

Strategic Trading in Informationally Complex Environments*

Nicolas S. Lambert[†]

Michael Ostrovsky[‡]

Mikhail Panov[§]

July 25, 2017

Abstract

We study trading behavior and the properties of prices in informationally complex markets. Our model is based on the single-period version of the linear-normal framework of Kyle (1985). We allow for essentially arbitrary correlations among the random variables involved in the model: the value of the traded asset, the signals of strategic traders and competitive market makers, and the demand from liquidity traders. We show that there always exists a unique linear equilibrium, characterize it analytically, and illustrate its properties with a number of applications. We then use this characterization to study the informational efficiency of prices as the number of strategic traders becomes large. If liquidity demand is positively correlated (or uncorrelated) with the asset value, then prices in large markets aggregate all available information. If liquidity demand is negatively correlated with the asset value, then prices in large markets aggregate all information except that contained in liquidity demand.

KEYWORDS: Information aggregation, rational expectations equilibrium, efficient market hypothesis, market microstructure, strategic trading.

*We thank Beth Allen, Mehmet Ekmekci, Tibor Heumann, Craig Holden, David Kreps, Pete Kyle, Andy Postlewaite, Sven Rady, Marzena Rostek, and the co-editor and five anonymous referees for helpful comments, suggestions, and discussions. We thank Giorgio Martini for research assistance. We are grateful to the National Science Foundation for financial support (grant CCF-1101209). Ostrovsky also thanks the Alfred P. Sloan Foundation for financial support.

[†]Stanford Graduate School of Business, Stanford, CA 94305, U.S.A.; nlambert@stanford.edu.

[‡]Stanford Graduate School of Business, Stanford, CA 94305, U.S.A.; ostrovsky@stanford.edu.

[§]New York University, New York, NY 10012, U.S.A.; panov@nyu.edu.

1 Introduction

Whether and how dispersed information enters into market prices is one of the central questions of information economics. A key difficulty in answering this question is the strategic behavior of informed traders. A trader who has private information about the value of an asset has an incentive to trade in the direction of that information. However, the more he trades, the more he reveals his information, and the more he moves the prices closer to the true value of an asset. Thus, to maximize his profits, an informed trader may stop short of fully revealing his information, and so the informational efficiency of market prices may fail.

In one important case, however, market prices may still accurately reflect dispersed information: the case in which the number of informed traders is large, and each of these traders is small. In such markets, each of the traders has limited impact on prices, but their aggregate behavior reflects the aggregate information available in the market. As a result, market prices are close to those that would prevail if all private information were publicly available.

Non-strategic explorations of this intuition go back to Hayek (1945), Grossman (1976), and Radner (1979). Subsequently, a line of research (which we discuss in more detail in Section 1.1) has considered strategic foundations for this intuition, studying the strategic behavior of informed agents in finite markets, and then considering the properties of prices as the number of agents becomes large. This stream of work, however, imposes very strict assumptions on how information is distributed among the agents, typically assuming that the signals are symmetrically distributed, or satisfy other related restrictions so that in equilibrium, the strategies of all informed traders are identical. In practice, however, the distribution of information in the economy can be much more complex. Some agents may be better informed than others. Different groups of agents may have access to different, potentially interdependent, sources of information. Some agents may be informed about the fundamental value of the security, while others may possess “technical” information about the market or other traders. And all such possibilities may be present in a market at the same time.

Our paper makes two main contributions.

First, we present a tractable framework that makes it possible to study trading in such informationally complex environments. Our model is based on the single-period version of the model of Kyle (1985). As in that paper, an important assumption that makes our model analytically tractable is the assumption of joint normality of random variables involved: the true value of the traded asset, the signals of strategic traders, the signals of competitive market makers, and the demand coming from liquidity traders. Beyond that assumption, however, we impose essentially no restrictions on the joint distribution of these variables, making it possible to model informationally rich situations such as those described above. In this framework, we show that there always exists a unique linear equilibrium, which can be computed in closed form.

Second, we explore the informational properties of equilibrium prices as the number of informed agents becomes large. We assume that there are several types of agents, with each agent of a given type receiving the same information (possibly affected by idiosyncratic noise), and fix the matrix

of correlations of signals across the types. We then allow the numbers of agents of every type to grow. We find that the properties of prices in large markets depend on the informativeness of the demand from liquidity traders. If liquidity demand is uncorrelated with the value of the asset or is positively correlated with it (conditional on other signals), then prices in large markets aggregate all available information. If liquidity demand is negatively correlated with asset value, then prices in large markets aggregate all available information except that contained in liquidity demand. Crucially, in both cases, as markets become large, the information possessed by the strategic traders is fully aggregated and fully incorporated into market prices, for very general (multidimensional and asymmetric) information structures.¹

We also illustrate our model with two sets of applications.² First, we consider a natural question of whether having more information is always advantageous for a strategic trader. The answer turns out to be subtle. In the context of a single market, if one trader is more informed than another, then the former trader indeed has a higher expected profit than the latter. However, if a strategic trader receives more information in one market than he does in another one (with other characteristics of those markets being the same), he may be worse off in the market in which he is more informed.

Second, we explore a question in the spirit of Bergemann and Morris (2013): how much can the outcomes in our model vary when the fundamentals of the economy are fixed, but the informational structure is not? We find that the information structure plays an important role in determining market outcomes: if it is allowed to vary without any restrictions, the resulting bounds on the outcomes are quite wide, even in the most restrictive case of markets in which liquidity demand is independent of all other variables in the model. This finding may at first glance seem at odds with our information aggregation result, which gives very sharp predictions on outcomes in large markets. Of course, there is no contradiction, since for the information aggregation result, we do place a restriction on the underlying informational structure: there are several groups of symmetrically informed agents, and these groups grow large.

We conclude the paper with two sets of results related to information aggregation in large markets.

First, we characterize the properties of prices in a “hybrid” case, in which some information is available only to a small number of traders (“scarce” information), while some other information is available to a large number of traders (“abundant” information). As the number of traders having access to abundant information grows, the equilibrium converges to the one that would obtain if these traders were not present in the market at all, and instead their information was observed by the market maker (but not by the remaining strategic traders).

Second, to investigate the driving force behind our information aggregation result, we consider a simpler model in which there are no liquidity traders, and in which the sensitivity of prices to

¹The presence of exogenous liquidity demand plays an important role in our results: it makes trading possible by providing a source of profits for the strategic traders. Our information aggregation results rely on a slightly stronger assumption that the variance of liquidity demand is positive conditional on the signals of the strategic traders and the market maker (see Section 5 and footnote 20).

²In addition to these applications, in Section 5 of the Online Appendix we present a number of further illustrative examples.

aggregate quantity is fixed (instead of being endogenously determined by a Bayesian market maker). We present the model in the language of Cournot competition, but note that it is isomorphic to a model of trading with a mechanical (rather than Bayesian) market maker. We find that in this simpler model, information dispersed among the strategic agents gets fully aggregated in the limit as their numbers grow—just as in the first model.

1.1 Related Literature

The literature on strategic foundations of information aggregation and revelation in markets goes back to Wilson (1977), who considers an auction-based model in which multiple partially informed agents bid on a single object. Other work in this tradition includes Milgrom (1981), Pesendorfer and Swinkels (1997), Kremer (2002), and Reny and Perry (2006). These papers find that under suitable conditions, information gets aggregated when the number of bidders becomes large. However, these results depend critically on strong symmetry assumptions on the bidders’ signals and strategies.

Another stream of literature, going back to Kyle (1989), considers equilibria in demand and supply functions, where bidders specify how many units of an asset they demand or supply at each price, and the market maker picks the price that clears the market.³ Most papers in this tradition also require a very high degree of symmetry among the traders, typically assuming that they are ex ante identical, receive symmetrically distributed information, and employ identical strategies in equilibrium.⁴

The stream of literature most closely related to our paper is the work building on Kyle (1985). In that literature, one or more strategic traders, fully or partially informed about the value of the traded asset, are present in the market. These strategic traders submit market orders to centralized market makers. There are also liquidity traders who submit exogenously determined market orders. The market makers set the price of the asset equal to their Bayesian estimate of its value based on the aggregate order flow. Our paper borrows much of its analytical framework from this literature. The key difference is that while many of the papers in this area consider both static and dynamic models of trading but place restrictive assumptions on the information structure, our paper places virtually no restrictions on the information structure (beyond joint normality), and focuses on the one-period model of trading and on the informational properties of prices as the number of strategic

³See Vives (2008) for a textbook treatment of that literature.

⁴Notable exceptions are recent papers by Rostek and Weretka (2012), who replace symmetry with a weaker assumption of “equicommonality” on the matrix of correlations of agents’ values; Rostek and Yoon (2014), who go beyond equicommonality and provide conditions on the (potentially asymmetric) matrix of correlations for the existence of linear equilibrium; Manzano and Vives (2016), who consider the case of two groups of traders, with traders belonging to the same group observing identical signals; and Babus and Kondor (2016), who assume a symmetric matrix of correlations of agents’ values, but allow for asymmetries in the graph of possible trading relationships. There are important differences between our model and the settings of those papers. First, in our model, while agents generally receive different signals, their valuations for the security are the same, while in the above papers, the valuations are allowed to differ. Correspondingly, while the focus in our paper is on whether prices fully aggregate and reveal information, in the above papers the focus is on whether prices are “privately revealing.” Second, the trading mechanisms are different: in our model, agents submit quantity orders, while in the above papers, agents submit demand and supply curves. So while the questions are related, our results and those of the above papers are not directly comparable.

traders becomes large.

In the original model of Kyle (1985), there is only one informed trader, who knows the value of the asset. Admati and Pfleiderer (1988), Holden and Subrahmanyam (1992), Foster and Viswanathan (1996), and Back et al. (2000) study generalizations of the dynamic model of Kyle (1985) in which multiple informed traders are either all fully informed about the asset value, or receive imperfect signals about it, in which case different traders may observe different signals, but the distribution of these signals across the traders is symmetric (as are the traders' strategies). Caballé and Krishnan (1994) and Pasquariello (2007) consider multi-asset versions of the one-period model with multiple traders, but still maintain the assumption of symmetry of information among the traders. Dropping the assumption of normality of the underlying random variables, Bagnoli et al. (2001) provide conditions for the existence and uniqueness of linear equilibria in one-period models with multiple strategic traders whose (possibly imperfect) signals about the value of the asset are distributed symmetrically.⁵

Several papers go beyond the case of fully symmetric distributions of strategic traders' signals. Foster and Viswanathan (1994) consider a dynamic model with two strategic traders in which one trader is strictly more informed than the other. Dridi and Germain (2009) study a one-period model in which the signals of strategic traders are independent conditionally on the true value of the security, but may have different precisions. Colla and Mele (2010) consider a dynamic model in which strategic traders are located on a circle, with the correlations of signals being stronger for traders who are closer to each other (in this model, as in the Rostek and Weretka (2012) model discussed above, all traders use identical strategies in equilibrium).

Bernhardt and Miao (2004) consider a dynamic model with a general information structure, allowing, as our paper does, for essentially arbitrary covariance matrices of traders' signals.⁶ However, while Bernhardt and Miao (2004) characterize necessary and sufficient conditions for linear equilibria (analogous to Steps 1 and 2 in the proof of Theorem 1 in our paper, but in a multi-period setting), and use these conditions to study the properties of such equilibria analytically and numerically in some specific examples, they do not provide general results on equilibrium existence or uniqueness and do not provide general closed-form equilibrium characterizations. Whether such results can be established for a general multi-period setting is an open question.

There are also a number of papers building on Kyle's (1985) one-period model in which the information structure is not limited to strategic traders observing signals about the asset value. In Jain and Mirman (1999), the market maker receives a separate informative signal about the value of the asset, in addition to observing the order flow. In Rochet and Vila (1994) and Foucault

⁵See also Nöldeke and Tröger (2001, 2006) for the analysis of the role of the normality assumption for the existence of linear equilibria in one-period Kyle (1985)-style models with multiple strategic traders who receive perfect signals about the value of the asset.

⁶There are several differences between the models. Bernhardt and Miao (2004) consider a model with multiple trading periods, while we restrict attention to one period. On the other hand, unlike Bernhardt and Miao (2004), we allow liquidity demand to be correlated with the asset value and the signals of informed traders. We also allow the market maker to observe informative signals in addition to the order flow. Finally, we do not impose any special structure on how the informed traders' signals are related to the value of the asset (and other random variables in the model), beyond joint normality.

and Lescourret (2003), some of the strategic traders observe signals about the amount of liquidity demand.⁷ These features of the information structure are naturally incorporated in our general model. Hence, our equilibrium existence and uniqueness result, as well as the characterization we derive, provide a unified approach with closed-form solutions to various models that include these features. In Section 5 of the Online Appendix, we provide several examples illustrating the flexibility of our general model, and its ability to naturally incorporate such features as the market maker receiving a signal about the value of the asset and the strategic traders observing signals about liquidity demand, among others.

In Section 7, we study information aggregation in a model of Cournot competition. The literature on information aggregation under Cournot competition as the number of firms becomes large goes back to Li (1985) and Palfrey (1985). These papers consider environments in which all firms' signals about the true state of the world are symmetrically distributed. In contrast, our information aggregation result holds for essentially arbitrary variance-covariance matrices of firms' signals.⁸ Our focus in Section 7 is on information aggregation as the number of firms becomes large, and the parallels between this information aggregation result and the main information aggregation result in the paper. Thus, we do not explore in depth the connections between equilibrium outcomes in Cournot competition (in which the slope of the demand curve is fixed) and in the model based on the framework of Kyle (in which the slope of the demand curve is determined endogenously) for a fixed, finite number of strategic traders. For the case of symmetric distributions of signals, these connections (along with the connections to equilibrium outcomes in a model of demand-function competition in the spirit of Kyle (1989)) are explored by Bergemann et al. (2015).

Finally, on a more conceptual level, our paper is related to the work of Bergemann and Morris (2013, 2016) on the analysis of games with rich information structures. Bergemann and Morris introduce the notion of Bayes-Correlated Equilibrium, which makes it possible to abstract away from the details of how information is distributed among the agents in the economy, and to make robust predictions that do not depend on those details. We discuss the connection to those papers in more detail in Section 4.2.

2 Model

There is a security traded in the market, whose value v is not initially known to market participants. There are n strategic traders, $i = 1, \dots, n$. Prior to trading, each strategic trader i privately observes a multidimensional signal $\theta_i \in \mathbb{R}^{k_i}$. For convenience, we denote by $\theta = (\theta_1; \theta_2; \dots; \theta_n)$ the vector⁹ summarizing the signals of all strategic traders. The dimensionality of vector θ is $K = \sum_{i=1}^n k_i$.

⁷Röell (1990) and Sarkar (1995) also consider related one-period models in which some agents observe signals about liquidity demand. Madrigal (1996) considers a dynamic model in which a speculator is informed about liquidity demand.

⁸As in Li (1985) and Palfrey (1985), we also assume that the firms' marginal costs of production are constant. Vives (1988) shows that full information aggregation in large Cournot markets is not necessarily obtained when marginal costs are increasing, even in the fully symmetric case.

⁹We denote the row vector with elements x_1, \dots, x_k by (x_1, \dots, x_k) , and the column vector with the same elements by $(x_1; \dots; x_k)$. All vectors are column vectors unless specified otherwise.

There is also a market maker, who privately observes signal $\theta_M \in \mathbb{R}^{k_M}$, $k_M \geq 0$ (when $k_M = 0$, the market maker does not receive any signals, as in the standard Kyle (1985) model).^{10,11} Finally, there are liquidity traders, whose exogenously given random demand u is in general not directly observed by either the strategic traders or the market maker.

The key assumption that makes the model analytically tractable is that all of the random variables mentioned above— v , θ , θ_M , and u —are jointly normally distributed. Specifically, we assume that the vector $\mu = (v; \theta; \theta_M; u)$ is drawn randomly from the multivariate normal distribution with expected value 0 and variance Ω . The assumption that the expected value of vector μ is equal to zero is simply a normalization that allows us to simplify the notation. We also assume that every variance-covariance matrix for signal θ_i of strategic trader i and the variance-covariance matrix of the market maker’s signal θ_M are full rank. This assumption is without loss of generality; it simply eliminates redundancies in each trader’s signals. Note that we do not place a full rank restriction on matrix Ω itself: for instance, two different strategic traders are allowed to have perfectly correlated signals. The only substantive restrictions that we place on matrix Ω are as follows.

Assumption 1 At least one strategic trader receives at least some information about the value of the security, beyond that contained in the market maker’s signal. Formally:

$$\text{Cov}(v, \theta | \theta_M) \neq 0. \tag{1}$$

Assumption 2 The market maker does not perfectly observe the demand from liquidity traders. Formally:

$$\text{Var}(u | \theta_M) > 0. \tag{2}$$

2.1 Trading and Payoffs

After observing his signal θ_i , each strategic trader i submits his demand $d_i(\theta_i)$ to the market. In addition, the realized demand from liquidity traders, u , is also submitted to the market. The market maker observes her signal θ_M and the total demand $D = \sum_{i=1}^n d_i(\theta_i) + u$, and subsequently sets the price of the security, $P(\theta_M, D)$, based on these observations. Securities are traded at this price $P(\theta_M, D)$ (with each strategic trader getting his demand $d_i(\theta_i)$, liquidity traders getting u , and the market maker taking the position of size $-D$ to clear the market). At a later time, the true value of the security is realized, and each strategic trader i obtains profit $\pi_i = d_i(\theta_i) \cdot (v - P(\theta_M, D))$.

2.2 Linear Equilibrium

Our solution concept is essentially the same as that in Kyle (1985): linear equilibrium. Definition 1 below formalizes the notion of equilibrium, while Definition 2 states what it means for an equilibrium to be linear.

¹⁰Strictly speaking, θ_i and θ_M are random variables whose realizations are in \mathbb{R}^{k_i} and \mathbb{R}^{k_M} .

¹¹The multidimensionality of the traders’ and the market maker’s signals is a key feature that allows our model to incorporate complex informational interdependencies discussed in the introduction.

Definition 1 A profile of demand functions $d_i(\cdot)$ and pricing rule $P(\cdot, \cdot)$ form an equilibrium if

- (i) on the equilibrium path, the price P set by the market maker is equal to the expected value of the security conditional on θ_M and D , given the primitives and the demand functions $d_i(\cdot)$; and
- (ii) for every trader i , for every realization of signal θ_i , the expected payoff from submitting demand $d_i(\theta_i)$ is at least as high as the expected payoff from submitting any alternative demand d'_i , given the realization of signal θ_i , the pricing rule $P(\cdot, \cdot)$ and the profile of strategies of other traders $(d_j(\cdot))_{j \neq i}$.¹²

Definition 2 Equilibrium $(\{d_i(\cdot)\}_{i=1, \dots, n}, P(\cdot, \cdot))$ is linear if functions d_i and pricing rule P are linear functions of their arguments, i.e., $d_i(\theta_i) = \alpha_i^T \theta_i$ for some $\alpha_i \in \mathbb{R}^{k_i}$ and $P(\theta_M, D) = \beta_M^T \theta_M + \beta_D D$ for some $\beta_M \in \mathbb{R}^{k_M}$ and $\beta_D \in \mathbb{R}$.¹³

3 Equilibrium Existence and Uniqueness

We can now state and prove our first main result.

Theorem 1 *There exists a unique linear equilibrium.*

The proof of Theorem 1 is in Appendix A. The notation used in the proof, as well as in some of the subsequent sections, is given in Section 3.1 below. The proof is constructive, yielding a closed-form characterization of the unique equilibrium. This characterization is presented in Section 3.2.

The proof consists of several steps. We first show that if all strategic traders follow linear strategies, then the pricing rule resulting from Bayesian updating is also linear; and that if all strategic traders other than trader i follow linear strategies, and the market maker is also using a linear pricing rule (with a positive coefficient β_D on aggregate demand D), then the unique best

¹²Our interpretation of condition (i) is similar to that of Kyle (1985): it is a reduced-form way of representing the outcome of Bertrand competition among multiple market makers. In that interpretation, Kyle (1985) assumes that all market makers observe the total order flow and nothing else. In our case, all market makers observe the total order flow D and the signal θ_M , and nothing else. (For an alternative way of modeling competition among liquidity-supplying market makers, in which they post price schedules and make positive profits in equilibrium, see Biais et al. (2000, 2013).)

Another, technical difference from the equilibrium notion of Kyle (1985) is that in our case, condition (i) is required to hold only on the equilibrium path. In the standard Kyle (1985) model and many of its generalizations, every observation of the market maker can be rationalized as being on the equilibrium path, and thus this qualifier is not needed. In our case, it is in general possible that for some strategy profiles $d_i(\cdot)$, only some realizations of aggregate demand D can be observed by the market maker if the strategic traders follow those strategies. In such cases, by analogy with perfect Bayesian equilibrium, our definition restricts the beliefs of the market maker on the equilibrium path, where they are pinned down by Bayes rule, and does not restrict them off the equilibrium path. For an example in which not all realizations of aggregate demand are observed in equilibrium, consider the following market. Value $v \sim N(0, 1)$. There is one strategic trader with signal θ_1 who observes the value perfectly: $\theta_1 = v$. The demand of liquidity traders is $u = -v$. Then in the unique linear equilibrium, the demand of the strategic trader is equal to the value of the security, and the aggregate demand is thus always equal to zero. See Section 1 of the Online Appendix for details.

¹³In principle, we could consider a more general definition that allows the strategies and the pricing rule to potentially have nonzero intercepts. However, in our setting, linear equilibria with nonzero intercepts do not exist. See Section 2 of the Online Appendix for a formal proof of this statement.

response of trader i is also linear. Next, we show that the best response conditions allow us to express all parameters of the pricing rule and the traders' strategies as functions of "market depth" $\gamma = 1/\beta_D$. Using that derivation, we show that the system of best response conditions can be reduced to a quadratic equation in γ . Finally, we prove that this quadratic equation has exactly one positive root, which concludes the proof.

3.1 Notation

We decompose the covariance matrix Ω of the vector $(v; \theta_1; \dots; \theta_n; \theta_M; u)$ as follows:

$$\begin{pmatrix} \sigma_{vv} & \Sigma_{v1} & \cdots & \Sigma_{vn} & \Sigma_{vM} & \sigma_{vu} \\ \Sigma_{1v} & \Sigma_{11} & \cdots & \Sigma_{1n} & \Sigma_{1M} & \Sigma_{1u} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ \Sigma_{nv} & \Sigma_{n1} & \cdots & \Sigma_{nn} & \Sigma_{nM} & \Sigma_{nu} \\ \Sigma_{Mv} & \Sigma_{M1} & \cdots & \Sigma_{Mn} & \Sigma_{MM} & \Sigma_{Mu} \\ \sigma_{uv} & \Sigma_{u1} & \cdots & \Sigma_{un} & \Sigma_{uM} & \sigma_{uu} \end{pmatrix}.$$

In this matrix, every σ represents a (scalar) variance or covariance of the asset value and/or the demand of liquidity traders, and every Σ represents a (generally non-scalar) covariance matrix of an element of vector $(v; \theta_1; \dots; \theta_n; \theta_M; u)$ with another element. We also introduce notation for the covariance matrices of the entire vector of traders' signals, $\theta = (\theta_1; \dots; \theta_n)$, with itself and with other elements of vector μ . Specifically:

$$\Sigma_{\theta\theta} = Var(\theta) = \begin{pmatrix} \Sigma_{11} & \cdots & \Sigma_{1n} \\ \vdots & \ddots & \vdots \\ \Sigma_{n1} & \cdots & \Sigma_{nn} \end{pmatrix}, \quad \Sigma_{\theta M} = Cov(\theta, \theta_M) = \begin{pmatrix} \Sigma_{1M} \\ \vdots \\ \Sigma_{nM} \end{pmatrix},$$

$$\Sigma_{\theta v} = Cov(\theta, v) = \begin{pmatrix} \Sigma_{1v} \\ \vdots \\ \Sigma_{nv} \end{pmatrix}, \quad \Sigma_{\theta u} = Cov(\theta, u) = \begin{pmatrix} \Sigma_{1u} \\ \vdots \\ \Sigma_{nu} \end{pmatrix}.$$

In addition, we use the following matrices:

$$\Sigma_{diag} = \begin{pmatrix} \Sigma_{11} & 0 & 0 & 0 \\ 0 & \Sigma_{22} & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \Sigma_{nn} \end{pmatrix},$$

$$\Lambda = \Sigma_{diag} + \Sigma_{\theta\theta} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T,$$

$$A_u = \Lambda^{-1}(\Sigma_{\theta u} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mu}),$$

$$A_v = \Lambda^{-1}(\Sigma_{\theta v} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mv}).$$

(We show in the proof of Theorem 1 that matrix Λ is invertible.)

3.2 Closed-Form Solution

The proof of Theorem 1 is constructive, producing the following expressions for the parameters of interest.

Depth $\gamma = -\left(b + \sqrt{b^2 - 4ac}\right)/2a$, where

$$\begin{aligned} a &= -A_v^T \Sigma_{diag} A_v, \\ b &= A_v^T (2\Sigma_{diag} + \Lambda) A_u + \Sigma_{uM} \Sigma_{MM}^{-1} \Sigma_{Mv} - \sigma_{uv}, \\ c &= \text{Var}(A_u^T \theta - u | \theta_M). \end{aligned}$$

(The proof shows that $a < 0$, $c > 0$, and thus $\gamma > 0$.) Equilibrium pricing rule and strategies are then as follows:

$$\begin{aligned} \beta_D &= \frac{1}{\gamma}, \\ \beta_M &= \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) - \beta_D \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u), \\ \alpha &= \frac{1}{\beta_D} A_v - A_u. \end{aligned}$$

These expressions are simplified in the case $k_M = 0$, when the market maker does not observe any private signals (other than the aggregate demand D).¹⁴ In that case,

$$\begin{aligned} a &= -A_v^T \Sigma_{diag} A_v, \\ b &= A_v^T (2\Sigma_{diag} + \Lambda) A_u - \sigma_{uv}, \\ c &= \text{Var}(A_u^T \theta - u), \end{aligned}$$

where

$$\begin{aligned} \Lambda &= \Sigma_{\theta\theta} + \Sigma_{diag}, \\ A_u &= \Lambda^{-1} \Sigma_{\theta u}, \\ A_v &= \Lambda^{-1} \Sigma_{\theta v}. \end{aligned}$$

These expressions are further simplified if, in addition, the demand from liquidity traders, u , is uncorrelated with the other random variables in the model. Then $b = 0$ and $\gamma = \sqrt{\frac{\sigma_{uu}}{A_v^T \Sigma_{diag} A_v}}$, and so

$$\beta_D = \sqrt{\frac{A_v^T \Sigma_{diag} A_v}{\sigma_{uu}}} \quad \text{and} \quad \alpha = \sqrt{\frac{\sigma_{uu}}{A_v^T \Sigma_{diag} A_v}} A_v.$$

Finally, the proof of Theorem 1 allows us to obtain convenient closed-form expressions for the expected profits of strategic traders and the expected losses of liquidity traders, in the general setting. Specifically, the expected profit of trader i is equal to $\beta_D \alpha_i^T \Sigma_{ii} \alpha_i$, and the expected loss of liquidity traders is equal to $\beta_D \alpha^T \Sigma_{diag} \alpha$.¹⁵

¹⁴Strictly speaking, our proof does not apply directly to the case $k_M = 0$ since, for example, it uses the inverse of the covariance matrix of θ_M . However, one can drop all terms related to θ_M from the proof and immediately obtain the proof for that case. Alternatively, one can consider a model in which the market maker observes a signal that is independent of all other random variables. The equilibrium in that model will be equivalent to one in which $k_M = 0$.

¹⁵The expressions in Step 2 of the proof of Theorem 1 imply that for any trader i , conditional on realization $\tilde{\theta}_i$ of

3.3 Discussion of the Proof of Theorem 1

The first part of the proof, which shows the linearity of best responses to linear strategies and linear pricing rules, is standard in the literature on linear-normal equilibria. The main novel contribution of the next part of the proof is to transform the potentially unwieldy, very general system of equations with $\sum_{i=1}^n k_i + k_M + 1$ unknowns into a manageable, analytically tractable set of expressions, as follows. First, for any fixed value of “market depth” $\gamma = 1/\beta_D$, the remaining $\sum_{i=1}^n k_i + k_M$ unknowns can be expressed in a convenient matrix form as a function of γ . Next, using that representation, we can show that γ must be a root of a quadratic equation, where, again, the coefficients have manageable, compact matrix representations in terms of the underlying primitives of the model. Of course, obtaining a quadratic equation on γ is not sufficient: a quadratic equation can have two roots (and the model can thus in principle suffer from equilibrium multiplicity) or zero roots (and the model can thus suffer from equilibrium non-existence). In our setting, we have an additional constraint that γ must be positive (because market sensitivity β_D must be positive). The last part of the proof shows that the quadratic equation obtained in the previous step is guaranteed to have exactly one positive root for all possible values of the primitives. The proof of this statement is fairly subtle, and relies on the compact and tractable matrix representations obtained in the previous steps. So while the general outline of the proof is parallel to those in the earlier literature, its main novelty is in the generality of the underlying model, and in showing that despite this generality, one does not need to worry about equilibrium existence or equilibrium selection issues, and moreover can use tractable, convenient closed-form expressions to characterize the equilibrium.

A natural question is to what extent this approach is applicable to the analysis of asymmetric linear equilibria in another canonical linear-normal setting: that of competition in demand/supply schedules, in the tradition of Kyle (1989), Vives (2011), and other related papers. Broadly speaking, with jointly normally distributed signals and CARA or quadratic utility functions, the first two steps of the proof “go through” and the linearity of best response functions is preserved. That is, if all traders other than trader i submit demand/supply schedules that are linear in market price p , the optimal demand/supply schedule that trader i will submit in response will also be linear in p . Moreover, this “best response” schedule can be characterized in compact closed form as a function of the primitives of the model and the parameters of the strategies of other traders—even for very general, multidimensional and asymmetric information structures and strategy profiles like those that we consider in the current paper. However, proving that the resulting system of equations has a solution, determining whether and when it is unique, and characterizing its properties, all become much more challenging, for several reasons. First, with small numbers of traders, a linear equilibrium may not exist for strategic reasons, even in the original symmetric model of Kyle (1989). Second, Bayesian inference by strategic traders becomes more complicated.

signal θ_i , the equilibrium expected profit is equal to $\beta_D \left(\alpha_i^T \tilde{\theta}_i \right)^2 = \beta_D \alpha_i^T \tilde{\theta}_i \tilde{\theta}_i^T \alpha_i$. Thus, the unconditional expected profit of trader i is equal to $\beta_D \alpha_i^T E[\theta_i \theta_i^T] \alpha_i = \beta_D \alpha_i^T \text{Var}(\theta_i) \alpha_i = \beta_D \alpha_i^T \Sigma_{ii} \alpha_i$. The expected loss of liquidity traders is equal to the sum of the expected profits of strategic traders, $\sum_{i=1}^n \beta_D \alpha_i^T \Sigma_{ii} \alpha_i = \beta_D \alpha^T \Sigma_{diag} \alpha$.

In the model of Section 2, trader i 's expectation of asset value, v , depends only on his signal θ_i , and the expected price of the asset depends linearly on the parameters of other players' strategies.¹⁶ By contrast, in models of price schedule competition, trader i conditions his expectation of the value of the security on realized price p , which in turn depends endogenously on the parameters of other players' strategies. While in the linear-normal world, this conditioning is analytically tractable, it results in a nonlinear function of the parameters of players' strategies, substantially complicating the analysis.¹⁷ Third, in the model of the current paper, as well as in symmetric equilibria of price schedule competition, each trader faces the same sensitivity of price to his own demand. By contrast, in asymmetric equilibria of price schedule competition, each trader in general faces a different residual demand curve from those that other traders face, with different slopes, complicating the analysis.

For these reasons, the analysis of demand and supply schedule competition has been largely restricted to various cases with symmetric equilibria. However, given the compact closed-form expressions for linear best responses, and the resulting system of well-behaved polynomial equations, there may be ways to obtain positive results for this framework despite the general lack of closed-form solutions of this system of equation. Some recent papers make progress on that front. Rostek and Yoon (2014) solve the resulting systems of equations numerically, for a variety of asymmetric examples. Manzano and Vives (2016) obtain closed-form solutions for a special case with two groups of traders, with traders belonging to the same group observing identical signals. What makes this special case tractable is that the second problem described above (dividing by the variance of p) disappears, because prices in this particular case are privately revealing: knowing his own signal (and thus the signals of other trader in his group) and price p , each trader can immediately infer the signal of the traders in the other group. And the third problem becomes manageable as well, as there are only two different price sensitivities that need to be dealt with. As a result, the system of equations can be reduced to an analytically tractable cubic equation in one variable.

4 Applications

Before proceeding to our second main result (Theorem 2 in Section 5), we illustrate our general framework with a number of applications. In Section 4.1, we consider the question of whether having more information is always advantageous for a strategic trader. In Section 4.2, we explore the range of possible outcomes in our model when the fundamentals of the economy are fixed, but the informational structure is allowed to vary, in the spirit of Bergemann and Morris (2013). We also provide a number of additional illustrative examples in Section 5 of the Online Appendix.

¹⁶The strategy of player j enters into player i 's expectation of the price of the asset as an additive term $-\beta_D \alpha_j^T E[\theta_j | \theta_i]$.

¹⁷Intuitively, conditioning on p involves multiplying the observed price by the inverse of its variance. The variance of p is a second-degree polynomial in the parameters of players' strategies.

4.1 The Value of Additional Information

In this subsection, we address the question of whether having more information is always advantageous for a strategic trader. In decision problems, the answer is of course immediate: getting more information is always weakly better than getting less, because the decision maker can always dismiss the additional information if he so desires. In strategic situations, however, that is not necessarily the case, as other players may adjust their behavior when a player receives extra information.

In the current setting, the answer to this question turns out to be subtle. We first show (Proposition 1) that in the context of a single market, if one trader is more informed than another one (in the sense that the first trader knows the information of the second one, and also possibly has some additional information), then the expected profit of the first trader is weakly higher than that of the latter. We then show (Example 1) that if a strategic trader receives more information in one market than he does in another one (with all other parameters in these two markets being the same), then the trader may be worse off in the market where he is more informed.

Proposition 1 *Consider a market in the general framework of Section 2, and suppose strategic trader A is more informed than strategic trader B : the latter observes a multidimensional signal θ_B , while (slightly abusing our notation) the former observes a multidimensional signal $(\theta_A; \theta_B)$. Then in the unique linear equilibrium, the expected profit of trader A is weakly higher than the expected profit of trader B .*

The proof of Proposition 1 is in Section 3 of the Online Appendix. The key ingredient of the proof is Lemma OA.1, which shows that if two strategic traders share some common information, then in equilibrium, they put exactly the same weight on this information (even if the additional signals that they observe are different). It then follows that in the context of Proposition 1, traders A and B make the same expected profits from the information contained in θ_B , and then trader A also makes additional nonnegative profits from the additional information contained in signal θ_A .

By contrast, the following example shows that becoming more informed is not necessarily good for a strategic trader: equilibrium effects may be negative and may outweigh the benefit from extra information.

Example 1 *The value of the security is $v \sim N(0,1)$. There are two strategic traders. Trader 1 observes a noisy estimate of v : $\theta_1 = v + \epsilon$, where $\epsilon \sim N(0,1)$ is a random variable independent of v . Trader 2 observes $\theta_2 = v$. Finally, there is demand from liquidity traders, $u \sim N(0,1)$, which is independent of all other random variables. The resulting covariance matrix is*

$$\Omega = \begin{pmatrix} 1 & 1 & 1 & 0 \\ 1 & 2 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Using the notation and closed-form characterization from the preceding section, we have $\Lambda = \Sigma_{\theta\theta} + \Sigma_{diag} = \begin{pmatrix} 4 & 1 \\ 1 & 2 \end{pmatrix}$, $\Lambda^{-1} = (1/7) \begin{pmatrix} 2 & -1 \\ -1 & 4 \end{pmatrix}$, and so

$$A_v = \Lambda^{-1}\Sigma_{\theta v} = \begin{pmatrix} \frac{1}{7} \\ \frac{3}{7} \end{pmatrix} \text{ and } \beta_D = \sqrt{\frac{A_v^T \Sigma_{diag} A_v}{\sigma_{uu}}} = \frac{\sqrt{11}}{7}.$$

The equilibrium strategies of traders 1 and 2, α_1 and α_2 , are thus given by

$$\begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} = \frac{1}{\beta_D} A_v = \begin{pmatrix} \frac{1}{11}\sqrt{11} \\ \frac{3}{11}\sqrt{11} \end{pmatrix}.$$

The equilibrium expected profit of every trader i is equal to $\beta_D \alpha_i^T \Sigma_{ii} \alpha_i$. Thus, in the current example, the expected profit of trader 1 is equal to

$$\pi_1 = \left(\sqrt{11}/7\right) \times \left(\sqrt{11}/11\right) \times 2 \times \left(\sqrt{11}/11\right) = 2\sqrt{11}/77 \approx 0.086,$$

and the expected profit of trader 2 is equal to

$$\pi_2 = \left(\sqrt{11}/7\right) \times \left(3\sqrt{11}/11\right) \times 1 \times \left(3\sqrt{11}/11\right) = 9\sqrt{11}/77 \approx 0.388.$$

Now consider a modified market, in which trader 2 becomes more informed: he observes both v and ϵ , i.e., $\theta_2 = (v; \epsilon)$. The resulting covariance matrix is now

$$\Omega = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 \\ 1 & 2 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

The corresponding auxiliary matrices are

$$\Lambda = \begin{pmatrix} 4 & 1 & 1 \\ 1 & 2 & 0 \\ 1 & 0 & 2 \end{pmatrix}, \quad \Lambda^{-1} = \frac{1}{12} \begin{pmatrix} 4 & -2 & -2 \\ -2 & 7 & 1 \\ -2 & 1 & 7 \end{pmatrix}, \quad \text{and } A_v = \Lambda^{-1}\Sigma_{\theta v} = \frac{1}{12} \begin{pmatrix} 2 \\ 5 \\ -1 \end{pmatrix},$$

and so

$$\beta_D = \sqrt{\frac{A_v^T \Sigma_{diag} A_v}{\sigma_{uu}}} = \frac{\sqrt{34}}{12}, \quad \alpha_1 = \begin{pmatrix} 2 \\ \sqrt{34} \end{pmatrix}, \quad \text{and } \alpha_2 = \begin{pmatrix} \frac{5}{\sqrt{34}} \\ \frac{-1}{\sqrt{34}} \end{pmatrix}.$$

Thus, in this modified market, the expected profit of trader 1 is equal to

$$\pi'_1 = \left(\sqrt{34}/12\right) \times \left(2/\sqrt{34}\right) \times 2 \times \left(2/\sqrt{34}\right) = \sqrt{34}/51 \approx 0.114$$

and the expected profit of trader 2 is equal to

$$\pi'_2 = \left(\sqrt{34}/12\right) \begin{pmatrix} 5/\sqrt{34} \\ -1/\sqrt{34} \end{pmatrix}^T \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 5/\sqrt{34} \\ -1/\sqrt{34} \end{pmatrix} = \frac{13\sqrt{34}}{204} \approx 0.372.$$

After getting more informed, trader 2 is worse off: $\pi'_2 \approx 0.372 < 0.388 \approx \pi_2$. Perhaps surprisingly, trader 1 (whose information did not change) is better off: $\pi'_1 \approx 0.114 > 0.086 \approx \pi_1$. Of

course, there is no contradiction here with Proposition 1: we still have $\pi'_2 > \pi'_1$, consistent with the fact that in the modified market, trader 2 (who observes v and ϵ separately) is better informed than trader 1 (who only observes $v + \epsilon$).

To understand these effects, note that when trader 2 observes trader 1's "mistake" ϵ , in equilibrium he will "undo" one half of that mistake: the weight trader 2 puts on ϵ is negative one half of the weight that trader 1 puts on his signal. This "undoing" allows trader 1 to trade more aggressively (increasing the equilibrium multiplier α_1 on his signal from ≈ 0.302 to ≈ 0.343), which in turn causes the market maker to also increase her price sensitivity β_D (from ≈ 0.474 to ≈ 0.486). The more aggressive behavior of trader 1 and the increased price sensitivity of the market maker have a negative impact on trader 2's profits, outweighing the benefits of trader 2's new information.

4.2 The Range of Possible Outcomes

As argued by Bergemann and Morris (2013), the information structure of players in a game may be hard to observe, and thus it is interesting to know what range of predictions one can obtain without making specific assumptions on that structure. In this section, we explore this question in the context of our model. For tractability, we restrict attention to the case in which the market maker does not observe any signals beyond the aggregate demand. For normalization, we fix the variance of the asset value, σ_{vv} , and the variance of liquidity demand, σ_{uu} , and also assume that the pooled information of all strategic traders is sufficient to know the asset value: $Var(v|\theta) = 0$.

We consider three classes of markets. The first (and smallest) class C_1 contains markets in which liquidity demand is independent of the value of the asset, and the strategic traders do not observe any information about liquidity demand (just like in the canonical Kyle (1985) model). The second, larger class C_2 contains markets in which liquidity demand is still independent of the value of the asset, but now the strategic traders may observe some information about liquidity demand. Finally, the third, most general class C_3 contains markets in which liquidity demand may be correlated with the value of the asset, and the strategic traders may observe some information about it.

Within those classes, we allow the information structure of the strategic traders to vary freely (subject to the constraint that as a group, they know the true value of the asset, v), and also allow the number of strategic traders to vary. For each class, we find the lower and the upper bounds on four outcomes of interest: the variance of market prices, $Var(p)$; the variance of aggregate market demand, $Var(D)$; the sensitivity of the market maker, β_D ; and the expected loss of liquidity traders, $-E[u(v - p)]$.

For each of the four variables, within each of the three classes, we find the infimum and the supremum of possible outcome values in the unique linear equilibrium, across all possible information structures and possible numbers of strategic traders. In the setting of Bergemann and Morris (2013), the range of possible Bayes Nash equilibrium outcomes of variables of interest, across information structures, corresponds to the range of possible Bayes correlated equilibrium outcomes of the corresponding basic game (Bergemann and Morris, 2016). Our characterization of the bounds

on outcome variables is similar in spirit, but note that the parallel is not exact, because the full strategic interaction in our model is dynamic: traders $i = 1, \dots, n$ move first (and simultaneously, just like in the settings of Bergemann and Morris), but then the market maker moves second after observing aggregate demand D .¹⁸ Of course, the dynamic game can be “collapsed” into the corresponding normal form, but then care needs to be taken in the interpretation of signal D of the market maker, which is not an exogenous signal about the state of the world or a signal coming from a correlation device, but is rather an endogenous signal that depends on the strategies of other players.

The results are summarized in the following proposition.

Proposition 2 *The bounds on the four outcome variables of interest, for three classes of markets, are as in the table below.*

	Class C_1	Class C_2	Class C_3
$Var(p)$	sup = σ_{vv} inf = 0	sup = σ_{vv} inf = 0	sup = σ_{vv} inf = 0
$Var(D)$	sup = ∞ inf = σ_{uu}	sup = ∞ inf = 0	sup = ∞ inf = 0
β_D	sup = $\frac{1}{2}\sqrt{\sigma_{vv}/\sigma_{uu}}$ inf = 0	sup = ∞ inf = 0	sup = ∞ inf = 0
$-E[u(v - p)]$	sup = $\frac{1}{2}\sqrt{\sigma_{vv}\sigma_{uu}}$ inf = 0	sup = $\frac{1}{2}\sqrt{\sigma_{vv}\sigma_{uu}}$ inf = 0	sup = $\sqrt{\sigma_{vv}\sigma_{uu}}$ inf = 0

The proof of Proposition 2 is in Section 4 of the Online Appendix. Several technical and substantive comments are in order about these bounds. First, while some of these bounds are in fact achieved in some markets (e.g., the highest possible loss of liquidity traders in classes C_1 and C_2 is achieved in the standard Kyle (1985) model with one strategic trader), others are not (e.g., the expected loss of liquidity traders is always positive, although it can be made arbitrarily small). Second, while these bounds are across markets with different numbers of strategic traders, the bounds across markets with a fixed number of traders mechanically get (weakly) wider as that number of traders increases (because the “extra” traders can be endowed with completely uninformative signals, thus not trading and not having any impact on the equilibrium outcomes). So the bounds in Proposition 2 can be alternatively viewed as limits on the bounds for markets with a fixed number of traders, as that number grows large (although again, some of these bounds are achieved in markets with a finite number of strategic traders; in some cases as low as one).

¹⁸Formally, in our model, the market maker sets the price non-strategically, to be equal to the expected value of the security based on her information. However, this behavior is identical to that of a fully strategic player who chooses p that maximizes the expected value of the expression $-(v - p)^2$. So with this re-interpretation, our model can be viewed as a standard two-period game, with traders $i = 1, \dots, n$ moving first and the market maker moving second.

On an economically substantive level, these bounds show the critical importance of the underlying information structure for the predictions of our model. For instance, without any assumptions on the information structure, prices can range from fully informative about the value of the security to completely uninformative,¹⁹ despite the fact that the players, as a group, have full information about v . This finding may at first glance seem at odds with the second main result of our paper, discussed in the next section: as markets grow large, the information of strategic traders gets aggregated and fully incorporated in market prices. Of course, there is no contradiction, but the contrast highlights the importance of the assumption on information structures that we impose in our information aggregation result: there is a finite number of groups of traders, symmetrically informed within each group, and the size of each group grows large. We now turn to the formal statement and proof of the result.

5 Information Aggregation in Large Markets

Consider a sequence of markets, indexed by $m = 1, 2, \dots$. Every market is in the general framework of Section 2. In every market, there are n groups of strategic traders, with at least one trader in each group. Index i , $1 \leq i \leq n$, now denotes a group of traders. The size of group i in market m is denoted by $\ell_i^{(m)}$. Every trader j in group i receives a k_i -dimensional signal $\theta_i + \xi_{i,j}$, where θ_i denotes the signal component common to all traders in group i and $\xi_{i,j}$ denotes the idiosyncratic component of trader j . We denote by $\theta = (\theta_1; \dots; \theta_n)$ the vector of common components of the signals, and denote by Ω the covariance matrix of vector $\mu = (v; \theta; \theta_M; u)$. The idiosyncratic components $\xi_{i,j}$ are distributed identically across the traders in group i , with each $\xi_{i,j}$ distributed according to a k_i -dimensional normal distribution with mean 0 and variance Σ_i^ξ . Every $\xi_{i,j}$ is independent of all other random variables in the model. We place no restrictions on matrices Σ_i^ξ . In particular, we allow for the case $\Sigma_i^\xi = 0$, in which all traders in group i receive the same signal θ_i .

We assume that Ω and $\Sigma_1^\xi, \dots, \Sigma_n^\xi$ are the same for all markets m . The number of traders in each group, however, changes with m : specifically, we assume that for every i , $\lim_{m \rightarrow \infty} \ell_i^{(m)} = \infty$, i.e., all groups become large as m becomes large. We do not impose any restrictions on the rates of growth of those groups: e.g., the sizes of some groups may grow much faster than those of other groups.

We slightly strengthen one of the two conditions on matrix Ω made in Section 2, replacing Assumption 2 with the following:²⁰

¹⁹Since by construction, $E[v|p] = p$, we have the following simple formula for the informativeness of the price: $Var(v - p) = \sigma_{vv} - Var(p)$.

²⁰Under the original Assumptions 1 and 2, information may not get aggregated as markets become large. To see that, consider a modification of the example introduced in footnote 12. Value $v \sim N(0, 1)$. There are m strategic traders with the same signal $\theta_1 = v$. The demand of liquidity traders is $u = -v$. Then in the unique linear equilibrium, the demand of each strategic trader is equal to θ_1/m , the aggregate demand of all strategic trader is equal to $\theta_1 = v = -u$, the aggregate demand of all traders is equal to zero, and thus the equilibrium price is also always equal to zero, for any m . Thus, there is no information aggregation of any kind in the limit as m becomes large. See Section 1 of the Online Appendix for a formal derivation of these results.

Assumption 2L $Var(u|\theta, \theta_M) > 0$.

It follows from Theorem 1 that for each m , there exists a unique linear equilibrium in the corresponding market. Let $p^{(m)}$ denote the random variable that is equal to the resulting price in the unique linear equilibrium of market m .

We can now state and prove our main result on information aggregation in large markets. If the demand from liquidity traders is positively correlated with the true value of the asset (conditional on other signals), then prices in large markets aggregate all available information: $p^{(m)}$ converges to $E[v|\theta, \theta_M, u]$. If liquidity demand is negatively correlated with the true value of the asset, then prices in large markets aggregate all available information *except* that contained in liquidity demand: $p^{(m)}$ converges to $E[v|\theta, \theta_M]$. If liquidity demand is uncorrelated with the true value of the asset, then both statements are true: $p^{(m)}$ converges to $E[v|\theta, \theta_M, u] = E[v|\theta, \theta_M]$.

Theorem 2

- If $Cov(u, v|\theta, \theta_M) \geq 0$, then $\lim_{m \rightarrow \infty} E \left[(p^{(m)} - E[v|\theta, \theta_M, u])^2 \right] = 0$.
- If $Cov(u, v|\theta, \theta_M) \leq 0$, then $\lim_{m \rightarrow \infty} E \left[(p^{(m)} - E[v|\theta, \theta_M])^2 \right] = 0$.

In Appendix B, we prove Theorem 2 for the special case in which the covariance matrix of random vector $(\theta; \theta_M; u)$ is full rank. This additional assumption guarantees that certain matrices remain invertible in the limit as m becomes large, which in turn allows us to give a direct proof of the theorem without technical complications. However, this special case rules out some interesting possibilities (e.g., one type of traders knowing strictly more than another type of traders), so in the Online Appendix (Section 8), we provide the full proof of Theorem 2, without this simplifying assumption.

The intuition for the information aggregation result is that, when the number of informed traders of each type is large, the information of each strategic trader has to be (almost) fully incorporated into the market price, since otherwise each trader of that type would be able to make a non-negligible profit, which cannot happen in equilibrium. Also, as the size of every group i grows, the idiosyncratic noise in the aggregate demand from that group vanishes, leaving only the “informative” part of the demand that is driven by the common component θ_i .²¹ The signal of the market maker gets incorporated into the market price by construction. Finally, with liquidity demand, the situation is more subtle. When liquidity demand is positively correlated with the asset value ($Cov(u, v|\theta, \theta_M) > 0$), equilibrium strategies and market depth adjust precisely in a

²¹The fact that the idiosyncratic components in signals have no impact on equilibrium outcomes in large markets is parallel to the results in McLean and Postlewaite (2002) and McLean et al. (2005) in which the agents with such idiosyncratic components in signals have non-redundant information, but become “informationally small” as markets become large: adding the information of an extra agent to the information of others does not significantly impact the Bayesian estimate of the value of the security. Note, however, that “informational smallness” by itself is not sufficient for our results. In an economy without idiosyncratic components in signals, agents become “informationally small” as soon as the size of each group i is at least two. However, information is generally not aggregated in our setting in finite markets, even if the size of each group is two or greater and all the agents in each group i receive the same signal θ_i .

way that makes liquidity demand get incorporated into the market price “correctly,” i.e., with the same weight as it would be incorporated into the market price by a Bayesian observer who was fully informed about all the random variables in the model (except value v). As a result, price $p^{(m)}$ converges to $E[v|\theta, \theta_M, u]$, and so all information available in the market is incorporated into the market price. However, when liquidity demand is negatively correlated with the value of the asset ($Cov(u, v|\theta, \theta_M) < 0$), this cannot happen. In equilibrium, aggregate demand always enters the market price with a positive sign (sensitivity β_D is positive). Thus, liquidity demand also enters the market price with a positive sign. However, a fully informed Bayesian observer would put a *negative* weight on liquidity demand—which cannot happen in any linear equilibrium, for any parameter values. So what happens instead as m becomes large is that the variance of the aggregate demand from informed traders grows to infinity (in contrast to the case $Cov(u, v|\theta, \theta_M) > 0$, in which it converges to a finite value). And thus, as m grows, liquidity demand u has less and less impact on the market price, and in the limit it has no impact at all: price $p^{(m)}$ converges to $E[v|\theta, \theta_M]$. The same happens in the case $Cov(u, v|\theta, \theta_M) = 0$, for the same reason, but in that case $E[v|\theta, \theta_M]$ is equal to $E[v|\theta, \theta_M, u]$, and so price $p^{(m)}$ does converge to the expected value of the asset given all the information available in the market.

Another way to get intuition about the result is to notice that as a particular group i becomes large, its aggregate behavior converges to that of a single agent who is trying to minimize the expected square of the difference between the true value of the asset and its market price, $E[(v - p)^2]$.²² By construction, the market maker is also trying to minimize $E[(v - p)^2]$ (subject to the constraint that the sensitivity of price to aggregate demand, β_D , is positive). Thus, as market size grows large, the system in essence behaves as a game with n partially informed traders (each corresponding to a particular group i and receiving the signal θ_i) and a market maker, all of whom have the same objective function: to minimize the expected square of the mispricing. The commonality of objective functions implies that the profile of policies by these $n + 1$ agents that minimizes the expected squared mispricing will be an equilibrium of this limit game. When $Cov(u, v|\theta, \theta_M)$ is positive, the profile of policies that minimizes the expected squared mispricing is the one that sets the price $p = E[v|\theta, \theta_M, u]$, incorporating all the information available in the market. When

²²To see this, fix a market and the corresponding equilibrium, and consider group i with ℓ traders, all of whom observe the same signal $\theta_i \in \mathbb{R}$ (the cases with multidimensional signals or idiosyncratic components are more notationally cumbersome, but the conclusions are the same). Suppose in equilibrium each of these traders, after observing realization $\tilde{\theta}_i$, submits demand $d^* = \alpha_i \tilde{\theta}_i$. Let $p_{-i} = p - \beta_D (\ell \alpha_i \theta_i)$ denote the random variable corresponding to what the price in the market would have been if all traders in group i demanded zero instead of submitting the demands prescribed by the equilibrium. Conditional on $\tilde{\theta}_i$, each agent in group i submits demand d that is maximizing $d \times \left(E[v - p_{-i}|\tilde{\theta}_i] - \beta_D(\ell - 1)d^* - \beta_D d \right)$, which implies $d = \frac{1}{2\beta_D} \left(E[v - p_{-i}|\tilde{\theta}_i] - \beta_D(\ell - 1)d^* \right)$. Since in equilibrium $d = d^*$, this in turn implies $\beta_D(\ell + 1)d^* = E[v - p_{-i}|\tilde{\theta}_i]$. The aggregate demand of group i is equal to ℓd^* , and so as $\ell \rightarrow \infty$, the aggregate demand of group i converges to $\frac{1}{\beta_D} E[v - p_{-i}|\tilde{\theta}_i]$. Now suppose we instead have a single trader i who observes a realization $\tilde{\theta}_i$ of signal θ_i and whose objective is to minimize $E[(v - p)^2]$. We have $E[(v - p)^2|\tilde{\theta}_i] = E[(v - p_{-i} - \beta_D d)^2|\tilde{\theta}_i] = E[(v - p_{-i})^2|\tilde{\theta}_i] - 2d\beta_D E[v - p_{-i}|\tilde{\theta}_i] + \beta_D^2 d^2$. The first term of this sum does not depend on d , and so the expression is minimized at $d = \frac{1}{\beta_D} E[v - p_{-i}|\tilde{\theta}_i]$ —which is precisely the quantity to which the aggregate demand of group i converges in the original game as $\ell \rightarrow \infty$.

$Cov(u, v|\theta, \theta_M)$ is negative or zero, setting the price at $p = E[v|\theta, \theta_M, u]$ is impossible, since that would require setting $\beta_D \leq 0$, which is not allowed. In fact, since β_D has to be positive, any profile of strategies by the $n + 1$ agents has to put positive weight on u in forming the price—which the $n + 1$ agents do not want to do. So they will want to set β_D to be infinitesimally small, and then adjust the strategies of the n partially informed traders accordingly, to get price p to be close to $E[v|\theta, \theta_M]$ —which provides the infimum of the square of the mispricing given the constraint $\beta_D > 0$. Note that this intuition also illustrates that mathematically, there is no asymmetry between the cases of $Cov(u, v|\theta, \theta_M) > 0$ and $Cov(u, v|\theta, \theta_M) < 0$, and the difference in predictions for those cases arises from the economic incentives of the agents. Namely, if in the original game the goal of the strategic traders was to *lose* as much money as possible, the limit game with the $n + 1$ traders would in fact be the same as in our original case, except that the constraint would be $\beta_D < 0$. And so all information would get aggregated in the case $Cov(u, v|\theta, \theta_M) < 0$ (and market depth would remain bounded), and only information contained in θ and θ_M , but not that contained in u , would get aggregated in the case $Cov(u, v|\theta, \theta_M) > 0$ (and market depth would go to infinity).

The information aggregation result in Theorem 2 raises some natural questions. The first one is to what extent it matters that the variance of liquidity traders’ demand $u^{(m)}$ remains constant as markets become large. What would happen if that variance also grew together with the number of strategic traders? Of course, the profits made by the strategic traders and their equilibrium strategies would be affected. It turns out, however, that equilibrium prices would remain unchanged. Specifically, for a given market, if liquidity demand were scaled by some factor ρ , the equilibrium strategies of all strategic traders would also get rescaled by the same factor ρ , the sensitivity of market maker’s pricing rule to the aggregate demand, β_D , would get rescaled by $1/\rho$, and the equilibrium prices would thus stay the same. Proposition OA.4 in Section 7 of the Online Appendix formally proves these statements. This result, in turn, immediately implies that the conclusion of Theorem 2 would not be affected if we allowed liquidity demand $u^{(m)}$ to scale as a function of m .

Another question is whether the presence of a Bayesian market maker is critical for information aggregation. Is it important that there is an agent in the economy who is accurately setting prices based on the information available to her? To answer this question, in Section 7 of the paper we consider a model of Cournot competition, which can be viewed as an analogue of the Kyle-style model of Section 2 with one key difference: the Bayesian market maker is replaced with a mechanical market maker whose sensitivity to aggregate demand, β , is exogenously fixed, instead of being determined endogenously in equilibrium. We find that the presence of a Bayesian market maker is not critical for information aggregation: Proposition 4 in Section 7.3 shows that as the number of firms grows, the outcome (equilibrium price and total quantity produced) of Cournot competition with information dispersed among the firms converges to that of Cournot competition in which all firms have access to all information.

Finally, a natural question is what happens if some groups remain “small,” while others grow “large.” The next section addresses this question.

6 Information in “Hybrid” Markets

In many situations, some “scarce” information about the value of a security is known by only a small number of traders, perhaps just one, while some other information, while not publicly available, may be more “abundant,” and may be observed by a large number of traders. In this section, we explore how these two types of information get incorporated into market prices in equilibrium.

It is intuitive that due to strategic considerations, “scarce” information will not be fully incorporated into market prices, and the traders possessing this information will make positive profits, while “abundant” information will be almost fully incorporated into market prices (and the traders possessing it will make vanishingly small profits). What is less immediate is the interplay between these two types of information, and how they get combined with the information observed directly by the market maker and the information contained in liquidity demand. In particular, a seemingly natural conjecture is that “abundant” information will enter the price essentially as a public signal, observed by everyone in the economy. Our last result shows that this is not the case: instead, “abundant” information, in the limit, enters into market prices in the same way as if it were directly observed by the market maker—but not by the strategic traders observing “scarce” information. As Examples OA.6 and OA.7 in Section 5.3 of the Online Appendix illustrate, this is substantively different from the case in which “abundant” information is observed by all the agents in the economy.

Formally, using the notation introduced in Section 5, suppose that for some $s \geq 1$, the sizes of the groups $i = 1, \dots, s < n$ remain constant as m varies, i.e., $\ell_i^{(m)} = \ell_i$ for some ℓ_i , while for $i = s + 1, \dots, n$, the size of group i grows to infinity, i.e., $\ell_i^{(m)} \rightarrow \infty$. We will refer to groups $i = 1, \dots, s$ as “small groups,” and to groups $i = s + 1, \dots, n$ as “large groups.” Every trader j of a small group i receives signal θ_i . Every trader j of a large group i receives signal $\theta_i + \xi_{i,j}$, where θ_i is the component common to all traders of group i , and $\xi_{i,j}$ is the idiosyncratic component of trader j , independently distributed according to a normal distribution with mean 0 and variance Σ_i^ξ .²³

Throughout this section, let θ_S be the vector of signals of the small groups, i.e., $\theta_S = (\theta_1; \dots; \theta_s)$, and let θ_L be the vector of common components of the signals of the large groups, i.e., $\theta_L = (\theta_{s+1}; \dots; \theta_n)$. We make two assumptions:

Assumption 1H $Cov(v, \theta_S | \theta_L, \theta_M) \neq 0$.

Assumption 2H $Var(u | \theta_L, \theta_M) > 0$.

The first assumption states that at least one of the small groups has some information about the asset value that is not observed by the market maker or the large groups. This assumption is analogous to Assumption 1 of Section 2, ensuring that some information about the value of the

²³Note that the assumption that all traders in the same small group i receive the same signal θ_i is without loss of generality: one small group of size ℓ_i in which traders also receive idiosyncratic components with nonzero variance can be represented as ℓ_i small groups of size one.

asset remains “scarce” even in the limit. The second assumption states that the information of the market maker and the joint information of large groups is not sufficient to fully learn liquidity demand. This assumption is analogous to Assumption 2L of Section 5.

Our next result shows that under Assumptions 1H and 2H, equilibrium prices in the above sequence of markets converge to the equilibrium price that would obtain in an alternative market, in which only the small groups of traders are present (with the same information as in the original markets, θ_S), and in which the market maker directly observes both her original signal θ_M and the common components of signals observed by the large groups of traders in the original markets, θ_L . Let $\{p^{(m)}\}$ denote the sequence of random variables that are equal to the prices in the linear equilibria of the original sequences of markets indexed by m . Let $p^{(alt)}$ denote the random variable that corresponds to the equilibrium price obtained in the alternative market.

Theorem 3 $\lim_{m \rightarrow \infty} E \left[(p^{(m)} - p^{(alt)})^2 \right] = 0.$

In Section 9 of the Online Appendix, we prove Theorem 3 for the special case in which the variance of vector $(\theta_S; \theta_L; \theta_M; u)$ is full rank. As in the case of Theorem 2, this assumption simplifies the argument by guaranteeing that certain matrices remain invertible in the limit. However, this assumption rules out some interesting possibilities (e.g., some small groups know some elements of the common components of signals of some large groups), and so in Section 10 of the Online Appendix, we provide the full proof of Theorem 3 without this simplifying assumption. The techniques used in the proofs are similar to those used in the proofs of Theorem 2, except that the presence of small groups requires a separate treatment, because their strategic incentives do not vanish in the limit. Also, note that unlike in Theorem 2, the result in Theorem 3 does not depend on the sign of the conditional covariance of liquidity demand and asset value. The reason for that is that in the large-market case, when $Cov(u, v | \theta, \theta_M) \leq 0$, as the market was getting larger, market maker’s sensitivity β_D was converging to zero, removing the impact of u on the market price. In the hybrid-market case, even as some groups become large, there are still some groups that remain small and whose traders thus possess “scarce” information which would have allowed them to make infinite profits if β_D converged to zero. So in the hybrid-market case, β_D remains bounded away from zero even in the limit, regardless of the sign of $Cov(u, v | \theta, \theta_M)$.

We conclude this section with a final observation. As Examples OA.6 and OA.7 in Section 5 of the Online Appendix show, the expected profit of an informed agent can be strictly higher when he observes the signal of the market maker than when he does not, because observing the information of the market maker allows the informed trader to better use the part of his information that is *not* known to the market maker. In the case of “hybrid” markets, Theorem 3 shows that equilibria converge to those that would obtain if the information of “large” groups was observed by the market maker, but *not publicly*, so that the “small” groups do not observe that information. This situation may create incentives for trading information. If some small-group traders were to obtain information from some of the large-group traders, those small-group traders could increase their expected profit by a non-negligible amount. At the same time, in the limit, large-group traders

make zero profits anyway, so they would not lose anything by sharing this information with the small-group traders. Thus, if trading information is allowed, the large-group information may end up being purchased by small-group traders, and thus the market, in the limit, may behave as if that information was observed publicly. We leave the formal analysis of this intuition to future research.

7 Cournot Competition

The model of Section 2 (to which we will refer as the Kyle model throughout this section) has many “moving parts.” In particular, it has three types of agents: fully optimizing strategic traders, mechanical liquidity traders, and Bayesian market makers. It is thus natural to ask which of these components are the driving forces behind our result on information aggregation (Theorem 2). Would the result break down without a market maker who explicitly sets prices to be equal to the expected value of the security? Is it essential that there are liquidity traders who in expectation lose money and by doing so “subsidize” trading and information discovery?

To shed light on these questions, in this section we consider a model that contains neither Bayesian market makers nor liquidity traders, but is otherwise closely related to the Kyle model. (As we explain in footnote 26 below, the model of this section can be equivalently viewed as a model of trading with a mechanical market maker whose sensitivity to demand is exogenously fixed.) The model we consider in this section is asymmetric Cournot competition, in which firms observe imperfect (and generally different) signals about the intercept of the market demand function (an analogue of the value of the security v in the Kyle model) and the question we address is, again, whether this asymmetric information gets aggregated as the market grows large.²⁴ We show that information does indeed get aggregated as the number of firms increases—the total quantity and price in the market converge to those that would obtain if all the firms had access to all available information.²⁵

²⁴Bergemann et al. (2015) also compare informational properties of trading under Cournot competition, in which the slope of the price response is exogenously fixed, and of trading in a setting in the spirit of Kyle (1985), in which the slope of the price response is endogenously determined by a Bayesian market maker (as well as in a setting of demand function competition in the spirit of Kyle (1989), which we do not consider). Their focus, however, is different from ours: while we study information aggregation in the limit as the number of strategic traders becomes large, Bergemann et al. (2015) keep the number of players fixed and study the spaces of possible equilibrium outcomes under general information structures, the equivalences of these outcomes under different equilibrium notions (Bayes Correlated Equilibrium and Bayes Nash Equilibrium), and the properties of these spaces (such as the first and second moments of the equilibrium distributions of players’ actions).

²⁵This result is closely related to the results of Li (1985) and Palfrey (1985), who also observe that under constant marginal costs of production, Cournot competition efficiently aggregates distributed information as the number of firms becomes large. The key difference between our result and those of Li (1985) and Palfrey (1985) is that we allow for an arbitrary matrix of correlations of the firms’ signals, while they require the signals to be symmetric. Note that the assumption of constant marginal costs is important for the results Li (1985) and Palfrey (1985), and thus also for our information aggregation result. As Vives (1988) shows, if production costs are quadratic (and so marginal costs are increasing in quantity instead of being constant), the market does not converge to the full-information outcome as the number of firms becomes large, even if the firms are ex ante identical.

7.1 Model

There are n firms in the market for a good. Each firm has a constant marginal cost of production c per unit of the good, and no fixed costs. The demand function for the good is not initially known to the firms. Rather, if the firms in aggregate produce Q units of the good, the resulting market price will be

$$p = v - \beta Q,$$

where $\beta > 0$ is the commonly known slope of the inverse demand function, and v is the uncertain intercept of that function.

Prior to making a production decision, each firm i observes a multidimensional signal $\theta_i \in \mathbb{R}^{k_i}$. Vector $\theta = (\theta_1; \dots; \theta_n)$ summarizes the signals of all firms. We assume that vector $(\theta; v)$ is drawn randomly from the multivariate normal distribution with expected value $(0; \bar{v})$ and covariance matrix Ω . We further assume, without loss of generality, that for every i , the covariance matrix of signal θ_i is full rank.

After observing its signal θ_i , each firm i simultaneously decides to produce quantity $q_i(\theta_i)$ of the good. The total amount produced is thus $Q = \sum_{i=1}^n q_i(\theta_i)$. The resulting market price is $p(v, Q) = v - \beta Q$. The realized payoff of firm i is $(p - c)q_i$.

7.2 Linear Equilibrium

As before, we restrict attention to linear equilibria, i.e., those of the form $q_i(\theta_i) = \alpha_i^T \theta_i + \delta_i$ for some profile of vectors $\alpha_i \in \mathbb{R}^{k_i}$ and $\delta_i \in \mathbb{R}$. We denote a linear equilibrium by these linear coefficients, and define $\alpha = (\alpha_1; \dots; \alpha_n)$ and $\delta = (\delta_1; \dots; \delta_n)$.

Proposition 3 *The Cournot competition game has a unique linear equilibrium.*

The proof of Proposition 3 is in Appendix C. The proof is, in essence, a substantially simplified version of the proof of Theorem 1, which only involves the analogues of Steps 2 and 3 from that proof. Step 1 (market maker's Bayesian updating) is not needed, because there is no market maker in the current model, and the price impact of each individual unit of supply, β , is exogenously fixed, instead of being endogenously determined by the market maker. Recall that Step 3 of the proof of Theorem 1 allowed us to express all the equilibrium strategies of the traders as a function of a single parameter—the inverse of the market maker's sensitivity to aggregate demand, β_D . Steps 4 and 5 then derived a quadratic equation in that parameter and showed that it has a unique positive root. In the Cournot competition setting, the sensitivity β is fixed exogenously, and so the analogue of Step 3 concludes the proof.²⁶

²⁶To see the parallels between the two models more directly, consider the version of the model in Section 2 in which the market maker does not observe any direct signals, and the version of the Cournot competition in which the marginal cost c is zero. In the former, the realized payoff of an individual trader i from submitting demand d_i is $d_i(v - p) = d_i(v - \beta_D(d_i + \sum_{j \neq i} d_j + u))$. In the latter, the realized payoff of an individual firm i from producing q_i units of the good is $q_i(v - \beta Q) = q_i(v - \beta(q_i + \sum_{j \neq i} q_j))$. So the Cournot competition setting can be viewed as a version of the Kyle model in which liquidity demand is fixed at zero, and the market maker is mechanical, with the sensitivity exogenously fixed at $\beta_D = \beta$, instead of being endogenously determined by the Bayes rule.

The closed-form solutions no longer involve the roots of a quadratic equation, and take the following form. Each firm i 's strategy is given by

$$q_i(\theta_i) = \alpha_i^T \theta_i + \delta_i,$$

where for each i ,

$$\delta_i = \beta^{-1} \frac{\bar{v} - c}{n + 1},$$

and vector α is given by

$$\alpha = \beta^{-1} (\Sigma_{\theta\theta} + \Sigma_{diag})^{-1} \Sigma_{\theta v},$$

where matrices $\Sigma_{\theta\theta}$, Σ_{diag} , and $\Sigma_{\theta v}$ are defined as before. The formula for vector α (as a function of price sensitivity β and the three matrices) is essentially the same as that in the Kyle model for the case in which the market maker does not observe any private signals (Section 3.2), except for the terms related to liquidity demand u that have no counterparts in the Cournot competition model. Of course, the key difference between the two formulas is that in the Kyle model, sensitivity β_D is derived endogenously, while in the Cournot competition model, sensitivity β is exogenously fixed.

7.3 Information Aggregation in Large Markets

We now turn to the behavior of markets with a large number of participants. Our modeling approach is analogous to that in Section 5. Specifically, consider a sequence of markets, indexed by $m = 1, 2, \dots$. The inverse demand function is the same in all markets m : $p(Q) = v - \beta Q$. In every market, there are n groups of firms, with at least one firm in each group. The groups are indexed by $i = 1, \dots, n$, and each group i in market m consists of $\ell_i^{(m)}$ firms, with $\ell_i^{(m)} \rightarrow \infty$ as $m \rightarrow \infty$.

Each firm j in group i receives signal $\theta_i + \xi_{i,j} \in \mathbb{R}^{k_i}$, where θ_i is the common signal component of all firms in group i , and $\xi_{i,j}$ is the idiosyncratic component of firm j . Random vector $(\theta_1; \dots; \theta_n; v)$ is distributed normally with mean $(0; \dots; 0; \bar{v})$ and variance-covariance matrix Ω . We also assume that the variance-covariance matrix of random vector $\theta = (\theta_1; \dots; \theta_n)$ is positive definite. Every $\xi_{i,j}$ is drawn from the normal distribution with mean zero and variance Σ_i^ξ , independently of all the other random variables in the model. We impose no restrictions on Σ_i^ξ , and in particular allow for the case $\Sigma_i^\xi = 0$, when all firms in group i observe identical signals.

As a benchmark, we also consider a sequence of alternative markets with the number of firms growing to infinity (for concreteness, let the number of firms in market m equal $N^{(m)} = \sum_{i=1}^n \ell_i^{(m)}$), but with a much simpler information structure: all ‘‘common components’’ of all signals are known to all firms. Formally, each firm j observes the same signal $\theta = (\theta_1; \dots; \theta_n)$. In this sequence of alternative markets, all information is shared by all firms, and as the number of firms increases, the outcomes (i.e., the total quantity produced and the equilibrium price) converge to the perfectly competitive equilibrium. Our next proposition shows that the outcomes in the original sequence of markets also converge to the same perfectly competitive outcome, thus aggregating all the information distributed among the firms. Formally, let $Q^{(m)}$ and $p^{(m)}$ denote the random variables corresponding to the total quantity produced and the price realized in the original market m , and

let $Q^{(alt,m)}$ and $p^{(alt,m)}$ denote the random variables corresponding to the total quantity produced and the price realized in the alternative market m where all the firms observe the same joint signal θ .

Proposition 4 $\lim_{m \rightarrow \infty} E[(Q^{(m)} - Q^{(alt,m)})^2] = 0$ and $\lim_{m \rightarrow \infty} E[(p^{(m)} - p^{(alt,m)})^2] = 0$.

The proof of Proposition 4 is in Appendix C. The proof proceeds by showing that in both sequences (original and alternative), for any realization of signals θ , the total quantity produced converges to the quantity produced in the perfectly competitive market with the intercept of the demand function equal to $E[v|\theta]$: $Q^*(\theta) = (E[v|\theta] - c)/\beta$. The result for the convergence of prices is then immediate.

Proposition 4 illustrates that the main driving force behind the information aggregation results in our paper is not the presence of a market maker who sets prices in an “intelligent” way, but rather the fact that the individual actions of informed players get aggregated (via aggregate demand in the Kyle model and via aggregate production in Cournot competition). The “aggregate action” of each group of players reflects that group’s common signal, and these aggregate actions of the groups are then further aggregated by the marketplace with the appropriate weights. This action aggregation feature is important for our results. In Section 11 of the Online Appendix, we provide a simple example of a Beauty Contest game in which dispersed information does not get aggregated in the limit, even though that game shares many of the features with the models considered above (normally distributed signals, linear best responses, and the uniqueness of linear equilibrium that can be characterized in closed form).

8 Concluding Remarks

Our paper leaves a number of open questions and directions for future research. One question is to what extent our analysis can be generalized to a dynamic setting, in which trading takes place over multiple periods, and each strategic trader takes into account the impact of his trading on his future arbitrage opportunities.

Second, the fact that our model admits explicit closed-form solutions for every profile of primitives makes it “embeddable” as part of richer settings and games. For instance, one can study pre-trading investment in costly acquisition of information (about the fundamentals of the traded security, about liquidity demand, or about the information of other strategic traders), mergers among the agents, or information sharing and trading among them. One can also consider the case of endogenous participation by liquidity traders, by considering a model with several different types of liquidity traders (e.g., retail investors, pension funds, insurance companies, etc.) whose demands may be differentially correlated with the value of the asset and/or with the informed traders’ or the market maker’s signals, and who choose to participate in the market only if their expected losses do not exceed certain thresholds.

Third, the tractability of our model may also extend, at least to some degree, to other related settings, such as those with risk-averse traders (with CARA utilities, to preserve the linear-quadratic

structure of the game), costly trading (with quadratic trading costs), multiple securities or trading venues (with traders having mean-variance preferences over payoffs), or partially informed liquidity traders.

Finally, a shared feature of the models considered in this paper (the model of trading in financial markets in Section 2 and the model of Cournot competition in Section 7) is that when making decisions, strategic players condition their behavior only on their own information, and not on market prices: in the model of financial trading, they submit market orders, and in the model of Cournot competition, they decide on the level of production before they get to observe any feedback from the market. A natural question is to what extent the types of rich informational asymmetries allowed in our models can also be considered in strategic models of markets in which players can condition their actions both on their own information and on endogenous market prices. As we discuss in Section 3.3, getting general closed-form solutions in such settings is challenging, both for economically fundamental strategic reasons (when the number of players is small), and for technical reasons (when the system of polynomial equations characterizing the equilibria of the market game does not have closed-form solutions). However, some of the results of our paper may continue to hold in such settings, perhaps under some additional assumptions. For example, one intuitive argument for our information aggregation result²⁷ is that a large group of symmetrically informed traders will behave in approximately the same way as a single trader who tries to minimize mispricing in the market. Thus, as all groups become large, the market, in effect, converges to a game in which all players pursue a common objective function: getting the price as close as possible to the asset value. This intuition may carry over to at least some of the settings in which players can condition their actions on market prices, and for large markets, existence results may be possible to obtain by means other than closed-form solutions. Thus, information aggregation results for general asymmetric information structures similar to the ones in Sections 2 and 7 may also hold in those settings.

We leave the exploration of these extensions and generalizations to future research.

Appendix A: Proof of Theorem 1

Step 1. Let $\alpha = (\alpha_1; \dots; \alpha_n)$ be a profile of linear strategies for the strategic traders. Each α_i in this profile is a vector $(\alpha_i^1; \dots; \alpha_i^{k_i}) \in \mathbb{R}^{k_i}$, corresponding to linear strategy

$$\begin{aligned} d_i(\theta_i) &= \alpha_i^1 \theta_i^1 + \dots + \alpha_i^{k_i} \theta_i^{k_i} \\ &= \alpha_i^T \theta_i, \end{aligned}$$

where $\theta_i^1, \dots, \theta_i^{k_i}$ are the elements of vector $\theta_i \in \mathbb{R}^{k_i}$.

Take any linear pricing rule $(\beta_M; \beta_D)$, $\beta_M \in \mathbb{R}^{k_M}$, $\beta_D \in \mathbb{R}$. Let vector $\beta = (\beta_M; \beta_D)$ summarize the pricing rule and let random vector $\eta = (\theta_M; D = \alpha^T \theta + u)$ denote the information available to the market maker when she sets the price. Then for this pricing rule to be consistent with profile

²⁷This argument is presented in more detail in the discussion that follows Theorem 2.

α , condition (i) of the definition of equilibrium requires that $\beta^T \eta = E[v|\eta]$, which is equivalent to the condition $Cov(v, \eta) = \beta^T Var(\eta)$.²⁸ Expressing $Cov(v, \eta)$ and $Var(\eta)$ using the notation from Section 3.1, we thus get the following equivalent characterization of condition (i) of the definition of equilibrium:

$$(\beta_M^T, \beta_D) \begin{pmatrix} \Sigma_{MM} & \Sigma_{\theta M}^T \alpha + \Sigma_{Mu} \\ \alpha^T \Sigma_{\theta M} + \Sigma_{Mu}^T & \alpha^T \Sigma_{\theta\theta} \alpha + 2\Sigma_{\theta u}^T \alpha + \sigma_{uu} \end{pmatrix} = (\Sigma_{vM}, \Sigma_{\theta v}^T \alpha + \sigma_{vu}). \quad (3)$$

Step 2. We now consider the optimization problem of a strategic trader i . Suppose he observes signal realization $\tilde{\theta}_i$ of signal θ_i , and subsequently submits demand d . Assuming that other traders $j \neq i$ follow linear strategies α_j , and that the market maker follows a linear pricing rule $(\beta_M; \beta_D)$, the expected profit of trader i from submitting demand d when observing realization $\tilde{\theta}_i$ is equal to

$$E \left[d \left(v - \beta_M^T \theta_M - \beta_D \left(d + \sum_{j \neq i} \alpha_j^T \theta_j + u \right) \right) \middle| \theta_i = \tilde{\theta}_i \right]. \quad (4)$$

Using the fact that d is a choice variable, and thus d and d^2 are constants from the point of view of taking expectations, we can rewrite equation (4) as

$$d \cdot E \left[v - \beta_M^T \theta_M - \beta_D \left(\sum_{j \neq i} \alpha_j^T \theta_j + u \right) \middle| \theta_i = \tilde{\theta}_i \right] - d^2 \cdot \beta_D. \quad (5)$$

Now, if $\beta_D < 0$, trader i can make an arbitrarily large expected profit, and no single d maximizes it—hence, β_D cannot be negative in equilibrium.

If $\beta_D = 0$, and $E \left[v - \beta_M^T \theta_M \middle| \theta_i = \tilde{\theta}_i \right] \neq 0$, then again trader i can make an arbitrarily large expected profit, and no single d maximizes it. But it follows from Assumption 1 in the model²⁹ that for at least one trader i , for at least some (in fact, for almost all) realizations $\tilde{\theta}_i$, we have $E \left[v - \beta_M^T \theta_M \middle| \theta_i = \tilde{\theta}_i \right] \neq 0$ —hence, β_D cannot be equal to zero in equilibrium.

Finally, if $\beta_D > 0$, then there is a unique d maximizing the expected profit:

$$d^* = \frac{1}{2\beta_D} E \left[v - \beta_M^T \theta_M - \beta_D \left(\sum_{j \neq i} \alpha_j^T \theta_j + u \right) \middle| \theta_i = \tilde{\theta}_i \right] \quad (6)$$

$$= \frac{1}{2\beta_D} \left(\Sigma_{iv}^T - \beta_M^T \Sigma_{iM}^T - \beta_D \left(\sum_{j \neq i} \alpha_j^T \Sigma_{ij}^T + \Sigma_{iu}^T \right) \right) \Sigma_{ii}^{-1} \tilde{\theta}_i, \quad (7)$$

where equation (7) is the standard projection/signal extraction formula for jointly normal variables. Note that d^* is a linear function of $\tilde{\theta}_i$, and vector α_i is uniquely determined by pricing rule $(\beta_M; \beta_D)$ and strategies α_j for $j \neq i$.

²⁸To see the equivalence, note first that $\beta^T \eta = E[v|\eta] \implies Cov(v, \eta) = Cov(E[v|\eta], \eta) = Cov(\beta^T \eta, \eta) = \beta^T Var(\eta)$. To go in the opposite direction, note that $Cov(v, \eta) = \beta^T Var(\eta) = Cov(\beta^T \eta, \eta) \implies Cov(v - \beta^T \eta, \eta) = 0$. Since variables $v - \beta^T \eta$ and η are jointly normal, $Cov(v - \beta^T \eta, \eta) = 0$ implies that they are independent, and thus for every realization $\tilde{\eta}$ of random variable η , $E[v - \beta^T \eta | \eta = \tilde{\eta}] = E[v - \beta^T \eta] = 0$, which implies that for every realization $\tilde{\eta}$, $E[v|\eta = \tilde{\eta}] = E[\beta^T \eta | \eta = \tilde{\eta}] = \beta^T \tilde{\eta}$.

²⁹Assumption 1 says that at least one strategic trader i has some useful information beyond that contained in the market maker's signal: $Cov(v, \theta | \theta_M) \neq 0$.

Step 3. We showed in Steps 1 and 2 that profile of strategies α and pricing rule $(\beta_M; \beta_D)$ form a linear equilibrium if and only if $\beta_D > 0$ and the following two conditions hold:

- (i) $(\beta_M^T, \beta_D) \begin{pmatrix} \Sigma_{MM} & \Sigma_{\theta M}^T \alpha + \Sigma_{Mu} \\ \alpha^T \Sigma_{\theta M} + \Sigma_{Mu}^T & \alpha^T \Sigma_{\theta\theta} \alpha + 2\Sigma_{\theta u}^T \alpha + \sigma_{uu} \end{pmatrix} = (\Sigma_{vM}, \Sigma_{\theta v}^T \alpha + \sigma_{vu})$;
- (ii) for all i , $\alpha_i^T = \frac{1}{2\beta_D} \left(\Sigma_{iv}^T - \beta_M^T \Sigma_{iM}^T - \beta_D \left(\sum_{j \neq i} \alpha_j^T \Sigma_{ij}^T + \Sigma_{iu}^T \right) \right) \Sigma_{ii}^{-1}$.

We will now show that there is a unique profile (α, β) satisfying these conditions, thus proving the existence and uniqueness of linear equilibrium.

First, we rewrite condition (ii), for all i , as:

$$2\Sigma_{ii}\alpha_i = \frac{1}{\beta_D} (\Sigma_{iv} - \Sigma_{iM}\beta_M) - \sum_{j \neq i} \Sigma_{ij}\alpha_j - \Sigma_{iu} \quad (8)$$

or equivalently

$$\Sigma_{ii}\alpha_i + \sum_{j=1}^n \Sigma_{ij}\alpha_j = \frac{1}{\beta_D} (\Sigma_{iv} - \Sigma_{iM}\beta_M) - \Sigma_{iu}. \quad (9)$$

“Stacking” equations (9) for all i one under another, and rewriting the resulting system of equations in matrix form using the notation defined in Section 3.1, we obtain the following condition (equivalent to condition (ii)):

$$(\Sigma_{diag} + \Sigma_{\theta\theta})\alpha = \gamma\Sigma_{\theta v} - \Sigma_{\theta M}\beta'_M - \Sigma_{\theta u}, \quad (10)$$

where for convenience we define $\gamma = 1/\beta_D$, $\beta'_M = \beta_M/\beta_D$.

Next, using this notation, and transposing the matrix equation in condition (i), that condition can be written as a system of two equations:

$$\Sigma_{MM}\beta'_M + \Sigma_{\theta M}^T \alpha + \Sigma_{Mu} = \gamma\Sigma_{Mv}, \quad (11)$$

$$\alpha^T \Sigma_{\theta M} \beta'_M + \Sigma_{uM} \beta'_M + \alpha^T \Sigma_{\theta\theta} \alpha + 2\Sigma_{\theta u}^T \alpha + \sigma_{uu} = \gamma(\Sigma_{\theta v}^T \alpha + \sigma_{vu}). \quad (12)$$

Step 4. We will now solve the system of equations (10), (11), and (12). Equation (11) allows us to express β'_M as a function of α and γ :

$$\beta'_M = \Sigma_{MM}^{-1} (\gamma\Sigma_{Mv} - \Sigma_{\theta M}^T \alpha - \Sigma_{Mu}). \quad (13)$$

We then plug this expression of β'_M into equation (10):

$$(\Sigma_{diag} + \Sigma_{\theta\theta})\alpha = \gamma\Sigma_{\theta v} - \Sigma_{\theta M}\Sigma_{MM}^{-1} (\gamma\Sigma_{Mv} - \Sigma_{\theta M}^T \alpha - \Sigma_{Mu}) - \Sigma_{\theta u},$$

or, isolating α on the left-hand side and collecting the terms with γ ,

$$(\Sigma_{diag} + \Sigma_{\theta\theta} - \Sigma_{\theta M}\Sigma_{MM}^{-1}\Sigma_{\theta M}^T)\alpha = (\Sigma_{\theta v} - \Sigma_{\theta M}\Sigma_{MM}^{-1}\Sigma_{Mv})\gamma - (\Sigma_{\theta u} - \Sigma_{\theta M}\Sigma_{MM}^{-1}\Sigma_{Mu}).$$

Note that

$$\begin{aligned} \Sigma_{\theta\theta} - \Sigma_{\theta M}\Sigma_{MM}^{-1}\Sigma_{\theta M}^T &= Var(\theta) - Cov(\theta, \theta_M)Var(\theta_M)^{-1}Cov(\theta_M, \theta) \\ &= Var(\theta|\theta_M), \end{aligned}$$

where the last equation follows from the standard projection formula for multivariate normal distributions. Thus, matrix $\Sigma_{\theta\theta} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T$ is positive semidefinite, and matrix $\Sigma_{diag} + \Sigma_{\theta\theta} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T$ is positive definite (and thus invertible). Letting

$$\begin{aligned}\Lambda &= \Sigma_{diag} + \Sigma_{\theta\theta} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T, \\ A_u &= \Lambda^{-1} (\Sigma_{\theta u} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mu}), \\ A_v &= \Lambda^{-1} (\Sigma_{\theta v} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mv}),\end{aligned}$$

we can express α as a linear function of γ :

$$\alpha = \gamma A_v - A_u.$$

Plugging this expression into (13), we can also express β'_M as a linear function of γ :

$$\begin{aligned}\beta'_M &= \Sigma_{MM}^{-1} (\gamma \Sigma_{Mv} - \Sigma_{\theta M}^T (\gamma A_v - A_u) - \Sigma_{Mu}) \\ &= \gamma \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) - \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u).\end{aligned}$$

Using these expressions, we can now rewrite equation (12) as a quadratic equation of just one scalar variable, γ :

$$a\gamma^2 + b\gamma + c = 0, \tag{14}$$

where

$$\begin{aligned}a &= A_v^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) + A_v^T \Sigma_{\theta\theta} A_v - \Sigma_{\theta v}^T A_v, \\ b &= -A_v^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u) - A_u^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) \\ &\quad + \Sigma_{uM} \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) - 2A_v^T \Sigma_{\theta\theta} A_u + 2\Sigma_{\theta u}^T A_v + \Sigma_{\theta v}^T A_u - \sigma_{vu}, \\ c &= A_u^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u) - \Sigma_{uM} \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u) \\ &\quad + A_u^T \Sigma_{\theta\theta} A_u - 2\Sigma_{\theta u}^T A_u + \sigma_{uu}.\end{aligned}$$

Therefore, finding a linear equilibrium is equivalent to finding a positive root of equation (14). To prove that this equation has a unique such root, we first simplify the expressions for a , b , and c . (For the proof, it is sufficient to simplify a and c , but getting a simplified expression for b is useful for deriving an explicit analytic characterization of the equilibrium.) Starting with a :

$$\begin{aligned}a &= A_v^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) + A_v^T \Sigma_{\theta\theta} A_v - \Sigma_{\theta v}^T A_v, \\ &= A_v^T [(\Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mv} - \Sigma_{\theta v}) + (\Sigma_{\theta\theta} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T) A_v] \\ &= A_v^T [(-\Lambda A_v) + (\Lambda - \Sigma_{diag}) A_v] \\ &= -A_v^T \Sigma_{diag} A_v.\end{aligned}$$

Next,

$$\begin{aligned}b &= -A_v^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u) - A_u^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) \\ &\quad + \Sigma_{uM} \Sigma_{MM}^{-1} (\Sigma_{Mv} - \Sigma_{\theta M}^T A_v) - 2A_v^T \Sigma_{\theta\theta} A_u + 2\Sigma_{\theta u}^T A_v + \Sigma_{\theta v}^T A_u - \sigma_{vu}, \\ &= 2A_v^T (\Sigma_{\theta u} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mu}) + A_u^T (\Sigma_{\theta v} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mv})\end{aligned}$$

$$\begin{aligned}
& +2A_v^T (\Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T - \Sigma_{\theta\theta}) A_u + \Sigma_{uM} \Sigma_{MM}^{-1} \Sigma_{Mv} - \sigma_{uv} \\
= & 2A_v^T \Lambda A_u + A_u^T \Lambda A_v \\
& +2A_v^T (\Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T - \Sigma_{\theta\theta}) A_u + \Sigma_{uM} \Sigma_{MM}^{-1} \Sigma_{Mv} - \sigma_{uv} \\
= & A_v^T (2\Sigma_{diag} + \Lambda) A_u + \Sigma_{uM} \Sigma_{MM}^{-1} \Sigma_{Mv} - \sigma_{uv}.
\end{aligned}$$

Finally,

$$\begin{aligned}
c & = A_u^T \Sigma_{\theta M} \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u) - \Sigma_{uM} \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u) \\
& + A_u^T \Sigma_{\theta\theta} A_u - 2\Sigma_{\theta u}^T A_u + \sigma_{uu} \\
& = -(\Sigma_{uM} - A_u^T \Sigma_{\theta M})^T \Sigma_{MM}^{-1} (\Sigma_{Mu} - \Sigma_{\theta M}^T A_u) \\
& + A_u^T \Sigma_{\theta\theta} A_u - 2\Sigma_{\theta u}^T A_u + \sigma_{uu} \\
& = \begin{pmatrix} A_u \\ -1 \end{pmatrix}^T C \begin{pmatrix} A_u \\ -1 \end{pmatrix},
\end{aligned}$$

where

$$\begin{aligned}
C & = \begin{pmatrix} \Sigma_{\theta\theta} & \Sigma_{\theta u} \\ \Sigma_{\theta u}^T & \sigma_{uu} \end{pmatrix} - \begin{pmatrix} \Sigma_{\theta M} \\ \Sigma_{uM} \end{pmatrix} \Sigma_{MM}^{-1} \begin{pmatrix} \Sigma_{\theta M} \\ \Sigma_{uM} \end{pmatrix}^T \\
& = Var((\theta; u)) - Cov((\theta; u), \theta_M) Var(\theta_M)^{-1} Cov(\theta_M, (\theta; u)) \\
& = Var((\theta; u) | \theta_M).
\end{aligned}$$

Thus,

$$c = Var(A_u^T \theta - u | \theta_M).$$

Step 5. We will now determine the signs of coefficients a and c .

Matrix Σ_{diag} is positive definite, by construction. Vector A_v is not equal to zero: matrix Λ^{-1} is positive definite, and vector $\Sigma_{\theta v} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mv} = Cov(\theta, v | \theta_M)$ is not equal to zero (by Assumption 1 of the model). Thus, $a = -A_v^T \Sigma_{diag} A_v < 0$.

To determine the sign of coefficient c , recall that we have shown in Step 4 that $c = Var(A_u^T \theta - u | \theta_M)$. So if we show that $c \neq 0$, it will immediately follow that $c > 0$.

If $A_u = 0$, then $c \neq 0$ follows from Assumption 2 of the model (which says that the market maker does not perfectly observe liquidity demand: $Var(u | \theta_M) > 0$).

Suppose $A_u \neq 0$. It is convenient to introduce an auxiliary random variable, ϕ , drawn randomly from the normal distribution with mean zero and covariance matrix Σ_{diag} , independent of all other random variables in the model. Note that matrix A_u now has a simple interpretation:

$$A_u = Var(\theta + \phi | \theta_M)^{-1} Cov(\theta, u | \theta_M) = Var(\theta + \phi | \theta_M)^{-1} Cov(\theta + \phi, u | \theta_M).$$

Let $\epsilon = u - A_u^T (\theta + \phi)$. Then $c = Var(\epsilon + A_u^T \phi | \theta_M)$. To show that $c \neq 0$, it is thus sufficient to show that $\epsilon + A_u^T \phi$ is not constant, conditional on θ_M . To show that, consider $Cov(\epsilon + A_u^T \phi, A_u^T (\theta + \phi) | \theta_M) = Cov(\epsilon, A_u^T (\theta + \phi) | \theta_M) + Cov(A_u^T \phi, A_u^T (\theta + \phi) | \theta_M)$.

First, $Cov(\epsilon, A_u^T (\theta + \phi) | \theta_M) = Cov(u - A_u^T (\theta + \phi), A_u^T (\theta + \phi) | \theta_M) = Cov(u, \theta + \phi | \theta_M) A_u -$

$$A_u^T \text{Var}(\theta + \phi | \theta_M) A_u = 0.$$

Second, $\text{Cov}(A_u^T \phi, A_u^T(\theta + \phi) | \theta_M) = \text{Var}(A_u^T \phi | \theta_M) = A_u^T \Sigma_{diag} A_u$, which is not equal to zero, because $A_u \neq 0$ and Σ_{diag} is positive definite. Therefore, $\text{Cov}(\epsilon + A_u^T \phi, A_u^T(\theta + \phi) | \theta_M) \neq 0$, and thus $\epsilon + A_u^T \phi$ is not constant conditional on θ_M , and so $c > 0$.

Thus, $a < 0$, $c > 0$, and hence equation (14) has exactly one positive root. Therefore, there exists a unique linear equilibrium.

Appendix B: Proof of Theorem 2 (Special Case)

Step 1. Consider first a specific market m , and, for convenience, drop superscript (m). We know there exists a unique linear equilibrium. In this equilibrium, any two strategic traders in the same group have the same linear strategy (otherwise, by swapping the strategies of these two traders, we would be able to obtain a different linear equilibrium). Denote by α_i the aggregate demand multiplier, in equilibrium, of group i ; i.e., given signal $\theta_i + \xi_{i,j}$ of trader j in group i , the trader submits demand $\frac{1}{\ell_i} \alpha_i^T (\theta_i + \xi_{i,j})$.

For the remainder of this proof, we define the variables $\xi_i = \frac{1}{\ell_i} \sum_j \xi_{i,j}$, $\xi = (\xi_1; \dots; \xi_n)$, and the matrices

$$\Sigma^\xi = \text{Var}(\xi) = \begin{pmatrix} \frac{1}{\ell_1} \Sigma_1^\xi & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{1}{\ell_n} \Sigma_n^\xi \end{pmatrix}, \quad \widehat{\Sigma}_{diag} = \begin{pmatrix} \frac{1}{\ell_1} (\Sigma_{11} + \Sigma_1^\xi) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{1}{\ell_n} (\Sigma_{nn} + \Sigma_n^\xi) \end{pmatrix}.$$

With this notation, the equilibrium condition (i) in Step 1 of the proof of Theorem 1—the market maker's inference given her information—becomes:

$$(\beta_M^T, \beta_D) \begin{pmatrix} \Sigma_{MM} & \Sigma_{\theta M}^T \alpha + \Sigma_{Mu} \\ \alpha^T \Sigma_{\theta M} + \Sigma_{Mu}^T & \alpha^T (\Sigma_{\theta\theta} + \Sigma^\xi) \alpha + 2 \Sigma_{\theta u}^T \alpha + \sigma_{uu} \end{pmatrix} = (\Sigma_{vM}, \Sigma_{\theta v}^T \alpha + \sigma_{vu})$$

where we observe that the only modification is in the variance of the overall demand, which is now written

$$\begin{aligned} \text{Var} \left(\sum_{i,j} \frac{\alpha_i^T}{\ell_i} (\theta_i + \xi_{i,j}) + u \right) &= \alpha^T \Sigma_{\theta\theta} \alpha + 2 \Sigma_{\theta u}^T \alpha + \sigma_{uu} + \sum_i \frac{\alpha_i^T \Sigma_i^\xi \alpha_i}{\ell_i} \\ &= \alpha^T (\Sigma_{\theta\theta} + \Sigma^\xi) \alpha + 2 \Sigma_{\theta u}^T \alpha + \sigma_{uu}. \end{aligned}$$

The equations (11) and (12) that capture condition (i) then become slightly different:

$$\Sigma_{MM} \beta'_M + \Sigma_{\theta M}^T \alpha + \Sigma_{Mu} = \gamma \Sigma_{Mv}, \quad (15)$$

$$\alpha^T \Sigma_{\theta M} \beta'_M + \Sigma_{uM} \beta'_M + \alpha^T (\Sigma_{\theta\theta} + \Sigma^\xi) \alpha + 2 \Sigma_{\theta u}^T \alpha + \sigma_{uu} = \gamma (\Sigma_{\theta v}^T \alpha + \sigma_{vu}). \quad (16)$$

The equilibrium condition (ii) in Step 3 of the proof of Theorem 1—the best response of strategic trader i —is also slightly different. In this new notation, it becomes:

$$\frac{1}{\ell_i} \alpha_i^T = \frac{1}{2\beta_D} \left(\Sigma_{iv}^T - \beta_M^T \Sigma_{iM}^T - \beta_D \left(\sum_{j \neq i} \alpha_j^T \Sigma_{ij}^T + \frac{\ell_i - 1}{\ell_i} \alpha_i^T \Sigma_{ii}^T + \Sigma_{iu}^T \right) \right) (\Sigma_{ii} + \Sigma_i^\xi)^{-1},$$

which is equivalent to:

$$\frac{\Sigma_{ii}}{\ell_i} \alpha_i + \frac{2\Sigma_i^\xi}{\ell_i} \alpha_i + \sum_j \Sigma_{ij} \alpha_j = \beta_D^{-1} [\Sigma_{iv} - \Sigma_{iM} \beta_M] - \Sigma_{iu}.$$

Similarly to equation (10) in the proof of Theorem 1, this can be rewritten as

$$(\widehat{\Sigma}_{diag} + \Sigma_{\theta\theta} + \Sigma^\xi) \alpha = \gamma \Sigma_{\theta v} - \Sigma_{\theta M} \beta'_M - \Sigma_{\theta u}, \quad (17)$$

where γ and β'_M are defined as before.

Next, again by analogy with the proof of Theorem 1, we define

$$\begin{aligned} \widehat{\Lambda} &= \widehat{\Sigma}_{diag} + \Sigma_{\theta\theta} + \Sigma^\xi - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T, \\ \widehat{A}_u &= \widehat{\Lambda}^{-1} (\Sigma_{\theta u} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mu}), \\ \widehat{A}_v &= \widehat{\Lambda}^{-1} (\Sigma_{\theta v} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mv}), \end{aligned}$$

and then finding a linear equilibrium is equivalent to solving the quadratic equation

$$a\gamma^2 + b\gamma + c = 0,$$

where

$$\begin{aligned} a &= -\widehat{A}_v^T \widehat{\Sigma}_{diag} \widehat{A}_v, \\ b &= \widehat{A}_v^T \left(2\widehat{\Sigma}_{diag} + \widehat{\Lambda} \right) \widehat{A}_u + \Sigma_{uM} \Sigma_{MM}^{-1} \Sigma_{Mv} - \sigma_{uv}, \\ c &= Var(\widehat{A}_u^T (\theta + \xi) - u | \theta_M) = Var(\widehat{A}_u^T \theta - u | \theta_M) + \widehat{A}_u^T \Sigma^\xi \widehat{A}_u. \end{aligned}$$

Since by definition $\gamma = 1/\beta_D$, solving the above quadratic equation is equivalent to solving the quadratic equation

$$c\beta_D^2 + b\beta_D + a = 0,$$

which turns out to be a more convenient characterization that we will proceed with. As before, we also have a simple expression for the vector of strategies α :

$$\alpha = \widehat{A}_v / \beta_D - \widehat{A}_u.$$

Step 2. Let us now consider the entire sequence of markets, and restore superscript (m) for the variables. From the simplifying assumption that $Var(\theta; \theta_M; u)$ is full rank, it follows that both $Var(\theta | \theta_M)$ and $Var(\theta_M | \theta)$ are full rank, and thus invertible.

As $m \rightarrow \infty$, $\widehat{\Sigma}_{diag}^{(m)} \rightarrow 0$ and $\Sigma^{\xi, (m)} \rightarrow 0$. Thus,

$$\begin{aligned} \widehat{\Lambda}^{(m)} &\rightarrow \Sigma_{\theta\theta} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{\theta M}^T = Var(\theta | \theta_M), \\ \widehat{A}_u^{(m)} &\rightarrow Var(\theta | \theta_M)^{-1} (\Sigma_{\theta u} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mu}) = Var(\theta | \theta_M)^{-1} Cov(\theta, u | \theta_M), \\ \widehat{A}_v^{(m)} &\rightarrow Var(\theta | \theta_M)^{-1} (\Sigma_{\theta v} - \Sigma_{\theta M} \Sigma_{MM}^{-1} \Sigma_{Mv}) = Var(\theta | \theta_M)^{-1} Cov(\theta, v | \theta_M). \end{aligned}$$

Therefore, using that $\widehat{A}_u^T \Sigma^{\xi, (m)} \widehat{A}_u \rightarrow 0$,

$$\begin{aligned} a^{(m)} &\rightarrow 0, \\ b^{(m)} &\rightarrow Cov(v, \theta | \theta_M) Var(\theta | \theta_M)^{-1} Var(\theta | \theta_M) Var(\theta | \theta_M)^{-1} Cov(\theta, u | \theta_M) + \Sigma_{uM} \Sigma_{MM}^{-1} \Sigma_{Mv} - \sigma_{uv} \end{aligned}$$

$$\begin{aligned}
&= \text{Cov}(v, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} \text{Cov}(\theta, u | \theta_M) - \text{Cov}(u, v | \theta_M) = -\text{Cov}(u, v | \theta, \theta_M), \\
c^{(m)} &\rightarrow \text{Var}(\text{Cov}(u, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} \theta - u | \theta_M) + 0 \\
&= \text{Var}(E[u | \theta, \theta_M] - u | \theta_M) = \text{Var}(u | \theta, \theta_M).
\end{aligned}$$

Note that these convergence results imply that $\beta_D^{(m)}$ converges to some finite value, since $\lim_{m \rightarrow \infty} c^{(m)} = \text{Var}(u | \theta, \theta_M) > 0$ (where the last inequality is due to Assumption 2L). If $\text{Cov}(u, v | \theta, \theta_M) > 0$, then $\lim_{m \rightarrow \infty} \beta_D^{(m)} = \text{Var}(u | \theta, \theta_M)^{-1} \text{Cov}(u, v | \theta, \theta_M)$. If $\text{Cov}(u, v | \theta, \theta_M) \leq 0$, then $\lim_{m \rightarrow \infty} \beta_D^{(m)} = 0$. We now consider the limiting behavior of price $p^{(m)}$ in these two cases separately.

Step 3, Case $\text{Cov}(u, v | \theta, \theta_M) > 0$. Note first that

$$\begin{aligned}
E[v | \theta, \theta_M, u] &= E[v | \theta_M] \\
&\quad + \text{Cov}(v, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} (\theta - E[\theta | \theta_M]) \\
&\quad + \text{Cov}(v, u | \theta, \theta_M) \text{Var}(u | \theta, \theta_M)^{-1} (u - E[u | \theta, \theta_M]) \\
&= E[v | \theta_M] \\
&\quad + \text{Cov}(v, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} (\theta - E[\theta | \theta_M]) \\
&\quad + \text{Cov}(v, u | \theta, \theta_M) \text{Var}(u | \theta, \theta_M)^{-1} \\
&\quad \quad \times (u - E[u | \theta_M] - \text{Cov}(u, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} (\theta - E[\theta | \theta_M])).
\end{aligned}$$

Thus, $E[v | \theta, \theta_M, u]$ is a linear function of θ , θ_M , and u :

$$E[v | \theta, \theta_M, u] = w_M^T \theta_M + w_\theta^T \theta + w_u u,$$

where weights w are as follows:

$$\begin{aligned}
w_M^T &= \text{Cov}(v, \theta_M) \text{Var}(\theta_M)^{-1} \\
&\quad - \text{Cov}(v, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} \text{Cov}(\theta, \theta_M) \text{Var}(\theta_M)^{-1} \\
&\quad - \text{Cov}(v, u | \theta, \theta_M) \text{Var}(u | \theta, \theta_M)^{-1} \text{Cov}(u, \theta_M) \text{Var}(\theta_M)^{-1} \\
&\quad + \text{Cov}(v, u | \theta, \theta_M) \text{Var}(u | \theta, \theta_M)^{-1} \text{Cov}(u, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} \text{Cov}(\theta, \theta_M) \text{Var}(\theta_M)^{-1}; \\
w_\theta^T &= \text{Cov}(v, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1} \\
&\quad - \text{Cov}(v, u | \theta, \theta_M) \text{Var}(u | \theta, \theta_M)^{-1} \text{Cov}(u, \theta | \theta_M) \text{Var}(\theta | \theta_M)^{-1}; \\
w_u &= \text{Cov}(v, u | \theta, \theta_M) \text{Var}(u | \theta, \theta_M)^{-1}.
\end{aligned}$$

Next, price $p^{(m)}(\theta, \xi^{(m)}, \theta_M, u)$ in market m can be expressed as

$$\begin{aligned}
p^{(m)}(\theta, \xi^{(m)}, \theta_M, u) &= \beta_M^{(m)T} \theta_M + \beta_D^{(m)} \left(\alpha^{(m)T} (\theta + \xi^{(m)}) + u \right) \\
&= \beta_M^{(m)T} \theta_M + \beta_D^{(m)} \alpha^{(m)T} (\theta + \xi^{(m)}) + \beta_D^{(m)} u.
\end{aligned}$$

To prove the statement of the theorem for this case, note that

$$E \left[\left(p^{(m)}(\theta, \xi^{(m)}, \theta_M, u) - E[v | \theta, \theta_M, u] \right)^2 \right] = (\beta_D^{(m)})^2 \alpha^{(m)T} \Sigma^{\xi, (m)} \alpha^{(m)}$$

$$+ \begin{pmatrix} \beta_M^{(m)} - w_M \\ \beta_D^{(m)} \alpha^{(m)} - w_\theta \\ \beta_D^{(m)} - w_u \end{pmatrix}^T \text{Var} \left(\begin{pmatrix} \theta_M \\ \theta \\ u \end{pmatrix} \right) \begin{pmatrix} \beta_M^{(m)} - w_M \\ \beta_D^{(m)} \alpha^{(m)} - w_\theta \\ \beta_D^{(m)} - w_u \end{pmatrix}.$$

Since $\Sigma^{\xi, (m)} \rightarrow 0$, it is enough to show that $\beta_D^{(m)} \rightarrow w_u$, $\beta_D^{(m)} \alpha^{(m)} \rightarrow w_\theta$, and $\beta_M^{(m)} \rightarrow w_M$.

The first convergence result is immediate:

$$\lim_{m \rightarrow \infty} \beta_D^{(m)} = \text{Var}(u|\theta, \theta_M)^{-1} \text{Cov}(u, v|\theta, \theta_M) = w_u.$$

Next:

$$\begin{aligned} \lim_{m \rightarrow \infty} \beta_D^{(m)} \alpha^{(m)} &= \lim_{m \rightarrow \infty} \widehat{A}_v^{(m)} - \beta_D^{(m)} \widehat{A}_u^{(m)} \\ &= \text{Var}(\theta|\theta_M)^{-1} \text{Cov}(\theta, v|\theta_M) \\ &\quad - \text{Var}(u|\theta, \theta_M)^{-1} \text{Cov}(u, v|\theta, \theta_M) \text{Var}(\theta|\theta_M)^{-1} \text{Cov}(\theta, u|\theta_M) \\ &= w_\theta. \end{aligned}$$

Finally:

$$\begin{aligned} \lim_{m \rightarrow \infty} \beta_M^{(m)} &= \lim_{m \rightarrow \infty} \Sigma_{MM}^{-1} \left(\Sigma_{Mv} - \Sigma_{\theta M}^T \widehat{A}_v^{(m)} \right) - \beta_D^{(m)} \Sigma_{MM}^{-1} \left(\Sigma_{Mu} - \Sigma_{\theta M}^T \widehat{A}_u^{(m)} \right) \\ &= \Sigma_{MM}^{-1} \left(\Sigma_{Mv} - \Sigma_{\theta M}^T \text{Var}(\theta|\theta_M)^{-1} \text{Cov}(\theta, v|\theta_M) \right) \\ &\quad - \text{Var}(u|\theta, \theta_M)^{-1} \text{Cov}(u, v|\theta, \theta_M) \Sigma_{MM}^{-1} \left(\Sigma_{Mu} - \Sigma_{\theta M}^T \text{Var}(\theta|\theta_M)^{-1} \text{Cov}(\theta, u|\theta_M) \right) \\ &= \text{Var}(\theta_M)^{-1} \text{Cov}(\theta_M, v) \\ &\quad - \text{Var}(\theta_M)^{-1} \text{Cov}(\theta, \theta_M)^T \text{Var}(\theta|\theta_M)^{-1} \text{Cov}(\theta, v|\theta_M) \\ &\quad - \text{Var}(u|\theta, \theta_M)^{-1} \text{Cov}(u, v|\theta, \theta_M) \text{Var}(\theta_M)^{-1} \text{Cov}(\theta_M, u) \\ &\quad + \text{Var}(u|\theta, \theta_M)^{-1} \text{Cov}(u, v|\theta, \theta_M) \text{Var}(\theta_M)^{-1} \text{Cov}(\theta, \theta_M)^T \text{Var}(\theta|\theta_M)^{-1} \text{Cov}(\theta, u|\theta_M) \\ &= w_M. \end{aligned}$$

Step 3, Case $\text{Cov}(u, v|\theta, \theta_M) \leq \mathbf{0}$. In this case, note that

$$\begin{aligned} E[v|\theta, \theta_M] &= E[v|\theta_M] \\ &\quad + \text{Cov}(v, \theta|\theta_M) \text{Var}(\theta|\theta_M)^{-1} (\theta - E[\theta|\theta_M]). \end{aligned}$$

Thus, $E[v|\theta, \theta_M]$ is a linear function of θ and θ_M :

$$E[v|\theta, \theta_M] = w_M^T \theta_M + w_\theta^T \theta,$$

where weights w are as follows:

$$\begin{aligned} w_M^T &= \text{Cov}(v, \theta_M) \text{Var}(\theta_M)^{-1} \\ &\quad - \text{Cov}(v, \theta|\theta_M) \text{Var}(\theta|\theta_M)^{-1} \text{Cov}(\theta, \theta_M) \text{Var}(\theta_M)^{-1}; \\ w_\theta^T &= \text{Cov}(v, \theta|\theta_M) \text{Var}(\theta|\theta_M)^{-1}. \end{aligned}$$

As before, price $p^{(m)}(\theta, \xi^{(m)}, \theta_M, u)$ in market m can be expressed as

$$p^{(m)}(\theta, \xi^{(m)}, \theta_M, u) = \beta_M^{(m)T} \theta_M + \beta_D^{(m)} \left(\alpha^{(m)T} (\theta + \xi^{(m)}) + u \right)$$

$$= \beta_M^{(m)T} \theta_M + \beta_D^{(m)} \alpha^{(m)T} (\theta + \xi^{(m)}) + \beta_D^{(m)} u.$$

As in Step 2, noting that $Var(\xi^{(m)}) \rightarrow 0$, to prove the theorem for this case, it is thus sufficient to show that as m grows, $\beta_D^{(m)} \rightarrow 0$, $\beta_D^{(m)} \alpha^{(m)} \rightarrow w_\theta$, and $\beta_M^{(m)} \rightarrow w_M$. The first convergence result, $\beta_D^{(m)} \rightarrow 0$, was proven at the end of Step 2 above. Next,

$$\begin{aligned} \lim_{m \rightarrow \infty} \beta_D^{(m)} \alpha^{(m)} &= \lim_{m \rightarrow \infty} \widehat{A}_v^{(m)} - \beta_D^{(m)} \widehat{A}_u^{(m)} \\ &= Var(\theta|\theta_M)^{-1} Cov(\theta, v|\theta_M) \\ &\quad - \left[\lim_{m \rightarrow \infty} \beta_D^{(m)} \right] Var(\theta|\theta_M)^{-1} Cov(\theta, u|\theta_M) \\ &= Var(\theta|\theta_M)^{-1} Cov(\theta, v|\theta_M) \\ &= w_\theta. \end{aligned}$$

Finally,

$$\begin{aligned} \lim_{m \rightarrow \infty} \beta_M^{(m)} &= \lim_{m \rightarrow \infty} \Sigma_{MM}^{-1} \left(\Sigma_{Mv} - \Sigma_{\theta M}^T \widehat{A}_v^{(m)} \right) - \beta_D^{(m)} \Sigma_{MM}^{-1} \left(\Sigma_{Mu} - \Sigma_{\theta M}^T \widehat{A}_u^{(m)} \right) \\ &= \lim_{m \rightarrow \infty} \Sigma_{MM}^{-1} \left(\Sigma_{Mv} - \Sigma_{\theta M}^T \widehat{A}_v^{(m)} \right) \\ &= Var(\theta_M)^{-1} Cov(\theta_M, v) \\ &\quad - Var(\theta_M)^{-1} Cov(\theta, \theta_M)^T Var(\theta|\theta_M)^{-1} Cov(\theta, v|\theta_M) \\ &= w_M. \end{aligned}$$

Appendix C: Proofs of Propositions in Section 7

C.1 Proof of Proposition 3

Fix a firm i , and suppose every firm $j \neq i$ plays according to a linear strategy

$$q_j(\theta_j) = \alpha_j^T \theta_j + \delta_j.$$

Suppose firm i observes realization $\tilde{\theta}_i$ of signal θ_i . The expected payoff of firm i from producing q units of the good is then equal to

$$E \left[q \left(v - \beta \left(q + \sum_{j \neq i} (\alpha_j^T \theta_j + \delta_j) \right) - c \right) \middle| \theta_i = \tilde{\theta}_i \right],$$

which can be rewritten as

$$q \cdot E \left[v - \beta \left(\sum_{j \neq i} (\alpha_j^T \theta_j + \delta_j) \right) - c \middle| \theta_i = \tilde{\theta}_i \right] - q^2 \cdot \beta.$$

Since by assumption $\beta > 0$, there is a unique q maximizing the expected profit:

$$q^* = \frac{1}{2\beta} \left(\bar{v} + \Sigma_{iv}^T \Sigma_{ii}^{-1} \theta_i - \beta \left(\sum_{j \neq i} (\alpha_j^T \Sigma_{ij}^T \Sigma_{ii}^{-1} \theta_i + \delta_j) \right) - c \right), \quad (18)$$

where we re-use our earlier notation for various covariance matrices.

Thus, if all firms other than i use strategies linear in their signals, firm i 's (unique) best response strategy is also linear in its signal. Moreover, the intercept and the slope of that strategy are uniquely determined. The intercept is given by

$$\delta_i = \frac{1}{2\beta} \left(\bar{v} - \beta \sum_{j \neq i} \delta_j - c \right), \quad (19)$$

and the slope is given by

$$\alpha_i^T = \frac{1}{2\beta} \left(\Sigma_{iv}^T \Sigma_{ii}^{-1} - \beta \sum_{j \neq i} \alpha_j^T \Sigma_{ij}^T \Sigma_{ii}^{-1} \right). \quad (20)$$

For the intercepts, multiplying both sides of equation (19) by 2β and moving one of the $\beta\delta_i$ terms under the summation sign, we get

$$\beta\delta_i = \bar{v} - \beta \sum_{j=1}^n \delta_j - c,$$

and so all δ_i are equal:

$$\delta_i = \frac{\bar{v} - c}{\beta(n+1)}.$$

For the slopes α_i , we follow manipulations analogous to those in the proof of Theorem 1: multiply both sides of equation (20) by $2\beta\Sigma_{ii}$ (on the right), move one of the $\beta\alpha_i^T\Sigma_{ii}$ terms under the summation sign, transpose the equation, and “stack” the resulting equations for all i . The resulting system of equation can be rewritten as

$$\beta\Sigma_{diag}\alpha = \Sigma_{iv} - \beta\Sigma_{\theta\theta}\alpha,$$

and so the vector of slopes α is given by the formula

$$\alpha = \frac{1}{\beta} (\Sigma_{\theta\theta} + \Sigma_{diag})^{-1} \Sigma_{iv},$$

because our assumptions imply that matrix $(\Sigma_{\theta\theta} + \Sigma_{diag})$ is invertible.

C.2 Proof of Proposition 4

Consider first the original sequence of markets and fix a particular market m (and for convenience, drop the superscript (m) for now). By Proposition 3, there exists a unique linear equilibrium. To explicitly characterize this equilibrium, we use the arguments and the notation almost identical to those in Step 1 of the proof of Theorem 2.

Specifically, by symmetry, any two firms in the same group use the same linear strategy in equilibrium. Denote by α_i the *aggregate* supply multiplier of group i , and by δ_i the *aggregate* intercept of group i . Thus, a specific firm j in group i , after observing its signal $\theta_i + \xi_{i,j}$, will produce quantity

$$\frac{1}{\ell_i} \alpha_i^T (\theta_i + \xi_{i,j}) + \frac{1}{\ell_i} \delta_i.$$

As in the proof of Theorem 2, let $\xi_i = \frac{1}{\ell_i} \sum_j \xi_{i,j}$ (the average idiosyncratic term in group i), let

$\xi = (\xi_1; \dots; \xi_n)$, and define matrices Σ^ξ and $\widehat{\Sigma}_{diag}$ as

$$\Sigma^\xi = Var(\xi) = \begin{pmatrix} \frac{1}{\ell_1} \Sigma_1^\xi & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{1}{\ell_n} \Sigma_n^\xi \end{pmatrix}, \quad \widehat{\Sigma}_{diag} = \begin{pmatrix} \frac{1}{\ell_1} (\Sigma_{11} + \Sigma_1^\xi) & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{1}{\ell_n} (\Sigma_{nn} + \Sigma_n^\xi) \end{pmatrix}.$$

It is immediate from Proposition 3 that $\delta_i = \frac{\ell_i(\bar{v}-c)}{\beta(N+1)}$. For vector α , writing down the first-order conditions for all firms j of all groups i , and combining them in the same way as in Step 1 of the proof of Theorem 2, we get the expression

$$\left(\widehat{\Sigma}_{diag} + \Sigma_{\theta\theta} + \Sigma^\xi \right) \alpha = \beta^{-1} \Sigma_{\theta v}. \quad (21)$$

(Note that equation (21) is almost identical to equation (17) in Step 1 of the proof of Theorem 2, except that the latter also contains the terms related to the signal observed by the market maker and the demand from liquidity traders.)

Let us now again write the market indices explicitly, so that

$$\alpha^{(m)} = \beta^{-1} \left(\Sigma_{\theta\theta} + \widehat{\Sigma}_{diag}^{(m)} + (\Sigma^\xi)^{(m)} \right)^{-1} \Sigma_{\theta v},$$

and

$$\delta^{(m)} = \frac{\ell_i(\bar{v}-c)}{\beta(N^{(m)}+1)}.$$

The total quantity produced in market m , as a function of θ and $\xi^{(m)}$, is then

$$\begin{aligned} Q^{(m)} &= \beta^{-1} \left(\left(\Sigma_{\theta\theta} + \widehat{\Sigma}_{diag}^{(m)} + (\Sigma^\xi)^{(m)} \right)^{-1} \Sigma_{\theta v} \right)^T \theta \\ &\quad + \beta^{-1} \left(\left(\Sigma_{\theta\theta} + \widehat{\Sigma}_{diag}^{(m)} + (\Sigma^\xi)^{(m)} \right)^{-1} \Sigma_{\theta v} \right)^T \xi^{(m)} \\ &\quad + \frac{N^{(m)}(\bar{v}-c)}{\beta(N^{(m)}+1)}. \end{aligned}$$

As m goes to infinity, $\frac{N^{(m)}(\bar{v}-c)}{\beta(N^{(m)}+1)}$ converges to $\beta^{-1}(\bar{v}-c)$, and matrices $\widehat{\Sigma}_{diag}^{(m)}$ and $(\Sigma^\xi)^{(m)}$ converge to zero. Moreover, $\xi^{(m)} \xrightarrow{L^2} 0$. Thus, as m goes to infinity,

$$Q^{(m)} \xrightarrow{L^2} \beta^{-1} \left((\Sigma_{\theta\theta}^{-1} \Sigma_{\theta v})^T \theta + (\bar{v}-c) \right) = \beta^{-1} (E[v|\theta] - c).$$

For the alternative sequence of markets, note that each alternative market m can be viewed as a special case of the “original” market, with just one group $i = 1$, and no idiosyncratic components of signals within the group (i.e., $\Sigma_i^\xi = 0$). Thus, the above derivation applies to this special case, and so for the alternative sequence, we also have

$$Q^{(alt,m)} \xrightarrow{L^2} \beta^{-1} (E[v|\theta] - c),$$

and so

$$Q^{(m)} - Q^{(alt,m)} \xrightarrow{L^2} 0.$$

Moreover, since $p^{(m)} = v - \beta Q^{(m)}$ and $p^{(alt,m)} = v - \beta Q^{(alt,m)}$, we immediately get

$$p^{(m)} - p^{(alt,m)} \xrightarrow{L^2} 0.$$

References

- Admati, A. R. and P. Pfleiderer (1988). A theory of intraday patterns: Volume and price variability. *Review of Financial Studies* 1(1), 3–40.
- Babus, A. and P. Kondor (2016). Trading and information diffusion in over-the-counter markets. Working paper.
- Back, K., C. H. Cao, and G. A. Willard (2000). Imperfect competition among informed traders. *Journal of Finance* 55(5), 2117–2155.
- Bagnoli, M., S. Viswanathan, and C. Holden (2001). On the existence of linear equilibria in models of market making. *Mathematical Finance* 11(1), 1–31.
- Bergemann, D., T. Heumann, and S. Morris (2015). Information and market power. Working paper.
- Bergemann, D. and S. Morris (2013). Robust predictions in games with incomplete information. *Econometrica* 81, 1251–1308.
- Bergemann, D. and S. Morris (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11, 487–522.
- Bernhardt, D. and J. Miao (2004). Informed trading when information becomes stale. *Journal of Finance* 59(1), 339–390.
- Biais, B., D. Martimort, and J.-C. Rochet (2000). Competing mechanisms in a common value environment. *Econometrica* 68(4), 799–837.
- Biais, B., D. Martimort, and J.-C. Rochet (2013). Corrigendum to “Competing mechanisms in a common value environment”. *Econometrica* 81(1), 393–406.
- Caballé, J. and M. Krishnan (1994). Imperfect competition in a multi-security market with risk neutrality. *Econometrica* 62(3), 695–704.
- Colla, P. and A. Mele (2010). Information linkages and correlated trading. *Review of Financial Studies* 23(1), 203–246.
- Dridi, R. and L. Germain (2009). Noise and competition in strategic oligopoly. *Journal of Financial Intermediation* 18(2), 311 – 327.
- Foster, F. D. and S. Viswanathan (1994). Strategic trading with asymmetrically informed traders and long-lived information. *Journal of Financial and Quantitative Analysis* 29, 499–518.
- Foster, F. D. and S. Viswanathan (1996). Strategic trading when agents forecast the forecasts of others. *Journal of Finance* 51(4), 1437–1478.
- Foucault, T. and L. Lescourret (2003). Information sharing, liquidity and transaction costs. Mimeo.

- Grossman, S. (1976). On the efficiency of competitive stock markets where traders have diverse information. *Journal of Finance* 31(2), 573–585.
- Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review* 35(4), 519–530.
- Holden, C. W. and A. Subrahmanyam (1992). Long-lived private information and imperfect competition. *The Journal of Finance* 47(1), 247–270.
- Jain, N. and L. J. Mirman (1999). Insider trading with correlated signals. *Economics Letters* 65(1), 105–113.
- Kremer, I. (2002). Information aggregation in common value auctions. *Econometrica* 70(4), 1675–1682.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica* 53(6), 1315–1336.
- Kyle, A. S. (1989). Informed speculation with imperfect competition. *Review of Economic Studies* 56(3), 317–355.
- Li, L. (1985). Cournot oligopoly with information sharing. *RAND Journal of Economics* 16(4), 521–536.
- Madrigal, V. (1996). Non-fundamental speculation. *Journal of Finance* 51(2), 553–578.
- Manzano, C. and X. Vives (2016). Market power and welfare in asymmetric divisible good auctions. Working paper.
- McLean, R., J. Peck, and A. Postlewaite (2005). On price-taking behavior in asymmetric information economies. *Essays in Dynamic General Equilibrium Theory: Festschrift for David Cass, Edited by A. Citanna, J. Donaldson, H. Polemarchkis, P. Siconolfi, and S. Spear, Springer Verlag*, 129–142.
- McLean, R. and A. Postlewaite (2002). Informational size and incentive compatibility. *Econometrica* 70(6), 2421–2453.
- Milgrom, P. R. (1981). Rational expectations, information acquisition, and competitive bidding. *Econometrica* 49(4), 921–943.
- Nöldeke, G. and T. Tröger (2001). Existence of linear equilibria in the Kyle model with multiple informed traders. *Economics Letters* 72(2), 159 – 164.
- Nöldeke, G. and T. Tröger (2006). A characterization of the distributions that imply existence of linear equilibria in the Kyle-model. *Annals of Finance* 2(1), 73 – 85.
- Palfrey, T. R. (1985). Uncertainty resolution, private information aggregation and the cournot competitive limit. *The Review of Economic Studies* 52(1), 69–83.
- Pasquariello, P. (2007). Imperfect competition, information heterogeneity, and financial contagion. *Review of Financial Studies* 20(2), 391–426.
- Pesendorfer, W. and J. M. Swinkels (1997). The loser’s curse and information aggregation in common value auctions. *Econometrica* 65(6), 1247–1282.
- Radner, R. (1979). Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica* 47(3), 655–678.

- Reny, P. J. and M. Perry (2006). Toward a strategic foundation for rational expectations equilibrium. *Econometrica* 74(5), 1231–1269.
- Rochet, J.-C. and J.-L. Vila (1994). Insider trading without normality. *Review of Economic Studies* 61(1), 131–152.
- Röell, A. (1990). Dual-capacity trading and the quality of the market. *Journal of Financial Intermediation* 1(2), 105–124.
- Rostek, M. and M. Weretka (2012). Price inference in small markets. *Econometrica* 80(2), 687–711.
- Rostek, M. and J. H. Yoon (2014). Inference design. Working paper. Available at <http://gtcenter.org/Archive/2015/Conf/Yoon1951.pdf>.
- Sarkar, A. (1995). Dual trading: Winners, losers, and market impact. *Journal of Financial Intermediation* 4(1), 77–93.
- Vives, X. (1988). Aggregation of information in large cournot markets. *Econometrica* 56(4), 851–876.
- Vives, X. (2008). *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton, NJ: Princeton University Press.
- Vives, X. (2011). Strategic supply function competition with private information. *Econometrica* 79(6), 1919–1966.
- Wilson, R. (1977). A bidding model of perfect competition. *Review of Economic Studies* 44(3), 511–518.