

Robust Price Formation*

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

November 18, 2014

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: Excess price volatility, price/trading-flow correlation, stochastic volatility 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. This means that dealers’ actual information structure (*i.e.*, what each dealer believes about the market fundamentals and about what other dealers believe, and so on) is quite complex and might vary over time and across markets.¹ However, because the information structure is largely dependent on the dealers’ subjective beliefs, it is not directly observable in practice. As a consequence, it is essentially impossible to assess the extent to which a given theoretical model’s assumptions on dealers’ information structure capture or not actual informational asymmetries. In existing market microstructure models, tractability imposes strong informational assumptions, and involves specific functional forms assumptions regarding the distribution of fundamentals and private signals.² Yet, because modeling dynamic interactions among asymmetrically informed dealers is a formidable task, the current theoretical literature is silent about the robustness of predictions of standard microstructure models to changes in the dealers’ information environment.³

The objectives of this paper are, first, to present a tractable price-formation theory delivering predictions that are robust to the specification of this information; and second, to provide the

¹See Bergemann and Morris (2013) for a definition of information structure in games of incomplete information.

²For instance, almost all models assume that trading prices are set by equally uninformed dealers to a level reflecting these dealers’ beliefs about fundamentals.

³At each 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.

market microstructure theorist with a practical toolkit of sufficient conditions allowing to check whether a candidate equilibrium is robust to changes in information structure, and of necessary conditions that allows to discard equilibria that are fragile. For this purpose, we consider a class of dynamic financial markets microstructure models in which risk-neutral liquidity providers (such as dealers or market-makers) interact with traders. For this class of models, we characterize equilibria that are robust to the extent that dealers' dynamic pricing strategies remain optimal no matter the private information and beliefs of a dealer about the economy fundamentals. In particular a robust equilibrium is a subgame perfect equilibrium *no matter the dealer's belief*. That is, among the many price formation strategies that form sub-game perfect equilibria for some specification of dealers' beliefs, we focus on those strategies that form a sub-game perfect equilibria no matter the specification of dealers' beliefs. For this reason we will often refer to robust equilibria as 'belief-free equilibria' (henceforth BFE).⁴

We believe that robust equilibria are interesting for a number of reasons. First, in terms of their scope: as we show, robustness implies that hardly any assumption on the dealers' information is called for. This is less restrictive than assuming that all dealers share the same exact beliefs about fundamentals. Also minimal assumptions are required concerning trading protocols, and the concept is tractable enough to allow for models that encompass many real-world trading protocols. Second, in terms of their ability to explain seemingly unrelated empirical findings. For a Glosten and Milgrom-type model, we show that the robust equilibria supports various stylized facts that cannot be explained using the textbook solution concept: positive profits for liquidity providers, stochastic price volatility, price bubbles, and bounded dealers' inventories. While each of these facts can be explained by existing models separately, we are aware of no single model delivering them simultaneously. Third, in terms of tractability: we actually focus on a subclass of robust equilibria that are arguably as tractable as the classical zero-profit models in market

⁴See 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.

microstructure.

In robust equilibria price formation strategies must be optimal no matter the dealer's information about the value of the fundamentals. Thus, the notion of robust equilibrium is more demanding than standard game-theoretic refinements. As a consequence, robust equilibria need not exist for static trading games. Nevertheless, we show that when moving to a dynamic framework, most common market microstructure games possess equilibria that are robust as long as the discount factor between two trading rounds is large enough.⁵ This condition is naturally met for market microstructure economies where trading frequency is intra-day.

BFE, i.e. robust equilibria impose the following testable restrictions on the equilibrium outcomes:

1) Dealers can gain or lose money over short periods, but their average long-run profit is strictly positive independently of the asset's fundamental value. This is in contrast with the classical prediction that dealers' expected per trading period profit is nil and demonstrates that nil expected profits of price makers is not a robust feature of price formation.⁶

In a BFE, dealers make positive profits through the intermediation of traders' order flow rather than by taking large positions on assets. Interestingly, these predictions are consistent with the empirical study of NYSE specialists by Colmerton-Forde, Hendershott, Jones, Moulton and M. Seasholes (2010). They find that "Specialists in aggregate lose money on [only] about 10% of the trading days" and that "specialists usually earn positive trading revenue on short-term (intraday) round-trip transactions." Our finding de-emphasizes the role of information about the asset fundamentals on dealers' behavior. In a BFE, what matters for a dealer are the levels of quotes that induce an abundant but balanced order flow from traders. Knowing these levels is

⁵In fact, multiple robust equilibria then exist, but recall that, because it is a refinement, this multiplicity is necessarily less severe than with other solution concepts.

⁶The zero profit conditions is often justified through the hypothesis of free-entry of new dealers. However, it requires the additional assumption that incumbents and entrant dealers have exactly the same information. Without this simplifying informational assumption, it is a priori unclear what free-entry would imply for dealer's profits.

necessary and sufficient for dealers in order to be able to realize many profitable round-trips. Hence dealers can ignore all information that does not affect traders' behavior.

2) Risk-neutral dealers tend to maintain balanced inventories. That is, dealers' inventories are mean reverting. This contrasts with the view that (absent risk aversion or institutional constraints on inventory size) inventory levels should not affect dealers' behavior.⁷ Thus, BFE provides an alternative explanation of the empirical evidence of market makers mean reverting their inventories.⁸

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 in the long run prices reflect all public information. That is, an asset trading price never stabilizes around its fundamental value. This no matter the amount of public information. In particular, 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. This generates patterns that are consistent with the phenomena of excess, and stochastic price volatility, two well known and empirically established properties of stock prices.

The intuition behind these results is as follows. First, dealers can always guarantee zero profit by abstaining from trading. Because in equilibrium dealers' strategies 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

⁷For 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.

⁸See for example Madhavan and Smidt (1993), Hansch, Naik, and Viswanathan (1999), Reiss and Werner (1998), and Naik and Yadav (2003).

belief about fundamentals. On the contrary, when the equilibrium strategy leads to sufficiently balanced inventories, fundamentals has little impact on dealers' profits. As a result, in a robust equilibrium, 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 flows from traders. Quotes react to the order flow, which 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 indefinitely, and while they point to the right exploiting strategy more often than not, with low probability they also lead to incorrect exploiting rounds. Hence, a Bayesian dealer could possibly believe that the current exploiting phase is incorrect. 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 exploring phases, the order flow is informative, leading to quotes that are highly sensitive to the volume of trade. In exploiting phases, trading flow is balanced and originates from liquidity traders; as a result, quote volatility is reduced. Transition to and from low to high volatility phases is stochastic and depends on the evolution of trading histories. Thus, stochastic price volatility results from the alternation of exploring and exploiting phases.

In the first part of the paper, we consider a broad 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 that we analyze is rich along

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 allowed.

We show that a dynamic trading game admits robust 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 whenever inter-dealer trading is allowed.

Despite the restrictiveness of robust equilibria, there remains considerable leeway in specifying belief-free equilibria. Rather than delineating precisely the scope of these equilibria, we take advantage of this leeway to focus on a tractable subset that accords with additional regularities documented in the empirical literature.

In particular, in the second part of the paper, we illustrate the functioning of such belief-free equilibria in the simple framework of the Glosten and Milgrom (1985) model where we make no assumption about each dealers' private information. We compare the robust equilibria with the zero-profit equilibrium (henceforth GME) described by Glosten and Milgrom (1985). Both GME and BFE explain correlation between trading flow and price changes, a pattern that is consistent with a wide body of empirical work.⁹ 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. Thus, while excess volatility and

⁹See for instance, Chordia, Roll and Subrahmanyam (2002) and Boehmer and Wu (2008) for security markets, Pasquariello and Vega (2005) for bond markets, Evans and Lyons (2002) for currency markets and Fleming, Kirby and Ostdiek (2006) for weather-sensitive commodity markets.

volatility clustering are trademarks in a BFE, they are not explained by the GME.

Another stylized fact that a BFE can explain but not the GME is that liquidity is negatively correlated with dealers' inventory, dealers' profits, and price volatility (as documented by Colmerton-Frode et al. (2010)).

The applications of repeated games to the market microstructure that are closest to our work are Dutta and Madhavan (1997) Benveniste et al. (1992) and Desgranges and Foucault (2005). 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.¹⁰

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 about 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, giving rise to a belief-free like equilibrium.

The paper is organized as follows. Section 1 presents the general framework of our model. Section 2 describes some salient properties of the single stage trading game. Section 3 defines BFE in the repeated game. Section 4 and 5 presents necessary and sufficient conditions, respectively, for a strategy to form a BFE of the repeated game. Section 6 presents an example based on the

¹⁰The same results hold if the frequency of trading is high compared to the frequency with which the state of nature changes.

model of Glosten and Milgrom and compares the BFE with the canonical equilibrium. Section 7 discusses extensions to imperfect monitoring about dealers' actions, non-stationary states of nature and dealers' strategies based on private information. Section 8 concludes. All proofs are in Appendix.

1 A Model of Price Formation

A risky asset is exchanged for money among short-lived risk averse traders and $n > 1$ long-lived risk neutral dealers (n is finite). Trading takes place over infinitely many periods $t = 1, 2, \dots$. At time 0 and once for all, Nature choses the state ω in the set Ω . We allow the state of Nature to affect the economy in two dimensions: the fundamental value of the asset and the composition of the population of traders. We denote with $W(\omega)$ the fundamental value of the asset in state ω . We assume that $W(\omega) = v(\omega) + e(\omega)$, where \tilde{v} and \tilde{e} are independently distributed and bounded. Furthermore we assume $E[\tilde{e}] = 0$. We will denote with \underline{e} and \bar{e} , the smallest and largest possible value of \tilde{e} . As in Back and Barush (2004), a public release of information takes place at a random time τ , 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 $W(\omega)$.

In every period a randomly selected trader comes to the market, trades and leaves the market. A trader is characterized by a triple $\theta = (g, y, c)$, its type, where $g : \mathbb{R} \rightarrow \mathbb{R}$ is the trader's utility function that we assume to be increasing and concave, y is his initial endowment of the risky asset and c his initial endowment of cash. We assume that θ belongs to a compact set Θ . We denote with $Z(\theta, \omega)$ the probability that in any given period t the trader is of type $\theta \in \Theta$ given that the state of nature is ω .

Traders know the realization of \tilde{v} but not the realization of \tilde{e} . There is no correlation between the distribution of traders type and \tilde{e} . This implies that information about the composition of

traders population is useless to learn $\tilde{\epsilon}$. We remain agnostic about dealers' information structure. Namely, we make no assumption about the private information of any given dealer regarding the true state ω . In a belief-free equilibrium, defined below, each dealer's strategy must be optimal no matter the true state ω .

2 Stage trading round

2.1 Trading mechanism

In every period t each dealer i chooses an action in the finite set A_i . A dealers' action profile a is an element in the set $A := \times_i A_i$. We denote with $\tilde{a} \in \Delta A$ a possibly correlated action profile for dealers and with $\tilde{a}(a)$ the probability that the outcome of the correlated action profile be $a \in A$ when dealers play \tilde{a} . The set of actions available to time t trader is S finite. To fix ideas, we call $a \in A$ a dealers' action profile, and $s \in S$ a trader's reaction. To any given action-reaction profile $(a, s) \in A \times S$ corresponds a transaction (possibly nil). Let $Q_i(a, s)$ and $P_i(a, s)$ denote the transfer to dealer i of the risky asset and money, respectively, resulting from the transaction. Transfers to the trader are denoted $Q_T(a, s)$ and $P_T(a, s)$. Thus, if the state of nature is ω , an action-reaction profile (a, s) translates into a change in the wealth of agent j given by

$$W(\omega)Q_j(a, s) + P_j(a, s). \tag{1}$$

We assume that transfers of cash and asset belong to finite grids G and G_q , respectively. Both grids G and G_q include the point 0. In what follows, if $-\frac{P_T(a,s)}{Q_T(a,s)} = p > 0$, we say that the *trader buys at price p* if $Q_T(a, s) > 0$ and that the *trader sells at price p* if $Q_T(a, s) < 0$. The trading protocol defined by the functions $Q(\cdot)$ and $P(\cdot)$ is such that no agent can be forced into trade. Namely, no matter what the other agents do, first each agent can chose not to trade. Second, for any positive price $p \in G$, and quantity $q \in G_q$, each agent j can chose his action so that, if

he trades, then he trades at price p and at most q shares of the asset.

2.1.1 Applications

Here, we illustrate how some of the trading protocols analyzed in the literature fit into our framework.

Quote driven markets (Glosten and Milgrom (1985)) The set of actions available to each dealer is the set of bid and ask quotes to be chosen on the price grid G . Formally $A_i = G \times G$. The trader observes dealer's quotes and chooses whether to buy one unit, sell one unit, or not trading, that is $G_q = \{-1, 0, 1\}$. Thus, $S = \{sell, no-trade, buy\}$ is the set of possible market orders that can be chosen by a trader. A trader's market order is executed against the best dealers' quotes.

Quote driven market (Biais, Martimort and Rochet (2000)) The set of actions available to each dealer is a schedule $T(\cdot)$, which specifies his willingness to trade $q \in G_q$ shares of the asset against transfer of $T(q)$, on the grid G . Thus $A_i = [G]^{M_q}$, where M_q denote the cardinality of G_q . The trader observes dealer's schedules and chooses how many shares to trade with each dealer. Thus, $S = [G_q]^n$ and a trader reaction s specifies how many shares the trader exchanges with each one of the n dealers.

Limit order markets There is no unanimous way to model limit order markets (see for example Foucault (1999), Goettler et al. (2005), Foucault et al. (2005), Rosu (2009) among others). We present one possible specification that captures the functioning of a limit order market. First, at the beginning of the period, each dealer submits a limit order where prices and quantities belong to the grids G and G_q . This generates a book of limit orders associated

to each price on the grid G . Second, time t trader chooses whether to submit a market order which trades against outstanding orders in the book, and/or a limit order at a specified price, which enters the book at that price. Third, dealers submit market orders which trades against the book. The set of actions for both dealers and traders is $A_i = S = G \times G_q \times G_q$. Namely, $\{(p_j, q_j), m_j\} \in A_j$ specifies agent j limit and market order. Note that because $0 \in G_q$, each market participant can choose not to submit a limit and/or market order.

2.2 Traders and Dealers

Traders In each period t , a trader is drawn from a population of traders, trades to adjust the composition of their portfolio and then leaves the market. Traders are investors who trade both for speculative and hedging reasons. Formally, time t trader comes to the market knowing the realization of \tilde{v} but having no information about \tilde{e} . He trades to maximize the expected utility of his post-trade wealth and then leaves the market. Namely if $v(\omega) = v$, in time t dealers action is $a \in A$ and the trader's type is $(g, y, c) \in \Theta$, then the trader's reaction will be

$$D(v, a) := \operatorname{argmax}_{s \in S} E [g((v + \tilde{e})(Q_T(a, s) + y) + P_T(a, s) + c)]$$

Let denote with $\theta(v, a, s) \subseteq \Theta$ the set of traders' types such that $D(v, a) = s$; where $s \in S$. Then,

$$F(\omega, a, s) := Z(\theta(v(\omega), a, s), \omega)$$

denotes the probability that time t trader chooses reaction s , given that dealers actions are a and that the state of Nature is ω .

Dealers There are n finite dealers they are long lived and are risk neutral. We make no assumption about the private information of any given dealer regarding the true state ω . However for each realization of ω we can compute dealers' payoffs in one trading round resulting from any

given action profile $a \in A$ set by the dealers. Let $u_i(\omega, a)$ denote dealer i 's expected payoff, or *reward*, in a *single period* t given the state ω and dealers' action profile a . Here, expectations are taken with respect to the possible trader's actions s given the state ω . Namely,

$$u_i(\omega, a) = \sum_{s \in S} F(\omega, a, s) (W(\omega)Q_i(a, s) + P_i(a, s)) \quad (2)$$

When considering a distribution over dealers' action profiles $\tilde{a} \in \Delta A$, dealer i expected payoff in state ω is

$$u_i(\omega, \tilde{a}) := \sum_{a \in A} \tilde{a}(a) u_i(\omega, a),$$

We also define the expected change in inventory for dealer i resulting from $\tilde{a} \in \Delta A$ in state ω as

$$Q_i(\omega, \tilde{a}) := \sum_{a \in A} \tilde{a}(a) \sum_{s \in S} F(\omega, a, s) Q_i(a, s)$$

2.3 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.¹¹ 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 is¹²

$$\sum_{t=1}^{\infty} (1 - \delta) \delta^t u_i(\omega, a^t). \quad (3)$$

¹¹Allowing for a stochastic discount factor complicates exposition but does not affect results as long as the discount factor remains close enough to one.

¹²Here, expectation is taken with respect to the possible realizations of traders' orders $\{s^t\}_{t=1}^\infty$, taking the state ω as given.

A public strategy profile (strategy henceforth) is a mapping $\sigma : \cup_t H^t \rightarrow \times_i \Delta A_i$ that associate to each public history h^t the action profile that dealers play at time t . For any given state ω , and any finite 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 that we denote $\tilde{a}_{(\omega, \sigma, h^t)} \in \Delta A$.¹³ 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 . Then we have

$$V_i(\omega, \sigma | h^t) = \sum_{a \in A} \tilde{a}_{(\omega, \sigma, h^t)}(a) u_i(\omega, a) = W(\omega) Q_i(\omega, \sigma | h^t) + P_i(\omega, \sigma | h^t),$$

where

$$Q_i(\omega, \sigma | h^t) := \sum_{a \in A} \tilde{a}_{(\omega, \sigma, h^t)}(a) \sum_{s \in S} Q_i(a, s) F(\omega, a, s)$$

is the expected change in dealer i ' asset inventory, and

$$P_i(\omega, \sigma | h^t) := \sum_{a \in A} \tilde{a}_{(\omega, \sigma, h^t)}(a) \sum_{s \in S} P_i(a, s) F(\omega, a, s)$$

is the expected change in dealer i 's cash holding.

We are interested in belief-free equilibria of the repeated game (hereafter, BFE) that, loosely speaking, are sub-game perfect equilibria that are robust to any specification of each dealer information about the true ω . Because dealers' beliefs might differ arbitrarily, in a BFE, dealer's strategy must be optimal no matter what the true realization of ω is. In other words, we look for a dealer's strategy profiles σ such that at any point in time t , after any history h^t , for any dealer i , it is optimal to chose action according to $\sigma_i(h^t)$ no matter what he believes about the

¹³Formally, an occupation measure of $a \in A$ is the discounted expected frequency with which action profile a will be played:

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

true state of nature ω . Formally,

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:*

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i} V_i(\omega, \sigma_i, \sigma_{-i}^* | h^t), \quad (4)$$

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

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.¹⁴ Thus, a BFE is a subgame-perfect equilibrium no matter the specific dealers' information structure. Second, a BFE is an equilibrium no matter whether dealers are Bayesian or not. Beside, 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 $\omega \in \Omega$.¹⁵

3 Learnable and non-learnable states

We start by defining what any dealer can learn about the true state ω no matter what his private information is. This corresponds to what can be learned about ω from observing how traders react to dealers' actions. Formally, the way the function F is affected by the state of Nature determines the information about the true ω that can be gathered from traders' behavior.

¹⁴To see this, note that in a perfect Bayesian equilibrium, dealers' strategies satisfy

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

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 also a perfect Bayesian equilibrium, but a perfect Bayesian equilibrium need not be belief-free.

¹⁵Unlike in Easley and O'Hara (2010), where some of the traders are ambiguity averse, here ambiguity aversion applies to dealers.

For example, if traders' behavior is identical in state ω and in state ω' , then it is impossible to tell apart these two states even after observing an infinite history of trades. Namely, F defines a partition, that we assume to be finite, $\hat{\Omega} := \{\hat{\omega}_1, \hat{\omega}_2, \dots, \hat{\omega}_k\}$ of the space Ω . Formally, partition $\hat{\Omega}$ satisfies the following two conditions:

1. Two states $\omega, \omega' \in \Omega$ belong to the same element $\hat{\omega} \in \hat{\Omega}$ if and only if, for any given $(a, s) \in A \times S$, it results

$$F(\omega, a, s) = F(\omega', a, s)$$

In words, traders behavior is identical in state ω and on state ω' . In this case we say that ω and ω' are *not statistically distinguishable* through traders' behavior.

2. Conversely, suppose that two states $\omega, \omega' \in \Omega$ satisfy $\omega \in \hat{\omega}$, $\omega' \in \hat{\omega}'$, where $\hat{\omega}, \hat{\omega}' \in \hat{\Omega}$ and $\hat{\omega} \neq \hat{\omega}'$. Then there exists a non-empty set $A(\hat{\omega}, \hat{\omega}') \subseteq A$ such that if $a \in A(\hat{\omega}, \hat{\omega}')$, then

$$F(\omega, a, s) \neq F(\omega', a, s)$$

for some $s \in S$.

In words, there are suitable choices of dealers' actions (i.e., for $a \in A(\hat{\omega}, \hat{\omega}')$), for which the distribution of traders' reactions s differs for $\omega \in \hat{\omega}$ and $\omega \in \hat{\omega}'$. In this case we say that ω and ω' are *statistically distinguishable* through traders' behavior.

In other words, a sufficiently long observation of how traders react to dealers' action allows to identify the element $\hat{\omega} \in \hat{\Omega}$ that contains the true state ω . However nothing in traders behavior will ever allow to say which element $\omega \in \hat{\omega}$ is the true state. In particular, because traders have no information about the \tilde{e} of the asset value, and the distribution of their types does not depends on \tilde{e} , their behavior does not depend on the actual value of \tilde{e} . Hence, F satisfies the following property:

Non-learnable States (NLS): *Traders behavior does not allow to learn the $e(\omega)$ component of $W(\omega)$.*

Beyond property **NLS**, the precise shape of F depends on the particular choice of the distribution Z over traders utility function and starting portfolios. However, because traders know the component $v(\omega)$ of the asset value $W(\omega)$, in general their order flow can be affected by the true value of $v(\omega)$. More precisely, we assume that there are suitable choices of dealers' actions, that induce a distribution of traders' reactions s that is measurable with respect to $v(\omega)$:

Learnable States (LS): *if $v(\omega) \neq v(\omega')$ then ω and ω' are statistically distinguishable.*

Property **LS** implies that it is possible to choose the sequence of dealers' actions such that a sufficiently long history of the traders' reactions allows to statistically learn the value of $v(\omega)$. Thus property **LS** implies that, contrary to what happens for $e(\omega)$, the $v(\omega)$ component of $W(\omega)$ is *statistically learnable* from traders' behavior.

4 Achievable payoffs for dealers

We now consider the payoffs that dealer can achieve in a single trading round when their actions are only based on what can be learned from traders' behavior. To this purpose we introduce a third property of the function F that follows from the fact that traders seek to maximize their post-trade utility function, they are risk averse, come to the market with some inventory of the risky asset and knowing $v(\omega)$. This implies that traders never buy the asset at a price that is too high nor sell at a price that is too low when compared to $v(\omega)$. However traders who initially have large enough short or long position in the risky asset will be willing to buy at price strictly larger than $v(\omega)$ and to sell at price strictly smaller than $v(\omega)$, respectively.¹⁶ In

¹⁶This guarantees that there is room for trade between dealers and traders. Otherwise the Milgrom and Stokey (1982) no-trade theorem applies.

other words, there always is a positive probability that time t trader is willing to buy or to sell the asset at price p , as long as p is not too far from $v(\omega)$. Formally,

Elastic Trader Demand (ETD): *There are $\bar{\rho} \geq \underline{\rho} > 0$ such that for all $\omega \in \Omega$,*

1. *If (a, s) are such that trader buys at price $p < v(\omega) + \underline{\rho}$, then $F(\omega, a, s) > 0$.*
2. *If (a, s) are such that trader sells at price $p > v(\omega) - \underline{\rho}$, then $F(\omega, a, s) > 0$.*
3. *If (a, s) are such that trader buys at price $p > v(\omega) + \bar{\rho}$, then $F(\omega, a, s) = 0$.*
4. *If (a, s) are such that trader sells at price $p < v(\omega) - \bar{\rho}$, then $F(\omega, a, s) = 0$.*

We are interested in distributions over dealers' quotes allowing them to achieve different profiles of payoffs in a single trading round. Namely, situations where all dealers make positive profits, situations where all dealers make negative profits, situation where a dealer cannot make positive profit once the other dealers know his beliefs on ω , and situations where dealers' profits differ. Properties **LS**, **NLS** and **ETD** guarantees that such distributions exist but might depend on the true value of the $v(\omega)$ component or more generally from the $\hat{\omega} \in \hat{\Omega}$ containing the true state ω .

Proposition 1 *If F satisfies properties **LS**, **NLS** and **ETD**, then for any given $\hat{\omega} \in \hat{\Omega}$,*

1. *Positive maximum payoffs: There exists a non-empty set $A^*(\hat{\omega}) \subseteq \Delta A$, if and only if $u_i(\omega, a) > 0$ for all $\tilde{a} \in A^*(\hat{\omega})$, $\omega \in \hat{\omega}$ 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 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,

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

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 $\omega \in \hat{\omega}$.

In words, the set $A^*(\hat{\omega})$ is the set of dealers action profile (possibly mixed) guaranteeing that each dealer i makes strictly positive profits in all states $\omega \in \hat{\omega}$. Let $u^* > 0$ denote a lower bound on payoffs from actions in $A^*(\hat{\omega})$. Roughly speaking, *Positive maximum payoffs* and *Negative minimum payoffs* 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, respectively. *Non-positive expected payoffs* guarantee that when fixing a dealer i and his beliefs about the true ω , the other dealers can guarantee that this dealer's expected payoff is non-positive. This is achieved for example if the other dealers provide the maximum liquidity at a price equal to the expected value of the asset for dealer i . *Non-equivalent payoffs* 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.

As we will show below, the fact that the stage trading game satisfy these four properties is key in constructing belief-free equilibria of the repeated game.

We conclude this section with an example of the Glosten and Milgrom (1985) economy to illustrate properties **NLS**, **LS** and **ETD** and Proposition 1

Example 1 Consider the Glosten and Milgrom trading mechanism described above, i.e., $A = G \times G$ and $S = \{\text{sell}, \text{no-trade}, \text{buy}\}$. Let α_i, β_i be the bid and ask quotes set by dealer i and let α and β be the best ask and best bid across dealers, respectively. Let $v(\omega) \in \{v_1, v_2\}$, with $v_1 < v_2$,

and $e(\omega) \in \{\underline{e}, \bar{e}\}$.

Traders All traders have the same utility function g that is strictly increasing and strictly concave, however they differ for the composition of their initial portfolio. The distribution Z of traders initial portfolio is state independent.

Then, it is possible to chose the distribution Z of traders initial portfolio so that traders behavior is described by the following function F :

$$F(\omega, a, \text{sell}) = \max \left\{ 0, \min \left\{ \frac{1}{2}, \frac{\beta - v(\omega) + \rho}{4\rho} \right\} \right\}, \quad (5)$$

$$F(\omega, a, \text{buy}) = \max \left\{ 0, \min \left\{ \frac{1}{2}, \frac{v(\omega) + \rho - \alpha}{4\rho} \right\} \right\}, \quad (6)$$

$$F(\omega, a, \text{no-trade}) = 1 - F(\omega, \beta, \text{sell}) - F(\omega, \alpha, \text{buy}). \quad (7)$$

where $\rho > \sqrt{2}(v_2 - v_1)$.¹⁷ Note that, whereas there are four possible states ω in Ω , one for each of the possible values of the asset, the partition $\hat{\Omega}$ has only two elements: $\hat{\omega}_1$ containing the two states where $v(\omega) = v_1$ and $\hat{\omega}_2$ containing the two states where $v(\omega) = v_2$.

Dealers Given a dealers quotes profile a , and a state ω , dealer i 's payoff is

$$u_i(\omega, a) = (W(\omega) - \beta_i)F(\omega, a, \text{sell})1_{\{\beta_i=\beta\}}\eta_\beta(a) + (\alpha_i - W(\omega))F(\omega, a, \text{buy})1_{\{\alpha_i=\alpha\}}\eta_\alpha(a)$$

where $\eta_\beta > 0$ and $\eta_\alpha > 0$ are tie-breaking rules applied in case more than one dealer sets the best bid or ask, respectively.

For an example of action profile in $a \in A^*(\hat{\omega})$, consider the case where all dealers set $\alpha_i = v(\hat{\omega}_1) + d$ and $\beta_i = v(\hat{\omega}_1) - d$ where $0 < d < \rho$. Then from equations (5) and (6) it follows that for any $\omega \in \hat{\omega}$ dealer i expected payoff is $u_i(\omega, a) = \eta_\beta(a)d(\rho - d)/4\rho + \eta_\alpha(a)d(\rho - d)/4\rho > 0$ no

¹⁷In words, the probability of a trader selling is increasing in the bid price β and decreasing in the $v(\omega)$ of the asset value. It is at most 1/2 and is nil if the bid price is smaller than $v(\omega) - \rho$. The probability of a trader buying is symmetric.

matter the value of $e(\omega)$.

For an example of action $\underline{a}_{-i}^i(\mu_{\hat{\omega}})$, consider that if dealer i has beliefs $\mu_{\hat{\omega}}$ then he values the asset exactly $W_{\mu_{\hat{\omega}}} := v(\hat{\omega}) + \sum_{\omega \in \hat{\omega}} \mu_{\hat{\omega}}(\omega)e(\omega)$. Thus if the other dealers bid and ask quotes are such that the best bid and the best ask are equal to $W_{\mu_{\hat{\omega}}}$, then dealer i expected profit cannot be strictly positive.

5 Necessary conditions for BFE price formation

The purpose of this section is to identify the features that are necessary for a strategy to form a BFE. This can be useful for sorting out the predictions of market-microstructure theory that are robust from those that rely on specific assumption about dealers information structure. Put it different, a strategy that forms a subgame perfect equilibrium for a specification of dealer's beliefs but that does not satisfy at least one of the properties described below is not robust to a changes in dealer's information structure.

First, note that a BFE must remain a sub-game perfect equilibrium even when dealers have no private information about ω , a situation we cannot exclude from our analysis. In this case, the only information about ω that can emerge is the one coming from traders' behavior. That is, dealers will be able to distinguish two states only if they are statistically distinguishable, i.e., if they belong to two different elements of the partition $\hat{\Omega}$. As a consequence, the equilibrium play cannot differ in two states ω and ω' belonging to the same element $\hat{\omega}$ of $\hat{\Omega}$. In particular, because traders' behavior provides no information about $e(\omega)$, the equilibrium distribution of trades cannot differ for two states ω and ω' that only differ for the value of \tilde{e} . Formally,

Lemma 1 (Measurability with respect to traders' behavior) *Let σ^* be a BFE, $\hat{\omega}$ an element of $\hat{\Omega}$ and h^t a finite history, then*

$$\tilde{a}_{(\omega, \sigma^*, h^t)} = \tilde{a}_{(\hat{\omega}, \sigma^*, h^t)},$$

for all states $\omega \in \hat{\omega}$.

Second, note that each dealer can guarantee zero profit by abstaining from trading. This implies that in a BFE, for each state $\omega \in \Omega$, each dealer's payoff cannot be strictly negative. Otherwise, for some beliefs, a dealer would prefer to deviate to the no-trade action.

Lemma 2 (Strictly positive dealers' payoffs) *Let σ^* be a BFE, $\hat{\omega}$ an element of $\hat{\Omega}$ and h^t a finite history, then for all states $\omega \in \hat{\omega}$,*

$$V_i(\omega, \sigma^* | h^t) \geq 0,$$

Moreover, if $Q_i(\hat{\omega}, \sigma^ | h^t) \neq 0$, then the weak inequality holds for at most one $\omega \in \hat{\omega}$.*

In other words, in a BFE dealers make non-negative profits state by state. In addition, when the equilibrium leads to changes in a dealer's inventory, he will make strictly positive profits in most states.

Third, Lemma 1 and 2 put a bound on the size of the inventories that dealers accumulate in a BFE. Namely, the same σ^* strategy can lead to positive payoffs in different states only if the equilibrium trading volume is relatively balanced and maintains dealers' inventories bounded. The intuition is simple. If on average dealers buy and sell the same quantity of the risky asset, then, because they are neither net buyers nor net sellers, their payoffs would not depend on $W(\omega)$ and would be positive as long as on average they sell the shares for more than what they paid them. By contrast, if for instance dealers take large positions in the asset, because of **ETD**, they would have to pay at least $v(\omega) - \bar{p}$ per share. But in a state where $e(\omega) < 0$ and small enough, dealers would lose money. Formally, let consider $Q(\omega, \sigma^* | h^t) := \sum_i Q_i(\omega, \sigma^* | h^t)$, i.e., the expected change in dealers aggregate inventory in state ω after history h^t in a BFE. This can be seen as the sum of two components: a negative component denoted $Q^-(\omega, \sigma^* | h^t) \leq 0$

that are transactions in which traders buy from dealers, and a positive component denoted $Q^+(\omega, \sigma^* | h^t) \geq 0$ that are transaction in which traders sell to dealers. The volume of trade between traders and dealers is $Q^+(\omega, \sigma^* | h^t) - Q^-(\omega, \sigma^* | h^t) \geq 0$. Then we have

Lemma 3 (Bounded dealers' inventories) *Let σ^* be a BFE and h^t a finite history, then for all states $\omega \in \Omega$,*

$$\frac{|Q(\omega, \sigma^* | h^t)|}{Q^+(\omega, \sigma^* | h^t) - Q^-(\omega, \sigma^* | h^t)} \leq \min \left\{ \frac{\bar{p}}{\underline{e}}, -\frac{\bar{p}}{\underline{e}} \right\} \quad (8)$$

Fourth, the fact that dealers' payoffs must be positive (Lemma 2) implies that if $\omega \in \hat{\omega}$, then a BFE must lead $a^t \in A^*(\hat{\omega})$ sufficiently often. When $a^t \in A^*(\hat{\omega})$, we say that the dealers are in a $\hat{\omega}$ -exploiting period. However, if for two elements $\hat{\omega}, \hat{\omega}' \in \hat{\Omega}$ one has $A^*(\hat{\omega}) \cap A^*(\hat{\omega}') = \emptyset$, then action profile leading to positive payoffs when $\omega \in \hat{\omega}$ would lead to negative payoffs if $\omega \in \hat{\omega}'$ and vice versa. In case dealers have no information about ω (a possibility we cannot rule out in our environment), it is the traders behavior that must allow the play to tell apart the two states and hence allow to use the appropriate exploiting actions. This is possible only if dealers actions belong to $A(\hat{\omega}, \hat{\omega}')$ sufficiently often. When $a^t \in A(\hat{\omega}, \hat{\omega}')$, we say that dealers are in an exploring period. Thus, both exploiting and exploring periods are necessary ingredients of a BFE and need to be played sufficiently often.

Lemma 4 (Frequent exploring) *Suppose that $A^*(\hat{\omega}) \cap A^*(\hat{\omega}') = \emptyset$, and let σ^* be a BFE and h^t a finite history. Then for any $\omega \in \Omega$ action profile in $A(\hat{\omega}, \hat{\omega}')$ are played with strictly positive occupation measure:*

$$\tilde{a}_{(\omega, \sigma^*, h^t)}(A(\hat{\omega}, \hat{\omega}')) > 0.$$

This lemma states that exploring periods must be relatively frequent, no matter the past history. As we illustrate below, because exploring period are associated with higher sensitivity of prices to trading volume, their recurrence makes prices more volatile than Bayesian beliefs on $\hat{\Omega}$.

6 Sufficient Conditions for BFE Pricing

In this section we show existence of BFE by constructing one. We first introduce the building blocks of our a candidate strategy. We then show how to combines these ingredients to obtains a BFE.

Equilibrium strategies ingredients: 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 $\hat{\omega}$. Fix $\varepsilon > 0$, and let us say that that the market measure π *points to* a state $\hat{\omega}$ if it attaches a probability $\pi(\hat{\omega}) \geq 1 - \varepsilon$ to state $\hat{\omega}$.

Let $\phi : \Pi \times A \times S \rightarrow \Pi$ 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 $\{a^\tau, s^\tau\}_{\tau=0}^{t-1}$ 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 $\sigma : \Pi \times S \rightarrow \Delta A$. Instead, a public strategy profile (strategy henceforth) is a mapping $\hat{\sigma} : \cup_t H^t \rightarrow \times_i \Delta A_i$.

Fix an arbitrary starting market measure $\pi^0 \in \Pi$, an updating rule ϕ and a partial strategy σ , and consider a situation where dealers use the partial strategy σ and the market measure evolves according to ϕ . We will say that the couple (ϕ, σ) is ε -*learning* if, over many periods, the market measure point at the $\hat{\omega}$ that contains the true state ω 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:

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

$$\Pr_{\omega, \sigma} \left[\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T 1_{\{\pi^t(\hat{\omega}) > 1 - \varepsilon\}} < 1 - \varepsilon \right] < \varepsilon, \quad (9)$$

We say that a couple (ϕ, σ) is ε -exploiting if whenever the market measure points at some $\hat{\omega}$, play is such that a dealers' payoffs are strictly positive in all states ω included in $\hat{\omega}$. Formally:

Definition 3 *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 $\Pr_\sigma [a^t \in A^*(\hat{\omega}) \mid h^t] > 1 - \varepsilon$.*

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

Theorem 1 *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 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 (ϕ, σ) that is both ε -exploring and ε -exploiting forms a strategy profile satisfying the necessary conditions for a BFE described in Section 5. Namely, first, the way dealers set their actions is clearly measurable with respect to traders' behavior (Lemma 1) because dealers' action at t only depend on π^t that is itself a function of the public history. Second, this strategy leads to positive payoffs (Lemma 2) because the market measure points frequently to the right $\hat{\omega}$ (ε -exploring) and when this happens the dealers' payoff is positive (ε -exploiting). Third the strategy generate exploring (Lemma 4). In fact, for (ϕ, σ) to be ε -exploring, it is necessary that, no matter the past history, the actions that allow to distinguish the true $\hat{\omega}$ from the others $\hat{\omega}' \in \hat{\Omega}$ are played with strictly positive frequency.¹⁸ Thus, when dealers behave according to a pair (ϕ, σ) that is both ε -exploring and ε -exploiting, they all achieve long term positive profits independently of the state ω . In the proof of Theorem 1 we show that dealers have no incentive to deviate. In fact, because of the properties detailed in Proposition 1, there are strategies that

¹⁸The fact that dealers' inventories are bounded (Lemma 3) is consequence of the fact that dealer's payoff remain positive for all values of $W(\omega)$.

will be 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.

7 BFE Market Making vs. zero expected profit equilibrium: An Example and some empirical implications

In this section we analyze a particular class of BFE for the specific quote driven market of Example 1. We compare this equilibrium with the canonical zero-profit equilibrium (*i.e.*, the GME) that can be obtained if we make the additional assumptions that all dealers are equally uninformed.

Let us first consider the canonical equilibrium. This equilibrium relies on the assumptions that all dealers start with a common prior belief that $\Pr[v(\omega) = v_2] = p^0$ and $E[\tilde{e}] = 0$. Then there is a perfect Bayesian equilibrium in which, in any period t : (i) each dealer's expected profit is nil and (ii), best bid and ask quotes in the quote-driven market satisfy

$$\alpha^t = \alpha(p^t) := E[v(\omega)|h^t, s^t = buy] = E[v|h^t] + \frac{\rho}{2} - \frac{1}{2}\sqrt{\rho^2 - 4\text{Var}[v|h^t]}, \quad (10)$$

$$\beta^t = \beta(p^t) := E[v(\omega)|h^t, s^t = sell] = E[v|h^t] - \frac{\rho}{2} + \frac{1}{2}\sqrt{\rho^2 - 4\text{Var}[v|h^t]}, \quad (11)$$

$$p^{t+1} = \phi_B(p^t, \alpha^t, \beta^t), \quad (12)$$

where $\phi_B(p^t, (\alpha^t, \beta^t), s^t)$ denotes the posterior probability that $v(\omega) = v_2$ resulting from the prior probability p^t and from the trader's reaction s^t to dealers' quotes α^t .¹⁹ The r.h.s. of equation (10) and (11) are obtained considering that the probability of a trader buying, selling or not trading are given by equation (6), (5), and (7), respectively. Expressions $E[v|h^t]$ and $\text{Var}[v|h^t]$

¹⁹In order to simplify the exposition and notation we neglect the rounding required from the fact that quotes belong to a grid.

are the expectation and the variance of $v(\omega)$, respectively, computed using the belief p^t that evolves according to (12). This equilibrium has the advantage of being “Markovian:” first, in every period t , best bid and ask quotes only depend on dealers’ common belief p^t ; second, next period dealers’ common posterior beliefs p^{t+1} only depend on the common time- t prior p^t and on (a^t, s^t) , dealers’ and trader’s actions at time t . Note, however that this quoting strategy cannot be a BFE because the zero expected profit conditions implies that there is at least one value of $W(\omega)$ for which dealers’ profits are strictly negative. Also, this 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 .²⁰

We now illustrate how a BFE can be achieved with strategies that have a Markov structure that is as simple as the one of the canonical equilibrium. To this purpose we build a market measure on the partition $\hat{\Omega}$, a probability updating rule ϕ and a strategy σ mapping the market measure into dealers quotes. The pair (ϕ, σ) will be ε -learning and ε -exploiting. Thus by Theorem 1, dealers quoting behavior resulting from (ϕ, σ) is BFE. We first provide a verbal description of the equilibrium.

Recall that for Example 1 we have $\hat{\Omega} = \{\hat{\omega}_1, \hat{\omega}_2\}$, where $\hat{\omega}_k$ is the set of states $\omega \in \Omega$ such that $v(\omega) = v_k$, $k = 1, 2$. Let us first define a market measure π on $\hat{\Omega}$ and a probability updating rule ϕ . 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. Afterwards, the market measure evolves according to the following updating rule $\phi : \Pi \times A \times S \rightarrow \Pi$:

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

where $\phi_B(\pi^t, a^t, s^t)$ is the Bayesian posterior computed from prior π^t . In words, $\phi(\pi^t, a^t, s^t)$

²⁰To 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 $\alpha_i^t < \alpha(p^t)$.

associates to a probability $\pi^t \in \Pi$ and an action-reaction profile (a^t, s^t) a probability π^{t+1} that is the point in Π that is closest to the Bayesian posterior computed using prior equal to π^t and the information provided by traders' reaction s^t to dealers action a^t . Note that because in a BFE we drop all assumptions about each dealer's information and belief, the market measure needs not reflect any of the dealer's belief. In particular, π^0 can be chosen arbitrarily in the interval $[\varepsilon, 1 - \varepsilon]$ and afterwards, π^t does not always evolve according to the Bayes' rule. Nevertheless, it will be the level of the market measure that determines the equilibrium quotes set by rational Bayesian dealers.

Namely, we say that for $\pi^t > 1 - \varepsilon$ (resp. for $\pi^t < \varepsilon$), the game is in a $\hat{\omega}_1$ -exploiting phase (resp. $\hat{\omega}_2$ -exploiting phase). For $\pi^t \in [\varepsilon, 1 - \varepsilon]$, the game is in the exploring phase. We can now describe the mapping σ associating to the value of the market measure dealers' quotes. Fix d such that $0 < d < \rho$. During a $\hat{\omega}$ -exploiting phase, the best ask and bid equilibrium quotes satisfy

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

$$\beta^t = v(\hat{\omega}) - d, \tag{15}$$

respectively. During an exploring phase the best ask and bid quotes satisfy

$$\alpha^t = \alpha(\pi^t) + d, \tag{16}$$

$$\beta^t = \beta(\pi^t) - d, \tag{17}$$

where $\alpha(\cdot)$ and $\beta(\cdot)$ are the functions defined in equations (10) and (11), respectively. Strictly positive payoff at the dealer's individual level can be achieved by introducing any sharing rule allowing each dealer to set the best bid and ask quotes a strictly positive fraction of the time.

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: dealers' quotes induce an informative flow of trades. 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, $a^t \in A^*(\hat{\omega}_1)$ (resp. $a^t \in A^*(\hat{\omega}_2)$). This guarantees that dealers gain the spread while keeping their inventory small in absolute value.

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 whose beliefs differ from the market measure must be given incentives to play along and wait for play to shift towards the exploiting phase corresponding to the asset value that he might believe in. 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 the market measure evolve such that two conditions are met. 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. This is satisfied by our choice of ϕ : in the $\hat{\omega}_k$ -exploiting phase quotes straddle $v(\hat{\omega}_k)$ and induce a balanced flow of trade only if the true state $\omega \in \hat{\omega}_k$. Thus, an unbalanced order flow changes the market measure and eventually leads to a new exploring phase. 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 $v(\hat{\omega})$ 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. To have sensitive market measure no matter the past history, we impose that π^t remains in the interval $\Pi = [\varepsilon/4, 1 - \varepsilon/4]$: it evolves as a Bayesian beliefs as long as posterior remains in Π . Because the resulting market measure is never too concentrated on a state, the time it requires to point to the correct $\hat{\omega}$ is bounded, no matter the past history.

Note that if dealers follow these strategies, at any time t dealers expect their future profits to be strictly positive independently of the true ω and independently of the past trading history. Thus, even if a dealer's beliefs differ from the market measure, it will not deviate because (using standard repeated-game logic) the “Non-positive expected payoffs” result of Proposition 1 guarantees that other dealers can ensure that the deviating dealer makes zero profits for a long enough but finite period. Thus a deviation cannot be profitable if the discount rate is sufficiently low.

7.1 Simulation and explanatory power of BFE and GME

To illustrate some salient differences between such a BFE and the GME, we simulate price behavior resulting from the GME and the BFE equilibria in our leading example.²¹

Excess price volatility: In a BFE quotes are intrinsically more volatile than in the GME. This is due to the fact that in the GME, dealers' quotes reflect dealers common Bayesian belief that will eventually attach 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 arbitrary close to 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. This is illustrated in Figure 1, which reports a simulation of the two equilibria

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

for $v(\omega) = v_1$. The sequence of potential buying and selling traders is the same for the two equilibria. The right panel of Figure 1 reports the evolution of quotes in the BFE. The left panel reports the evolution quotes in GME.

Figure 1: Evolution of bid (red dots) and ask (blue dots) quotes in the GME (left panel) and in the BFE (right panel). Ask quotes and bid quotes are in blue, and magenta, respectively.

Volatility clustering: The recurrence of exploring and exploiting phases gives rise to price volatility clusters. In an exploring phase, dealers attract informative unbalanced traders' order flow, whereas in exploiting phase dealers make profits from the intermediation of relatively 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 1. 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 phase is illustrated in Figure 5 which shows how the market measure reacts to the trading volume in a exploiting phase (left panel) and in an exploring phase (right panel). 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 tend to be preceded by consistent traders' orders imbalance and significant changes in dealers' inventory. The transition from high to low volatility phases follows

the fading of traders' orders imbalance and a stabilization of dealers' inventory. Note that in the BFE, the volatility regime shift is completely endogenous and occurs in the absence of news. Not so in the GME which predicts that, in the absence of news, price volatility is bound to fade.

Quote volatility vs. trading flow, bid-ask spread and profits: This Markov BFE has some interesting implications regarding the correlation between price volatility, liquidity (as measured by bid-ask spread), dealers' aggregate inventory and profits. In an exploring phase, orders are more informative, in comparison to exploiting phases, exploring phases are associated with larger bid-ask spread, price volatility and aggregate inventory and with lower profits. This is consistent with the empirical regularities observed in Comerton-Forde et al. (2010): liquidity is negatively correlated with dealers' profits and inventories and with price volatility.

News and volatility: Whereas this benchmark model does not allow for exogenous shocks in information, it is straightforward to extend the model to allow for exogenous arrival of public news about the asset fundamental. Our Markov BFE can easily account for this extension by having the market measure depending on all public information, *i.e.*, public news as well as order flow. Unexpected news arriving when the market is an exploiting phase, move the market measure and may trigger an exploring phase. As a result, just following the news, price volatility increases and this may generate price overshooting and/or undershooting with respect to the level of quotes that will be reached once a new exploiting phase starts.

Dealers' profits: One of the necessary conditions for an equilibrium to be belief-free is that dealers' long term profits are strictly positive state by state. In the BFE this is achieved by maintaining a spread that is larger than the one predicted in the GME. 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 the GME the average dealers' aggregate per-period profit quickly converges to 0, in the BFE it is of the same magnitude as d (see Figure 2). Note that a dealer's *ex post* profit also depends on the value of $e(\omega) \in \{\underline{e}, \bar{e}\}$. Figure 3 represents the *ex post* cumulative profit for $e(\omega) = \underline{e}$ and $e(\omega) = \bar{e}$. In the GME (left panel of Figure 3), the dealers'

cumulative profit remains negative for at least one realization of $e(\omega)$. In the BFE, the dealers' cumulative profit eventually becomes positive no matter the realized $e(\omega)$ (right panel of Figure 3).

Figure 2: Evolution of the average per-period profit taking $e(\omega) = 0$ in GM (red dashed line) and in the BFE (blue solid line).

Dealers' inventories: A strictly positive profit in all states can only be achieved when aggregate inventory does not explode. For this reason, in a BFE, dealers' inventory must remain bounded. Not so in the GME. For example, in the simulation, $v(\omega) = v_1$, hence, traders tend to sell more than buy the asset. The right panel of Figure 4 reports the evolution of aggregate inventory in GME (red line) and in BFE (blue line). In the GME dealers' aggregate inventory tends to explode. Quite to the opposite, in the BFE dealers' aggregate inventory remain more balanced thanks to the bias in quotes c^t . Equation (8) provides the upper and lower bound for the ratio between the average change in dealers' inventory and the average trading volume. The larger $\bar{e} - \underline{e}$ and the smaller ρ , the smaller should be this ratio, on average. The variable $\bar{e} - \underline{e}$ can be interpreted as the asset intrinsic uncertainty, *i.e.*, the remaining uncertainty after

Figure 3: Cumulative dealers' profits for $e(\omega) = 0$ (dashed line), $e(\omega) = \underline{e}$ (blue line) and $e(\omega) = \bar{e}$ (red line).

incorporating traders' information.²² The parameter ρ is a measure of traders' willingness to trade for hedging rather than speculative reasons. In standard market microstructure parlance, ρ increase with liquidity traders' activity. Thus, the BFE predicts that dealers' inventories are more balanced in the presence of intrinsic uncertainty and for companies that are seldom traded by institutional investors.

8 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 and 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 such features.

A restriction of our model is that dealers' actions are observable. This might not be real-

²²Some possible proxies for the presence of intrinsic uncertainty are growth companies vs. utility companies, youth of the firm the firm's sector, product market innovations, R&D investments, business sensitivity to exogenous risks such as weather or other natural risks, and foreign country risk.

Figure 4: Red dashed lines represent the GME and blue solid lines the 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.

istic 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) 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 easily be 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 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 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

Figure 5: Market measure and dealers' inventory in an exploiting phase (left panel) and an exploring phases (right panel).

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 trade itself does not affect fundamentals. The verification that σ is a belief-free equilibrium follows exactly the same steps as in the main proof.

A third restriction is that long-term market participants do not take advantage of their private information, if any. This is an implication of our definition of belief-free equilibrium, which requires robustness to any possible information structure. What really matters for dealers 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. The fact that, in our equilibrium, dealers do not take advantage of their private information might be counter-intuitive, but there is no difficulty in re-defining our model to accommodate for such behavior without abandoning the belief-free assumption altogether. 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, with respect to 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. Whereas these type of belief-free equilibria can be

characterized, they would be less robust to changes in the information structure and hence they can help explaining price formation only for those actual cases where only few specific forms of information asymmetries are likely. We believe that such an extension raises interesting questions and technical challenges that motivate further study.

9 Conclusion

This paper considers market microstructure models in which long-lived dealers interact with short-lived traders. We characterize 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 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, 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 reasonable 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, or that minimize the expected

time required for the market measure to point at the true state, or even equilibria 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 Proposition 1

1. *Positive maximum payoffs*: Take any $\hat{\omega}$, it is sufficient to show that the set $A^*(\omega)$ is not empty. Fix dealer i and consider the following two action profiles $a(i)$ and $a'(i)$ in which all dealers different from i set the no-trade action. In $a(i)$ dealer i sets his action so that if a trader trades he can only buy at price strictly larger than $v(\hat{\omega}) + \underline{\rho}$ and he cannot sell. Conversely, in $a'(i)$ dealer i sets his action so that a trade can only consists in the trader selling at price strictly smaller than $v(\hat{\omega}) + \underline{\rho}$. Because of **ETD** and **NLS** we have that for all $\omega \in \hat{\omega}$, the expected asset transfer to dealer i are equals to $Q(\hat{\omega}, a(i)) > 0$ and $Q(\hat{\omega}, a'(i)) < 0$ for action $a(i)$ and $a'(i)$ respectively. Now let $q = Q(\hat{\omega}, a'(i)) / (Q(\hat{\omega}, a'(i)) - Q(\hat{\omega}, a(i))) \in [0, 1]$ and consider $\tilde{a}(i)$ obtained by playing $a(i)$ with probability q and $a'(i)$ with probability $1 - q$, this translates in dealer i expected profit of at least $Q(\hat{\omega}, a(i))2\underline{\rho} > 0$ no matter the value of $e(\omega)$. In facts in expectation he buys $qQ(\hat{\omega}, a(i))$ shares for a price less than $v(\hat{\omega}) - \underline{\rho}$ and he sells the same quantity for at least $v(\hat{\omega}) + \underline{\rho}$ 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 $u_i(\omega, \tilde{a}) = Q(\hat{\omega}, a(i))2\underline{\rho}/n > 0$ for every i and every $\omega \in \hat{\omega}$, no matter the value of the $e(\omega)$ component.
2. *Negative minimum payoffs*: Fix dealer i and consider the the action $\underline{a}(i)$ in which all dealers different from i set the no-trade action. In $\underline{a}(i)$ dealer i sets his action so that if a trader trades he can only buy at price strictly smaller than $v(\hat{\omega}) - \underline{\rho}$ and he cannot sell. Because of **ETD**, there will be trader willing to sell at such price, implying that dealer' i payoff is negative, no matter the true value of ω and hence for all $\omega \in \hat{\omega}$. Consider the random strategy $\underline{a}(\omega)$ obtained by first selecting a dealer i with probability $1/n$ and then playing $\underline{a}(i)$ satisfies the requirement. 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 probability measure $\mu_{\hat{\omega}} \in \Delta\hat{\omega}$. 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 on the price grid G that are closets to $W_{\mu_{\hat{\omega}}}$, with $p_1 \leq W_{\mu_{\hat{\omega}}} \leq p_2$. Let define $\underline{a}_{-i}^i(\mu_{\hat{\omega}})$ as follows. Each dealer $j \neq i$ set an action such that any other market participant can buy and sell up to the maximum tradable quantity at price $p : W_{\mu_{\hat{\omega}}}$. Let consider dealer i expected payoff when his belief

that the state is ω is equal to $\mu_{\hat{\omega}}(\omega)$. His expected payoff fro playing a_i when the other dealers play $\underline{a}_{-i}(\mu_{\hat{\omega}})$ is

$$\begin{aligned} \sum_{\omega \in \hat{\omega}} \mu_{\hat{\omega}}(\omega) u_i(\omega, a_i, \underline{a}_{-i}(\mu_{\hat{\omega}})) &= \sum_{\omega \in \hat{\omega}} \mu_{\hat{\omega}}(\omega) (v(\omega) + e(\omega)) Q_i(\omega, a_i, \underline{a}_{-i}) + P_i(\omega, a_i, \underline{a}_{-i}) \\ &= W_{\mu_{\hat{\omega}}} Q_i(\hat{\omega}, a_i, \underline{a}_{-i}) + P_i(\hat{\omega}, a_i, \underline{a}_{-i}) \end{aligned} \quad (18)$$

where the second equality follows from the fact that for any $\omega \in \hat{\omega}$ and $a \in A$, property **NLS** implies that $v(\omega) = v(\hat{\omega})$, $Q_i(\omega, a) = Q_i(\hat{\omega}, a)$ 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 exactly $W_{\mu_{\hat{\omega}}}$ buys a quantity $Q_i(\hat{\omega}, a_i, \underline{a}_{-i})$ of the asset in exchange for $P_i(\hat{\omega}, a_i, \underline{a}_{-i})$. 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 dealer actions are such that he can only trade at price $p = W_{\mu_{\hat{\omega}}}$, implying that dealer i 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 $p = W_{\mu_{\hat{\omega}}}$ from the other dealers he will trade with dealer i only if he propose equal or better trading condition, 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 case dealer i payoff of expression (18) cannot be strictly positive. Note that if because of the price grid trade cannot occur at price exactly equal to $W_{\mu_{\hat{\omega}}}$ then it is possible for dealer i to make some strictly positive payoff. However this payoff can be made arbitrarily small for a dense enough price grid.

4. *Non-equivalent payoffs*: Consider the strategy $\tilde{a}(i)$ defined in point 1. above. When dealer play $\tilde{a}(i)$, dealer i payoff is positive whereas all other dealers payoff is nil. Let $a^i(\hat{\omega})$ 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 payoff is nil whereas all other dealers payoff is 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 Lemma 1

A BFE must be an equilibrium even when dealers have no private information about ω . In this case the distribution over history can be different in two states ω and ω' only if traders' behavior differ in those two states. ■

Proof of Lemma 2

To see that the equilibrium payoff cannot be negative, suppose that for some state ω and 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 deviating and playing the no-trade action guaranteeing him a nil payoff, thus σ^* cannot be a BFE. To see that the equilibrium payoff is strictly positive in most states, take any $\hat{\omega} \in \hat{\Omega}$. Note that for all $\omega \in \hat{\omega}$ $v(\omega)$ is the same because of **NLS** and the equilibrium play is the same because of Lemma 1. Thus we have

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

Suppose $Q_i(\hat{\omega}, \sigma^*|h^t) \neq 0$, and take two different states $\omega, \omega' \in \hat{\omega}$, because $e(\omega) \neq e(\omega')$, then it must be $V_i(\omega, \sigma^*|h^t) \neq V_i(\omega', \sigma^*|h^t)$, so $V_i(\omega, \sigma^*|h^t) = 0$ for at most one state $\omega \in \hat{\omega}$. ■

Proof of Lemma 3

Fix $\hat{\omega} \in \hat{\Omega}$. Property **LS** implies that $v(\omega) = v(\hat{\omega})$ for all $\omega \in \hat{\omega}$ whereas property **NLS** implies that knowing $v(\omega)$ does not allow to infer anything about $e(\omega)$ whose maximum and minimum possible values are $\bar{e} > 0$, respectively. Because of Lemma 1 for any state $\omega \in \hat{\omega}$ we have that the change in dealers aggregate inventory is the same $Q(\hat{\omega}, \sigma^*|h^t)$. This change in inventory is mirrored by the change in traders aggregate inventory, *i.e.*, $Q(\hat{\omega}, \sigma^*|h^t) = -Q_T(\hat{\omega}, \sigma^*|h^t)$. Let consider dealers' aggregate payoff. Because of property **ETD**, traders will never buy for more than $v(\hat{\omega}) + \bar{\rho}$ neither sell for less $v(\hat{\omega}) - \bar{\rho}$, hence dealers' aggregate payoff cannot be greater than:

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

that is non-negative only if

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

Lemma 2 requires each dealer's payoff to be non-negative and hence a fortiori their aggregate payoff. Hence the above expression needs be positive for all realization of $e(\omega)$. Then the result follows from the fact Q, Q^+ and Q^- do not depend on $e(\omega)$ and that $e(\omega) \in [\underline{e}, \bar{e}]$. ■

Proof of Lemma 4

Suppose that, after some history h^t dealers actions never belong to $A(\hat{\omega}, \hat{\omega}')$, that is, in all

period $t' > t$ the dealers action profile $a^{t'} \notin A(\hat{\omega}, \hat{\omega}')$, implying $F(\hat{\omega}, a^{t'}, s) = F(\hat{\omega}', a^{t'}, s)$ for all $s \in S$. In other words, observation of traders behavior will not tell apart $\hat{\omega}$ from $\hat{\omega}'$. Now, measurability of the equilibrium play with respect to traders behavior (Lemma 1) implies that the equilibrium occupation measure after h^t is the same no matter whether the true state ω is in $\hat{\omega}$ or $\hat{\omega}'$. Let denote $\tilde{a}(h^t)$ this occupation measure. Because $A^*(\hat{\omega}) \cap A^*(\hat{\omega}') = \emptyset$, one must have either $\tilde{a} \notin A^*(\hat{\omega})$ or $\tilde{a} \notin A^*(\hat{\omega}')$ or both. But this implies that after history h^t dealers continuation payoffs are strictly negative for some ω . Thus, the continuation strategy cannot be a BFE because it would contradicting Lemma 2.

■

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

$$\begin{aligned} V_i^\delta(\omega, \sigma | h^\tau) &> (1 - \delta) \sum_{t=\tau}^{\infty} \delta^{t-\tau} (q^t((1 - \varepsilon)u^* - \varepsilon\bar{u}) - (1 - q^t)\bar{u}) \\ &= (1 - \varepsilon)(u^* + \bar{u})(1 - \delta) \sum_{t=\tau}^{\infty} \delta^t q^{t-\tau} - \bar{u}. \end{aligned} \quad (19)$$

Now condition 1 of Definition 2, implies that

$$\Pr_{\omega, \sigma} \left[\lim_{\delta \rightarrow 1} (1 - \delta) \sum_{t=\tau}^{\infty} \delta^t q^{t-\tau} > 1 - \varepsilon \right] > 1 - \varepsilon. \quad (20)$$

Hence we have that

$$\lim_{\delta \rightarrow 1} V_i^\delta(\omega, \sigma | h^\tau) > (1 - \varepsilon)^3(u^* + \bar{u}) - (1 + \varepsilon)\bar{u}. \quad (21)$$

As the r.h.s. is strictly positive for $\varepsilon = 0$, it is also positive for all ε smaller than some $\bar{\varepsilon} > 0$. Continuity of V_i^δ in δ implies there exists $\underline{\delta} < 1$ such that for $\varepsilon < \bar{\varepsilon}$, 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 5 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

$$0 < u_i(\omega, \tilde{a}^i(\hat{\omega})) < u_i(\omega, \tilde{a}^j(\hat{\omega})) < u(\omega, a). \quad (22)$$

for all $i \neq j$.

Proof. Consider the convex combination

$$\tilde{a}^i(\hat{\omega}) := \beta_1(\hat{\omega})\beta_2(\hat{\omega})\underline{a}(\hat{\omega}) + \beta_1(\hat{\omega})(1 - \beta_2(\hat{\omega}))a^i(\hat{\omega}) + (1 - \beta_1(\hat{\omega}))a, \quad (23)$$

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 $\{\tilde{a}^i(\hat{\omega})\}_{i=1, \dots, n}$ also satisfies Assumption **B-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 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 $\sigma^{i,\varepsilon}(h^t) = (1 - \varepsilon)\tilde{a}^i(\hat{\omega}) + (\varepsilon/L)\sum_{a \in A_L} a$. 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 $\hat{\omega} \in \hat{\Omega}$. 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 $\{a_i^1, \dots, a_i^{m^\eta}\}$ 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 $\underline{\sigma}^{i,\varepsilon}(h^t) = (1 - \varepsilon)\underline{a}_{-i}^i(\hat{\omega}_i)(h^t) + (\varepsilon/L)\sum_{a \in A_L} a$, 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

$$\lim_{\delta \rightarrow 1, \varepsilon \rightarrow 0} V_i^\delta(\omega, \sigma_i, \underline{\sigma}_{-i}^{i,\varepsilon}|h^t) \leq 0. \quad (24)$$

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 $\underline{\delta} \leq \bar{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$, players do not gain from deviating.

Note that this construction yields a 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] Benveniste, L., A. Marcus, and W. Wilhelm, 1992. "What's so Special about the Specialist?," *Journal of Financial Economics*, **32**, 61–86.
- [3] Bergemann, D. and S. Morris, 2013. "Robust Predictions in Games with Incomplete Information," *Econometrica*, **81**, 1251–1308.
- [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] Colmerton-Frode C, T. Hendershott, C. Jones, P. Moulton and M. Seasholes, 2010. “Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues,” *Journal of Finance*, **65**, 295–331.
- [15] Desgranges, G. and T. Foucault, 2005. “Reputation-Based Pricing and Price Improvements,” *Journal of Economics and Business*, **57**, 493–527.
- [16] De Frutos, M. and C. Manzano, 2005. “Trade Disclosure and Price Dispersion,” *Journal of Financial Markets*, **8**, 183–216.
- [17] De Meyer, B., 2010. “Price Dynamics on a Stock Market with Asymmetric Information,” *Games and Economic Behavior*, **69**, 42–71.
- [18] Du, S. and H. Zhu, 2012. “Ex-post Equilibria in Double Auctions of Divisible Assets,” *mimeo*.
- [19] Dutta, P. and A. Madhavan, 1997. “Competition and Collusion in Dealers Markets,” *Journal of Finance*, **52**, 1817–1846.
- [20] Easley, D. and M. O’Hara, 1987. “Price, Trade Size and Information in Security Markets,” *Journal of Financial Economics*, **19**, 60–90.
- [21] Easley, D. and M. O’Hara, 2010. “Microstructure and Ambiguity,” *Journal of Finance*, **65**, 2289–2316.

- [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. "Folk Theorem for Repeated Games with Discounting or with Incomplete Information," *Econometrica*, **54**, 533–554.
- [26] Fudenberg, D. and Y. Yamamoto, 2010. "Repeated Games Where the Payoffs and Monitoring Structure are Unknown," *Econometrica*, **78**, 1673–1710.
- [27] Fudenberg, D. and Y. Yamamoto, 2011. "Learning from Private Information in Noisy Repeated Games," *Journal of Economic Theory*, **146**, 1733–1769.
- [28] 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.
- [29] Goettler R., C. Parlour and U. Raja, 2009. "Equilibrium in a Dynamic Limit Order Market," *Journal of Finance*, **60**, 2149–2192.
- [30] Goettler R., C. Parlour and U. Raja, 2009. "Informed Traders in Limit Order Markets," *Journal of Financial Economics*, **93**, 67–87.
- [31] Gromb, D. and D. Vayanos, 2002. "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs," *Journal of Financial Economics*, **66**, 361–407.
- [32] 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.
- [33] Harsanyi, J., 1967–1968. "Games of Incomplete Information Played by Bayesian Players, Parts I-III. " *Management Science*, **14**, 159–182, 320–334, 486–502.
- [34] Hasbrouck, J. and G. Sofianos, 1993. "The Trades of Market Makers: An Empirical Analysis of NYSE Specialists," *Journal of Finance*, **48**, 1565–1593.

- [35] Ho, T. and H. Stoll, 1981. “Optimal Dealer Pricing under Transactions and Return Uncertainty,” *Journal of Finance*, **9**, 47–73.
- [36] Hörner, J. and S. Lovo, 2009. “Belief-free Equilibria in Games with Incomplete Information,” *Econometrica*, **77**, 453–487.
- [37] 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.
- [38] Leach, C. and A. Madhavan, 1993. “Price Experimentation and Security Market Structure,” *Review of Financial Studies*, **6**, 375–404.
- [39] Madhavan, A., 2000. “Market Microstructure: A Survey,” *Journal of Financial Markets*, **3**, 205–258.
- [40] Madhavan, A., and S. Smidt, 1993. “An Analysis of Daily Changes in Specialist Inventories and Quotations,” *Journal of Finance*, **48**, 1595–1628.
- [41] Milgrom, P., and N. Stokey, 1982. “Information, Trade and Common Knowledge,” *Journal of Economic Theory*, **26**, 17–27.
- [42] 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.
- [43] 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.
- [44] Pasquariello, P. and C. Vega, 2005. “Informed and Strategic Order Flow in the Bond Market,” *mimeo*.
- [45] Reiss, P. and I. Werner, 1998. “Does Risk Sharing Motivate Interdealer Trading?,” *Journal of Finance*, **53**, 1657–1703.
- [46] Reiss, P. and I. Werner, 2005. “Anonymity, Adverse Selection, and the Sorting of Interdealer Trades,” *Review of Financial Studies*, **18**, 559–636.
- [47] Roëll, A., 1988. “Regulating Information Disclosure Among Stock Exchange Market Makers,” FMG Discussion Papers Series 51, London School of Economics.