat{\omega}), \dots, a^n(\hat{\omega})\} \in [\Delta A]^n$ such that $u_i(\omega, a^i(\hat{\omega})) < u_i(\omega, a^j(\hat{\omega}))$ for all $i \neq j$ and $\omega \in \hat{\omega}$.

Let $u^* = \min_i \min_{\omega} \min_{A^*_+(\hat{\omega}(\omega))} u_i(a, \omega) > 0$ denote a lower bound on payoffs from actions in $A^*(\hat{\omega})$. Roughly speaking, the properties *Positive maximum payoff* and *Negative minimum payoff* guarantee that for each statistically distinguishable state $\hat{\omega}$, there are action profiles providing each dealer with at least $u^* > 0$ and action profiles leading to strictly negative payoffs. *Non-positive expected payoffs* guarantees that it is possible to punish each dealer in each state. *Non-equivalent payoffs* states that for each $\hat{\omega}$, one can find as many action profiles as there are dealers, such that dealer i 's least favorite action profile is the i -th profile.

Lemma 8 *If Θ and Z are such that properties **RFT** are satisfied, then the dealers' reward function is regular.*

A dealer's payoff function in state ω , given a strategy profile σ and history h^t , can be decomposed as

$$V_i(\omega, \sigma \mid h^t) = w(\omega)Q_i(\omega, h^t) + P_i(\omega, h^t),$$

where, as in Section 3.2, $Q_i(\omega, h^t)$ is the expected discounted change in inventory for dealer i . In turn, this can be decomposed as the difference between the expected purchase volume $Q_i^+(\omega, h^t)$ and the sales volume $Q_i^-(\omega, h^t)$ for dealer i . The quantity $P_i(\omega, h^t)$ is the expected discounted change in cash for dealer i . Let $Q^+(\omega, h^t) := \sum_i Q_i^+(\omega, h^t)$ and $Q^-(\omega, h^t) := \sum_i Q_i^-(\omega, h^t)$ be dealers' aggregate expected purchase and sales volume.¹⁹

The definition of BFE for this general model does not change from that in Section 3.1. Hence, we focus directly on the necessary and sufficient conditions for a dealer's strategy to form a BFE when the reward function is regular.

¹⁹Note that since the state can affect both $v(\omega)$ and the distribution of dealer types, Q_i and P_i do not depend on $v(\omega)$ only, as was the case for the simpler model of Section 1. We do not exclude the possibility that there are states in which \tilde{v} is the same but \tilde{Z} differs.

5.3 Necessary conditions

In this section we show that, qualitatively, the necessary conditions for a strategy profile to be a BFE in the baseline model (Section 3.2) also apply to our general set-up.

Proposition 2 *Let σ form a BFE and $\hat{\omega}$ be an element of $\hat{\Omega}$. Then, for all $\omega \in \hat{\omega}$ and all histories h^t , we have the following:*

1. *(Measurability with respect to traders' behavior) A time- t trader's equilibrium behavior corresponds to the function F given in (20). The equilibrium occupation measure $\nu_{(\omega, \sigma, h^t)} \in \Delta(A \times S)$ is the same for all $\omega \in \hat{\omega}$.*
2. *(Positive dealer's payoffs) For each $\omega \in \hat{\omega}$, each dealer's continuation payoff $V_i(\omega, \sigma \mid h^t)$ is non-negative. Moreover, if $Q_i(\hat{\omega}, h^t) \neq 0$, then dealer i 's continuation payoff is nil for at most one $\omega \in \hat{\omega}$.*
3. *(Bounded dealers' aggregate inventory) Dealers' aggregate inventory and trading volume satisfy*

$$|Q^+(\omega, h^t) - Q^-(\omega, h^t)| \leq (Q^+(\omega, h^t) + Q^-(\omega, h^t)) \frac{\bar{\rho}}{\bar{e}}. \quad (22)$$
4. *(Average information efficiency) There is a level of $l(\hat{\omega})$ and of $\psi(\bar{e})$, decreasing in \bar{e} , such that the average transaction price is $\psi(\bar{e})$ -close to $l(\hat{\omega})$.*
5. *(Excess volatility) The expected time required for the transaction prices to be $\psi(\bar{e})$ -close to $l(\hat{\omega})$ is bounded, regardless of the past history.*

The interpretation of these properties is the same as those in Lemmas 3- 7.

5.4 Sufficient conditions

Equilibrium strategy ingredients: The following generalizes the construction of Section 3.3.

We begin by defining a *market measure* π . Let $\Pi \subseteq \Delta\hat{\Omega}$ be a closed set of probability distributions over $\hat{\Omega}$ and π denote an element in Π . Let $\pi(\hat{\omega})$ denote the probability that π attaches to $\hat{\omega}$. Let $\psi : \Pi \times A \times S \rightarrow \Pi$ be a probability updating rule, *i.e.*, $\pi^{t+1} = \psi(\pi^t, a^t, s^t)$. We are interested in simple strategies such that, on the equilibrium path and in each period t , dealers' actions depend on π^t (and possibly on s^{t-1}) only. We want the partial strategy to be ε -learning and ε -exploiting in the sense of Definitions 2 and 3. To be ε -exploiting, it is necessary that, independent of the level of the market measure π , the actions that allow one to distinguish the true $\hat{\omega}$ from the other state in $\hat{\Omega}$, must be played with strictly positive frequency. Formally, (ψ, σ) must be *exploratory* in the sense that $\forall \omega \in \Omega, \forall \hat{\omega}' \in \hat{\Omega}$ such that $\hat{\omega}' \neq \hat{\omega}(\omega)$, and for any $\pi^0 \in \Pi$,

$$\Pr_{\omega, \sigma} \left[\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T 1_{\{a^t \in A(\hat{\omega}', \hat{\omega}(\omega))\}} > 0 \right] = 1. \quad (23)$$

With this modification, one can construct a BFE qualitatively similar to that in Section 2. We have the following:

Theorem 1 *Suppose dealers' reward functions are regular. There exists $\bar{\varepsilon} > 0$ such that for any $\varepsilon < \bar{\varepsilon}$, if (ψ, σ) is ε -learning and ε -exploiting, then there exists $\underline{\delta} < 1$ such that the outcome induced by σ is a BFE outcome, for all $\delta \in (\underline{\delta}, 1)$.*

That is, there exists a BFE $\hat{\sigma}$ that specifies the same action profile as the partial strategy σ , after any history after which no player has deviated.

6 Discussion

Our environment is restrictive in several dimensions. In particular, dealers' actions are observed by all other dealers. Furthermore, the state of the world that determines the fundamentals is fixed permanently at time 0. Further, in equilibrium, long-term market participants are unable to exploit their private information. Finally, we assume that the set of states, the set of actions

and the asset's fundamental value are bounded. Here, we sketch how the model can be extended and the analysis adapted to accommodate such features.

The assumption that dealers' actions are observable might not be realistic for some opaque markets, for instance, when dealers' quotes are anonymous or when some dealers privately execute the orders of some of their clients. Imperfect monitoring of actions makes it more difficult to detect a dealer's deviation from the mutually profitable collusive-type strategy. This reduces the threat of punishment and complicates the implementation of collusive behavior. However, this does not eliminate the dealers' ability to sustain a BFE, as long as equilibrium strategies are built in a way that make deviations (at least statistically) detectable. For example, Christie and Shultz (1994) and Christie, Harris and Schultz (1994) document how Nasdaq dealers used to quote only on even-eight quotes. Deviations from such a collusive scheme can be easily detected even when quotes are anonymous. More generally, imperfect monitoring of players' actions is not an issue for the existence of a BFE (as demonstrated in Fudenberg and Yamamoto (2011)). The intuition is relatively simple. Consider what is publicly observed by dealers. This could be their public quotes or the previous trades. Then, it is possible to build a strategy mapping this public history onto dealers' actions that guarantees positive profits. The construction is amended to recognize that the only deviations that can be punished are those that impact the public history and are hence detectable. That is, imperfect monitoring of dealers' actions might impose further restrictions on the type of equilibrium strategies that can be sustained in a BFE.

Allowing for fluctuations in the value of the asset raises no difficulty, *as long as* these fluctuations take place at a much slower rate than does the learning process. That is, in the definition of (ψ, σ) to have ε -learning, we must account for the fact that $\hat{\omega}(\omega_t)$ depends on time t . Hence, the learning requirement is considerably stronger. We must think of learning the fundamental value as occurring at another time scale than the fluctuations of the value itself—perhaps the learning occurs within a day of trading, an interval of time over which the fluctuations in the fundamental value are sufficiently small to be considered negligible. If trading periods are at high frequency

(say, milliseconds), fundamentals hardly change from one such period to the next. Of course, we have in mind that the flow of trade itself does not affect fundamentals. The verification that σ is a BFE follows exactly the same steps as in the main proof.

A third restriction is that long-term market participants do not exploit their private information, if any. This is an implication of our definition of BFE, which requires optimality for any possible information structure. What matters is that dealers can identify the set of quotes that balance supply and demand coming from the mass of investors. As these quotes can be ultimately learned from the observation of the trading flow, dealers' private information is not crucial. The fact that, in our equilibrium, dealers do not exploit their private information might be counter-intuitive, but there is no difficulty in re-defining our model to accommodate such behavior without abandoning the belief-free assumption. Rather than taking the asset value as a primitive that determines a distribution over the players' private signals, one can regard the players' private signals as a primitive that determines the asset's value. In that case, we can re-define a strategy profile to be belief-free if it is the case that, for every player, given his private signal, his strategy (that can depend on his private signal) is optimal independent of the other players' possible strategies. That is, given a player's signal, there is a set of signal profiles of his opponents that are consistent with his; for each such signal profile, his opponents play some strategy profile. Belief-freeness requires the player's strategy to be optimal in response to all these profiles. In fact, it is clear that we do not need to impose that the players' combined signals specify the value of the asset. Rather, it specifies a set of possible values, with respect to all of which the best-reply property must hold. This provides a natural extension of the definition of BFE that allows dealers to exploit their private information. Whereas such belief-free equilibria can be characterized, they are unlikely to exist under very broad conditions, and hence they can help to explain price formation only for those actual cases in which only few specific forms of information asymmetries are likely. We believe that such an extension raises interesting questions and technical challenges that should motivate further study.

While the assumption that the asset value \tilde{W} is bounded might be realistic, it is not innocuous. This is because a dealer who believes that the asset’s fundamental value is arbitrarily large can expect to make arbitrarily large profits from a single deviation. In this case, equilibria may not exist for fixed $\delta < 1$. However, the boundedness of action profiles is not restrictive, if the asset’s value is bounded, in the sense that equilibrium quantities and prices are necessarily bounded.

7 Conclusion

This paper considers market microstructure models in which long-lived dealers interact with short-lived traders. We characterize equilibrium price formation strategies that do not depend on dealers’ beliefs about fundamentals. Belief-free equilibria feature two key ingredients. First, dealers collectively learn the value of those fundamentals that affect traders’ demand. Second, for any given value of these fundamentals, dealers generate positive profits from intermediating traders’ demand. This has three implications that contrast with those delivered by canonical microstructure models relying on the assumption of equally uninformed competitive dealers. First, dealers’ long-term profit is strictly positive, independent of the asset’s fundamental value. This profit is obtained through intermediating traders’ demand. Second, the trading price need not reflect any of the dealers’ beliefs and is generally more volatile than prices that reflect the evolution of Bayesian beliefs. Third, dealers’ inventories tend to be balanced even in the absence of risk aversion or institutional constraints. Given that belief-free equilibrium is more stringent than traditional solution concepts, it might be surprising that so much flexibility remains—in particular, the equilibrium is not unique. Hence, we have focused on a belief-free equilibrium with a simple Markovian structure. When applied to a version of the Glosten and Milgrom model, it explains well-documented stylized empirical facts. For specific microstructure games, it might then be sensible to focus on belief-free equilibria that satisfy further criteria. For example, depending on the specific trading model considered, one could analyze equilibria that maximize

the dealers' aggregate payoff, that minimize the expected time required for the market measure to point at the true state, that even minimize the aggregate cost of learning, or more generally, strategies that form a belief-free equilibrium for the largest range of dealers' discount rates.

Appendix A

Proof of Lemma 1

We provide the proof for $F(\omega, a, 1)$. The proof for $F(\omega, a, -1)$ is analogous.

If the state is ω , the profile of dealer quotes is a and no information about \tilde{e} is available, a trader with inventory y prefers selling to not trading if

$$(y - 1)v(\omega) + \beta(a) - \frac{\gamma}{2}(y - 1)^2\sigma^2 \geq yv(\omega) - \frac{\gamma}{2}y^2\sigma^2.$$

This is true for y not smaller than the threshold $\frac{1}{2} + \frac{v(\omega) - \beta(a)}{\gamma\sigma^2}$. Thus, the probability that trader t sells is the probability that his inventory y is above this threshold, that is, $1 - Z\left(\frac{1}{2} + \frac{v(\omega) - \beta(a)}{\gamma\sigma^2}\right)$. The expression for F follows from the fact that $Z(\cdot)$ is the distribution function of a uniform distribution on the interval $[-\phi/2, \phi/2]$. ■

Proof of Lemma 2

We provide the proof for the bid price. The proof for the ask price is analogous. Consider a dealer whose beliefs are such that $\Pr(\tilde{v} = v_2) = p$ and $E[\tilde{e}] = 0$. His expected profit from buying the asset at price $\beta \in [v_1, v_2]$ is

$$p(v_2 - \beta)F(v_2, a, 1) + (1 - p)(v_1 - \beta)F(v_1, a, 1),$$

where a is such that $\beta(a) = \beta$. The expression for $\beta(p)$ is obtained by first replacing $F(\cdot)$ with the expression given in Lemma 1, considering that $\beta \geq v_1$ and the assumption $\rho > v_2 - v_1$ and, second, finding the β that makes the dealer's expected payoff equal to 0. ■

Proof of Lemma 3

A BFE must be an equilibrium even when dealers have no private information about ω . In this case, first, the behavior of traders is given by Lemma 1. Second, suppose that $v(\omega) = v(\omega')$,

then traders' behavior is the same in states ω and ω' . Thus, there is no agent in the economy who can distinguish between these two states, and hence, all agents' behavior, quotes and orders, must be the same in the two states. ■

Proof of Lemma 4

To see that the equilibrium payoff cannot be negative, suppose that for some state ω , dealer i and history h^t , we have $V_i(\omega, \sigma^* | h^t) < 0$. Then, if dealer i believes that the true state is ω , he would be better off by setting $\beta_i^\tau < v_1 - \rho$ and $\alpha_i^\tau > v_2 + \rho$ in all periods $\tau > t$. This guarantees that he never trades and hence obtains zero continuation payoff. Thus, σ^* cannot be a BFE. To see that if a dealer trades, his equilibrium payoff is nil in at most two states and is strictly positive in all other states, take any $v \in \{v_1, v_2\}$. Note that for all ω such that $v(\omega) = v$, because of Lemma 3, one has

$$V_i(\omega, \sigma^* | h^t) = (v + e(\omega))Q_i(v, \sigma^* | h^t) + P_i(v, \sigma^* | h^t).$$

Suppose that $Q_i(v, \sigma^* | h^t) \neq 0$ and take two different states ω, ω' with $v(\omega) = v(\omega') = v$ and $e(\omega) \neq e(\omega')$. Then, it must be that $V_i(\omega, \sigma^* | h^t) \neq V_i(\omega', \sigma^* | h^t)$. Because $V_i(\omega, \sigma^* | h^t)$ cannot be strictly negative, it can be nil for at most one of the states $\omega \in \hat{\omega}_{v_1}$ and is strictly positive for all other states in $\hat{\omega}_{v_1}$. The same argument applies for $\hat{\omega}_{v_2}$. ■

Proof of Lemma 5

From Lemma 1 and Lemma 3, we have that dealers cannot purchase the asset for less than $v(\omega) - \rho$ or sell it for more than $v(\omega) + \rho$. Fix $v \in \{v_1, v_2\}$. For all ω satisfying $v(\omega) = v$, we have

$$\begin{aligned} V_i(\omega, \sigma | h^t) &\leq Q_i^+(v, h^t)(W(\omega) - v + \rho) + Q_i^-(v, h^t)(v + \rho - W(\omega)) \\ &= (Q_i^+(v, h^t)(e(\omega) + \rho) + Q_i^-(v, h^t)(\rho - e(\omega))), \end{aligned}$$

which is non-negative only if

$$e(\omega)(Q_i^+(v, h^t) - Q_i^-(v, h^t)) > -\rho(Q_i^+(v, h^t) + Q_i^-(v, h^t)).$$

Because in a BFE we cannot exclude the case in which dealer i believes that $e(\omega) = \bar{e}$ or the case where he believes $e(\omega) = -\bar{e}$, the above inequality must be satisfied for the two extreme levels of \tilde{e} and, as a consequence, for all possible values of $e(\omega) \in [-\bar{e}, \bar{e}]$. This implies inequality (12). ■

Proof of Lemma 6

Let $\tilde{a} \in A^*(v)$. Then, for all dealers i and all states ω such that $v(\omega) = v$, we have $u_i(\omega, \tilde{a}) \geq 0$, implying

$$\sum_i u_i(\omega, \tilde{a}) = \sum_{a \in A} \tilde{a}(a) (F(\omega, a, 1)(W(\omega) - \beta(a)) + F(\omega, a, -1)(\alpha(a) - W(\omega))) \geq 0.$$

Fix a state ω , and consider a quote profile a . Observe that if $\beta(a) > v(\omega) + \rho \frac{\phi+1}{\phi-1}$, then the trader sells with probability 1, but by setting quotes such that $\beta(a) = v(\omega) + \rho \frac{\phi+1}{\phi-1}$ trader would sell with the same probability at a strictly lower price. If $\beta(a) < v(\omega) - \rho$, the trader does not sell and dealers would have the same reward from setting $\beta(a) = v(\omega) - \rho$. A similar reasoning applies to the ask side. Hence, in what follows, we focus on a such that $\beta(a) \in [v(\omega) - \rho, v(\omega) + \rho \frac{\phi+1}{\phi-1}]$ and $\alpha(a) \in [v(\omega) - \rho \frac{\phi+1}{\phi-1}, v(\omega) + \rho]$. Then, we have

$$\begin{aligned} \sum_i u_i(\omega, a) &= F(\omega, a, 1)(W(\omega) - \beta(a)) + F(\omega, a, -1)(\alpha(a) - W(\omega)) \\ &< e(\omega)(F(\omega, a, 1) - F(\omega, a, -1)) + \rho(F(\omega, a, 1) + F(\omega, a, -1)) \\ &< e(\omega)(F(\omega, a, 1) - F(\omega, a, -1)) + 2\rho \\ &= e(\omega) \frac{\phi-1}{\rho\phi} \left(\frac{\alpha(a) + \beta(a)}{2} - v(\omega) \right) + 2\rho, \end{aligned}$$

where the first inequality follows from the fact that traders do not sell for less than $v(\omega) - \rho$ or

buy for more than $v(\omega) + \rho$. The last equality comes from Lemma 3 and the expression for F in Lemma 1.

Thus, $\tilde{a} \in A^*(v)$ implies that for *all possible values* of $e(\omega) \in [-\bar{e}, \bar{e}]$, one has that

$$0 \leq \sum_i u_i(\omega, \tilde{a}) < e(\omega) \frac{\phi - 1}{\rho\phi} \left(\frac{\alpha(\tilde{a}) + \beta(\tilde{a})}{2} - v \right) + 2\rho.$$

This leads to

$$\left| \frac{\alpha(\tilde{a}) + \beta(\tilde{a})}{2} - v \right| < \frac{2\rho^2\phi}{(\phi - 1)\bar{e}},$$

that is, the average mid-quote $\frac{\alpha(\tilde{a}) + \beta(\tilde{a})}{2}$ cannot differ from $v(\omega)$ for more than $\frac{2\rho^2\phi}{(\phi - 1)\bar{e}}$, regardless of what ω is. ■

Proof of Lemma 7

Consider two states ω and ω' with $v(\omega) = v_1$ and $v(\omega') = v_2$. Observe that if $\bar{e} > \hat{e}$, then there is no \tilde{a} such that the mid-quote is $\frac{2\rho^2\phi}{(\phi - 1)\bar{e}}$ -close to both v_1 and v_2 . Thus, $A^*(v_1) \cap A^*(v_2) = \emptyset$. Note that, by definition of A^* , if a quote profile $\tilde{a} \notin A^*(v)$, then there is a state ω , with $v(\omega) = v$, and a dealer i , such that $u_i(\omega, \tilde{a}) < 0$. Thus, if $\bar{e} > \hat{e}$ and $\tilde{a} \in A^*(v)$, then there is $l < 0$, a dealer i and a state ω , with $v(\omega) \neq v$, such that $u_i(\omega, \tilde{a}) < l$. From Lemma 6, we also know that $\tilde{a} \in A^*(v)$ requires

$$\left| \frac{\alpha(\tilde{a}) + \beta(\tilde{a})}{2} - v \right| < \frac{2\rho^2\phi}{(\phi - 1)\bar{e}}.$$

Take any history h^t and let $t + \tau_1$ the first time that $\left| \frac{\alpha(a^{t+\tau_1}) + \beta(a^{t+\tau_1})}{2} - v_1 \right| < \frac{2\rho^2\phi}{(\phi - 1)\bar{e}}$. Then, there is a dealer i and a state ω , with $v(\omega) = v_1$, such that in this state, this dealer's continuation payoff is not larger than $(1 - \delta^{\tau_1})l + \delta^{\tau_1}\bar{u}$, where $\bar{u} > 0$, finite, is such that $|u_i(\omega, a)| < \bar{u}$ for all dealers i , $\omega \in \Omega$, and $a \in \Delta A$. This is the maximum profit a dealer can possibly make in one trading round and is bounded by $\bar{e} + \rho$. Let $z(\tau \mid \omega, h^t)$ denote the equilibrium probability that $\tau_1 = \tau$ conditional on the true state being ω and the past history being h^t . Because we cannot exclude the possibility that dealer i believes that the state is ω , this dealer's equilibrium

continuation payoff satisfies

$$V_i(\omega, \sigma, h^t) \leq \sum_{\tau=0}^{\infty} z(\tau \mid \omega, h^t) ((1 - \delta^\tau)l + \delta^\tau \bar{u}) = l + (\bar{u} - l) \sum_{\tau=0}^{\infty} z(\tau \mid \omega, h^t) \delta^\tau.$$

Let $\kappa := -l/(\bar{u} - l) \in (0, 1)$. From $V_i(\omega, \sigma, h^t) \geq 0$, we have that for any finite T ,

$$\begin{aligned} \kappa &\leq \sum_{\tau=0}^{\infty} z(\tau \mid \omega, h^t) \delta^\tau \leq \sum_{\tau=0}^{T-1} z(\tau \mid \omega, h^t) \delta^\tau + \sum_{\tau=T}^{\infty} z(\tau \mid \omega, h^t) \delta^\tau \\ &\leq \sum_{\tau=0}^{T-1} z(\tau \mid \omega, h^t) + \delta^T \sum_{\tau=T}^{\infty} z(\tau \mid \omega, h^t) = \Pr(\tilde{\tau}(\omega, h^t) < T)(1 - \delta^T) + \delta^T, \end{aligned}$$

because $\sum_{\tau=0}^{T-1} z(\tau \mid \omega, h^t) = \Pr(\tilde{\tau}(\omega, h^t) < T)$. This implies $\Pr(\tilde{\tau}(\omega, h^t) < T) < \frac{\kappa - \delta^T}{1 - \delta^T}$. ■

■

Proof of Proposition 1

Let us first show that

Lemma 9 *The proposed Markov strategy is ε -learning.*

Proof. Let ε_0 be the parameter that determines the size of the set $\Pi := [\varepsilon_0/4, 1 - \varepsilon_0/4]$ in the definition of the strategy. We will show that for any given $\varepsilon > 0$, one can set ε_0 such that the strategy is ε -learning.

To fix notations, we assume that the true state is $\hat{\omega}_1$ and denote $\Pr(\cdot) = \Pr(\cdot \mid \hat{\omega}_1, \sigma)$. The (Bayesian) updated probability of state $\hat{\omega}_2$ is denoted p^t and evolves according to the recursion $p^{t+1} = \psi_B(p^t, a^t, s^t)$, where ψ_B denotes Bayesian updating. The probability that the market measure assigns to $\hat{\omega}_2$ at time t is denoted π^t and evolves according to the following updating rule:

$$\pi^{t+1} \arg \min_{\pi \in \Pi} \|\pi - \psi_B(\pi^t, a^t, s^t)\|. \quad (24)$$

We will prove the following.

Lemma 10 *For all $\varepsilon > 0$, there exists a choice of $0 < \varepsilon_0 < \varepsilon$, such that for all sufficiently high values of δ ,*

$$\Pr \left[\sum_t (1 - \delta) \delta^{t-1} 1_{\{\pi^t > \varepsilon\}} > \varepsilon \right] < \varepsilon.$$

Intuitively, this follows from the fact that Bayesian beliefs converge to the truth. Specifically, we gather the following properties from standard Bayesian statistics (see, *e.g.*, Chamley (2004)).

P1: The belief p^t converges to 0 almost surely under $\Pr(\cdot \mid \hat{\omega}_1, \sigma)$.

P2: The likelihood ratio $\frac{p^t}{1-p^t}$ is a martingale under $\Pr(\cdot \mid \hat{\omega}_1, \sigma)$.

First, we use these properties to estimate $E[\sum_t (1 - \delta) \delta^{t-1} \pi^t]$.

Claim 1 *There is a constant $c > 1$ such that for all sufficiently high values of δ ,*

$$E \left[\sum_t (1 - \delta) \delta^{t-1} \pi^t \right] \leq \frac{c\varepsilon_0}{4 - c\varepsilon_0} + \varepsilon_0/4.$$

Proof. From Property P1, we obtain that after finitely many steps, at some stage t , $p^t \leq \varepsilon_0/4$ and thus $\pi^t = \varepsilon_0/4$. Precisely, let $\hat{T} = \inf\{t : p^t \leq \varepsilon_0/4\}$. Property 1 implies that $\Pr(\hat{T} > T)$ converges to 0 as T goes to infinity. Thus, for some T_0 , $\Pr(\hat{T} > T_0) \leq \varepsilon_0/4$. Let us fix such a T_0 , and write

$$E \left[\sum_t (1 - \delta) \delta^{t-1} \pi^t \right] = \sum_T \Pr(\hat{T} = T) E \left[\sum_t (1 - \delta) \delta^{t-1} \pi^t \mid \hat{T} = T \right].$$

From the choice of T_0 ,

$$E \left[\sum_t (1 - \delta) \delta^{t-1} \pi^t \right] \leq \sum_{T \leq T_0} \Pr(\hat{T} = T) E \left[\sum_t (1 - \delta) \delta^{t-1} \pi^t \mid \hat{T} = T \right] + \varepsilon_0/4.$$

For each $T \leq T_0$,

$$E \left[\sum_t (1 - \delta) \delta^{t-1} \pi^t \mid \hat{T} = T \right] \leq 1 - \delta^T + \delta^T X,$$

where $X := E_{\varepsilon_0} [\sum_t (1 - \delta) \delta^{t-1} \pi^t]$ is the expected discounted average of the π^t 's, assuming that this (Markov) process starts from $\pi^1 = \varepsilon_0/4$. We can choose a δ sufficiently high that for all $T \leq T_0$, $1 - \delta^T \leq \varepsilon_0/4$. Then,

$$E \left[\sum_t (1 - \delta) \delta^{t-1} \pi^t \right] \leq \sum_{T \leq T_0} \Pr(\hat{T} = T) (\varepsilon_0/4 + \delta^T X) + \varepsilon_0/4 \leq X + \varepsilon_0/2.$$

Now, let us estimate X . Starting from $\pi^1 = \varepsilon_0/4$, with some probability α , the Bayesian updating decreases; then, $\pi^2 = \varepsilon_0$, and the process is restarted.

With the complementary probability $1 - \alpha$, the Bayesian updating increases and $\pi^2 = p^2$. The next argument establishes that Bayesian updating leads to very small changes. Recall that from Bayes' rule,

$$p^{t+1} = \frac{p^t F(v_2, a^t, s^t)}{p^t F(v_2, a^t, s^t) + (1 - p^t) F(v_1, a^t, s^t)},$$

and thus

$$\frac{p^{t+1}}{p^t} = \frac{F(v_2, a^t, s^t)}{p^t F(v_2, a^t, s^t) + (1 - p^t) F(v_1, a^t, s^t)}.$$

Using the expressions for $F(\cdot)$ given in Lemma 1, and observing that according to the proposed strategy $v_1 \leq \beta^t \leq \alpha^t \leq v_2$, we have that $\frac{p^{t+1}}{p^t} < 1 + \frac{v_2 - v_1}{\rho} \in (1, 2)$ because $\rho > v_2 - v_1$. There is thus a constant $c := 1 + \frac{v_2 - v_1}{\rho} > 1$ such that $\pi^2 \leq c\pi^1 = c\varepsilon_0/4$. We use now property P2, which states that for the Bayesian process, $\frac{p^t}{1 - p^t}$ is a martingale. In particular, $E \left[\frac{p^t}{1 - p^t} \right] = \frac{p^1}{1 - p^1}$. Because $p^t \leq \frac{p^t}{1 - p^t}$, we find that $E[p^t] \leq \frac{p^1}{1 - p^1}$. We conclude that in the case in which the Bayesian process starts from $p^2 \leq c\varepsilon_0/4$, the expectation remains below $\frac{c\varepsilon_0/4}{1 - c\varepsilon_0/4}$. This holds as long as the process π^t follows the Bayesian trajectory. If it declines to $\varepsilon_0/4$ again, then this still holds since $\varepsilon_0/4 \leq \frac{c\varepsilon_0/4}{1 - c\varepsilon_0/4}$. If it increases to $1 - \varepsilon_0/4$, then $\pi^t \leq p^t$, and thus, $E[\pi^t] \leq E[p^t] \leq \frac{c\varepsilon_0/4}{1 - c\varepsilon_0/4}$; therefore, this holds again. Hence, we have that if we start from $\pi_0 = \varepsilon_0/4$, then $E[\pi^t] \leq \frac{c\varepsilon_0/4}{1 - c\varepsilon_0/4}$ for all $t > 0$.

Now, X satisfies the following recursion:

$$\begin{aligned}
X &= (1 - \delta)\varepsilon_0/4 + \delta \left(\alpha X + (1 - \alpha) \mathbb{E}_{\pi^2} \left[\sum_t (1 - \delta) \delta^{t-1} \pi^{t+1} \right] \right) \\
&\leq (1 - \delta)\varepsilon_0/4 + \delta \left(\alpha X + (1 - \alpha) \sum_t (1 - \delta) \delta^{t-1} \frac{c\varepsilon_0}{4 - c\varepsilon_0} \right) \\
&= (1 - \delta)\varepsilon_0/4 + \delta \left(\alpha X + (1 - \alpha) \frac{c\varepsilon_0}{4 - c\varepsilon_0} \right).
\end{aligned}$$

Solving for X , this yields

$$X \leq \frac{1 - \delta}{1 - \delta + \delta(1 - \alpha)} \varepsilon_0/4 + \frac{\delta(1 - \alpha)}{1 - \delta + \delta(1 - \alpha)} \frac{c\varepsilon_0}{4 - c\varepsilon_0} \leq \frac{c\varepsilon_0}{4 - c\varepsilon_0}.$$

We then obtain $\mathbb{E}[\sum_t (1 - \delta) \delta^{t-1} \pi^t] \leq \frac{c\varepsilon_0}{4 - c\varepsilon_0} + \varepsilon_0/2$. ■

■

ε -exploiting: If at time t , the market measure points at state $\hat{\omega}_2$, then $\alpha(a^t) = v_2 + d$ and $\alpha(a^t) = v_2 - d$. The dealers' aggregate payoff is

$$\begin{aligned}
\sum_i u_i(\omega, a) &= F(\omega, a, 1)(W(\omega) - \beta(a)) + F(\omega, a, -1)(\alpha(a^t) - W(\omega)) \\
&= (F(\omega, a, 1) - F(\omega, a, -1))e(\omega) + (F(\omega, a, 1) + F(\omega, a, -1))d = \frac{2d(\rho - d)}{\phi\gamma\sigma^2},
\end{aligned}$$

which is strictly positive because $0 < d < \rho$, and it does not depend on $e(\omega)$. The strategy is such that this aggregate profit is shared across dealers. Namely dealer i 's profit is $\theta_i \frac{2d(\rho - d)}{\phi\gamma\sigma^2}$, which is strictly positive in all states ω such that $v(\omega) = v_2$. Hence, $a^t \in A_+^*(v_2)$. The same argument applies to $\hat{\omega}_1$.

We can now show that if δ is large enough, then (σ, ϕ) describes the on-equilibrium-path behavior of a BFE. For this purpose, we first prove that on the equilibrium path, each dealer's payoff is strictly positive.

Let $\bar{u} > 0$, finite, be such that $|u_i(\omega, a)| < \bar{u}$ for all dealers i , $\omega \in \Omega$, and $a \in \Delta A$. A dealer's reward cannot be smaller than $-\bar{u}$. Consider a dealer i . Because (σ, ϕ) is ε -exploiting, dealer i 's reward when the market measure points to the true $\hat{\omega}$ is $u^* = \theta_i \frac{2d(\rho-d)}{\phi\gamma\sigma^2} > 0$.

Let $\omega \in \hat{\omega}$ be the true state and $\pi^t(\hat{\omega})$ the market measure of $\hat{\omega}$ at time t . Let q^t be the probability that at time t , the market measure satisfies $\pi^t(\hat{\omega}) > 1 - \varepsilon$. Because (ϕ, σ) is ε -exploiting, Definition 3 implies that with probability q_t , dealer i 's stage- t reward is at least $(1 - \varepsilon)u^* - \varepsilon\bar{u}$. Then, at time $\tau \geq 0$, dealer i 's payoff satisfies

$$\begin{aligned} V_i^\delta(\omega, \sigma \mid h^\tau) &> (1 - \delta) \sum_{t=\tau}^{\infty} \delta^{t-\tau} (q^t((1 - \varepsilon)u^* - \varepsilon\bar{u}) - (1 - q^t)\bar{u}) \\ &= (1 - \varepsilon)(u^* + \bar{u})(1 - \delta) \sum_{t=\tau}^{\infty} \delta^t q^{t-\tau} - \bar{u}. \end{aligned} \quad (25)$$

Because (ϕ, σ) is ε -exploring, from Definition 2, we have that

$$\Pr_{\omega, \sigma} \left[\lim_{\delta \rightarrow 1} (1 - \delta) \sum_{t=\tau}^{\infty} \delta^t q^{t-\tau} > 1 - \varepsilon \right] > 1 - \varepsilon. \quad (26)$$

Hence, we have that

$$\lim_{\delta \rightarrow 1} V_i^\delta(\omega, \sigma \mid h^\tau) > (1 - \varepsilon)^3(u^* + \bar{u}) - (1 + \varepsilon)\bar{u}. \quad (27)$$

As the r.h.s. is strictly positive for $\varepsilon = 0$, it is also positive for all ε smaller than some $\bar{\varepsilon} > 0$. The continuity of V_i^δ in δ implies that there exists $\underline{\delta} < 1$ such that for $\varepsilon < \bar{\varepsilon}$, $\delta > \underline{\delta}$, dealer i 's continuation payoff $V_i^\delta(\omega, \sigma \mid h^\tau)$ is strictly positive.

We now turn to the off-equilibrium-path behavior. To do so, we need two ingredients. First, let $\sigma^i : \Pi \rightarrow \Delta A$ be a partial strategy that differs from σ only in terms of dealers' sharing rule. Namely, compared to partial strategy σ , in partial strategy σ^i , dealer i receives a strictly smaller share of the dealers' aggregate payoff, whereas all other dealers obtain a strictly larger share.

Formally, $\theta_i^i < \theta_i$, whereas $\theta_j^i > \theta_j$ for $j \neq i$. Second, let us define a punishment strategy. Observe that for any given dealer i and any distribution $\mu \in \Delta\Omega$, there exists $\underline{a}_{-i}^i(\mu) \in \times_{j \neq i} \Delta A_j$ such that

$$\max_{a_i} \sum_{\omega \in \Omega} \mu(\omega) u_i(\omega, a_i, \underline{a}_{-i}^i(\mu)) \leq 0.$$

Namely, let $v_\mu := E_\mu[\tilde{W}]$ be the expected value of the asset given the belief μ . Let $\underline{a}_{-i}^i(\mu)$ consist in dealers other than i setting $\alpha = \beta = v_{\mu\omega}$. Then, dealer i cannot hope to trade at a strictly positive expected profit given his beliefs μ . This guarantees that we can extend Blackwell (1956)'s approachability argument to the discounted case: for any $\eta > 0$, there is $\delta^\eta < 1$, $m^\eta < \infty$ and m^η -period strategy \underline{a}_{-i} for player $-i$ such that if $\delta > \delta^\eta$, for any sequence $\{a_i^1, \dots, a_i^{m^\eta}\}$, player i 's discounted payoff during these m^η periods is smaller than η in each $\omega \in \Omega$. This strategy is the ingredient for the punishment partial strategy $\underline{\sigma}^i$. Given any σ_i , any ω , and any history h^t , the continuation payoff $V_i^\delta(\omega, \sigma_i, \underline{\sigma}_{-i}^i | h^t)$ is such that

$$\lim_{\delta \rightarrow 1, \varepsilon \rightarrow 0} V_i^\delta(\omega, \sigma_i, \underline{\sigma}_{-i}^i | h^t) \leq 0. \quad (28)$$

Let $\underline{\sigma}^i$ consist of player i playing a fixed action \underline{a}_i and the other players using strategies $\underline{\sigma}_{-i}^i$.

The remainder of the proof is standard; see Fudenberg and Maskin (1986). Given the partial strategy σ , define a strategy $\hat{\sigma}$ as follows. As long as no player unilaterally deviates, actions are specified by σ . Once a player (say i) unilaterally deviates, play proceeds according to $\underline{\sigma}^i$ for T periods (for some $\varepsilon > 0$, $T \in \mathbb{N}$ to be specified). If during this i -punishment phase, some player (say j) unilaterally deviates from $\underline{\sigma}^i$, play switches to the j punishment phase, in which $\underline{\sigma}^j$ is played for T periods. If T periods elapse without unilateral deviation during the i -punishment phase, play is then given by σ^i , where the punishing dealers obtain a larger share of the future profits at the expense of the punished dealer i . If there is a unilateral deviation from $\sigma^{i,\varepsilon}$ by j , play switches to the j -punishment phase, etc. It is now standard to show that for a sufficiently

large T and sufficiently small ε , there exists $\underline{\delta} \in (\underline{\delta}, 1)$ such that for all $\delta \in (\underline{\delta}, 1)$, players do not gain from deviating.

Note that this construction yields a BFE: the strategy is optimal irrespective of dealers' beliefs about ω on and off the equilibrium path.

■

Appendix B

Proof of Lemma 8

1. *Positive maximum payoffs:* Take any given $\hat{\omega}$, and let us show that the set $A^*(\hat{\omega})_+$ is not empty. Fix dealer i and consider the following two action profiles $a(i)$ and $a'(i)$ in which all dealers other than i select the no-trade action. In $a(i)$, dealer i chooses his action such that if a trader trades, he can only buy at price $v(\hat{\omega}) + \rho$, with $0 < \rho < \underline{\rho}$ and he cannot sell. Conversely, in $a'(i)$, dealer i selects his action such that a trade can only consist of the trader selling at price $v(\hat{\omega}) - \rho$. Because of **RTF** 3. and 4., we have that for all $\omega \in \hat{\omega}$, the expected asset transfer to dealer i is equal to $q_i(\hat{\omega}, a(i)) < 0$ and $q_i(\hat{\omega}, a'(i)) > 0$ for action $a(i)$ and $a'(i)$, respectively. Now, let $z_i := q_i(\hat{\omega}, a'(i)) / (q_i(\hat{\omega}, a'(i)) - q_i(\hat{\omega}, a(i))) \in [0, 1]$ and consider the $\tilde{a}(i)$ obtained by playing $a(i)$ with probability z_i and $a'(i)$ with probability $1 - z_i$. This translates into an expected profit for dealer i of $2\rho z_i q_i(\hat{\omega}, a'(i)) > 0$ regardless of the value of $e(\omega)$. In fact, in expectation, he buys $z_i q(\hat{\omega}, a'(i))$ shares for $v(\hat{\omega}) - \rho$, and he sells the same quantity for $v(\hat{\omega}) + \rho$ per share. Now, consider the random strategy \tilde{a} obtained by first selecting a dealer i with probability $1/n$ and then playing $\tilde{a}(i)$. This guarantees that $u_i(\omega, \tilde{a}) = 2\rho z_i q_i(\hat{\omega}, a'(i))/n > 0$ for every i and every $\omega \in \hat{\omega}$, irrespective of the value of the $e(\omega)$ component. Thus, $\tilde{a} \in A_+^*(\hat{\omega})$.
2. *Negative minimum payoffs:* Fix dealer i and consider the action $\underline{a}(i)$ in which all dealers other than i select the no-trade action. In $\underline{a}(i)$, dealer i chooses his action such that if a

trader trades, he can only buy at a price strictly smaller than $v(\hat{\omega}) - \bar{e}$ and he cannot sell. Because of **RTF** 3., there will be a trader willing to buy at this price, implying that dealer i 's payoff is negative regardless of the true value of ω and hence for all $\omega \in \hat{\omega}$. Consider the random strategy $\underline{a}(\hat{\omega})$ obtained by first selecting a dealer i with probability $1/n$ and then playing $\underline{a}(i)$. Clearly, $u_i(\omega, \underline{a}(\hat{\omega})) < 0$ for all $\omega \in \hat{\omega}$ and dealer i .

3. *Non-positive expected payoffs:* Fix dealer i and a distribution $\mu_{\hat{\omega}} \in \Delta\hat{\omega}$. Let $W_{\mu_{\hat{\omega}}} := v(\hat{\omega}) + \sum_{\omega \in \hat{\omega}} \mu_{\hat{\omega}}(\omega) e(\omega)$ be the expected fundamental value of the asset computed using probability measure $\mu_{\hat{\omega}}$. Let p_1 and p_2 be the two points in the set G that are closest to $W_{\mu_{\hat{\omega}}}$, with $p_1 \leq W_{\mu_{\hat{\omega}}} \leq p_2$. Define $\underline{a}_{-i}^i(\mu_{\hat{\omega}})$ as follows. Each dealer $j \neq i$ selects an action such that any other market participant can buy and sell up to the maximum tradable quantity at price $p = W_{\mu_{\hat{\omega}}}$. Consider dealer i 's expected payoff when his belief that the state is ω is equal to $\mu_{\hat{\omega}}(\omega)$. His expected payoff from playing a_i when the other dealers play $\underline{a}_{-i}^i(\mu_{\hat{\omega}})$ is

$$\begin{aligned} \sum_{\omega \in \hat{\omega}} \mu_{\hat{\omega}}(\omega) u_i(\omega, a_i, \underline{a}_{-i}^i(\mu_{\hat{\omega}})) &= \sum_{\omega \in \hat{\omega}} \mu_{\hat{\omega}}(\omega) (v(\omega) + e(\omega)) q_i(\omega, a_i, \underline{a}_{-i}^i) + p_i(\omega, a_i, \underline{a}_{-i}^i) \\ &= W_{\mu_{\hat{\omega}}} q_i(\hat{\omega}, a_i, \underline{a}_{-i}^i) + p_i(\hat{\omega}, a_i, \underline{a}_{-i}^i), \end{aligned} \quad (29)$$

where the second equality follows from the fact that by the definition of $\hat{\omega}$, for any $\omega \in \hat{\omega}$ and $a \in A$, one has $v(\omega) = v(\hat{\omega})$, $q_i(\omega, a) = q_i(\hat{\omega}, a)$ and $p_i(\omega, a) = p_i(\hat{\omega}, a)$. The last expression can be interpreted as the payoff of a dealer who values the asset at exactly $W_{\mu_{\hat{\omega}}}$ and buys a quantity $q_i(\hat{\omega}, a_i, \underline{a}_{-i}^i)$ of the asset in exchange for $p_i(\hat{\omega}, a_i, \underline{a}_{-i}^i)$. To see that this expression cannot be strictly positive, note first that if a_i is such that dealer i trades with some other dealer, the other dealers' actions are such that he can only trade at price $x = W_{\mu_{\hat{\omega}}}$, implying that dealer i 's profit is nil. Suppose that a_i is such that dealer i trades with the trader. Because the trader can trade any quantity at price $x = W_{\mu_{\hat{\omega}}}$ from the

other dealers, he will trade with dealer i only if the latter offers better trading conditions, that is, only if he can buy from dealer i for less than $W_{\mu_{\hat{\omega}}}$ or sell for more than $W_{\mu_{\hat{\omega}}}$. In both cases, dealer i 's payoff from expression (29) cannot be strictly positive.

4. *Non-equivalent payoffs:* Consider the strategy $\tilde{a}(i)$ defined in point 1 above. When dealers play $\tilde{a}(i)$, dealer i 's payoff is positive, whereas all other dealers' payoffs are zero. Let $a^i(\hat{\omega})$ be obtained by first selecting a dealer $j \neq i$ with probability $1/(n-1)$ and then playing $\tilde{a}(j)$. Because in this strategy, dealer i 's payoff is nil, whereas all other dealers' payoffs are strictly positive, we have $u_i(\omega, a^i(\hat{\omega})) < u_i(\omega, a^j(\hat{\omega}))$ for all $i \neq j$ and $\omega \in \hat{\omega}$.

■

Proof of Proposition 2

The proof of 1. and 2. is analogous to the proof of Lemmas 3 and 4 and is omitted.

3. Because of property **RTF**, traders will never buy for more than $v(\hat{\omega}) + \bar{\rho}$ or sell for less than $v(\hat{\omega}) - \bar{\rho}$. Hence, the dealers' aggregate payoff cannot be greater than

$$(v(\hat{\omega}) + e(\omega) - (v(\hat{\omega}) - \bar{\rho}))Q^+(\hat{\omega}, \sigma^* | h^t) + (v(\hat{\omega}) + \bar{\rho} - (v(\hat{\omega}) + e(\omega)))Q^-(\hat{\omega}, \sigma^* | h^t),$$

which is non-negative only if

$$\bar{\rho}(Q^+(\hat{\omega}, \sigma^* | h^t) - Q^-(\hat{\omega}, \sigma^* | h^t)) > (Q^+(\hat{\omega}, \sigma^* | h^t) + Q^-(\hat{\omega}, \sigma^* | h^t))e(\omega).$$

Because each dealer's payoff is non-negative (see 2.), so is their aggregate payoff. Therefore, the above expression must be positive for all realizations of $e(\omega)$. Then, the result follows from the fact that Q , Q^+ and Q^- do not depend on $e(\omega)$ and that $e(\omega) \in [-\bar{e}, \bar{e}]$.

4. The proof follows the same logic as that for the baseline model. Because of the **RTF** properties, for any $\hat{\omega} \in \hat{\Omega}$, there is a transaction price $l(\omega)$ that induces a trader's order flow

that is positive and balanced. Thus, to maintain a balanced inventory, trading quotes cannot systematically differ from $l(\hat{\omega})$.

5. The proof follows the same logic as that for the base model. If for some history, the time required in state $\hat{\omega}$ for the average trading quote to be close to $l(\omega)$ is too long, then in this transition period, dealers' inventory would expand dramatically. This would contradict result 3 of this proposition.

■

Proof of Theorem 1

Fix a game and a profile (ϕ, σ) satisfying the assumptions of the theorem, and let ω be the true state. Let $\bar{u} := \max_{\omega, a, i} |u_i(\omega, a)|$. Consider the play on the equilibrium path. Let q^t be the probability that at time t the market measure satisfies $\pi^t(\hat{\omega}(\omega)) > 1 - \varepsilon$. Thus, following point 2 in Definition 2 and the definitions of u^* and \bar{u} , with probability q_t , dealer i 's stage- t payoff is at least $(1 - \varepsilon)u^* - \varepsilon\bar{u}$. Then, at time $\tau \geq 0$, dealer i 's payoff satisfies

$$\begin{aligned} V_i^\delta(\omega, \sigma \mid h^\tau) &> (1 - \delta) \sum_{t=\tau}^{\infty} \delta^{t-\tau} (q^t((1 - \varepsilon)u^* - \varepsilon\bar{u}) - (1 - q^t)\bar{u}) \\ &= (1 - \varepsilon)(u^* + \bar{u})(1 - \delta) \sum_{t=\tau}^{\infty} \delta^t q^{t-\tau} - \bar{u}. \end{aligned} \quad (30)$$

Condition 1 of Definition 2 implies that

$$\Pr_{\omega, \sigma} \left[\lim_{\delta \rightarrow 1} (1 - \delta) \sum_{t=\tau}^{\infty} \delta^t q^{t-\tau} > 1 - \varepsilon \right] > 1 - \varepsilon. \quad (31)$$

Hence, we have that

$$\lim_{\delta \rightarrow 1} V_i^\delta(\omega, \sigma \mid h^\tau) > (1 - \varepsilon)^3(u^* + \bar{u}) - (1 + \varepsilon)\bar{u}. \quad (32)$$

As the r.h.s. is strictly positive for $\varepsilon = 0$, it is also positive for all ε smaller than some $\bar{\varepsilon} > 0$. The

continuity of V_i^δ in δ implies that there exists $\underline{\delta} < 1$ such that for $\varepsilon < \bar{\varepsilon}$, dealer i 's continuation payoff $V_i^\delta(\omega, \sigma \mid h^\tau)$ is strictly positive.

The next step is to show that dealers have no profitable deviations. For this purpose, we first establish a simple lemma.

Lemma 11 *For any given $\hat{\omega} \in \hat{\Omega}$, all $\omega \in \hat{\omega}$ and any player i , and any $a \in A^*(\hat{\omega})$, there exist n action profiles $\{\tilde{a}^1(\hat{\omega}), \dots, \tilde{a}^n(\hat{\omega})\} \in [\Delta A]^n$ such that*

$$0 < u_i(\omega, \tilde{a}^i(\hat{\omega})) < u_i(\omega, \tilde{a}^j(\hat{\omega})) < u(\omega, a), \quad (33)$$

for all $i \neq j$.

Proof. Consider the convex combination

$$\tilde{a}^i(\hat{\omega}) := \beta_1(\hat{\omega})\beta_2(\hat{\omega})\underline{a}(\hat{\omega}) + \beta_1(\hat{\omega})(1 - \beta_2(\hat{\omega}))a^i(\hat{\omega}) + (1 - \beta_1(\hat{\omega}))a, \quad (34)$$

for some $\beta_1(\hat{\omega}), \beta_2(\hat{\omega}) \in [0, 1]$, where $\underline{a}(\hat{\omega})$ satisfies Definition 4.2 and $a^i(\hat{\omega})$ is as in Definition 4.4. Note that $\{\tilde{a}^i(\hat{\omega})\}_{i=1, \dots, n}$ also satisfies Definition 4.4, as long as $\beta_1(\hat{\omega}) > 0$, $\beta_2(\hat{\omega}) < 1$. Furthermore, because $u(\omega, \underline{a}(\hat{\omega})) < 0$, we can set $\beta_2(\hat{\omega})$ close enough to one and $\beta_1(\hat{\omega})$ close enough to zero to guarantee that all payoffs are between 0 and $u(\omega, a)$. ■

We may now define n partial strategy profiles $\sigma^{i, \varepsilon}$ as follows. Let A_L denote a set of learning action profiles satisfying $A(\hat{\omega}, \hat{\omega}') \cap A_L \neq \emptyset$ for each pair $\hat{\omega} \neq \hat{\omega}'$. Let L denote the cardinality of A_L and $D_{\hat{\omega}}$ denote the Dirac measure attaching probability 1 to $\hat{\omega}$. If h^t is such that $\|\pi^t - D_{\hat{\omega}}\| < \varepsilon$, then let $\sigma^{i, \varepsilon}(h^t) = (1 - \varepsilon)\tilde{a}^i(\hat{\omega}) + (\varepsilon/L)\sum_{a \in A_L} a$. For all other h^t , let $\sigma^{i, \varepsilon}(h^t) = (1/L)\sum_{a \in A_L} a$.

In addition, define n partial “punishment” strategies $\underline{\sigma}^{i, \varepsilon}$ as follows. Fix any $\hat{\omega} \in \hat{\Omega}$. Lemma 8 guarantees that there are strategies satisfying Definition 4.3, and hence, we can extend the Blackwell (1956) approachability argument to the discounted case: for any $\eta > 0$ there is $\delta^\eta < 1$,

$m^\eta < \infty$ and m^η -period strategy $\underline{a}_{-i}(\hat{\omega})$ for player $-i$ such that if $\delta > \delta^\eta$, for any sequence $\{a_i^1, \dots, a_i^{m^\eta}\}$, player i 's discounted payoff during these m^η periods is smaller than η for each $\omega \in \hat{\omega}$. This Blackwell strategy is then an ingredient of the punishment partial strategy $\underline{\sigma}^{i,\varepsilon}$. If h^t is such that, for some $\hat{\omega}_i$, π^t assigns probability no more than ε to states outside of $\hat{\omega}_i$ but a probability of at least ε to all $\omega \in \hat{\omega}_i$, then $\underline{\sigma}^{i,\varepsilon}(h^t) = (1 - \varepsilon) \underline{a}^i(\hat{\omega}_i)(h^t) + (\varepsilon/L) \sum_{a \in A_L} a$, where $\underline{a}_{-i}^i(\hat{\omega}_i)(h^t)$ as defined above and $\underline{a}_{-i}^i(\hat{\omega}_i)$ is some fixed action. Note that, for $\varepsilon > 0$, each of these strategies is exploratory. Furthermore, given any σ_i , any ω , and any history h^t , the continuation payoff $V_i^\delta(\omega, \sigma_i, \underline{\sigma}_{-i}^{i,\varepsilon} | h^t)$ is such that

$$\lim_{\delta \rightarrow 1, \varepsilon \rightarrow 0} V_i^\delta(\omega, \sigma_i, \underline{\sigma}_{-i}^{i,\varepsilon} | h^t) \leq 0. \quad (35)$$

From here, the proof is standard; see Fudenberg and Maskin (1986). Given the partial strategy σ , define a strategy $\hat{\sigma}$ as follows. As long as no player unilaterally deviates, actions are specified by σ . Once a player (say i) unilaterally deviates, play proceeds according to $\underline{\sigma}^{i,\varepsilon}$ for T periods (for some $\varepsilon > 0$, $T \in \mathbb{N}$ to be specified). If during this i -punishment phase, some player (say j) unilaterally deviates from $\underline{\sigma}^{i,\varepsilon}$, play switches to the j -punishment phase, in which $\underline{\sigma}^{j,\varepsilon}$ is played for T periods. If T periods elapse without unilateral deviations during the i -punishment phase, play is then given by $\sigma^{i,\varepsilon}$. If there is a unilateral deviation from $\sigma^{i,\varepsilon}$ by j , play switches to the j -punishment phase, etc. It is now standard to show that, for a large enough T and small enough ε , there exists $\underline{\delta} \leq \bar{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$, players do not gain from deviating.

Note that this construction yields a BFE: the strategies are optimal irrespective of dealers' beliefs about ω on and off the equilibrium path.

■

References

- [1] Atkeson, A., A. Eisfeldt, and P. Weill, 2013. “The Market for OTC Derivatives,” *mimeo*.
- [2] Back, K., and S. Baruch, 2004. “Information in Security Markets: Kyle meets Glosten and Milgrom,” *Econometrica*, **72**, 433–465.
- [3] Benveniste, L., A. Marcus, and W. Wilhelm, 1992. “What’s so Special about the Specialist?,” *Journal of Financial Economics*, **32**, 61–86.
- [4] Bessembinder, H., and W. Maxwell, 2008. “Markets: Transparency and Corporate Bond Market,” *Journal of Economic Perspectives*, **22**, 217–234.
- [5] Biais B., D. Martimort, and J-C. Rochet, 2000. “Competing Mechanisms in a Common Value Environment,” *Econometrica*, **68**, 799–837.
- [6] Blackwell, D., 1956. “An Analog of the Minimax Theorem for Vector Payoffs,” *Pacific Journal of Mathematics*, **6**, 1–8.
- [7] Bloomfield, R., and M. O’Hara, 2000. “Can Transparent Markets Survive?,” *Journal of Financial Economics*, **55**, 425–459.
- [8] Boehmer, E., and J. Wu, 2008. “Order Flow and Prices,” *mimeo*.
- [9] Boulatov, A., and T. George, 2013. “Securities Trading when Liquidity Providers are Informed,” *Finance and Accounting Memos*, forthcoming.
- [10] Calcagno, R., and S. Lovo, 2006. “Bid-ask Price Competition with Asymmetric Information between Market-makers,” *Review of Economic Studies*, **73**, 329–355.
- [11] Calvet, L., and A. Fisher, 2004. “How to Forecast Long-Run Volatility: Regime-Switching and the Estimation of Multifractal Processes,” *Journal of Financial Econometrics*, **2**, 49–83.
- [12] Carlin, B., M. Lobo, and S. Viswanathan, 2007. “Episodic Liquidity Crises: Cooperative and Predatory Trading,” *Journal of Finance*, **62**, 2235–2274.
- [13] Chamley, C., 2004. *Rational Herds*, Cambridge University Press.

- [14] Chordia, T., R. Roll, and A. Subrahmanyam, 2002. "Order imbalance, liquidity and market returns," *Journal of Financial Economics*, **65**, 111–130.
- [15] Christie, W., and P. Schultz, 1994. "Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?," *Journal of Finance*, **49**, 1813–1840.
- [16] Christie, W., J. Harris, and P. Schultz, 1994. "Why did NASDAQ Market Makers Stop Avoiding Odd-Eighth Quotes?," *Journal of Finance*, **49**, 1841–1860.
- [17] Desgranges, G., and T. Foucault, 2005. "Reputation-Based Pricing and Price Improvements," *Journal of Economics and Business*, **57**, 493–527.
- [18] De Frutos, M., and C. Manzano, 2005. "Trade Disclosure and Price Dispersion," *Journal of Financial Markets*, **8**, 183–216.
- [19] De Meyer, B., 2010. "Price Dynamics on a Stock Market with Asymmetric Information," *Games and Economic Behavior*, **69**, 42–71.
- [20] Du, S., and H. Zhu, 2012. "Ex-post Equilibria in Double Auctions of Divisible Assets," *mimeo*.
- [21] Dutta, P., and A. Madhavan, 1997. "Competition and Collusion in Dealers Markets," *Journal of Finance*, **52**, 1817–1846.
- [22] Ellis, K., R. Michaely, and M. O'Hara, 2002. "The Making of a Dealer Market: From Entry to Equilibrium in the Trading of Nasdaq Stocks," *Journal of Finance*, **57**, 2289–2316.
- [23] Evans, M., and R. Lyons, 2002. "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy*, **110**, 170–180.
- [24] Fleming, J., C. Kirby, and B. Ostdiek, 2006. "Information, Trading, and Volatility: Evidence from Weather-Sensitive Markets," *Journal of Finance*, **61**, 2899–2930.
- [25] Fudenberg, D., and E. Maskin, 1986. "The Folk Theorem for Repeated Games with Discounting or with Incomplete Information," *Econometrica*, **54**, 533–554.
- [26] Foucault, T., 1999. "Order Flow Composition and Trading Costs in a Dynamic Limit Order Market," *Journal of Financial Markets*, **2**, 99–134.

- [27] Foucault, T., O. Kadan, and E. Kandel, 2005. “Limit Order Book as a Market for Liquidity,” *Review of Financial Studies*, **18**, 1171–1217.
- [28] Fudenberg, D., and Y. Yamamoto, 2010. “Repeated Games Where the Payoffs and Monitoring Structure are Unknown,” *Econometrica*, **78**, 1673–1710.
- [29] Fudenberg, D., and Y. Yamamoto, 2011. “Learning from Private Information in Noisy Repeated Games,” *Journal of Economic Theory*, **146**, 1733–1769.
- [30] Glosten, L., and P. Milgrom, 1985. “Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders,” *Journal of Financial Economics*, **14**, 71–100.
- [31] Goldstein, M., and E. Hotchkiss, 2007. “Dealer Behavior and the Trading of Newly Issued Corporate Bonds.” Available at SSRN: <http://ssrn.com/abstract=1022356>.
- [32] Goettler, R., C. Parlour, and U. Rajan, 2005. “Equilibrium in a Dynamic Limit Order Market,” *Journal of Finance*, **60**, 2149–2192.
- [33] Goettler, R., C. Parlour, and U. Rajan, 2009. “Informed Traders in Limit Order Markets,” *Journal of Financial Economics*, **93**, 67–87.
- [34] Green, R., B. Hollifield, and N. Schüroff, 2007. “Dealer intermediation and price behavior in the aftermarket for new bond issues,” *Journal of Financial Economics*, **86**, 643–682.
- [35] Green, R., D. Li, and N. Schüroff, 2010. “Price Discovery in Illiquid Markets: Do Financial Asset Prices Rise Faster Than They Fall?,” *Journal of Finance*, **65**, 1669–1702.
- [36] Hansch, O., N. Naik, and S. Viswanathan, 1998. “Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange,” *Journal of Finance*, **53**, 1623–1656.
- [37] Harris, L., and M. Piwowar, 2006. “Secondary Trading Costs in the Municipal Bond Market,” *Journal of Finance*, **61**, 1361–1397.
- [38] Hasbrouck, J., and G. Sofianos, 1993. “The Trades of Market Makers: An Empirical Analysis of NYSE Specialists,” *Journal of Finance*, **48**, 1565–1593.
- [39] Hendershott T., and A. Menkveld, 2014. “Price Pressure,” *Journal of Financial Economics*, **114**, 405–423.

- [40] Heston, S., 1993. “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options,” *Review of Financial Studies*, **6**, 327–343.
- [41] Hörner, J., and S. Lovo, 2009. “Belief-free Equilibria in Games with Incomplete Information,” *Econometrica*, **77**, 453–487.
- [42] Hörner, J., S. Lovo and T. Tomala, 2011. “Belief-free Equilibria in Games with Incomplete Information: Characterization and Existence,” *Journal of Economic Theory*, **146**, 1770–1795.
- [43] Huang, R., and H. Stoll, 1996. “Dealer Versus Auction Markets: A Paired Comparison of Execution Costs on Nasdaq and the NYSE,” *Journal of Financial Economics*, **41**, 313–357.
- [44] International Swaps and Derivative Association, 2010. *ISDA Market Surveys 1995-2010*, Available at <http://www2.isda.org/attachment/MjAyNW==/ISDA%20Market%20Surveys%202010-1995.docx>.
- [45] International Swaps and Derivative Association, 2013. *2013 Operations Benchmarking Survey*, Available at <http://www2.isda.org/attachment/NTUzOQ==/OBS%202013%20FINAL%200425.pdf>.
- [46] Kandel, E., and L. Marx, 1997. “Nasdaq Market Structure and Spread Patterns,” *Journal of Financial Economics*, **45**, 61–89.
- [47] Kyle A., 1985. “Continuous Auctions and Insider Trading,” *Econometrica*, **53**, 1315–1335.
- [48] LeRoy, S., and L. Porter, 1981. “The Present Value Relation: Tests based on Implied Variance Bonds,” *Econometrica*, **49**, 555–574.
- [49] Li, D., and N. Schüroff, 2014. “Dealer Networks,” Available at SSRN: <http://ssrn.com/abstract=2023201>.
- [50] Madhavan, A., and S. Smidt, 1993. “An Analysis of Daily Changes in Specialist Inventories and Quotations,” *Journal of Finance*, **48**, 1595–1628.
- [51] Milgrom, P., and N. Stokey, 1982. “Information, Trade and Common Knowledge,” *Journal of Economic Theory*, **26**, 17–27.

- [52] Moussa Saley, H., and B. De Meyer, 2003. “On the Strategic Origin of Brownian Motion in Finance,” *International Journal of Game Theory*, **31**, 285–319.
- [53] Naik, N., and P. Yadav, 2003. “Do Dealer Firms Manage Inventory on a Stock-by-Stock or a Portfolio Basis?,” *Journal of Financial Economics*, **69**, 325–353.
- [54] Pasquariello, P., and C. Vega, 2005. “Informed and Strategic Order Flow in the Bond Market,” *mimeo*.
- [55] Rosu I., 2009. “A Dynamic Model of the Limit Order Book,” *The Review of Financial Studies*, **22**, 4601–4641.
- [56] Reiss, P., and I. Werner, 1998. “Does Risk Sharing Motivate Interdealer Trading?,” *Journal of Finance*, **53**, 1657–1703.
- [57] Roëll, A., 1988. “Regulating Information Disclosure Among Stock Exchange Market Makers,” FMG Discussion Papers Series 51, London School of Economics.
- [58] Shiller R., 1981. “Do Stock Prices Move too much to be Justified by Subsequent Changes in Dividends?,” *The American Economic Review*, **71**, 421–426.
- [59] Weston, J., 2000. “Competition on the Nasdaq and the Impact of Recent Market Reforms,” *Journal of Finance*, **55**, 2565–2598.