Belief-free Price Formation*

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

December 11, 2012

Abstract

We analyze security price formation in a dynamic setting in which long-lived dealers repeatedly compete for trading with potentially informed retail traders. For a class of market microstructure models, we characterize equilibria in which dealers' dynamic pricing strategies are optimal no matter the private information each dealer may possess. In a generalized version of the Glosten and Milgrom model, these equilibria deliver price dynamics reminiscent of well-known stylized facts: price/trading-flow correlation, volatility clustering, price bubble and inventory/inter-dealer trading correlation.

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

JEL Codes: G1, G12, C72, C73.

*We thank Vincent Fardeau, Thierry Foucault and Ioanid Rosu for useful comments, as well as seminar audience at HEC, EDHEC, Warwick Business School, SEAT 2012 meeting in Brisbane and EFA 2012 Meeting in Copenhagen. This paper was previously circulated under the title "Belief-free Market Making." S. Lovo and T. Tomala gratefully acknowledge financial support from the HEC Foundation. J. Hörner gratefully acknowledges financial support from NSF Grant SES 092098.

[†]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 interact in a market that is open to short-lived market participants ("traders"). We characterize equilibria that are robust to any form of asymmetry of information among dealers.

It has been claimed in the market-microstructure empirical literature (see for instance Ellis, Michaely, and O'Hara (2002)) that dealers have access to different sources of information and that they need not be well aware of other dealers' sources of private information. However, in existing market microstructure models, tractability imposes strong informational assumptions, and specific functional assumptions regarding the distribution of fundamentals and private signals.[1] Because modeling dynamic interaction among asymmetrically informed dealers can be a formidable task, the theory is silent about the robustness of canonical microstructure theory predictions to changes in the dealers' information environment.[2] In practice, dealers' actual information structures are not directly observable, so that it is usually impossible to assess the extent to which a given model's assumptions on information structure reflect "real world" informational asymmetries.

The objective of this paper is to provide a tractable price-formation theory delivering predictions that are robust to details in the information structure. To this purpose, we consider a class of dynamic financial markets microstructure models in which risk-neutral financial intermediaries (such as dealers or market-makers) interact with traders. For this class of models, we characterize equilibria in which dealers' dynamic pricing strategies remain optimal no matter the private information of a dealer about the economy fundamentals (so called "belief-free equilibria",

[1] For instance, almost all models assume that trading prices are set by equally uninformed dealers to a level reflecting these dealers' beliefs on fundamentals.

[2] At any point in time each dealer anticipates how its behavior affects its current expected payoff as well as each competing dealer's posterior beliefs and future behavior. The problem is even more complex if a dealer is not certain about its competitors' prior beliefs.

BFE henceforth).[3]

There is no *a priori* reason why belief-free equilibria are more compelling than, for example, the "classical" zero-profit equilibrium. Nevertheless, we believe that these equilibria are interesting for a number of reasons. First, in terms of their scope: as we show, in a belief-free equilibrium, hardly any assumption on the dealers' information is called for. This seems more realistic than assuming that all dealers share the same exact beliefs about fundamentals. Also minimal assumptions are required concerning trading protocols, and the model is flexible enough to encompass many real-world trading protocols. Second, in terms of their ability to explain seemingly unrelated empirical findings: we present an example that supports various stylized facts: price-volume correlation, volatility clustering, price bubbles, and inventory-trading correlations. While each of these facts can be explained by some models, none delivers them simultaneously. Third, in terms of tractability: we actually focus on a subset of belief-free equilibria that are arguably as tractable as the classical zero-profit models in market microstructure. Finally, in the presence of multiple equilibria, it might be sensible for dealers to coordinate on equilibria generating positive profits. This is a feature of all belief-free equilibria.

All BFE equilibria enjoy the following properties: 1) Dealers can gain or lose money in the short run, but their long-run profit is strictly positive independently of the asset's fundamental value. This contrasts with the traditional prediction that dealers' expected per trading period profit is nil; 2) Risk-neutral dealers tend to maintain balanced inventories and make profits trough the intermediation of traders' order flow. This contrasts with the view that (absent risk aversion or institutional constraints on inventory size) inventory levels should not affect dealers' behavior.[4] Also, our finding de-emphasizes the role of information (about the asset) on dealers' behavior. In a BFE, what matters for a dealer is the level of quotes that induce an abundant but balanced

[3]See 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.

[4]For instance, in Ho and Stoll (1981) balanced inventory results from dealers' risk-aversion, whereas in Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2009) it results from the dealers' institutional inability to take a position beyond a certain size. Our model displays neither factor.

order flow from traders. Hence dealers can ignore all information that does not affect traders' behavior. 3) As in the canonical market microstructure models, movement in asset quotes are caused by the public information provided by the trading order flow. However, unlike in models in which quotes reflect beliefs about fundamentals, it is not the case that given a sufficient amount of public information, an asset trading price eventually stabilizes around its fundamental value. In other words, equilibrium quotes need not reflect any of the dealers' (Bayesian) belief, and price sensitivity to trading volume does not fade away as public information accumulates. Thus, long-term price volatility remains large even without exogenous shocks on fundamentals.

The rationale behind these results is as follows. First, dealers can always guarantee zero profit by abstaining from trading. Because in equilibrium dealers' strategy must be optimal no matter the dealer's belief about fundamentals, each dealer's equilibrium long-term profit must be positive for each possible value of fundamentals. Second, given the range of possible asset values, a strategy leading to a sufficiently unbalanced inventory would correspond to a negative value portfolio for some level of the asset fundamentals, and hence for some level of a dealer's belief about fundamentals. On the contrary, when the equilibrium strategy leads to sufficiently balanced inventories, the asset fundamental value will have little impact on dealers' profits. As a result, in a BFE, dealers' long term profit must mainly result from intermediation of traders' order flow. This is achieved through (what we refer to as) "exploiting periods" during which dealers set quotes prompting a balanced order flow and make positive profit from the bid-ask spread. Third, because the specific strategies that dealers adopt during exploiting periods depend on the fundamentals, dealers' equilibrium strategies must also display "exploring periods." During an exploring period, dealers' quotes prompt informative order flow from traders. Quotes react to the order flow, which then eventually provides enough information about the quoting strategy to be followed during exploiting periods. Because dealers might lose money during exploring periods, exploring phases cannot last very long, and while they point to the right exploiting strategy most often than not, with low probability they also lead to incorrect exploiting rounds.

Hence, a Bayesian dealer could possibly disagree with the consensus view during an exploiting phase. For such a dealer not to deviate, it must be that he expects the flow of public information to correct this view rapidly. Hence, unlike Bayesian beliefs that take arbitrarily long to budge once they are sufficiently degenerate, belief-free equilibrium prices must be sensitive to the order flow at all times –hence, they cannot simply reflect Bayesian beliefs about fundamentals. As a result, exploiting phases must always alternate with exploring phases, and quote sensitivity to order flow cannot fade away.

In the first part of the paper, we illustrate the functioning of such belief-free equilibria in the simple framework of the Glosten and Milgrom (1985) model, modified in two respects. First, to the standard quote driven market open to dealers making a market to traders, we add an inter-dealer market only accessible to dealers. Second, we make no assumption about dealers' private information.

The notion of belief-free equilibrium is more demanding than standard game-theoretic refinements. Therefore, it imposes stronger restrictions on equilibrium outcomes. Despite this, there remains considerable leeway in their specification. This flexibility suffices to explain some regularities documented in the empirical literature. First, because the flow of trade and any other relevant public information is used by dealers as a coordination device, movements in asset quotes are caused by the public information provided by the order flow or possibly the exogenous arrival of news. This is consistent with a wide body of empirical work, spanning from security markets (see for instance Chordia, Roll and Subrahmanyam (2002) and Boehmer and Wu (2008)), bond markets (Pasquariello and Vega (2005)), currency markets (Evans and Lyons (2002)), weather-sensitive commodity markets (Fleming, Kirby and Ostdiek (2006)), etc. Second, the alternation of exploring phases and exploiting phases implies that periods of high price volatility follow periods of low volatility, that is, high-volatility events tend to cluster in time.[5] Exploring phases attract informed traders, leading to quotes that are highly sensitive to

[5]See Cont (2001) for a discussion of volatility clustering and other stylized facts.

the volume of trade. In exploiting phases, trading flow is balanced and originates from liquidity traders; as a result quote volatility is reduced. Third, the inter-dealer market is used as a tool to redistribute among dealers the profits and losses they make against traders. As a consequence, trades in the quote driven market are a predictor of the trade a dealer subsequently makes in the inter-dealer market. That is, a dealer buying (selling) in the quote driven market is more likely to later sell (resp. buy) to other dealers in the inter-dealer market. This provides an explanation to the finding that dealers use the inter-dealer market to re-balance their inventory (see for instance Hasbrouck and Sofianos (1993), Reiss and Werner (1998) and (2005), Hansch, Naik and Viswanathan (1998), Evans and Lyons (2002)). Broadly speaking, dealers prefer sharing the profits that result from imperfect competition to competing them away.[6] As a result, whereas in the short-run dealers might gain or lose money, in the long run, they achieve positive profits. In fact, long-run profits are positive not only on average but also *ex post*, i.e., independently of the fundamental value of the asset.

Compared to the zero-profit equilibrium (GME henceforth) described by Glosten and Milgrom (1985), both explain correlation between trading volume and price changes. However, whereas in a GME dealers set quotes equal to the asset's expected value given past and current public information, so that expected per-period profits are zero, neither property holds in a BFE. These traditional predictions are not robust to changes in the dealers' information structure. A GME does not predict volatility clustering, nor does it explain inter-dealer trading. Furthermore, while a GME only applies to situations in which dealers are symmetrically uninformed, a BFE remains an equilibrium no matter the extent of information asymmetries among dealers.

In the second part of the paper, we consider a broader class of market microstructure models, in which some long-run market participants ("dealers") repeatedly interact in a market that is open to short-run market participants ("traders"). The class of models we analyze is broad along

[6] While our model focuses on the market for a single asset, its underlying logic of dealers' shared profits carries over to markets for different assets. This explains the finding by Ellis et al. (2002) that each market has a dominant dealer who makes most of the profit.

a number of dimensions. First it encompasses different trading protocols. Second, it comprises both fundamental uncertainty (i.e., uncertainty about the fundamental value of the asset) and non-fundamental uncertainty (for instance, uncertainty about the fraction of informed traders in the economy, the precision of their signals or traders' preferences). Third, within a given trading protocol and type of uncertainty, all specifications of asymmetries of information among dealers are covered.

We show that a dynamic trading game admits belief-free equilibria as long as the static game describing one trading round satisfies four simple conditions. Loosely speaking, for any given value of the fundamentals that is statistically learnable from the traders' behavior: first, there exists a way for dealers to earn a positive profit; second, there also exists a way to lose money; third, dealers have a way to "punish" a dealer in case of an observable deviation. The fourth condition is more technical but obtains in particular whenever inter-dealer trading is allowed.

Whenever the trading game satisfies these conditions, and discounting between consecutive trading rounds is sufficiently low, a "folk theorem" type result holds: for any given candidate dealers' strategy profile displaying the exploring and the exploiting features (as in the illustrative example), there is a belief-free equilibrium whose outcome coincides with the candidate profile.

What matters for our construction is that long-run market participants can use public information to coordinate on mutually profitable actions. Financial intermediaries that repeatedly meet in a stock market possess this ability. Implicit collusion in the stock market has been documented by Christie and Schultz (1994), Christie, Harris and Schultz (1994) and Ellis, et al. (2002). Evidence of long-term relationships is reported by Battalio, Ellul and Jennings (2007). In the market microstructure literature, Dutta and Madhavan (1997) model implicit collusion among dealers, while Benveniste et al. (1992) and Desgranges and Foucault (2005) analyze long-term relationships. These papers assume either no informational asymmetry, or short-lived informational asymmetries. Here instead, the state of nature is chosen once and for all, so that a dealer owning some private information might possibly take advantage of it over a long horizon.[7]

Few theoretical papers analyze the effect of asymmetric information among dealers. Even fewer do so within a dynamic framework. Some static examples in which dealers, or more generally liquidity providers, are asymmetrically informed are Roël (1988), Bloomfield and O'Hara (2000), de Frutos and Manzano (2005) and Boulatov and George (2010). 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 findings are sensitive to the precise assumptions regarding the dealers' information. Du and Zhu (2012) results 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 leading to a belief-free equilibrium.

The paper is organized as follows. Section 1 presents an example based on the model of Glosten and Milgrom. Section 2 develops the main model and states the central result. Section 3 discusses extensions to imperfect monitoring about dealers' actions, non-stationary states of nature and dealers' strategies based on private information. Section 4 concludes. All proofs are in Appendix.

1 A Model of Price Formation

In this section, we illustrate the definition, the logic and the main features of a belief-free equilibrium. The purpose of this section is not to construct a model that replicates all institutional features of a specific existing market, but rather to illustrate the logic of belief-free equilibria in a simple, well-known financial market microstructure framework à la Glosten and Milgrom (1985). In Section 3, we illustrate how the equilibrium strategy can be modified to account for

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

some institutional features of real markets that are ignored in the baseline model.

Set-up: A risky asset is exchanged for money among short-lived traders and $n > 1$ long-lived dealers (n is finite). Trading takes place over infinitely many periods $t = 1, 2, \ldots$ At time 0, Nature chooses the state ω in the set Ω (with at least four elements). The asset's fundamental value is $W(\omega) = v(\omega) + \psi(\omega)$, where $v(\omega) \in \{v_1, v_2\}$ and $\psi(\omega) \in \{\underline{e}, \overline{e}\}$, with $\underline{e} < 0 < \overline{e}$ and $0 < v_1 < v_2$. Thus, $W(\omega)$ takes values in $\{v_1 + \underline{e}, v_1 + \overline{e}, v_2 + \underline{e}, v_2 + \overline{e}\}$. As in Back and Barush (2004), a public release of information takes place at a random time θ, and conditional on it not having occurred yet, the probability that it occurs in the next period is constant. After the public announcement, all dealers' positions are liquidated at price $W(\omega)$.

The stage trading game: Each trading round t unfolds as follows. First, all dealers simultaneously choose their actions in an inter-dealer market and post their bid and ask quotes in a quote driven market. The inter-dealer market is closed to traders.[8] Second, a trader randomly arrives in the quote driven market, observes dealers' quotes, decides whether to trade or not one unit of the asset with dealers and then leaves the market.

As far as the inter-dealer market is concerned, it is sufficient to focus on its reduced form.[9] We denote by A_i^{ID} the finite set of actions available to dealer i in the inter-dealer market in any given period t. Let $A^{ID} := \times_i A_i^{ID}$ denote the set of dealers' action profiles in the inter-dealer market. Given $a^{ID} \in A^{ID}$, let $Q_i^{ID}(a^{ID}) \in \mathbb{R}$ and $P_i^{ID}(a^{ID}) \in \mathbb{R}$ denote the resulting net transfers of the asset and cash, respectively, from other dealers to dealer i. Note that

$$\sum_i Q_i^{ID}(a^{ID}) = \sum_i P_i^{ID}(a^{ID}) = 0, \tag{1}$$

meaning that the inter-dealer market leads to cash and asset redistribution among dealers. We

[8] We can dispense with the introduction of an inter-dealer market. In this case, transactions across dealers occur through the quote driven market.

[9] See for instance Evans and Lyons (2002) for a specification of FX inter-dealer market.

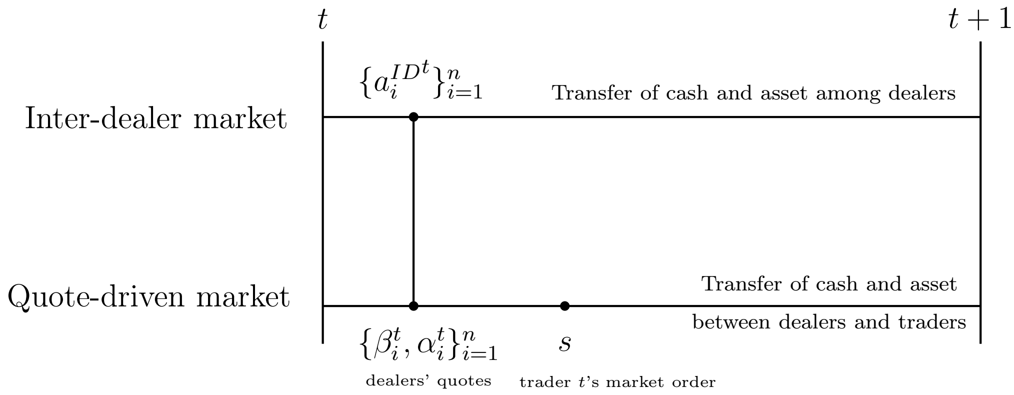

Figure 1: Stage trading round.

assume that there exists a *no trade action* $\underline{a}_i^{ID} \in A_i^{ID}$ allowing a dealer i to abstain from trading in this market. Also, assume that there are $\overline{Q} > 0$ and $\overline{P} > 0$, large but finite, such that any vector of (Q^{ID}, P^{ID}) of inter-dealer transfers satisfying (1) and $|Q_i^{ID}| \leq \overline{Q}$, $|P_i^{ID}| \leq \overline{P}$ for all i can be attained with an appropriate (possibly mixed) action profile $a^{ID} \in \Delta A^{ID}$.[10]

Let α_i^t and β_i^t be dealer i's bid and ask quotes posted in the quote driven market in period t, respectively. Quotes belong to a finite grid G of non-negative prices whose largest (smallest) element is larger than $v_2 + \overline{e}$ (resp. smaller than $v_1 - \underline{e}$). Let $\beta^t := \max_i \beta_i^t$ and $\alpha^t := \min_i \alpha_i^t$ denote the best bid and ask quotes in period t. Overall, an action profile for dealers in one trading round specifies for each dealer the action that it takes in the inter-dealer market as well as the bid and ask quotes it posts in the quote driven market. Formally, the set of dealers' action profiles is $A := A^{ID} \times G^n \times G^n$. We make no assumption about the private information of any given dealer regarding the true state ω, nor about the distribution of $W(\omega)$. In a belief-free equilibrium, defined below, each dealer's strategy must be optimal for each realization of the state ω, so the profile constitutes a (subgame-perfect) equilibrium independently of the presence or the extent of informational asymmetries among dealers and of dealers' beliefs about $W(\omega)$.

[10]Here and in what follows, ΔB denotes the set of distributions over a finite set B.

Assumptions: Traders have information about ω, but they also come to the market for liquidity reasons unrelated to ω. Namely, we posit that traders' information is about the $v(\omega)$ component of $W(\omega)$, but not the $\psi(\omega)$ component. Let $a^t \in A$ be any given dealers' action profile. Then $F(\omega, a^t, s) \in [0, 1]$ denotes the probability that, after observing a^t, trader t chooses action

$$s \in S := \{\textit{buy, sell, no trade, buy and sell}\},$$

given that the state is ω. Because traders have no information about the value of the $\psi(\omega)$ component of the asset fundamental, their behavior cannot depend on ψ. Thus, with some abuse of notation, assume:

Non-learnable States (NLS): *If $v(\omega) = v(\omega')$, then $F(\omega, a^t, s) = F(\omega', a^t, s) =: F(v(\omega), a^t, s)$ for all $a^t \in A$ and $s \in S$.*

By contrast, traders' demand may depend on the realization of $v(\omega)$. That is, there exists a subset $A(v_1, v_2) \subseteq A$ of dealers' action profiles for which traders' behavior is sensitive to the value of $v(\omega)$. Formally:

Learnable States (LS): *There is a non-empty set $A(v_1, v_2) \subseteq A$ such that, if $a^t \in A(v_1, v_2)$, then*

$$F(v_1, a^t, s) \neq F(v_2, a^t, s),$$

for some $s \in S$.

Assumption **LS** states that because some traders might have private information about the $v(\omega)$ component of the asset value $W(\omega)$, this component affects both the asset's liquidation value and traders' behavior. More precisely, there are suitable choices of dealers' actions (i.e., for $a^t \in A(v_1, v_2)$), for which the distribution of traders' reactions s is measurable with respect to $v(\omega)$. In general, this obtains as long as the bid-ask spread $\alpha^t - \beta^t$ is neither too large to induce no trading, nor negative so as to induce arbitrage trading. Thus, as long as a^t is in $A(v_1, v_2)$, a sufficiently long history of the traders' order flow allows to statistically tell apart whether $v(\omega) = v_1$ or $v(\omega) = v_2$.

To the contrary, Assumption **NLS** means that the $\psi(\omega)$-component of the asset value $W(\omega)$ affects the asset liquidation value, but because no trader is informed of it, the way traders react to dealers' quotes does not depend on the true $\psi(\omega)$. As a consequence, trading flow does not allow to tell apart whether $\psi(\omega) = \underline{e}$, or $\psi(\omega) = \overline{e}$. That is to say, **LS** implies that $v(\omega)$ is *statistically learnable* from traders' behavior, whereas **NLS** implies that $\psi(\omega)$ is not.

We also assume that traders never buy the asset at a price that is too high nor sell at a price that is too low. The range of prices at which traders do trade depends on traders' private information and hence on the level of $v(\omega)$. Formally,

Elastic Trader Demand (ETD): *There is* $\rho > 0$ *such that*

$$\begin{aligned} F(v(\omega), a^t, sell) = 0 \quad &\text{for} \quad \beta^t < v(\omega) - \rho \\ F(v(\omega), a^t, buy) = 0 \quad &\text{for} \quad \alpha^t > v(\omega) + \rho, \end{aligned}$$

for all $\omega \in \Omega$.

For concreteness, assume that the trader's buy and sell order are executed against the best ask, α^t, and bid, β^t, quotes, respectively. Let $u_i(\omega, a^t)$ denote dealer i's expected payoff, or *reward*, in period t given the state ω and dealers' action profile a^t.[11] Dealers' aggregate expected

[11]Here, expectations are taken with respect to the possible trader's orders (i.e., buy, sell and no trade) given the fundamentals ω. For instance, if orders are executed by the dealers setting the best quotes,

$$\begin{aligned} u_i(\omega, a^t) &= (W(\omega) - \beta_i^t) F(\omega, a^t, sell) 1_{\{\beta_i^t = \beta^t\}} \eta_\beta(a^t) + (\alpha_i^t - W(\omega)) F(\omega, a^t, buy) 1_{\{\alpha_i^t = \alpha^t\}} \eta_\alpha(a^t) \\ &+ W(\omega) q_i^{ID}(a^t) + c_i^{ID}(a^t), \end{aligned}$$

where η_β and η_α are tie-breaking rules applied in case more than one dealer sets the best bid or ask, respectively. However equation (2) holds no matter the identity of the dealer executing a trader's order as long as this order is executed at the best price.

payoff in period t can be written as

$$U(\omega, a^t) := \sum_i u_i(\omega, a^t) = (W(\omega) - \beta^t)F(v(\omega), a^t, sell) + (\alpha^t - W(\omega))F(v(\omega), a^t, buy). \qquad (2)$$

Condition **ETD** and expression (2) imply that dealers' aggregate trading stage payoff $U(\omega, a^t)$ is positive independently of the non-learnable component $\psi(\omega)$ only if a^t induces a bounded change in dealers' aggregate inventory. That is, only if [12]

$$-\frac{\rho}{\bar{e}} < \frac{F(v(\omega), a^t, sell) - F(v(\omega), a^t, buy)}{F(v(\omega), a^t, sell) + F(v(\omega), a^t, buy)} < -\frac{\rho}{\underline{e}}. \qquad (3)$$

In other words, only a relatively balanced traders' order flow can guarantee that dealer's aggregate profit be non-negative for some realization of the non-learnable component $\psi(\omega)$. On the other hand, a positive bid-ask spread inducing a non-nil but relatively balanced traders' order flow guarantees that dealers' aggregate profit is positive independently of the asset value $W(\omega)$. Namely we assume:

Positive Payoffs (PP): *For any given $\omega \in \Omega$, there is a non-empty set set $A^\star(\omega) \subset A$ such that if $a^t \in A^\star(\omega)$, then $\beta^t < \alpha^t$ and $F(v(\omega), a^t, sell) = F(v(\omega), a^t, buy) > 0$.*

Note that for $a^t \in A^\star(\omega)$, we have

$$U(\omega, a^t) = (\alpha^t - \beta^t)F(v(\omega), a^t, sell) > 0. \qquad (4)$$

Thus, **PP** implies that, if the state is ω, and dealers pick their action in $A^\star(\omega)$, they make strictly positive aggregate profit from pure intermediation, i.e., without taking a net position in the asset. This implies that given the true state ω, dealers' aggregate profits from setting $a \in A^\star(\omega)$ remain strictly positive independently of the asset's fundamental value $W(\omega)$.

[12] Inequality (3) is obtained noting that condition **ETD** implies $U(\omega, a^t) \leq (W(\omega) - v(\omega) + \rho)F(v(\omega), a^t, sell) + (v(\omega) + \rho - W(\omega))F(v(\omega), a^t, buy)$.

Consider the partition $\hat{\Omega} = \{\hat{\omega}_1, \hat{\omega}_2\}$ of the set Ω, where $\hat{\omega}_1 := \{\omega | v(\omega) = v_1\}$ and $\hat{\omega}_2 := \Omega \backslash \hat{\omega}_1$. Fix $j \in \{1, 2\}$. Note that because of **NLS**, the traders' behavior does not change with $\omega' \in \hat{\omega}_j$. It follows from the assumptions that for any given $\hat{\omega} \in \hat{\Omega}$, the *static trading game* satisfies the following four properties:

1. *Positive maximum payoffs*: there is a dealers' action profiles for which their trading round payoff is strictly positive, independently of $\omega' \in \hat{\omega}$. For instance, for $a \in A^{\star}(\omega')$, expression (4) implies $U(\omega', a^t) > 0$ for all $\omega' \in \hat{\omega}$ and one can choose the inter-dealer market action profile so that each dealer gets a strictly positive share of the aggregate payoff $U(\omega', a^t)$.

2. *Negative minimum payoffs*: There is a dealers' action profile for which their trading round payoff is negative independently of $\omega' \in \hat{\omega}$. There are many ways dealers' can make an aggregate loss that can be "shared" among dealers in the inter-dealer market.[13]

3. *Non-positive expected payoffs*: For any given dealer i, and any probability distribution $\mu \in \Delta\hat{\omega}$, the other dealers have some action profiles $\underline{a}_{-i}(\mu)$ forcing dealer i's trading round expected payoff to be non-positive (the expectation is w.r.t. the distribution μ). Formally,

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

   For example, if dealers other than dealer i adopt the no-trade actions $\underline{a}_{-i}^{ID}$ in the inter-dealer market and then set quotes such that $\min_{k \neq i} \alpha_k = \max_{k \neq i} \beta_k = v(\hat{\omega}) + \sum_{\omega \in \hat{\omega}_j} \mu(\omega)\psi(\omega)$, then dealer i's trading round expected payoff (computed according to μ) cannot be positive.

4. *Non-equivalent payoffs*: There exist feasible payoffs in which any given dealer has a lower profit than any other dealer's. For instance, fix dealer i and take an action profile $a \in A^{\star}(\omega')$

[13] For example, by setting a^t such that $v + \bar{e} < \beta^t = \alpha^t$ and $F(\omega', a^t, sell) > F(\omega', a^t, buy)$, it results $U(\omega', q^t) \leq (v + \bar{e} - \beta^t)(F(\omega', a^t, sell) - F(\omega', a^t, buy)) < 0$ for all $\omega' \in \hat{\omega}_j$.

resulting in a strictly positive aggregate payoff that is then shared in the inter-dealer market leaving to dealer i strictly less than what each other dealer gets.

The repeated game: We can now move to the repeated game. The stage game payoffs (or rewards) of the dealers are discounted at the common factor $\delta < 1$ and the (overall) game payoff is the average discounted sum of rewards. The discount factor δ accounts both for the dealers' time preference and for the possibility that the public information gets released in the current period.[14] In each period, dealers' actions and traders' reactions are observed by all dealers. Let H^t denote the set of public histories $h^t = \{a^\tau, s^\tau\}_{\tau=0}^{t-1}$. Given some sequence of action profiles $\{a^t\}_{t=1}^{\infty}$ by the dealers, dealer i's expected payoff in state ω is[15]

$$\sum_{t=1}^{\infty}(1-\delta)\delta^t u_i(\omega, a^t). \tag{5}$$

A public strategy profile (strategy henceforth) is a mapping $\sigma : \cup_t H^t \to \times_i \Delta A_i$. A strategy σ and a state ω induce a probability distribution over histories in the standard fashion. Let $V_i(\omega, \sigma | h^t)$ denote dealer i's expected continuation payoff after observing the public history h^t given state ω and strategy profile σ.

Definition 1 *A* belief-free equilibrium *(hereafter, BFE) 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:*

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i} V_i(\omega, \sigma_i, \sigma_{-i}^* | h^t), \tag{6}$$

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

[14]Allowing for a stochastic discount factor complicates exposition but does not affect results as long as the discount factor remains close enough to one.

[15]Here, expectation is taken with respect to the possible realizations of traders' orders $\{s^t\}_{t=1}^{\infty}$, taking the state ω as given.

Some remarks are in order. First, a BFE is a subgame-perfect equilibrium given any initial prior distribution of dealers' belief about ω and any additional private information a dealer might possess.[16] Second, a BFE is an equilibrium even if dealers are ambiguity averse, as long as ambiguity pertains to the distribution of the possible states of nature $\omega \in \Omega$.[17]

The equilibrium: Our purpose is to construct a particular class of equilibria. Here, we describe the logic underlying their structure for the simple model presented above. To begin with, let us consider the canonical equilibrium (i.e., the GME) that can be obtained if we make the additional assumptions that all dealers are equally uninformed and start from the common prior $p^0 = \Pr(v(\omega) = v_2)$ and that $E[\psi(\omega)] = 0$. Then there is a perfect Bayesian equilibrium in which, in any period t: (i) each dealer's expected profit is nil; (ii) there is no trade in the inter-dealer market; (iii) best bid and ask quotes in the quote driven market satisfy

$$\alpha^t = \alpha(p^t) \quad := \quad E\left[v(\omega)|h^{t-1}, s^t = buy\right], \tag{7}$$

$$\beta^t = \beta(p^t) \quad := \quad E\left[v(\omega)|h^{t-1}, s^t = sell\right], \tag{8}$$

$$p^{t+1} \quad = \quad \phi_B(p^t, a^t, s^t), \tag{9}$$

where $\phi_B(p^t, (\alpha^t, \beta^t), s^t)$ denotes the posterior probability that $v(\omega) = v_2$ resulting from the prior p^t and from the trader's reaction s^t to dealers' quotes a^t.[18] This equilibrium has the advantage of being Markovian: first, in every period t, best bid and ask quotes only depend on dealers'

[16] To see this, note that in a perfect Bayesian equilibrium, dealers' strategies satisfy

$$\sigma_i^* \in \text{argmax}_{\sigma_i} E\left[V_i(\omega, \sigma_i, \sigma_{-i}^*|h^t)|I_i\right],$$

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.

[17] Unlike in Easley and O'Hara (2010), where some of the traders are ambiguity averse, here ambiguity aversion applies to dealers.

[18] In order to simplify the exposition and notation we neglect the rounding required from the fact that quotes belong to a grid.

common belief p^t; second, next period dealers' common posterior beliefs p^{t+1} only depend on the common prior p^t and on (a^t, s^t), dealers' and trader's actions at time t. However this quoting strategy is not an equilibrium as soon as there is at least one dealer whose belief that $v(\omega) = v_2$ is not p^t.[19]

In a BFE, the same dynamic quoting strategy must be optimal no matter the belief a dealer might have about the true ω. Because dealers beliefs might differ arbitrarily, a dealer's strategy must be optimal no matter what the true realization of ω is. We first briefly illustrate how, if dealers are patient enough, this can be achieved with strategies that have a Markov structure that is as simple as the one of the canonical equilibrium. Initially, dealers post quotes that depend on an arbitrary given distribution on the possible values of $v(\omega)$, we will call this distribution the *market measure*. Assumption **LS** guarantees that these quotes can be chosen in $A(v_1, v_2)$ so that the resulting flow of trade provides information about the true $v(\omega)$. This information affects the value of the market measure and hence the evolution of quotes. This is what we call an *exploring phase*. None of the dealers' beliefs need reflect the market measure, but nevertheless no dealer deviates. This is because the "Non-positive expected payoffs" property guarantees that other dealers can ensure that the deviating dealer makes zero profits with sufficiently low discount rate (using standard repeated-game logic). The order flow eventually conveys sufficient information about the true value of $v(\omega)$ and then dealers switch to an *exploiting phase* where quotes belong to $A^\star(\omega)$, that is: best bid and ask quotes induce a balanced order flow and provide dealers with a flow of aggregate profit that is strictly positive no matter the value of $W(\omega)$. Each dealer obtains a strictly positive share of this profit through the redistribution of cash and asset taking place in the inter-dealer market. In response to the order flow during an exploiting phase, however, play can revert to the exploring phase, and so on. The reason why an exploiting phase cannot last forever is that a dealer who disagrees with the consensus asset value must be given incentives to

[19] To see this, note that if at some time t, dealer i's belief that $v(\omega) = v_2$ is $p_i^t \neq p^t$, then dealer i has a profitable deviation that consists in setting either $\beta_i^t > \beta(p^t)$ or $\alpha^t < \alpha(p^t)$.

play along and wait for play to shift towards the asset value that he might believe in. To preserve the Markovian structure, the level of current quotes and the transition from one phase into the other must only depend on the current level of the market measure (and possibly on the level of dealers aggregate inventory). At the same time, no matter the current level of the market measure and a dealer belief about ω, the dealer must expect that the play will shift toward the correct exploiting phase within a bounded period of time. Otherwise, even a patient dealer would prefer to deviate and generate extra profits in the current trading round (even if held down to zero profits afterwards), rather than to make losses during the long transition period required for the market measure to adjust to what it thinks the right exploiting phase is. This is possible only if, first, during an exploiting phase the market measure "transition rule" attaches decreasing probability to states that are unlikely in view of the flow of information provided by traders' orders; second, during an exploiting phase the market measure is not too persistent, but instead is sensitive to the new public information provided by traders' orders. Bayesian updating, for instance, would not satisfy these two properties: while it allows to pin down the true $v(\omega)$ almost surely eventually, it is too persistent for our purpose: once the market measure is sufficiently concentrated on a state, it takes arbitrarily long for a Bayesian belief to budge.

A formal description: To define equilibrium play more precisely, consider the partition $\hat{\Omega} = \{\hat{\omega}_1, \hat{\omega}_2\}$ and fix some small $\varepsilon > 0$ and some arbitrary $\pi^0 \in \Pi := [\varepsilon/4, 1-\varepsilon/4]$ as the initial weight the market measure assigns to $\hat{\omega}_2$. The following updating rule $\phi : \Pi \times A \times S \to \Pi$ is an example of rule that allows to identify the true state but is less persistent than a Bayesian updating rule:

$$\pi^{t+1} = \phi(\pi^t, a^t, s^t) := \arg\min_{\pi\in\Pi} \left\|\pi - \phi_B(\pi^t, a^t, s^t)\right\|. \tag{10}$$

The (on-path) equilibrium play can then be seen as the alternation of two type of phases: *exploring phases* and *exploiting phases*. Whenever $\pi^t \in [\varepsilon, 1-\varepsilon]$, the game is in an exploring phase and dealers' actions are such that $a^t \in A(v_1, v_2)$. This guarantees that dealers' quotes induce an informative flow of trades in the quote driven market. Thus, as time passes the market measure attaches more and more weight to the true $\hat{\omega}$. An exploiting phase is defined to start as soon as the market measure attaches enough weight to a particular state. Namely, whenever $\pi^t < \varepsilon$ (resp., $\pi^t > 1 - \varepsilon$), the game is in the $\hat{\omega}_1$-exploiting phase (resp. $\hat{\omega}_2$-exploiting phase). In this phase, $a^t \in A^\star(\hat{\omega}_1)$ (resp. $a^t \in A^\star(\hat{\omega}_2)$). This guarantees that dealers gain the spread without taking a net permanent position in the asset. In a BFE, a dealer's payoff must be positive no matter its belief about $\psi(\omega)$, a component that, because of **NLS**, cannot be integrated by the market measure. Assumption **PP** guarantees that for any given $v(\omega)$, dealer's can maintain balanced aggregate positions that generate strictly positive aggregate profits that do not depend on the asset true value W. However, the recurrence of exploring phases can lead dealers to accumulate relatively unbalanced portfolios leading to average inventories that would not satisfy condition (3) and hence to dealers' aggregate long term profits that are negative for some value of $\psi(\omega)$. This can be easily avoided by appropriately biasing dealers' quotes on the basis of the current level of their aggregate inventory, so as to induce traders to absorb dealers' excessive inventory .

Finally to make sure that each single dealer makes strictly positive profits, dealers can use the inter-dealer market at time t to share the dealers' aggregate positions and trades resulting from the orders at time $t - 1$. A sharing rule and resulting inter-dealer trade can be easily set so that at beginning of each round, each dealer i gets a strictly positive fraction of dealers' aggregate profit or loss resulting from the trade in the previous round made in the quote driven market.

Overall, we have defined a *partial strategy* profile, i.e., a mapping $\sigma : \Pi \times S \to \Delta A$, so that dealers' actions in the inter-dealer market and dealers' quotes at time t only depend on s^{t-1} and on π^t and possible on the level of dealers' aggregate inventory. This partial strategy together with the market measure updating rule ψ defined in (10) satisfies the following two properties:

It is ε-*learning*: For any current level of the market measure $\pi^t \in \Pi$, the expected time that it takes for the market measure to assign at least probability $1-\varepsilon$ to the true state is bounded, uniformly in π^t.

It is ε-*exploiting*: Whenever the market measure assigns at least probability $1-\varepsilon$ to the true state $\hat{\omega}$, each dealer's trading round payoff is strictly positive.

In Section 2, we consider a generalization of this microstructure model, and show that, as long as the general trading game satisfies Properties 1–4, one can define partial strategies that are ε-learning and ε-exploiting. The main result establishes that any such strategy profile defines a behavior that coincides with a belief-free equilibrium outcome, provided the dealers are patient enough.

We conclude this section with an example to illustrate some salient differences between such a BFE and the GME.

1.1 BFE Market Making vs. zero expected profit equilibrium: An Example

We consider here a specification of the model outlined above to compare dealers' quotes in a BFE to the quote resulting from the canonical belief-based equilibrium in which dealers make zero expected profit in every period. This is the equilibrium characterized by Glosten and Milgrom (1985), referred to as GME.

We assume that half of the population of traders is composed of potential buyers. The remainder are potential sellers. Traders trade both for liquidity and speculative reasons. They are informed about the component $v(\omega)$, but not about the $\psi(\omega)$ component. The following specification of the function F for $\beta^t \leq \alpha^t$, satisfies Assumptions **NLS**, **LS**, **PP** and **ETD**:[20]

$$\begin{aligned}
F(v(\omega), a^t, sell) &= \max\left\{0, \min\left\{\frac{1}{2}, \frac{\beta^t - v(\omega) + \rho}{4\rho}\right\}\right\}, && (11)\\
F(v(\omega), a^t, buy) &= \max\left\{0, \min\left\{\frac{1}{2}, \frac{v(\omega) + \rho - \alpha^t}{4\rho}\right\}\right\}, && (12)\\
F(v(\omega), a^t, no\ trade) &= 1 - F(\omega, \beta^t, sell) - F(\omega, \alpha^t, buy). && (13)
\end{aligned}$$

First, let us consider the GME. In this case, two additional assumptions are required: first, dealers are equally uninformed with belief p^t, and second, $E[\psi(\omega)] = 0$. Then, for $\rho > \sqrt{2}(v_2 - v_1)$, we can express time t ask and bid quotes resulting from dealers' zero-profit condition as

$$\begin{aligned}
\alpha(p^t) &:= E[v|h^t] + \frac{\rho}{2} - \frac{1}{2}\sqrt{\rho^2 - 4\mathrm{Var}[v|h^t]}, && (14)\\
\beta(p^t) &:= E[v|h^t] - \frac{\rho}{2} + \frac{1}{2}\sqrt{\rho^2 - 4\mathrm{Var}[v|h^t]}. && (15)
\end{aligned}$$

where $E[v|h^t]$ and $\mathrm{Var}[v|h^t]$ are the expectation and the variance of $v(\omega)$, respectively, computed using the common belief p^t that evolves according to (9).

Now, let us consider the following BFE. Fix some small $\varepsilon > 0$ and $\pi^0 \in \Pi$ and let π^t evolve according to equation (10), so that $\pi^t \in \Pi$ for all t. We say that for $\pi > 1 - \varepsilon$ (resp. for $\pi < \varepsilon$), the game is in v_1-exploiting phase (resp. v_2-exploiting phase). Fix d and c such that $0 < c < d < \rho - c$. Let inv^t denote dealers aggregate inventory at t and let $c^t := c(inv^t)$, where $c(.)$ is decreasing satisfying $c(0) = 0$. During a v-exploiting phase, the best ask and bid equilibrium quotes satisfy

$$\begin{aligned}
\alpha^t &= v + d + c^t, && (16)\\
\beta^t &= v - d + c^t, && (17)
\end{aligned}$$

[20]Namely, $A^\star(\omega)$ is the set of $a \in A$ such that $(\alpha^t, \beta^t) = (v(\omega) + \gamma, v(\omega) - \gamma)$ with $\gamma \in (0, \rho)$. The set $A(v_1, v_2)$ satisfies $\alpha^t \geq \beta^t$ with $\alpha^t < v_2 + \rho$ or $\beta^t > v_1 - \rho$, so that $F(v_1, a^t, s) \neq F(v_2, a^t, s)$ for some $s \in S$. Furthermore $F(v(\omega), a^t, sell)$ and $F(v(\omega), a^t, buy)$ are nil for $\beta^t < v(\omega) - \rho$ and $\alpha^t < v(\omega) + \rho$, respectively

respectively. For $\pi^t \in [\varepsilon, 1-\varepsilon]$, the game is in the exploring phase and best ask and bid quotes satisfy

$$\alpha^t = \alpha(\pi^t) + d + c^t, \tag{18}$$

$$\beta^t = \beta(\pi^t) - d + c^t. \tag{19}$$

Note that for a level of the market measure in the BFE that is identical to the Bayesian belief in the GME, best quotes in the two equilibria differ by at most $d + c^t$. However, the evolution of dealers' belief can substantially differ across the two equilibria. As a result, the same history of past trades may lead to sharply different quotes in GME and in BFE.

Namely, in a BFE quotes are intrinsically more volatile than in the GME. This is due to the fact that in the GME, dealers' eventual belief will attach probability arbitrarily close to 1 to the true value of v. This cannot happen for the BFE market measure, which can never be too concentrated on a given state and hence remains unstable. Thus, independently of the previous history of trade, and on dealers' actual beliefs about $v(\omega)$, the market measure and quotes will remain sensitive to the trading volume. As a result, exploring phases are recurrent. This is illustrated in Figure 2 that reports a simulation of the two equilibria for for $v(\omega) = v_1$ (blue line for BFE and red dashed line for GME).[21] The sequence of potential buying and selling traders is the same for the two equilibria. The right panel of Figure 2 reports the evolution of dealers' aggregate inventory for the two equilibria. The left panel reports the evolution of the market measure in a BFE and of the Bayesian belief in the GME. Because in the simulation $v(\omega) = v_1$, traders tend to sell more than buy the asset. As a result in GME dealers' aggregate inventory tends to explode. Not so in the BFE where dealer's aggregate inventory remain more balanced thanks to the bias in quotes c^t. Exploiting phases correspond to the periods in which the market

[21]The parameters used for this simulation are: $v_1 = \$15$, $v_2 = \$18$, $\rho = 15$, $d = \$0.05$, $\varepsilon = 0.05$, $p^0 = \pi^0 = 0.5$, $\overline{e} = -\underline{e} = 3$ and $c(inv^t) = -0.02 inv^t$. The Figure reports time series of 3000 trades.

measure is below the threshold ε (left graphic, solid magenta line). In these regions volatility and sensitivity to the volume of trade are low. Exploring phases occur when the market measure is above ε and display higher volatility and sensitivity to the trading volume. This is illustrated in Figure 3: the negative relation between the evolution of dealers' aggregate inventory and the market measure is stronger in exploring phases than in exploiting phases, whereas it is negligible in GME.

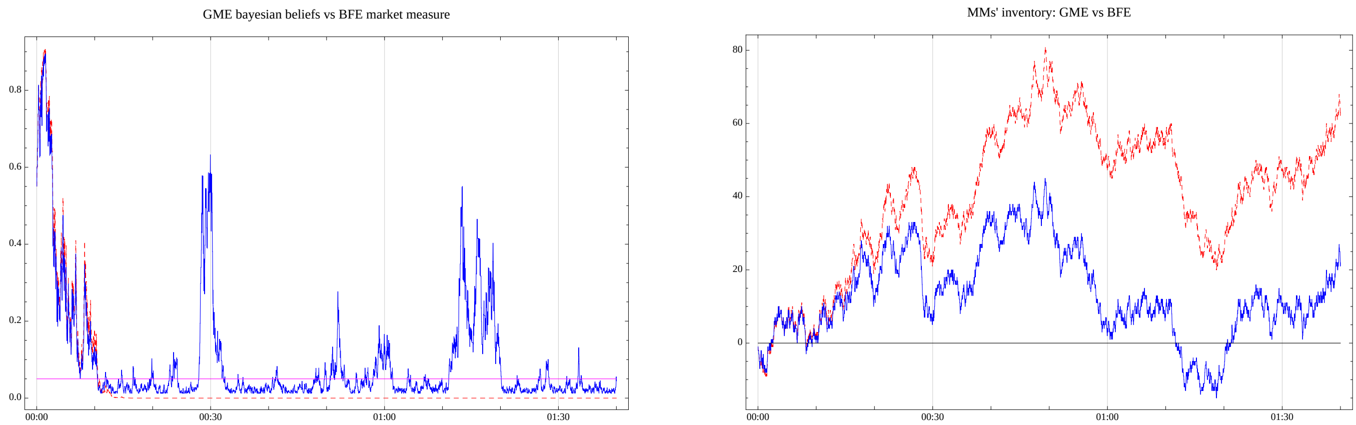

Figure 2: Red dashed line are used for the GME and blue solid lines for BFE. The left panel displays the evolution of the BFE market measure and of the GME Bayesian belief. The right panel reports the evolution of dealers aggregate inventory in the GME and in the BFE.

Figure 4 represents the evolution of quotes in the GME and in the BFE. In both cases, the flow of trade provides enough information about $v(\omega)$ allowing the Bayesian dealers' belief to eventually converge to the truth about $v(\omega)$. In the GME, this leads to a vanishing volatility and bid ask spread with quotes that remain arbitrary close to v_1. This is illustrated in the left panel of Figure 4.

In the BFE, quotes keep moving with the market measure depicted in Figure 2. Note that the spread remains bounded away from 0 even when the market measure is relatively concentrated, leading to an exploiting phase (see Figure 4, right panel). As a result, while in the GME the average dealers' aggregate per-period profit quickly converges to 0, it is of the same magnitude

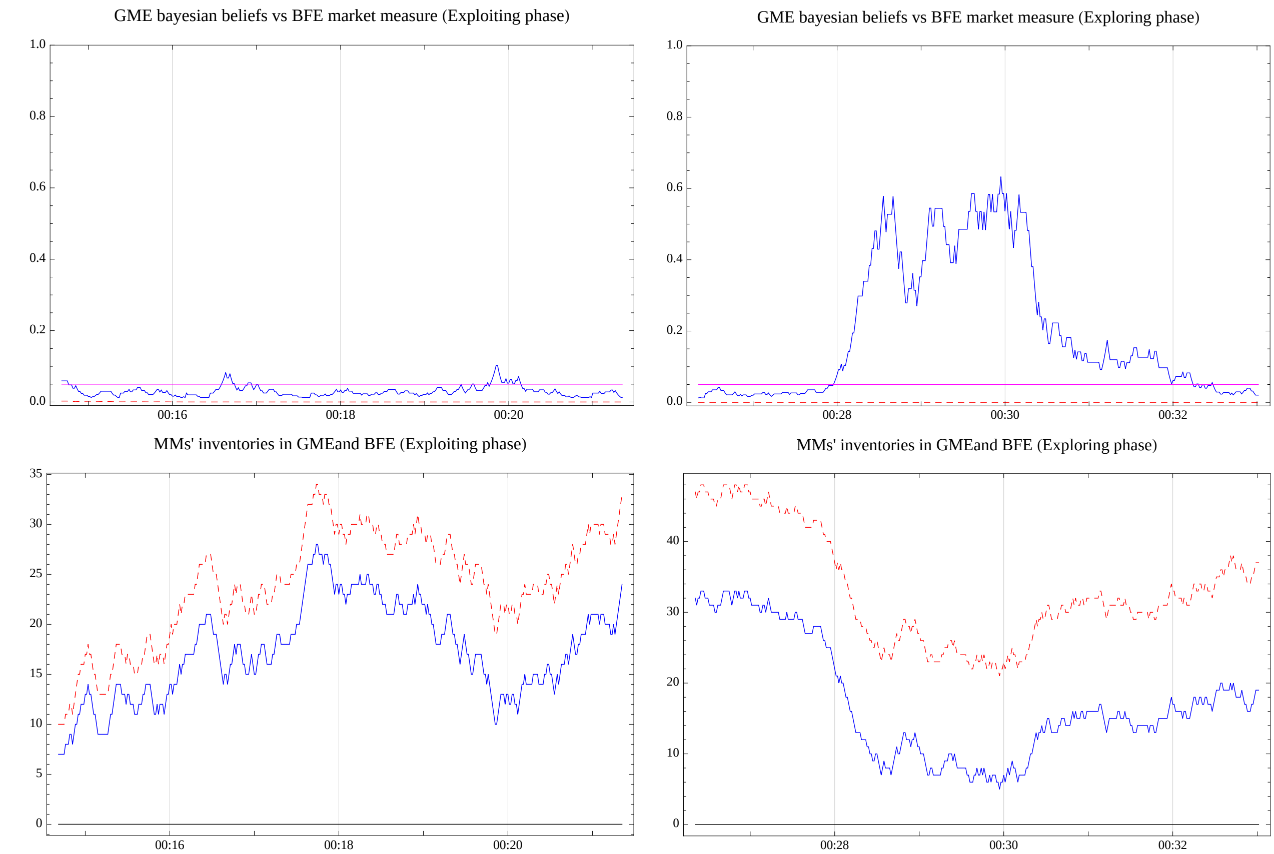

Figure 3: Market measure and dealers' inventory in an exploiting phase (left panel) and an exploring phases (right panel).

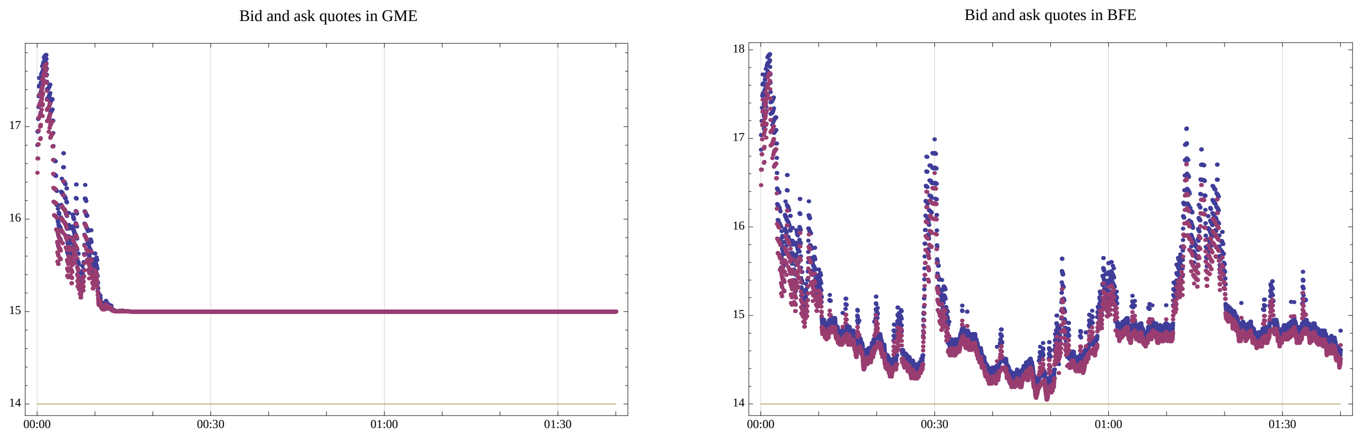

Figure 4: Evolution of bid and ask quotes in the GME (left panel) and in the BFE (right panel). Ask quotes and bid quotes are in blue, and magenta, respectively.

as d in the BFE (See Figure 5).

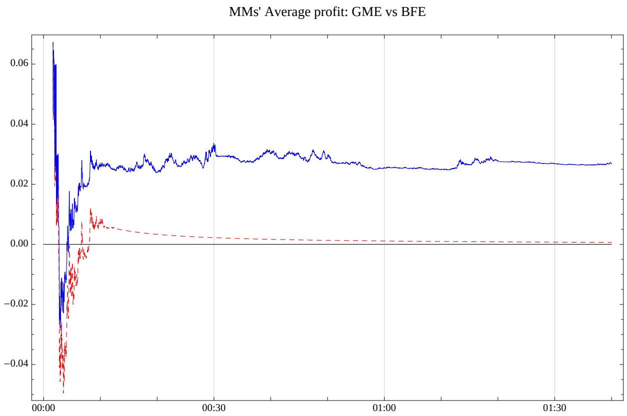

Figure 5: Evolution of the average per-period profit taking $\psi(\omega) = 0$ in GM (red dashed line) and in the BFE (blue solid line).

Note that a dealer's *ex post* profit also depends on the value of $\psi(\omega) \in \{\underline{e}, \overline{e}\}$. Figure 6 represents the *ex post* cumulative profit for $\psi(\omega) = \underline{e}$ and $\psi(\omega) = \overline{e}$. Note that in the GME (left panel of Figure 6), the dealers' cumulative profit remains negative for at least one realization of $\psi(\omega)$. In the BFE, the dealers' cumulative profit eventually becomes positive no matter the realized $\psi(\omega)$. This can be achieved also thanks to the fact that dealers aggregate inventory does not explode.

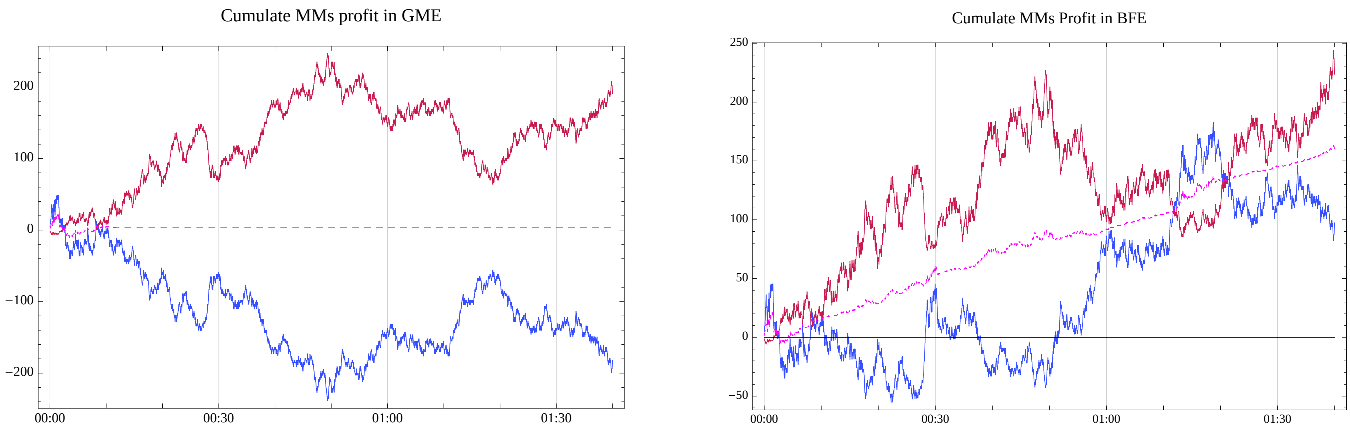

Figure 6: Cumulative dealers' profits for $\psi(\omega) = 0$ (dashed line), $\psi(\omega) = \underline{e}$ (blue line) and $\psi(\omega) = \overline{e}$ (red line).

We stress that GME is a perfect Bayesian equilibrium of the economy considered in this example *only if* dealers are equally uninformed, have common initial prior $\Pr[v(\omega) = v_2] = p^0$ and $E[\psi(\omega)] = 0$. The BFE remains a perfect Bayesian equilibrium even when dealers are asymmetrically informed and *no matter* dealers' prior beliefs about ω.

2 A General Model of Market Microstructure

Set-up: At time 0, Nature chooses a state ω in an arbitrary set Ω. In each period $t = 1, 2, \ldots,$ first n dealers choose an action profile in the finite set A. Second, traders react with an action in the finite set S. Depending on the specific trading mechanism, a dealer's action might specify bid and ask quotes and maximum quantities dealers stand ready to trade, but it can also correspond to a limit order or a market order in a limit order market, and/or to the trade in an inter-dealer parallel market. Traders' reactions can consist of market orders specifying the quantities to be traded at dealers' best quotes, and/or limit orders that will compete with existing limit orders.

Let $W(\omega)$ be the fundamental value of the risky asset in state ω. Dealers' actions and traders' reactions are publicly observable. For a given action profile $a \in A$ of dealers' actions and traders reaction $s \in S$, let $Q_i(a, s)$ and $P_i(a, s)$ denote the resulting amount of the risky asset and money, respectively, that other market participants transfer to dealer i. Then, if the state of nature is ω, dealer i's stage payoff is

$$u_i(\omega, a, s) := W(\omega)Q_i(a, s) - P_i(a, s). \tag{20}$$

Let $F : \Omega \times A \to \Delta S$ summarize traders' behavior. Namely, $F(\omega, a, s)$ is the probability that, in a given trading round, the traders' reaction is s, given that the state of nature is ω and dealers' actions are a. Thus, given state ω and dealers action profile a, we obtain dealer i's expected stage payoff by taking the expectations with respect to the possible reactions s:

$$u_i(\omega, a) := \sum_{s\in S} u_i(\omega, a, s)F(\omega, a, s). \tag{21}$$

Let $\overline{u} := \max_{\omega,a,i} |u_i(\omega, a)|$. Stage game payoffs are discounted at the common discount factor $\delta < 1$.[22]

Note that uncertainty about the state of nature ω can be fundamental and/or non-fundamental. Fundamental uncertainty pertains to the liquidation value of the asset. It affects dealers' payoffs directly through $W(\omega)$ and, because of the presence of traders with private information about $W(\omega)$, indirectly via $F(\omega, \cdot)$. The non-fundamental uncertainty relates to the behavior of traders, for example, it might pertain to their risk aversion, their inventory or the amount and quality of the information they possess regarding the asset value. The non-fundamental uncertainty solely affects $F(\omega, \cdot)$.

We define the information that can be gathered by observing the realization of the public signal s. For any pair of states $\omega, \omega' \in \Omega$, we denote by $A(\omega, \omega') \subseteq \Delta A$ the set of actions profiles a satisfying $F(\omega, a) \neq F(\omega', a)$. A state ω can be statistically distinguished from ω' only if $A(\omega, \omega') \neq \emptyset$. Let $\hat{\Omega}$ be the partition over Ω induced by the function F. 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 $\hat{\omega}(\omega)$ the element of $\hat{\Omega}$ containing $\omega \in \Omega$.

Note that the fact that the value of the asset differs in two states does not imply that the two states are statistically distinguishable. This happens for instance when no trader can tell those two states apart.[23] Similarly, even if the value of the asset is the same in two states, the two states could be distinguishable because the traders' reaction could be different, possibly because ω affects the traders' risk aversion.

[22] As for the example of 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.

[23] For instance, in the model of Section 1, states that differ only for the $\psi(\omega)$ component are indistinguishable.

We make the following Assumptions on $u_i(\omega, a)$, that subsume our earlier assumptions and Properties 1-4.

Assumption B: For any given $\hat{\omega} \in \hat{\Omega}$:

1. *Positive maximum payoffs*: There exists a non-empty set $A^\star(\hat{\omega}) \subseteq A$, such that $u_i(\omega, a) > 0$ for all $a \in A^\star(\hat{\omega})$, $\omega \in \hat{\omega}$ 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 given 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,

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

4. *Non-equivalent payoffs*: There exist n action profiles $\{a^1(\hat{\omega}), \ldots, a^n(\hat{\omega})\} \in [\Delta A]^n$ 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 $u^\star = \min_i \min_\omega \min_{A^\star(\hat{\omega}(\omega))} u_i(a, \omega) > 0$ denote a lower bound on payoffs from actions in $A^\star(\hat{\omega})$. Roughly speaking, Assumptions **B**-1 and **B**-2 guarantee that for each statistically distinguishable state $\hat{\omega}$, there are action profiles providing each dealer with at least $u^\star > 0$ and action profiles leading to strictly negative payoffs, respectively. Assumption **B**-3 guarantees that it is possible to punish each dealer in each state. Assumption **B**-4 states that for each $\hat{\omega}$ one can find as many action profiles as there are dealers such that dealer i prefers all the other $n-1$ action profiles to the i-th action profile.

Note that whenever the public information does not allow to pin down the exact value of $W(\omega)$, as long as traders demand is elastic, and thus they are not willing to trade the asset at a too high expected loss, an action in $A^\star(\hat{\omega})$ cannot conduce to a too imbalance change in dealers aggregate inventory.

Equilibrium strategies ingredients: The following generalizes the construction of the example of Section 2.

We start 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 $\phi : \Pi \times A \times S \to \Pi$ be a probability updating rule, i.e. $\pi^{t+1} = \phi(\pi^t, a^t, s^t)$. Thus, π^t can be recursively computed from the map ϕ, given the sequence (a^τ, s^τ) of actions and signals, and the initial value π^0. 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. That is, given ϕ, we define a *partial strategy* to be a map $\sigma : \Pi \times S \to \Delta A$. Instead, a public strategy profile (strategy henceforth) is a mapping $\hat{\sigma} : \cup_t H^t \to \times_i \Delta A_i$, where H^t is the set of histories $h^t = \{a^\tau, s^\tau\}_{\tau=0}^{t-1}$ specifying dealers' actions and traders' reactions until time t.

For a given updating rule ϕ and a partial strategy σ, we have the following definitions.

Definition 2 *1. The pair (ϕ, σ) is ε-*learning*, for $\varepsilon > 0$, if for any $\omega \in \Omega$ and any $\pi^0 \in \Pi$,*

$$\Pr_{\omega,\sigma}\left[\liminf_{T\to\infty} \frac{1}{T}\sum_{t=0}^{T} 1_{\{\pi^t(\hat{\omega}(\omega))>1-\varepsilon\}} < 1-\varepsilon\right] < \varepsilon, \tag{22}$$

*2. The pair (ϕ, 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 $\Pr_\sigma\left[a^t \in A^\star(\hat{\omega})|h^t\right] > 1-\varepsilon$.*

Let us say that that the market measure π^t *points to* a state $\hat{\omega}$ if $\pi^t(\hat{\omega}) \geq 1-\varepsilon$. Then we can interpret the two definitions as follows: ε-learning means that, over many periods, the market measure will not point at the $\hat{\omega}$ that contains the true state ω with a frequency that is smaller than ε. In other words, the market measure is rarely far from the truth, in terms of long-run frequency. The ε-exploiting property guarantees that whenever the market measure points at a some $\hat{\omega}$, play is such that a dealers' payoff is strictly positive in all states ω included in $\hat{\omega}$.

For (ϕ, σ) to be ε-learning, note that it is necessary that dealers' actions do not block the flow of information coming from traders' reactions. That is to say, no matter the level of the market measure π, the actions that allow to distinguish the true $\hat{\omega}$ from the other $\hat{\omega}' \in \hat{\Omega}$ must be played with strictly positive frequency. Formally, (ϕ, 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 $\pi^0 \in \Pi$,

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

Note also that if, as in the illustrative example, the pair (ϕ, σ) is such that the expected time required for the market measure to point at the true $\hat{\omega}$ is finite, then (ϕ, σ) satisfies (22).

Then we have the following:

Theorem 1 *Suppose that Assumption **B** is satisfied. Then there exists $\bar{\varepsilon} > 0$ such that for any $\varepsilon < \bar{\varepsilon}$, if (ϕ, σ) is ε-learning and ε-exploiting, then 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 belief-free equilibrium $\hat{\sigma}$ that specifies the same action profile as the partial strategy σ, after any history after which no player has deviated.

3 Extensions

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 once for all at time 0. Also, long-term market participants do not take advantage of their private information. Here, we sketch how the model can be extended and the analysis adapted to deal with these modifications.

A restriction of our model is that dealers' actions are observable. This might not be the case for some opaque markets as for instance when dealers' quotes are anonymous. Imperfect monitoring of actions makes it more difficult to detect a dealer's deviating from the mutually profitable collusive-type strategy. This reduces the threat of punishment and complicates implementing collusive-like behaviors. 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 detectable. For example, Christie and Shultz (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)). However, imperfect monitoring of dealers' actions might impose further restriction 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 that (ϕ, σ) be ε-learning, we must now 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 as the fluctuations of the value itself –perhaps 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 trades itself does not affect fundamentals. The verification that σ is a belief-free equilibrium follows exactly the same steps as in the main proof. As for the corresponding payoffs, if for every partial strategy, the pair consisting of the asset's value and the public signal follows an irreducible Markov chain, then computing the limiting payoff as $\delta \to 1$ is straightforward, by integration with respect to the invariant distribution. However, there is clearly a tension between the assumption that the value of the asset changes slowly enough for learning to occur fast, and the assumption that market participants are patient enough for the long-run distribution to be the only payoff-relevant aspect of the Markov chain. Therefore, if we view δ (or rather, $1-\delta$) as capturing the delay between consecutive trades rather than intrinsic patience by market participants, it makes little sense to assume that the fluctuations occur at a rate that is independent of delay –higher δ should correspond to slower fluctuations of the fundamentals. In that case, computing equilibrium payoffs as $\delta \to 1$ is more difficult.

A third important restriction is that long term market participants do not take advantage of their private information, if any. One can view this as an implication of our definition of belief-free equilibrium: dealers' beliefs do not affect their actions because what really matters is identifying 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. However, in our equilibrium, dealers cannot take advantage of the private information they might possess. Still, there is no difficulty in redefining belief-free equilibrium appropriately. Rather than taking the asset value as a primitive that determines a distribution over the players' private signals, one can think of 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 independently 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 against all these profiles. In fact, it is clear that we do not need to impose that the players' combined signals pin down the value of the asset. Rather, it pins down a set of possible values, for all of which the best-reply property must hold.

This provides a natural extension of the definition of belief-free equilibrium that allows dealers to take advantage of their private information. We believe that such an extension raises interesting questions and technical challenges that motivate further study.

4 Conclusion

This paper considers market microstructure models in which long-lived dealers interact with short-lived traders. We characterized equilibrium price formation strategies that are robust to changes in 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 the intermediation of traders' demand. This has has three robust 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 independently of the asset's fundamental value. This profit is obtained through intermediation of traders' demand. Second, trading price need not reflect any of the dealers' belief, 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 constraint. Given that belief-free equilibrium is more stringent than traditional solution concepts, it might be surprising that so much flexibility remains –in particular, equilibrium is not unique. Hence, we have focused on a belief-free equilibrium with a simple Markovian structure. When applied to a version of Glosten and Milgrom model, this explains well-documented stylized empirical facts. For specific microstructure games, it might then be reasonable to focus on belief-free equilibria that satisfy further criteria. For example, depending on the specific trading model considered, one could analyze strategies that maximize dealers' aggregate payoff or that minimize the expected time required by the market measure to point at the true state, or even strategies that minimize the aggregate cost of learning, or more generally strategies that form a belief-free equilibrium for the lowest possible level of dealers' patience.

Appendix

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. 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 $u^\star$ and $\overline{u}$, with probability q_t, dealer i stage t payoff is at least $(1-\varepsilon)u^\star - \varepsilon\overline{u}$. Then, at time $\tau \geq 0$, dealer i's payoff satisfies

$$\begin{aligned} V_i^\delta(\omega, \sigma | h^\tau) &> (1-\delta)\sum_{t=\tau}^{\infty} \delta^{t-\tau}\left(q^t((1-\varepsilon)u^\star - \varepsilon\overline{u}) - (1-q^t)\overline{u}\right) \\ &= (1-\varepsilon)(u^\star + \overline{u})(1-\delta)\sum_{t=\tau}^{\infty}\delta^t q^{t-\tau} - \overline{u}. \end{aligned} \tag{24}$$

Now condition 1 of Definition 2, implies that

$$\Pr_{\omega,\sigma}\left[\lim_{\delta\to 1}(1-\delta)\sum_{t=\tau}^{\infty}\delta^t q^{t-\tau} > 1-\varepsilon\right] > 1-\varepsilon. \tag{25}$$

Hence we have that

$$\lim_{\delta\to 1} V_i^\delta(\omega, \sigma | h^\tau) > (1-\varepsilon)^3(u^\star + \overline{u}) - (1+\varepsilon)\overline{u}. \tag{26}$$

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

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

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

$$0 < u_i(\omega, \tilde{a}^i(\hat{\omega})) < u_i(\omega, \tilde{a}^j(\hat{\omega})) < u(\omega, a). \tag{27}$$

for all $i \neq j$.

Proof. Consider the convex combination

$$\tilde{a}^i(\hat{\omega}) := \beta_1(\hat{\omega})\beta_2(\hat{\omega})\underline{a}(\hat{\omega}) + \beta_1(\hat{\omega})\left(1 - \beta_2(\hat{\omega})\right) a^i(\hat{\omega}) + \left(1 - \beta_1(\hat{\omega})\right) a, \tag{28}$$

for some $\beta_1(\hat{\omega}), \beta_2(\hat{\omega}) \in [0,1]$, where $\underline{a}(\hat{\omega})$ satisfies Assumption **B**-2, and $a^i(\hat{\omega})$ is as in Assumption **B**-4. Note that $\{\tilde{a}^i(\hat{\omega})\}_{i=1,\ldots,n}$ also satisfies Assumption **B**-4, as long as $\beta_1(\hat{\omega}) > 0$, $\beta_2(\hat{\omega}) < 1$. Furthermore, because $u\left(\omega, \underline{a}(\hat{\omega})\right) < 0$, we can pick $\beta_2(\hat{\omega})$ 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 couple $\hat{\omega} \neq \hat{\omega}'$. 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 $\sigma^{i,\varepsilon}\left(h^t\right) = (1-\varepsilon)\,\tilde{a}^i(\hat{\omega}) + (\varepsilon/L)\,\Sigma_{a \in A_L} a$ For all other h^t, let $\sigma^{i,\varepsilon}\left(h^t\right) = (1/L)\,\Sigma_{a \in A_L} a$.

In addition, define n partial "punishment" strategies $\underline{\sigma}^{i,\varepsilon}$ as follows. Fix any $\hat{\omega} \in \hat{\Omega}$. Condition **B**-3 guarantees that we can extend the Blackwell (1956) approachability argument to the discounted case: for any $\eta > 0$ there is $\delta^\eta < 0$, $m^\eta < \infty$ and m^η-period strategy $\underline{a}_{-i}(\hat{\omega})$ for player $-i$ such that if $\delta > \delta^\eta$, for any sequence $\{a_i^1, \ldots, a_i^{m^\eta}\}$ player i discounted payoff during these m^η periods is smaller than η in each $\omega \in \hat{\omega}$. This Blackwell strategy is then an ingredient for 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 probability at least ε to all $\omega \in \hat{\omega}_i$, then $\underline{\sigma}^{i,\varepsilon}\left(h^t\right) = (1-\varepsilon)\,\underline{a}^i(\hat{\omega}_i)(h^t) + (\varepsilon/L)\,\Sigma_{a \in A_L} a$, 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 $V_i^\delta\left(\omega, \sigma_i, \underline{\sigma}^{i,\varepsilon}_{-i}|h^t\right)$ is such that

$$\lim_{\delta \to 1, \varepsilon \to 0} V_i^\delta\left(\omega, \sigma_i, \underline{\sigma}^{i,\varepsilon}_{-i}|h^t\right) \leq 0. \tag{29}$$

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. As soon as 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 T large enough, and ε small enough, there exists $\underline{\underline{\delta}} \leq \underline{\delta} < 1$ such that for all $\delta \in \left(\underline{\underline{\delta}}, 1\right)$, players do not gain from deviating.

Note that this construction yields a belief-free equilibrium: The strategy are optimal irrespective of dealers' beliefs about ω on and off the equilibrium path.

■

References

[1] Back, K. and S. Baruch, 2004. "Information in Security Markets: Kyle meets Glosten and Milgrom," *Econometrica*, **72**, 433–465.

[2] Battalio, R., A. Ellul, and R. Jennings, 2007. "Reputation Effects in Trading on the New York Stock Exchange," *Journal of Finance*, **62**, 1243–1271.

[3] Benveniste, L., A. Marcus, and W. Wilhelm, 1992. "What's so Special about the Specialist?," *Journal of Financial Economics*, **32**, 61–86.

[4] Blackwell, D., 1956. "An Analog of the Minimax Theorem for Vector Payoffs," *Pacific Journal of Mathematics* **6**, 1–8.

[5] Bloomfield, R. and M. O'Hara, 2000. "Can Transparent Markets Survive?," *Journal of Financial Economics*, **55**, 425–459.

[6] Boehmer, E. and J. Wu, 2008. "Order Flow and Prices," *mimeo.*

[7] Boulatov, A. and T. George, 2010. "Securities Trading when Liquidity Providers are Informed," *mimeo.*

[8] Brunnermeier, M., L. Pedersen, 2009. "Market Liquidity and Funding Liquidity, " *Review of Financial Studies*, **22**, 2201–2238.

[9] Calcagno, R. and S. Lovo, 2006. "Bid-ask Price Competition with Asymmetric Information between Market-makers," *Review of Economic Studies*, **73**, 329–355.

[10] Chordia, T., R. Roll, and A. Subrahmanyam, 2002. "Order imbalance, liquidity and market returns," *Journal of Financial Economics*, **65**, 111–130.

[11] Christie, W. and P. Schultz, 1994. "Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?," *Journal of Finance*, **49**, 1813–1840.

[12] Christie, W., J. Harris and P. Schultz, 1994. "Why Did NASDAQ Market Makers Stop Avoiding Odd-Eighth Quotes?," *Journal of Finance*, **49**, 1841–1860.

[13] Cont, R., 2001. "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues," *Quantitative Finance*, **1**, 223–236.

[14] Desgranges, G. and T. Foucault, 2005. "Reputation-Based Pricing and Price Improvements," *Journal of Economics and Business*, **57**, 493–527.

[15] De Frutos, M. and C. Manzano, 2005. "Trade Disclosure and Price Dispersion," *Journal of Financial Markets*, **8**, 183–216.

[16] De Meyer, B., 2010. "Price Dynamics on a Stock Market with Asymmetric Information," *Games and Economic Behavior*, **69**, 42–71.

[17] Du,S. and H. Zhu, 2012. "Ex-post Equilibria in Double Auctions of Divisible Assets," *mimeo*

[18] Dutta, P. and A. Madhavan, 1997. "Competition and Collusion in Dealers Markets," *Journal of Finance*, **52**, 1817–1846.

[19] Easley, D. and M. O'Hara, 2010. "Microstructure and Ambiguity," *Journal of Finance*, **65**, 2289–2316.

[20] 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.

[21] Evans, M. and R. Lyons, 2002. "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy*, **110**, 170–180.

[22] Fleming, J., C. Kirby and B. Ostdiek, 2006. "Information, Trading, and Volatility: Evidence from Weather-Sensitive Markets," *Journal of Finance*, **61**, 2899–2930.

[23] Fudenberg, D. and E. Maskin, 1986. "Folk Theorem for Repeated Games with Discounting or with Incomplete Information," *Econometrica*, **54**, 533–554.

[24] Fudenberg, D. and Y. Yamamoto, 2010. "Repeated Games Where the Payoffs and Monitoring Structure are Unknown," *Econometrica*, **78**, 1673–1710.

[25] Fudenberg, D. and Y. Yamamoto, 2011. "Learning from Private Information in Noisy Repeated Games," *Journal of Economic Theory*, **146**, 1733–1769.

[26] 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.

[27] Gromb, D. and D. Vayanos, 2002. "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs,"*Journal of Financial Economics*, **66**, 361–407.

[28] 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.

[29] Harsanyi, J., 1967–1968. "Games of Incomplete Information Played by Bayesian Players, Parts I-III. "*Management Science*, **14**, 159–182, 320–334, 486–502.

[30] Hasbrouck, J. and G. Sofianos, 1993. "The Trades of Market Makers: An Empirical Analysis of NYSE Specialists," *Journal of Finance*, **48**, 1565–1593.

[31] Ho, T. and H. Stoll, 1981. "Optimal Dealer Pricing under Transactions and Return Uncertainty," *Journal of Finance*, **9**, 47–73.

[32] Hörner, J. and S. Lovo, 2009. "Belief-free Equilibria in Games with Incomplete Information," *Econometrica*, **77**, 453–487.

[33] Hörner, J., S. Lovo and T. Tomala, 2009. "Belief-free Equilibria in Games with Incomplete Information: Characterization and Existence," *Journal of Economic Theory*, **146**, 1770–1795.

[34] Leach, C. and A. Madhavan, 1993. "Price Experimentation and Security Market Structure," *Review of Financial Studies*, **6**, 375–404.

[35] Madhavan, A., 2000. "Market Microstructure: A Survey," *Journal of Financial Markets*, **3**, 205–258.

[36] 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.

[37] Pasquariello, P. and C. Vega, 2005. "Informed and Strategic Order Flow in the Bond Market," *mimeo.*

[38] Reiss, P. and I. Werner, 1998. "Does Risk Sharing Motivate Interdealer Trading?," *Journal of Finance*, **53**, 1657–1703.

[39] Reiss, P. and I. Werner, 2005. "Anonymity, Adverse Selection, and the Sorting of Interdealer Trades," *Review of Financial Studies*, **18**, 559–636.

[40] Roëll, A., 1988. "Regulating Information Disclosure Among Stock Exchange Market Makers," FMG Discussion Papers Series 51, London School of Economics.