

Information Aggregation and Allocative Efficiency in Smooth Markets

Krishnamurthy Iyer

School of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853,
kriyer@cornell.edu

Ramesh Johari

Department of Management Science and Engineering, Stanford University, Stanford, California 94305,
ramesh.johari@stanford.edu

Ciamac C. Moallemi

Graduate School of Business, Columbia University, New York, New York 10027,
ciamac@gsb.columbia.edu

Recent years have seen extensive investigation of the information aggregation properties of markets. However, relatively little is known about conditions under which a market will aggregate the private information of rational risk-averse traders who optimize their portfolios over time; in particular, what features of a market encourage traders to ultimately reveal their private information through trades? We consider a market model involving finitely many informed risk-averse traders interacting with a market maker. Our main result identifies a basic *asymptotic smoothness* condition on prices in the market that ensures information is aggregated as long as portfolios converge; furthermore, under this assumption, the allocation achieved is ex post Pareto efficient. Asymptotic smoothness is fairly mild: it requires that, eventually, infinitesimal purchases or sales should see the same per-unit price. Notably, we demonstrate that, under some mild conditions, algorithmic markets based on cost functions (or, equivalently, markets based on market scoring rules) aggregate the information of traders.

Keywords: information aggregation; perfect Bayesian equilibrium; risk aversion

History: Received August 16, 2011; accepted February 3, 2014, by Brad M. Barber, finance. Published online in *Articles in Advance* July 15, 2014.

1. Introduction

Recent years have seen a surge of development and interest in *prediction markets*. These are typically online markets where the assets pay a fixed amount if a given event occurs. The basic goal of a prediction market is to determine the likelihood that a given event will happen: informally, if the price of an event is low relative to the likelihood that it will occur, then at least some traders should find it profitable to buy—and thus the price will go up. In this way, the expectation is that prices in prediction markets *aggregate* the information available to individual traders (Pennock and Sami 2007).

A wide range of prediction markets have emerged in the last decade. Two of the most well-known prediction markets include Betfair and Intrade, which run prediction markets on a wide range of outcomes (including political races, economic outcomes, etc.); these are both markets where individuals trade securities with real currency. Other markets, such as the Hollywood Stock Exchange, Crowdcast, and Inkling, provide prediction markets based on various forms of virtual currency. In all cases, the goal is to leverage market mechanisms

to aggregate information from diverse traders into an efficient prediction.

This basic goal has led to extensive investigation of the information aggregation properties of prediction markets (see, e.g., Healy et al. 2010, Jian and Sami 2012). More broadly, however, the notion that information is aggregated in market prices has been a long-standing topic of study in economics (see, e.g., Hayek 1945, Radner 1979, Glosten and Milgrom 1985, Kyle 1985).¹ Roughly speaking, the idea that market prices should reflect the information collectively possessed by market participants is one element of the *efficient markets hypothesis* (Lo 2008). A second, closely related objective is that markets that function well should achieve Pareto-efficient allocations: it should not be possible to reallocate assets between agents so that some agent's expected utility is improved without reducing the expected utility of another agent. Taken together, these goals—*information aggregation*

¹ Various strands of economics literature consider aggregation of information in dynamic economic environments; see also, e.g., the literature on social learning (Banerjee 1992, Bikhchandani et al. 1992).

and *allocative efficiency*—underlie much of the appeal of decentralized markets in settings where traders possess asymmetric information.

Despite the central connection between information aggregation, Pareto efficiency, and market behavior, significant open problems remain. In particular, relatively little is known about conditions under which a market will aggregate the private information of rational risk-averse traders who optimize their portfolios over time, or whether the market will achieve efficient allocations. The central goal of this paper is to investigate conditions under which these results are obtained.

Our main contributions are as follows.

1. *Smoothness and information aggregation*: Our main result identifies a basic *smoothness* condition on prices in a complete market that ensures information will be aggregated. Formally, we show that in any perfect Bayesian equilibrium, if the portfolios held by traders converge over time, and the limiting price charged by the market maker is *continuously differentiable* at zero with respect to the quantity traded, then information is aggregated. This is an interesting result, because, in particular, it suggests the per-unit price for a small purchase and a small sale should be essentially the same—that is, there should be no bid–ask spread near zero quantity.

We note that without such a smoothness condition, it is possible that information aggregation may fail. Consider a strategic market maker that seeks exclusively to maximize trading profits. If such a market maker is sufficiently concerned that traders possess significantly superior private information regarding the value of a traded asset, then prices may be set in a manner that completely precludes trading and thus the dissemination of information (Glosten and Milgrom 1985, Glosten 1994). In settings where market makers faced with privately informed traders impose bid–ask spreads for infinitesimal trades that are uniformly bounded away from zero, our results do not apply. This includes many financial markets, for example, where a minimum “tick size” is often required by regulators or market operators. We emphasize, however, that the smoothness condition is needed only in the limit: the market maker may set positive bid–ask spreads after any finite sequence of trades; we only require the spreads to vanish in the limit.

Furthermore, there are many settings where a market maker is not averse to trading with traders with superior private information. In a prediction market, for example, the foremost goal of a market maker is to aggregate diffuse private information. The information so aggregated may be useful for forming more accurate forecasts of future uncertainties and aid in making better decisions. For example, prediction markets have been used to forecast project completion times (see,

e.g., Cowgill et al. 2009, Othman and Sandholm 2013). Knowing that a project would not be completed in time would allow a market maker to take remedial actions well in advance. Such a market maker may be willing to incur losses so as to induce trade with privately informed traders in order to achieve this end. Our main contribution is to demonstrate that a smoothness assumption on the market ensures sufficient trade and thereby guarantees the diffusion of the information available to the individual traders.

Notably, many algorithmic market makers used in prediction markets *are* smooth in the sense we require; in particular, we give conditions under which *cost function* market makers (or, equivalently, market makers based on market scoring rules) satisfy the smoothness requirement. In addition to virtual currency markets such as Inkling that use algorithmic market makers, others have been established using real currency as well. For example, Yoonew ran a prediction market that could be used by fans of sports teams to obtain tickets for championship games at discount prices; the market essentially allowed individuals to bid on tickets contingent on their team actually making it to the championship game. This market ran on a market scoring rule known as the *logarithmic market scoring rule*.

2. *Markets with risk-averse traders*: The market model we consider involves finitely many informed *risk-averse* traders interacting with a market maker. This is distinct from prior work in this area (Chen et al. 2010, Ostrovsky 2012), which primarily considered information aggregation among traders that are risk neutral.

This modeling choice is significant for two reasons. First, in many markets, risk aversion is important. For example, traders may bid on contracts in prediction markets as insurance to hedge against risks inherent in other elements of their portfolio. Second, if all traders and the market maker are risk neutral and strategic, then in general, the *no-trade theorem* applies and precludes trade (Milgrom and Stokey 1982). Informally, the problem is that two rational risk-neutral traders cannot take opposing positions in the market and both expect to be better off. However, with risk aversion, trading can occur even if all traders are rational: traders may trade purely on the motive of hedging, if the initial allocation of risk is not *ex ante* Pareto efficient. More generally, even in settings where no-trade theorems do not apply (e.g., when market makers are not strategic), the transfer of risk between agents with different risk aversion is an important motivation for trade.

3. *Allocative efficiency*: We show that regardless of the level of risk aversion of the traders, the final allocation and prices together constitute a *competitive equilibrium*; thus, in particular, the final portfolios of the traders are *ex post* Pareto efficient. Note that as a consequence of

this result, when traders are risk averse, prices may not reflect the true posterior probabilities of events occurring. This is because competitive equilibrium prices must also reflect the marginal expected utility of the traders; thus, in general, prices are *risk-adjusted* probabilities. If even one trader in the market is risk neutral, however, then prices are accurate posterior probabilities.

We note that though we demonstrate ex post Pareto efficiency, the market mechanism we describe will not generally be ex ante Pareto efficient. That is, we cannot exclude the possibility of other mechanisms that yield a higher ex ante expected utility to each agent. Furthermore, we also note that our notion of ex post efficiency excludes the market maker. This is reasonable, as we do not model the market makers as utility maximizers, and hence we do not have a satisfying method to incorporate their welfare. However, we note that the trade in the market is not primarily driven by any inefficiency in the initial allocation of the traders. Even if the initial allocation among the traders were ex ante efficient, an algorithmic market maker may simply trade securities at a discount to initiate and facilitate trade. Intuitively, since the supply of the securities is not constrained because of the presence of the market maker, there may be trade even if the initial allocation were efficient among the traders.

The most closely related paper to our work is by Ostrovsky (2012). He shows that if the securities under consideration are “separable” in an appropriate sense, and all traders are risk neutral, then information is aggregated in markets based on the competitive dealer model of Kyle (1985), as well as in markets based on market scoring rules. The markets we consider are complete with respect to payoff-relevant uncertainty (i.e., there exists one contingent contract for each possible payoff-relevant event that can occur); and with the partition model for signal structure that Ostrovsky considers, a complete market is always separable.² Our main innovation is in studying markets with risk-averse traders and establishing the aforementioned smoothness condition as essential to information aggregation.

The remainder of the paper is organized as follows. In the next section we define our basic model, including the game played by market participants, as well as perfect Bayesian equilibrium for this game. In §3, we formally define information aggregation. In §4, we claim that if the portfolios of traders converge and the market is asymptotically smooth, then information

is aggregated. In §5, we provide insight into ex post Pareto efficiency of the market. In §6, we study a class of markets based on cost functions and show information aggregation under the assumption of portfolio convergence.

2. Model

In this section we describe the operation of the market, as well as our equilibrium notion for the resulting game, perfect Bayesian equilibrium.

2.1. Market Operation

We consider a market consisting of n traders and organized by a market maker. Trading takes place in the market sequentially at an infinite sequence of times $t \in \{1, 2, \dots\}$. In particular, at time t , the trader i_t engages in trade with the market maker, and the sequence $\{i_t\}$ is known³ a priori to all the traders and the market maker. We further assume that each trader visits the market infinitely often. Here, note that t indexes the set of trading opportunities, not calendar time. We imagine a market in which trading occurs over a short, finite calendar time horizon by traders who make decisions rapidly. However, so as to not a priori limit the number of trading opportunities available to an agent, we allow each trader to access the market an infinite number of times.

The uncertainty in the value of future securities is captured by a random variable $\tilde{\omega}$ taking values in the finite set $\Omega = \{1, \dots, m\}$; we refer to the random variable $\tilde{\omega}$ as the *payoff-relevant state* of the world, and we assume that all traders have a common prior distribution for $\tilde{\omega}$. We assume that the payoff-relevant state is only revealed after all trades are completed. Furthermore, we assume the market for securities over payoff-relevant uncertainty, i.e., possible realizations of $\tilde{\omega}$, is *complete*. That is, we assume traders can trade in any of m securities labeled $1, \dots, m$; one share of security ω pays \$1 in state ω and nothing otherwise. These are often referred to as *Arrow securities*. Note that we do not make the assumption of complete markets over *all* sources of uncertainty; in particular, traders do not trade over *payoff-irrelevant* events that do not directly affect the trader’s preferences.

Suppose that trader $i_t = i$ visits the market at time t to trade with the market maker. Let $y_t \in \mathbb{R}^m$ denote the corresponding trade, where the component $y_t(\omega)$ is the

² Although in the case of finite signal space our model of signals can be reformulated as a partition model, our results hold in more general signal spaces. Moreover, Ostrovsky assumes that upon pooling of all the signals of the traders, there is no further uncertainty in the market, whereas we allow for payoff-relevant uncertainty even after the signals of the traders are pooled.

³ We note here that our results can be extended to allow for the random arrivals of traders, as long as (i) the sequence of arrivals is independent of the rest of the randomness in the system, and (ii) the identity of each visiting trader is common knowledge subsequent to each trade. Without the latter requirement, the operation of the market would create additional private information for each trader (e.g., namely, the times at which the trader traded), and this may interfere with the aggregation of the information in the private signals. See Footnote 8 for a related discussion.

quantity of security ω bought by the trader i at time t . The history h_t at time t as observed by the traders (and the market maker) consists of all the trades until time t . In other words,

$$h_t \triangleq (y_1, y_2, \dots, y_{t-1}).$$

As a matter of convention, we let $h_1 \triangleq \emptyset$ denote the null history at time 1. We denote by H_t the set of all possible histories up to time t , and we let $H_f \triangleq \bigcup_{t \geq 1} H_t$ denote the set of all possible finite histories. Finally, let $h_\infty \triangleq (y_1, y_2, \dots)$ denote the infinite history, or the path taken by the market, and let H_∞ denote the set of all possible infinite histories.

The portfolio of a trader i at time t consists of the different quantities of each security she holds. Let $w_{i,t}(\omega)$ denote the quantity of security ω held by trader i at time t ; we refer to the vector $w_{i,t}$ as the *portfolio* of trader i at time t . We assume that the initial portfolio of each trader is common knowledge among the traders and the market maker.

Observe that if a trader holds the portfolio $\mathbf{1} \triangleq (1, \dots, 1)$, i.e., one unit of each security, then the trader receives a payoff of \$1 regardless of the realized state. For this reason we refer to the $\mathbf{1}$ portfolio as *money*, or the *risk-free* asset, and throughout the paper we interpret monetary payment of $\$x$ to or from a trader as credits or debits of x units of the $\mathbf{1}$ portfolio. Equivalently, we can imagine all trades as being simultaneously settled upon realization of the payoff-relevant state.

The market maker determines the price for trades of different quantities of the securities; this price may depend on the history. In particular, we let $K(h_t, y)$ represent the price charged for a portfolio y after history h_t ; thus the trader's net trade at time t is $y_t - K(h_t, y_t)\mathbf{1}$. We assume the functional form of K is known to all traders a priori.⁴ To reflect the fact that the trade in the market is voluntary, we assume that the pricing function K does not penalize any trader for not participating in the market. More precisely, we assume that after any history h_t , the pricing function satisfies $K(h_t, \mathbf{0}) = 0$, where $\mathbf{0} \triangleq (0, \dots, 0)$.

2.2. The Game and Equilibrium

In this section our main goal is to define our equilibrium concept, *perfect Bayesian equilibrium* (PBE). Informally, PBE requires that traders' strategies maximize their utility given their beliefs over any uncertain elements of the model, and their beliefs are consistent with the strategies adopted by other traders in equilibrium.

⁴ We note that this does not preclude a strategic market maker: our model allows for the possibility that K is simply the equilibrium pricing strategy of the market maker. Our only requirement is that the price schedule $K(h_t, \cdot)$ be known to the trader arriving at time t ; in other words, the market maker's pricing strategy should be a deterministic function of history.

In our model, uncertainty arises because the payoff-relevant state $\tilde{\omega}$ is unknown; however, we assume traders are informed and receive signals regarding the true state. In this section we define signals, beliefs, strategies, utilities, and, ultimately, the concept of PBE.

It should be noted that, to be precise, these constructs should be defined in a measure-theoretic framework. However, for clarity of exposition, we suppress measure-theoretic details in the main text; in Appendix C we provide a formal measure-theoretic description of each of the elements introduced here.

2.2.1. Signals. Before trading begins, each trader i receives a private signal $\tilde{s}_i \in \Sigma_i$. We refer to $\tilde{s} \triangleq (\tilde{s}_1, \dots, \tilde{s}_n)$ as the joint signal and $\Sigma \triangleq \Sigma_1 \times \dots \times \Sigma_n$ as the joint signal space.⁵ Let \mathbf{P} denote the joint prior distribution of $\tilde{\omega}$ and $\tilde{s}_1, \dots, \tilde{s}_n$. We assume that this joint distribution \mathbf{P} is common knowledge among the traders.

An important class of signal structures that we focus on in this paper is one where the traders' private signals satisfy *conditional independence*.⁶ Formally, this requires that the private signals are independent conditional on the payoff-relevant state $\tilde{\omega}$. (We revisit and relax this assumption in Appendix B.)

2.2.2. Beliefs. Let $\mathcal{S} \triangleq H_\infty \times \Omega \times \Sigma$. Observe that an element of \mathcal{S} captures all uncertainty in the model: the path of the market, the payoff-relevant state, and the signals of all the traders. All uncertainty can therefore be represented by probability distributions over this space. In particular, we assume that after each history h_t , and having observed signal s_t , trader i 's *belief* $\nu_i(h_t, s_t)$ is a probability distribution over \mathcal{S} . This represents trader i 's forecast of the future actions by traders, the payoff-relevant state $\tilde{\omega}$, and the signal vector \tilde{s} .

2.2.3. Strategies. Trader i is said to follow the strategy δ_i if, after a history $h \in H_f$, the trader selects a trading decision according to the distribution specified by $\delta_i(h, s_t)$, where s_t denotes the private signal received by the trader. (Of course, trader i can only trade at those times where $i_t = i$.) Let $y_t(\delta_i, h_t)$ denote the trade specified by the strategy δ_i at time t . If δ_i is a mixed strategy, then $y_t(\delta_i, h_t)$ is the realized trade when trader i chooses her trade according to δ_i . Note that the distribution of $y_t(\delta_i, h_t)$ at time t depends only

⁵ A special case of our general model is the *finite partition model* of signals. This case arises when each Σ_i is a partition of the payoff-relevant state space Ω , and the private signal \tilde{s}_i denotes the partition element in which $\tilde{\omega}$ is contained. Conversely, it can be shown that any general signal structure where the signal space Σ_i of trader i is finite can be reformulated as a finite partition model. However, such a reformulation involves the introduction of additional components in the state, altering the state space Ω and possibly rendering a complete market incomplete.

⁶ Note that, in particular, the finite partition model satisfies conditional independence.

on the realized history h_t and the private signal \tilde{s}_i . When the context is clear, we omit the explicit history dependence. Let $\delta = (\delta_1, \dots, \delta_n)$ denote the collective strategy profile, where each trader i follows the strategy δ_i . Also, for each i , let $\delta_{-i} = (\delta_j)_{(j \neq i)}$ denote the strategy of every trader except i .

Note that the strategy profile δ , together with the common prior distribution \mathbf{P} over the payoff-relevant state of the world and the signals, defines a probability measure Q^δ over the space \mathcal{S} .

2.2.4. Utility and the Value Function. We assume that if trader i is holding a portfolio w_i , and the payoff-relevant state $\tilde{\omega} = \omega$ is realized, then trader i obtains a utility $u_i(\omega, w_i(\omega))$. The utility function $u_i: \Omega \times \mathbb{R} \rightarrow \mathbb{R}$ of each trader is common knowledge among the traders and the market maker. We consider traders who are risk averse, and who, all else being equal, prefer greater wealth to less wealth. These two characteristics of the traders are captured by assuming that, for each trader i and each payoff-relevant state ω , the function $u_i(\omega, \cdot)$ is strictly increasing and concave. We further make the assumption in our analysis that each function $u_i(\omega, \cdot)$ is differentiable. Note that all these assumptions hold trivially for risk-neutral traders, whose utility functions are given by $u_i(\omega, x) = x$ for all $\omega \in \Omega$, $x \in \mathbb{R}$. Observe that it is important to allow for each trader's utility to depend on the payoff-relevant state. This flexibility allows for the modeling of agents who may trade for hedging purposes, for example.

Although we assume that traders act as expected utility maximizers, as trading proceeds over an infinite horizon, defining this objective precisely is somewhat subtle. In particular, if the history up to time t is h_t , the strategy profile is δ , and trader i 's belief is ν_i , then trader i 's *value function* at time t is defined to be

$$J_{i,t}(\delta) \triangleq \mathbf{E}_{\nu_i(h_t, s_i), \delta} \left[\liminf_{T \rightarrow \infty} u_i(\tilde{\omega}, w_{i,T}(\tilde{\omega})) \right]. \quad (1)$$

We assume that traders act at each time period to maximize their value function.⁷

The objective (1) has several notable features. First, it is defined through a limiting portfolio of the trader. To motivate this, suppose for a moment that the portfolio of the trader converges over time. Then, the quantity $\lim_{T \rightarrow \infty} u_i(\tilde{\omega}, w_{i,T}(\tilde{\omega}))$ is the final utility of the trader i at the conclusion of trading in the market. Thus the value function captures the expected utility of the eventual portfolio of the trader. Defining the value function in terms of the limit inferior is a generalization that allows us to capture a broader class of possible

⁷ Note that the value function is defined as the expectation of a limit inferior of the utility; in what follows, we ultimately will focus on equilibria where portfolios converge along the equilibrium path, and in these cases, the value function reduces to the expectation of the limiting utility.

trader behaviors that is not restricted to only those cases where the portfolios are required to converge asymptotically. This approach is common, for example, in the literature on infinite horizon dynamic programming problems. Second, there is no discounting in the objective (1). This is because all intermediate transfers in the market occur through the risk-free asset. Thus *all* payments, whether from intermediate transfers or security payoffs, occur at the termination of the game. Moreover, as discussed earlier, we imagine the market as consisting of rapid trading over a short, finite calendar time horizon. In such a setting, it is not necessary to discount either wealth or utility. Finally, note that the value function is not separable over the trades that occur at different time periods; rather, it is derived from the limiting wealth and aggregates the impact of the actions taken at all times in a nonlinear way. Hence, analysis of the value function does not immediately decompose into the separate analysis of individual trades.

2.2.5. Perfect Bayesian Equilibrium. Finally, we describe the notion of PBE for the market defined above. (For a general definition of PBE, see Fudenberg and Levine 1983, Fudenberg and Tirole 1991.) We say the strategy profile δ and the beliefs ν_i for each i constitute a PBE if the following two conditions hold:

1. For each i , for each $h_t \in H_t$ and all $s_i \in \Sigma_i$, and for any strategy δ'_i , the value function $J_{i,t}(\delta'_i, \delta_{-i})$ exists, and we have

$$\delta_i \in \arg \max_{\delta'_i} J_{i,t}(\delta'_i, \delta_{-i}).$$

2. For each i and for each s_i , the belief $\nu_i(h_t, s_i)$ after any history $h_t \in H_t$ is derived from the belief $\nu_i(h_{t-1}, s_i)$ after history h_{t-1} through Bayes' rule whenever possible.

The first condition requires that after any history, the action specified by the strategy δ_i for trader i is optimal, holding fixed the strategies of all the other traders. The second condition requires that after any history, the traders update their beliefs through Bayes' rule whenever possible. Note that, along an equilibrium path, we have $\nu_i(h_t, s_i)(\cdot) = Q^\delta(\cdot | h_t, \tilde{s}_i = s_i)$ for all $s_i \in \Sigma_i$.

Whenever we consider a PBE associated with a particular strategy profile δ , we append δ as a superscript to all the relevant quantities. Thus, for example, $w_{i,t}^\delta$ denotes the portfolio of trader i at time t in the PBE with strategy profile δ .

3. Information Aggregation

In this section, we define the notion of information aggregation. Informally, information aggregation occurs when the private signals observed by individual traders

are aggregated into a common market belief. To give a precise definition, we begin by defining some notation.

Fix a PBE with strategy profile δ and beliefs $\{\nu_i\}$. Recall that Q^δ is the joint probability measure induced on the space \mathcal{S} by the strategy profile δ and the common prior distribution \mathbf{P} over the payoff-relevant state and the signals. For much of our analysis, we require notation for the belief of an uninformed market observer who is able to observe market transactions, is aware of all common knowledge, and shares the same prior distribution as market participants. In particular, we define φ_t , the *common belief of the market* over the payoff-relevant state and private signals at time t , by

$$\varphi_t(\cdot) \triangleq Q^\delta(\cdot | h_t).$$

The notion of a common belief is an important tool in analyzing dynamic games with asymmetric information in a number of different contexts—for example, in models of rational learning (Blume and Easley 1998), herding (Smith and Sorensen 2006), and reputation (Mailath and Samuelson 2006). A similar idea was employed by Ostrovsky (2012) in his work on information aggregation.

Furthermore, trader i has beliefs that are informed by both the history of trading and the private signal observed. We define $\varphi_{i,t}$ to be the belief of trader i over the payoff-relevant state and signals at time t , and in a PBE, this is given by

$$\varphi_{i,t}(\cdot) \triangleq Q^\delta(\cdot | h_t, \tilde{s}_i = s_i).$$

We note that in our definition φ_t and $\varphi_{i,t}$ are beliefs *only* over the state and signals, i.e., distributions over $\Omega \times \Sigma$. Furthermore, note that φ_t and $\varphi_{i,t}$ have implicit history dependence. Trivially, $\varphi_{i,t}$ can be obtained from φ_t by conditioning on trader i 's signal: $\varphi_{i,t}(\cdot) = \varphi_t(\cdot | \tilde{s}_i = s_i)$.⁸

The following result is common in analysis of PBE in infinite horizon games: it establishes that beliefs converge to well-defined limits. The proof, which is omitted, involves writing the probabilities as a Doob martingale and then applying the martingale convergence theorem.

LEMMA 1. *Under the measure Q^δ , almost surely, the sequence of beliefs φ_t (respectively, $\varphi_{i,t}$) converges weakly to a probability distribution φ_∞ (respectively, $\varphi_{i,\infty}$), where $\varphi_\infty(\cdot) = Q^\delta(\cdot | h_\infty)$, and $\varphi_{i,\infty}(\cdot) = Q^\delta(\cdot | h_\infty, \tilde{s}_i = s_i)$.*

We are now ready to make the main definition.

⁸Note that this is a consequence of the fact that, at any time, all information in the market is common knowledge, except the traders' private signals. This will no longer be true, for example, if the arrivals of traders are random, without revealing the trader's identity subsequent to each trade, as discussed in Footnote 3.

DEFINITION 1 (INFORMATION AGGREGATION). The market aggregates the information of the traders in a PBE δ if, almost surely, for all $\omega \in \Omega$,

$$\varphi_\infty(\tilde{\omega} = \omega) = \mathbf{P}(\tilde{\omega} = \omega | \tilde{s}).$$

On the left-hand side, we have the posterior common belief of the market after the infinite trading history has been observed. On the right-hand side, we have the posterior distribution of the payoff-relevant state if all traders' signals could be pooled. Thus, information aggregation requires that, via the trading history, the common market belief completely pools the private signals of the traders.

Note that the preceding definition does not require that the prices of the securities reflect the posterior beliefs of the traders. Our definition only requires that an uninformed outsider sharing the common knowledge and prior distribution of the traders, and having knowledge of the sequence of trades conducted, should be able to infer the relevant information in the joint private signal \tilde{s} . Subsequently, we show a stronger result: in fact, the asymptotic portfolios and prices together constitute a *competitive equilibrium* of the limiting economy, where traders maximize expected utility given the joint signal \tilde{s} . This observation allows us to show that an uninformed outsider can infer the information in the joint signal under much milder requirements. Furthermore, we demonstrate the relationship between prices and the utility functions of the traders.

The following lemma relates information aggregation to a condition on the limiting common belief of the market. We make use of this relation in the proof of our main theorem in the next section.

LEMMA 2. *In any PBE, information aggregation holds if and only if, almost surely, the state $\tilde{\omega}$ is independent of the joint signal \tilde{s} under the limiting belief φ_∞ .*

PROOF. Observe that after any history h_t in the market, we have $\varphi_t(\tilde{\omega} = \omega | \tilde{s}) = \mathbf{P}(\tilde{\omega} = \omega | \tilde{s})$. This is because the history cannot contain any more information about $\tilde{\omega}$ than the joint signal \tilde{s} . On taking the limit $t \rightarrow \infty$, we obtain from Lemma 1 that $\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}) = \mathbf{P}(\tilde{\omega} = \omega | \tilde{s})$. Thus, we obtain from the definition that information aggregation is equivalent to requiring the state $\tilde{\omega}$ to be independent of the joint signal \tilde{s} under the limiting belief φ_∞ . \square

4. Asymptotic Smoothness

A central theme of our paper is that if the prices set by the market maker are sufficiently "smooth" with respect to small purchases or sales by traders, then information will be aggregated. Formally, we introduce the following condition.

ASSUMPTION 1 (ASYMPTOTIC SMOOTHNESS). Consider a strategy profile δ and the associated induced distribution Q^δ . We assume that there exists an open neighborhood \mathcal{N} of zero such that for almost all (under Q^δ) histories h_∞ , for all $y \in \mathcal{N}$, the limit

$$K(h_\infty, y) \triangleq \lim_{t \rightarrow \infty} K(h_t, y)$$

exists, and it is finite and continuously differentiable in y .

The preceding condition essentially requires that, asymptotically, if a trader buys or sells an infinitesimal portfolio, the marginal price is the same. This rules out the possibility of a nonzero bid–ask spread for infinitesimal trades. Note that such nonzero bid–ask spreads have been observed in prior models, e.g., in the limit-order book model with adverse selection as studied by Glosten (1994). A central observation in those models is that if the initial bid–ask spread is too wide, then trading may not take place simply because informed traders will not find it profitable to participate. Using this insight it is straightforward to construct examples of markets that are not asymptotically smooth and in which information aggregation does not occur.

Although we have stated asymptotic smoothness as a property of the pricing function K , it may alternatively be viewed as the property of the underlying equilibrium. In other words, one might consider equilibria where the behavior of the market maker arises endogenously. In such settings, if the induced pricing function K that arises from the market maker’s activity is asymptotically smooth, then our results apply.

Additionally, we make the following assumption on the equilibria under consideration.

ASSUMPTION 2 (PORTFOLIO CONVERGENCE). Under the strategy profile δ , the portfolio of each trader i converges as trading proceeds in the market. More precisely, for each trader i , the limit

$$w_{i,\infty}^\delta(\omega) \triangleq \lim_{t \rightarrow \infty} w_{i,t}^\delta(\omega)$$

exists, and it is finite for all $\omega \in \Omega$, almost surely.

Portfolio convergence precludes those instances where the market maker’s pricing function is too weak to restrict traders from taking arbitrarily large positions. The assumption is violated, for example, in instances where a trader possesses a sequence of actions that yield unbounded utility. Furthermore, the assumption also captures the idea that the traders are unwilling to accept unbounded losses in *any* state, even states on which their belief assigns zero weight. In these respects, one can view Assumption 2 as excluding pathological equilibria.

The following theorem, which is a key step to the main result of this section, shows that under asymptotic smoothness and portfolio convergence, all the information contained in any single trader’s signal is eventually revealed through the trades in the market.

THEOREM 1. In any PBE where asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2) hold, we obtain that, almost surely, the true state $\tilde{\omega}$ is independent, under the limiting common belief φ_∞ , of each trader i ’s private information \tilde{s}_i .

Before we prove the theorem, we provide some intuition for the result. In a PBE, each trader starts with private information that may not be reflected in the pricing function of the market maker. As long as the information of any trader is not fully incorporated into the pricing function, asymptotic smoothness (Assumption 1) suggests that, eventually, the trader should be able to profit from trading infinitesimal quantities of the appropriate securities. Portfolio convergence in the PBE (Assumption 2), on the other hand, suggests that, eventually, the trader ceases to exploit any such opportunities. We therefore expect that all the information of an individual trader will ultimately be incorporated into the prices and thus into the limiting common belief. This implies that any further uncertainty in the private information \tilde{s}_i of trader i is independent of the payoff-relevant state $\tilde{\omega}$. Observe that our result holds even in the case where the utility of traders does not depend on the state. This is because we assume utility that is strictly increasing in wealth. Hence, traders with private information who are indifferent to the ultimate realization of the state still have incentive to trade in a manner that reveals their private information.

The proof of Theorem 1 uses the fact that in a PBE, a unilateral deviation by any trader should result in lower utility for that trader. Using this fact, we show that in the limit, the trader i ’s belief over $\tilde{\omega}$, given the infinite history h_∞ , is independent of her private signal \tilde{s}_i .

PROOF OF THEOREM 1. Given portfolio convergence and the continuity of the utility function, the value function of a trader i after the history h_t can be written as

$$J_{i,t}(\delta) = \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)}[u_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))].$$

Now, consider a unilateral deviation for trader i after any history h_t , where the trader, instead of following the strategy δ_i , decides to trade a fixed quantity $z \in \mathbb{R}^m$ and never trade in the market thereafter. Let $w_{i,t}^z$ denote the payoff vector of trader i after such a trade; i.e.,

$$w_{i,t}^z \triangleq w_{i,t-1}^\delta + z - K(h_t, z)\mathbf{1}.$$

As δ is a PBE strategy profile, any deviation after history h_t cannot lead to a higher utility for the trader. Thus, almost surely,

$$\mathbf{E}_{\nu_i(h_t, \tilde{s}_i)}[u_i(\tilde{\omega}, w_{i,t}^z(\tilde{\omega}))] \leq \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)}[u_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))] \quad (2)$$

for all $z \in \mathbb{R}^m$ and for all t . Next, observe that

$$\begin{aligned} & \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)}[u_i(\tilde{\omega}, w_{i,t}^z(\tilde{\omega}))] \\ &= \sum_{\omega \in \Omega} \varphi_{i,t}(\tilde{\omega} = \omega) u_i(\omega, w_{i,t}^\delta(\omega) + z(\omega) - K(h_t, z)). \end{aligned}$$

As the pricing function satisfies asymptotic smoothness (Assumption 1), we have $K(h_t, z) \rightarrow K(h_\infty, z)$ almost surely for all $z \in \mathcal{N}$. Moreover, from portfolio convergence (Assumption 2), we have $w_{i,t}^\delta(\omega) \rightarrow w_{i,\infty}^\delta(\omega)$ almost surely for each $\omega \in \Omega$. Furthermore, we know from Lemma 1 that for each $\omega \in \Omega$, the trader’s belief $\varphi_{i,t}(\tilde{\omega} = \omega)$ converges to $\varphi_{i,\infty}(\tilde{\omega} = \omega)$ almost surely. Using these facts and the continuity of the utility function, we obtain almost surely for all $z \in \mathcal{N}$ that

$$\begin{aligned} & \lim_{t \rightarrow \infty} \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)} [u_i(\tilde{\omega}, w_{i,t}^z(\tilde{\omega}))] \\ &= \lim_{t \rightarrow \infty} \sum_{\omega \in \Omega} \varphi_{i,t}(\tilde{\omega} = \omega) u_i(\omega, w_{i,t}^\delta(\omega) + z(\omega) - K(h_t, z)) \\ &= \sum_{\omega \in \Omega} \varphi_{i,\infty}(\tilde{\omega} = \omega) u_i(\omega, w_{i,\infty}^\delta(\omega) + z(\omega) - K(h_\infty, z)) \\ &= \mathbf{E}_{\varphi_{i,\infty}} [u_i(\tilde{\omega}, w_{i,\infty}^z(\tilde{\omega}))], \end{aligned} \tag{3}$$

where we have defined $w_{i,\infty}^z$ as

$$w_{i,\infty}^z(\omega) \triangleq w_{i,\infty}^\delta(\omega) + z(\omega) - K(h_\infty, z) \quad \text{for all } \omega \in \Omega.$$

On the other hand, almost surely,

$$\begin{aligned} & \lim_{t \rightarrow \infty} \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)} [u_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))] \\ &= \lim_{t \rightarrow \infty} \mathbf{E}_{Q^\delta} [u_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega})) | h_t, \tilde{s}_i] \\ &= \mathbf{E}_{Q^\delta} [u_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega})) | h_\infty, \tilde{s}_i] \\ &= \mathbf{E}_{\varphi_{i,\infty}} [u_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))]. \end{aligned} \tag{4}$$

Here, the first equality follows since ν_i is the belief of trader i in a PBE. The second equality follows by a backward martingale convergence argument. The third equality follows by the definition of $\varphi_{i,\infty}$.

Thus, from Equations (2)–(4), we obtain almost surely for all $z \in \mathcal{N}$ that

$$\mathbf{E}_{\varphi_{i,\infty}} [u_i(\tilde{\omega}, w_{i,\infty}^z(\tilde{\omega}))] \leq \mathbf{E}_{\varphi_{i,\infty}} [u_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))].$$

Since $w_{i,\infty}^0 = w_{i,\infty}^\delta$, we conclude that, almost surely,

$$0 \in \arg \max_{z \in \mathcal{N}} \mathbf{E}_{\varphi_{i,\infty}} [u_i(\tilde{\omega}, w_{i,\infty}^z(\tilde{\omega}))]. \tag{5}$$

By asymptotic smoothness, the pricing function $K(h_\infty, y)$ is differentiable at $y = \mathbf{0}$. Moreover, the traders have differentiable utility functions. This implies that the following first-order necessary condition holds:

$$\begin{aligned} & \varphi_{i,\infty}(\tilde{\omega} = \omega) u'_i(\omega, w_{i,\infty}^\delta(\omega)) \\ &= \nabla_{y(\omega)} K(h_\infty, \mathbf{0}) \mathbf{E}_{\varphi_{i,\infty}} [u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))] \\ & \quad \text{for all } \omega \in \Omega. \end{aligned} \tag{6}$$

Note that (6) holds almost surely over possible values of the infinite history h_∞ and the private signal \tilde{s}_i . Furthermore, observe that we are able to differentiate through the expectation since the expectation is only over the finitely many possible realizations of $\tilde{\omega}$.

Recall that $\varphi_{i,\infty}(\tilde{\omega} = \omega) = \varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)$. Using this fact and rearranging the above equation, we have almost surely that

$$\frac{\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)}{\mathbf{E}_{\varphi_\infty} [u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega})) | \tilde{s}_i]} = \frac{\nabla_{y(\omega)} K(h_\infty, \mathbf{0})}{u'_i(\omega, w_{i,\infty}^\delta(\omega))} \quad \text{for all } \omega \in \Omega. \tag{7}$$

The preceding equation is true almost surely over the possibilities for the infinite history h_∞ and the private signal \tilde{s}_i . However, observe that once h_∞ is known, the right-hand side of (7) is fixed and does not further depend on \tilde{s}_i . In particular, this implies that the quantity on the left-hand side is almost surely a constant for all realized values $\tilde{s}_i = s_i$ that are possible given a fixed infinite history h_∞ , i.e., for all values of \tilde{s}_i in the support of φ_∞ . On summing the left-hand side over all values of $\omega \in \Omega$, and noting that the numerator sums to 1, we obtain that the denominator $\mathbf{E}_{\varphi_\infty} [u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega})) | \tilde{s}_i = s_i]$ also does not depend on s_i in the support of φ_∞ . This, in turn, implies that for all $\omega \in \Omega$, the numerator $\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)$ itself does not depend on the realization of \tilde{s}_i in the support of φ_∞ ; i.e., $\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i) = \varphi_\infty(\tilde{\omega} = \omega)$ almost surely. From this argument we conclude that under the limiting common belief φ_∞ , the payoff-relevant state $\tilde{\omega}$ is independent of the private signal \tilde{s}_i of each trader i . \square

We have shown that in a PBE, asymptotic smoothness and portfolio convergence together suffice to ensure that the payoff-relevant state $\tilde{\omega}$ and each trader i ’s private information \tilde{s}_i are pairwise independent. However, note that the preceding theorem does not imply information aggregation; recall from Lemma 2 that for information aggregation to hold, we require the payoff-relevant state $\tilde{\omega}$ to be independent of the joint signal \tilde{s} . In particular, any information that is encoded collectively in the signals of multiple traders may still remain unrevealed through the trades in the market. Nevertheless, if we obtain a condition on the signal structure that ensures independence of the joint signal \tilde{s} and the payoff-relevant state $\tilde{\omega}$ from pairwise independence for each trader i , then using Lemma 2, we obtain information aggregation. The following theorem, which is the main result of this section, states that one such condition on the signal structure is conditional independence. In Appendix B, we show that a broader class of signal structures—namely, those satisfying *hierarchical conditional independence*—satisfy the requirement.

THEOREM 2. *Suppose the signal structure satisfies conditional independence. Then, in any PBE where asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2) hold, the market aggregates the information of the traders.*

PROOF. See Appendix A. \square

We thus obtain that in any PBE of the dynamic market where the signal structure is conditionally independent, if asymptotic smoothness and portfolio convergence hold, then the market aggregates the private information of the traders. The crucial step in our proof is relation (7). Using this equation, we conclude that, almost surely, conditional on the history, each trader's signal is independent of the payoff-relevant state. Note that we make use of the completeness of the market over payoff-relevant states in writing this equation. Without market completeness, we cannot infer that Equation (7) holds for each $\omega \in \Omega$, leading to a breakdown of our proof technique. Intuitively, in the absence of market completeness, traders may not be able to trade so as to fully exploit their private information, and hence their private information may not be aggregated.

Before concluding this section, we briefly discuss how our result for the dynamic market is related to information aggregation in static exchange economies. The standard tool to study information aggregation in such economies is a rational expectations equilibrium (REE) (Radner 1979), and with complete markets, information aggregation corresponds to the existence of a *fully revealing* REE. Under general signal structure, DeMarzo and Skiadas (1998, 1999) show the existence of a fully revealing REE in a restricted class of static exchange economies that satisfy an additional condition that they term as *quasicompleteness*. Furthermore, they introduce a *separably oriented* condition⁹ under which the fully revealing REE is the unique REE. Thus, in a quasi-complete exchange economy satisfying the separably oriented condition, information aggregation holds.

To see how our work is related to the results of DeMarzo and Skiadas (1998, 1999), we introduce a limiting exchange economy that operates after the dynamic market has closed. Formally, let δ be a PBE of the dynamic market where asymptotic smoothness and portfolio convergence hold. Consider an exchange economy where the traders' initial endowments are their limiting portfolios $\{w_{i,\infty}^\delta: i \leq n\}$, and their (common) prior belief is given by the limiting common market belief φ_∞ . Each trader i , on receiving the signal \tilde{s}_i , updates her belief to the limiting belief $\varphi_{i,\infty}$. Without loss of generality, we assume that for all $\omega \in \Omega$, we have $\varphi_\infty(\tilde{\omega} = \omega) > 0$ (otherwise, we just redefine Ω accordingly).

Using (6), it is straightforward to show that for such an economy, the no-trade allocation and the price vector $\nabla K(h_\infty, \mathbf{0})$ together constitute an REE. Furthermore, one can also show that the limiting exchange economy satisfies the quasicompleteness

condition (see Theorem 7 in Appendix D). From the results of DeMarzo and Skiadas (1998), we obtain that if, almost surely, the limiting exchange economy is separably oriented, then the aforementioned REE is unique and *fully revealing*. Since the market is complete, this implies that information aggregation holds in the PBE. To summarize, the results of DeMarzo and Skiadas imply that in our setting, for a PBE δ satisfying asymptotic smoothness and portfolio convergence, if, almost surely, the limiting exchange economy is separably oriented, then information aggregation holds.

It turns out that the converse statement is also true: for a PBE δ satisfying asymptotic smoothness and portfolio convergence, if information aggregation holds then, almost surely, the limiting exchange economy is separably oriented (see Theorem 8 in Appendix D). Thus, from the preceding discussion, we obtain that in complete markets, information aggregation is equivalent to the condition that (almost surely) the limiting economy is separably oriented. As a consequence, Theorem 2 can be understood as identifying conditions on the dynamic market—namely, conditional independence, asymptotic smoothness, and portfolio convergence—that ensure, almost surely, the existence of a limiting exchange economy that is quasicomplete, separably oriented, and for which no-trade is the unique, fully revealing REE.

5. Allocative Efficiency and Competitive Equilibrium

One of the primary reasons for organizing an information market is to ensure efficient allocation of risk among the participating traders. Moreover, efficiency is important to ensure that traders do not collude among themselves to leave the market and trade through other mechanisms. In this section we ask the following: in what sense is an asymptotically smooth market efficient? We focus on the notion of ex post Pareto efficiency, which requires that the set of limiting portfolios of the traders be Pareto efficient if the information of all the traders is revealed.

DEFINITION 2 (EX POST PARETO EFFICIENCY). An allocation of securities among the traders is ex post Pareto efficient if, holding the total quantity of securities fixed, no trader's expected utility under the pooled information posterior $\mathbf{P}(\cdot | \tilde{s})$ can be increased without decreasing some other trader's expected utility.

Ex post Pareto efficiency is tied closely to the concept of *competitive equilibrium*. Formally, consider an exchange economy among the traders where each trader maximizes expected utility with respect to $\mathbf{P}(\cdot | \tilde{s})$, and trader i has an initial portfolio given by w_i . A *competitive equilibrium* for this economy is specified by an allocation

⁹ Ostrovsky (2012) defines and studies a closely related notion of separable securities.

w_i^* for each trader i and a price $p(\omega)$ for each security ω such that the following two conditions hold:

1. *Market clearance*: The net portfolio of all the traders in the allocation $\{w_i^*: 1 \leq i \leq n\}$ should be the same as that in the initial allocation; that is,

$$\sum_{i=1}^n w_i^*(\omega) = \sum_{i=1}^n w_i(\omega) \quad \text{for all } \omega \in \Omega. \quad (8)$$

2. *Utility maximization*: For each trader i , the allocation w_i^* of the security to trader i should maximize her expected utility with respect to the pooled information posterior $\mathbf{P}(\cdot | \tilde{s})$, given the price vector p . In other words, for each trader i , the allocation w_i^* is a solution to the optimization problem

$$\begin{aligned} & \text{maximize } \mathbf{E}_p[u_i(\tilde{\omega}, w_i^*(\tilde{\omega})) | \tilde{s}] \\ & \text{subject to } \sum_{\omega \in \Omega} p(\omega) w_i^*(\omega) = \sum_{\omega \in \Omega} p(\omega) w_i(\omega). \end{aligned} \quad (9)$$

The first fundamental theorem of welfare economics states that the allocation in any competitive equilibrium is ex post Pareto optimal (Mas-Colell et al. 1995). Thus, it suffices for our purposes to show that the limiting allocation of securities among the traders, together with the limiting prices in the market, constitutes a competitive equilibrium under pooled information.

Observe that if information is aggregated, then every trader effectively chooses to stop trading given the joint signal \tilde{s} ; in particular, no trader can profitably improve her expected utility at the limiting portfolio. Informally, we might therefore expect that the limiting allocation, together with the limiting prices, must be a competitive equilibrium of the economy where traders maximize their expected utility conditional on \tilde{s} . The following theorem establishes formally.

THEOREM 3. *Suppose the signal structure satisfies (hierarchical) conditional independence. Consider a PBE with strategy profile δ satisfying asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2). Then, the collection of traders' limiting portfolios $\{w_{i,\infty}^\delta: 1 \leq i \leq n\}$ and the price vector $\phi \triangleq \nabla_y K(h_\infty, \mathbf{0})$ together constitute a competitive equilibrium of an exchange economy with pooled information, where each trader i 's initial portfolio is specified by $w_{i,\infty}^\delta$.*

PROOF. Observe that the market clearance condition (8) holds trivially in this case, since the final and initial allocation are the same. So for the allocation $\{w_{i,\infty}^\delta: 1 \leq i \leq n\}$ and the price vector ϕ to constitute a competitive equilibrium, we require that the utility maximization condition holds for all traders. As the maximization problem (9) is a convex program, sufficient conditions for an allocation $\{w_i: 1 \leq i \leq n\}$ to be optimal can be written as

$$\mathbf{P}(\tilde{\omega} = \omega | \tilde{s}) u'_i(\omega, w_i(\omega)) = \lambda p(\omega) \quad \text{for all } \omega \in \Omega,$$

where $\lambda \in \mathbb{R}$ is a Lagrange multiplier.

Now we know from Theorem 2 that the market aggregates the information of the traders. Thus, we obtain $\mathbf{P}(\tilde{\omega} = \omega | \tilde{s}) = \varphi_\infty(\tilde{\omega} = \omega) = \varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)$. Then, the sufficiency conditions become

$$\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i) u'_i(\omega, w_i(\omega)) = \lambda p(\omega).$$

It is now easily seen from (7) that the sufficiency conditions are satisfied by the allocation $w_i \triangleq w_{i,\infty}^\delta$ for the price vector $p = \phi \triangleq \nabla_y K(h_\infty, \mathbf{0})$ by setting the Lagrange multiplier $\lambda \triangleq \mathbf{E}_{\varphi_\infty}[u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))]$. This implies that the allocation $w_{i,\infty}^\delta$ is utility maximizing for trader i with the price vector ϕ . \square

The preceding argument proves that the limiting portfolio of each trader is ex post Pareto optimal, in the sense that if the traders were subsequently allowed to trade among themselves and circumvent the market maker, they would not have any incentive to do so. In this respect, asymptotically smooth markets exhibit allocative efficiency.

We conclude with a brief discussion regarding the information content of prices. Suppose the signal structure satisfies (hierarchical) conditional independence. Let $\phi \triangleq \nabla_y K(h_\infty, \mathbf{0})$ denote the eventual price vector in a PBE with strategy profile δ satisfying asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2). Because the limiting allocation is a competitive equilibrium, it is clear that prices must reflect *risk-adjusted* probabilities; this is captured by (7), which reveals that ϕ is the posterior state distribution, shaded by a factor proportional to the marginal utility in each state. In particular, if at least one trader is risk neutral, then ϕ will be equal to the posterior state distribution.

Although prices may not exactly reflect the posterior state distribution in general, by the definition of information aggregation, an observer of the market having the same knowledge as market participants can infer the information contained in the traders' private signals. However, this is too stringent a requirement. Using the fact that the limiting economy is in competitive equilibrium, we identify in the following theorem a milder requirement under which an observer can infer the posterior state distribution. The proof is straightforward and follows from a standard result for static economies in competitive equilibrium (see, e.g., Cochrane 2005).

THEOREM 4. *Suppose the signal structure satisfies (hierarchical) conditional independence. Consider a PBE with strategy profile δ satisfying asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2). An observer with access to the limiting price vector $\phi \triangleq \nabla_y K(h_\infty, \mathbf{0})$ and the limiting portfolio of a single trader i , $w_{i,\infty}^\delta$, along with her utility function u_i , can infer the posterior distribution φ_∞ of $\tilde{\omega}$.*

PROOF. Note that, as was argued at the end of the proof of Theorem 2, under these hypotheses, \tilde{s}_i is independent of $\tilde{\omega}$ under the measure φ_∞ . The first-order condition (7) in the proof of Theorem 2 can be rewritten as, for all $\omega \in \Omega$,

$$\frac{\varphi_\infty(\tilde{\omega} = \omega)}{\mathbf{E}_{\varphi_\infty}[u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))]} = \frac{\nabla_{y(\omega)}K(h_\infty, \mathbf{0})}{u'_i(\omega, w_{i,\infty}^\delta(\omega))} \triangleq R_i(\omega),$$

almost surely. We then have, for all $\omega \in \Omega$, almost surely,

$$\varphi_\infty(\tilde{\omega} = \omega) = R_i(\omega)\mathbf{E}_{\varphi_\infty}[u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))].$$

As $\varphi_\infty(\tilde{\omega} = \omega)$ is a probability distribution, we obtain for all $\omega \in \Omega$

$$\varphi_\infty(\tilde{\omega} = \omega) = \frac{R_i(\omega)}{\sum_{\omega' \in \Omega} R_i(\omega')}.$$

Thus, an uninformed observer with access to $\nabla_{y(\omega)}K(h_\infty, \mathbf{0})$ and $\{u'_i(\omega, w_{i,\infty}^\delta(\omega)) : \omega \in \Omega\}$ can infer the common belief φ_∞ . \square

6. Cost Function Market Maker

In this section we focus our attention on a class of algorithmic market makers defined by *cost functions*. In such a market, the price seen by a trader is set by a fixed rule that depends only on the total outstanding number of shares sold in the market. Cost function market makers encompass a wide class of markets; of particular significance is the fact that market scoring rules can be reformulated as cost functions (Chen and Pennock 2007, Agrawal et al. 2011). The ease of organizing a market based on a cost function has led to its use in many real settings, especially in combinatorial prediction markets (Pennock and Sami 2007); see also Jian and Sami (2012) for a recent experimental analysis of such markets. Our main observation is that for such markets, the asymptotic smoothness condition developed in §4 is implied by the portfolio convergence assumption (Assumption 2). Thus, in markets based on cost functions, information aggregation holds in any PBE where the portfolios of the traders converge.

Let $q_t(\omega)$ denote the total quantity of the security ω sold by the market maker until time t . We have the following relation between q_t and the trades $\{y_\tau : 1 \leq \tau \leq t\}$ that have occurred up to time t :

$$q_t = \sum_{\tau=1}^t y_\tau.$$

A cost function market maker is defined by a continuously differentiable function $C: \mathbb{R}^m \rightarrow \mathbb{R}$, satisfying the following condition:

$$C(q + a\mathbf{1}) = C(q) + a \quad \text{for all } a \in \mathbb{R}. \quad (10)$$

After any history h_t , the market maker prices the trade $y \in \mathbb{R}^m$ at time t according to

$$K(h_t, y) \triangleq C(q_{t-1} + y) - C(q_{t-1}).$$

Market makers based on cost functions are a special case of the general class of market makers studied in this paper. These market makers are distinguished by the fact that their pricing function depends on the history only through the total quantity of securities sold up to current time. Furthermore, observe that the total revenue obtained by the market maker at time t is given by

$$\sum_{\tau=1}^t K(h_\tau, y_\tau) = \sum_{\tau=1}^t C(q_\tau) - C(q_{\tau-1}) = C(q_t) - C(\mathbf{0}).$$

Thus, at any time, the total revenue of the market maker is also dependent only on the total quantity of securities sold, and not on the actual sequence of trades leading to the final position. Also, note that the condition imposed by (10) on the cost function ensures that the traders cannot profit via an exchange of money with the market maker. We note in passing that a major attraction of cost function market makers is that for such market makers, exogenous guarantees can typically be provided that ensure the loss to the market maker is *bounded* independent of the sequence of trades in the market; see Pennock and Sami (2007) for further details and specific examples.

The following theorem shows that for cost function market makers in markets with (hierarchically) conditionally independent signal structure, under the portfolio convergence assumption, the information of the traders is aggregated. Thus in such markets, we do not explicitly require the asymptotic smoothness assumption on the pricing function to obtain information aggregation.

THEOREM 5. *In any PBE of a cost function based market with a (hierarchically) conditionally independent signal structure, if portfolio convergence (Assumption 2) holds, then the market aggregates the information of the traders.*

PROOF. The proof involves showing that any PBE of a cost function-based market satisfying portfolio convergence also satisfies asymptotic smoothness. To that end, we first define the function $\Gamma: \mathbb{R}^m \rightarrow \mathbb{R}^m$ by

$$\Gamma(q) \triangleq q - \left(\frac{1}{m} \sum_{\omega \in \Omega} q(\omega)\right)\mathbf{1}.$$

Note that $\Gamma(q)$ is the projection of q onto the subspace of vectors in \mathbb{R}^m with components that sum to zero. For a PBE with strategy profile δ , we have, for all $y \in \mathbb{R}^m$,

$$\begin{aligned} K(h_t, y) &= C(q_{t-1}^\delta + y) - C(q_{t-1}^\delta) \\ &= C(q_{t-1}^\delta + y) - \left(\frac{1}{m} \sum_{\omega \in \Omega} q_{t-1}^\delta(\omega)\right) \end{aligned}$$

$$\begin{aligned} & - \left(C(q_{t-1}^\delta) - \left(\frac{1}{m} \sum_{\omega \in \Omega} q_{t-1}^\delta(\omega) \right) \right) \\ & = C(\Gamma(q_{t-1}^\delta) + y) - C(\Gamma(q_{t-1}^\delta)). \end{aligned}$$

Here, in the final equality, we have used the fact that $C(q + a\mathbf{1}) = C(q) + a$ for all $q \in \mathbb{R}^m$.

Next, as the total quantity of the securities sold by the market maker at time t equals the sum of the net change in the portfolio of all the traders and the total revenue obtained by the market maker, we have

$$q_t^\delta = \sum_{i=1}^n (w_{i,t}^\delta - w_{i,0}^\delta) + \left(\sum_{\tau=1}^t K(h_\tau, y_\tau^\delta) \right) \mathbf{1}. \quad (11)$$

This implies that

$$\begin{aligned} & \Gamma(q_t^\delta) \\ & = \sum_{i=1}^n \left(w_{i,t}^\delta - w_{i,0}^\delta - \left(\frac{1}{m} \sum_{\omega \in \Omega} (w_{i,t}^\delta(\omega) - w_{i,0}^\delta(\omega)) \right) \right) \mathbf{1}. \quad (12) \end{aligned}$$

As the PBE under consideration satisfies portfolio convergence, we infer from the above equation that there exists $q_\infty^\delta \in \mathbb{R}^m$ such that $\Gamma(q_t^\delta) \rightarrow q_\infty^\delta$ almost surely. Since $K(h_t, z) = C(\Gamma(q_{t-1}^\delta) + z) - C(\Gamma(q_{t-1}^\delta))$, and since the cost function is continuous and differentiable, it follows that, almost surely, $K(h_t, z) \rightarrow K(h_\infty, z) \triangleq C(q_\infty^\delta + z) - C(q_\infty^\delta)$ for all $z \in \mathbb{R}^m$, where the function $K(h_\infty, \cdot)$ is continuously differentiable. This proves that asymptotic smoothness holds in this case, implying through Theorem 2 that information aggregation obtains. \square

Thus, we obtain that for any algorithmic market maker using a cost function, if the signal structure is (hierarchically) conditionally independent, then portfolio convergence (Assumption 2) suffices to ensure that the private information of the traders gets aggregated.

7. Conclusion

In this paper we have studied the information aggregation properties of market equilibria that satisfy two conditions: *asymptotic smoothness* and *portfolio convergence*. We study a fairly general model consisting of finitely many rational, risk-averse traders. In our main results, we demonstrate that these two conditions lead to aggregation of information in the market as well as ex post Pareto efficiency of the allocation.

As noted in §1, information aggregation has long been a central object of study in the literature on markets in general and financial markets in particular. However, a significant additional motivation for our analysis comes from the recent rise of prediction markets. In particular, prediction markets based on algorithmic market makers will often possess the asymptotic smoothness property essentially for free (provided the algorithmic pricing rule is appropriately smooth).

Our results support the premise that such market makers should lead to information aggregation (provided that trading asymptotically ceases). In addition, a key observation of our work is that such markets can also serve as a mechanism to efficiently allocate risk in settings where traders may be risk averse.

Acknowledgments

The authors are grateful to the anonymous referees, the associate editor, and the department editor for helpful comments on the paper. The authors also thank Michael Ostrovsky and David Pennock for valuable discussions. This research has been supported in part by the National Science Foundation [Grants CMMI-123502, CMMI-0948434, CNS-0904609, CCF-0832820, and CNS-0644114] and by a Stanford Graduate Fellowship.

Appendix A. Proof of Theorem 2

To prove Theorem 2, we need the following two lemmas. The first lemma makes the intuitive observation that under Q^δ , even after observing past history, the signals are independent conditional on the payoff-relevant state. (Note that this result holds even after infinite history; i.e., $t = \infty$.)

LEMMA 3. *Under Q^δ , conditional on the history h_t and payoff-relevant state $\tilde{\omega}$, the private signals $\tilde{s}_1, \dots, \tilde{s}_n$ are independent for any $1 \leq t \leq \infty$.*

PROOF. The proof follows directly from Lemma 5 and the observation that conditionally independent signals are hierarchically conditionally independent with respect to the graph $G = (V, E)$, where $V = \{0, \dots, n\}$ and $E = \{(0, i): 1 \leq i \leq n\}$. \square

The second lemma shows that with conditionally independent signal structure, pairwise independence implies joint independence.

LEMMA 4. *Let X and $\{Y_i; i = 1, \dots, n\}$ be random variables such that (1) conditional on X , the random variables $\{Y_i\}$ are independent, and (2) X is independent of Y_i for each i . Then X and (Y_1, \dots, Y_n) are independent.*

PROOF. Let g and f_1, \dots, f_n be bounded functions. As Y_1, \dots, Y_n are independent conditional on X , we have

$$\mathbf{E} \left[\prod_i f_i(Y_i) \mid X \right] = \prod_i \mathbf{E}[f_i(Y_i) \mid X].$$

Moreover, we know that X and Y_i are independent for all i . This implies that $\mathbf{E}[f_i(Y_i) \mid X] = \mathbf{E}[f_i(Y_i)]$. Thus we obtain

$$\mathbf{E} \left[\prod_i f_i(Y_i) \mid X \right] = \prod_i \mathbf{E}[f_i(Y_i)].$$

We then reason as follows:

$$\begin{aligned} \mathbf{E} \left[g(X) \prod_i f_i(Y_i) \right] &= \mathbf{E} \left[g(X) \mathbf{E} \left[\prod_i f_i(Y_i) \mid X \right] \right] \\ &= \mathbf{E}[g(X)] \prod_i \mathbf{E}[f_i(Y_i)]. \end{aligned}$$

As this is true for all bounded functions g, f_1, \dots, f_n , we infer that X and (Y_1, \dots, Y_n) are independent. \square

We are now ready to prove Theorem 2.

PROOF OF THEOREM 2. By Lemma 3, under measure φ_∞ , the private signals $\tilde{s}_1, \dots, \tilde{s}_n$ are independent given the payoff-relevant state $\tilde{\omega}$. From Theorem 1, we know that for each trader i , the payoff-relevant state $\tilde{\omega}$ is independent of the private signal \tilde{s}_i , under the limiting common belief. Together with Lemma 4, this implies that under measure φ_∞ , the payoff-relevant state $\tilde{\omega}$ and the joint signal of the traders $\tilde{s} \triangleq (\tilde{s}_1, \dots, \tilde{s}_n)$ are independent. The result then follows from Lemma 2. \square

Appendix B. Hierarchically Conditionally Independent Signals

In the main paper, we considered a model where signals are *conditionally independent*; that is, conditional on the true payoff-relevant state $\tilde{\omega}$, the signals of the traders are independent. Conditionally independent signals are plausible if each trader independently conducts her own research about the payoff-relevant state $\tilde{\omega}$. To achieve conditional independence, we must assume that this process of obtaining information is inherently noisy, with independent noise across traders. Though traders may differ in the accuracy of their signals, a key assumption is that each trader conducts her research in complete isolation: signal information is not shared across traders, except through the market.

As a consequence, this model of obtaining information ignores the possibility that some traders might obtain their information from others in the market. For example, consider a scenario where some subset of traders first carries out research to estimate the state $\tilde{\omega}$. Other traders then contact this “primary” set; for these traders, their signals are inherently correlated with the signals of those primary traders they contacted. For example, an individual investor may obtain information secondhand from the research department of their brokerage. Our signal structure does not allow these correlation relationships, and thus our proofs of information aggregation no longer apply directly.

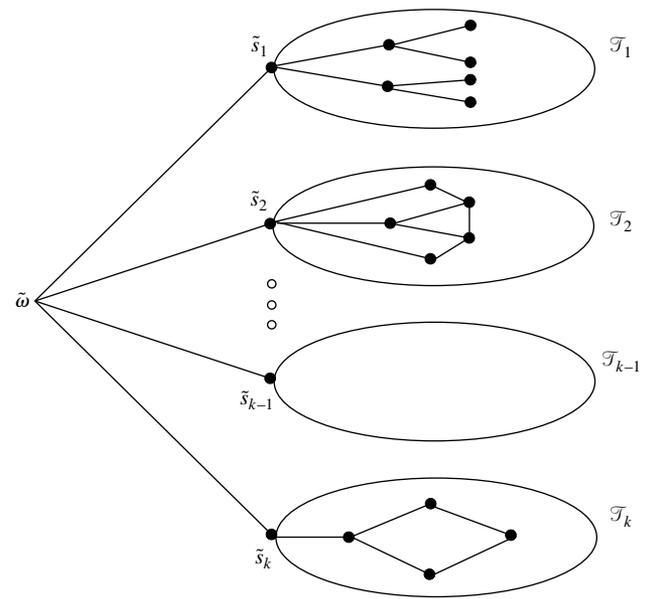
We refer to these more general correlation relationships as *hierarchical conditional independence*. The central probabilistic construct we require to capture this signal model is a *Markov random field*. A Markov random field captures correlation structure between a collection of random variables through an underlying graph.

Formally, a Markov random field is a probability measure associated with an undirected graph $G = (V, E)$, with vertex set V and edge set E . For any vertex $v \in V$, define $\partial v \triangleq \{u \in V : uv \in E\}$ to be the set of adjacent vertices, and let $\bar{\partial} v \triangleq \partial v \cup \{v\}$. Finally, let $G \setminus S$ denote the graph that remains if the set of nodes S is removed from V , and any edges with endpoints in S are removed from E .

DEFINITION 3. A collection of random variables $\{X_u : u \in V\}$ is a Markov random field with respect to a graph $G = (V, E)$ if for all $S, A, B \subset V$, where A is disconnected from B in $G \setminus S$, $\{X_v : v \in A\}$ is conditionally independent of $\{X_v : v \in B\}$ given $\{X_v : v \in S\}$.

We now give the definition of hierarchically conditionally independent signals. For ease of notation, define $\tilde{s}_0 \triangleq \tilde{\omega}$. We have the following definition.

Figure B.1 Hierarchically Conditionally Independent Signals



DEFINITION 4. The traders’ signals are said to be *hierarchically conditionally independent* if the set of random variables $\{\tilde{s}_0, \dots, \tilde{s}_n\}$ forms a Markov random field with respect to a graph G on the set $\{0, \dots, n\}$, with no simple cycles containing the node 0.

The key condition in the preceding definition is that the graph G contains no simple cycles with the node 0; see Figure B.1. Since \tilde{s}_0 is the payoff-relevant state $\tilde{\omega}$, this assumption implies that we can partition traders into subsets as follows. Let $\mathcal{T}_0 = \{1, \dots, k\}$ be the traders directly connected to node 0 in G ; informally, these are the traders who obtain information about the state by directly doing independent research. For each i , let \mathcal{T}_i be the set of traders who obtain their information (directly or indirectly) from trader i (including trader i herself); assume every trader lies in one such set. Hierarchical conditional independence then requires that *the sets $\mathcal{T}_1, \dots, \mathcal{T}_k$ are mutually disjoint*. In other words, each trader in \mathcal{T}_0 is a primary source of information for some subset of other traders, but no trader can have two primary sources.

Hierarchically conditionally independent signals are a generalization of conditionally independent signals. This can be seen at once by noticing that, for conditionally independent signals, the set $\{\tilde{s}_0, \dots, \tilde{s}_n\}$ forms a Markov random field with respect to a tree on $\{0, \dots, n\}$, rooted at 0, with an edge $(0, i)$ for all $i \in \{1, \dots, n\}$. At the same time, hierarchically conditionally independent signals are not fully general. For example, this model precludes a trader obtaining information from two different primary sources, e.g., an online investing site *and* her brokerage.

We now show that our main result continues to hold if the traders’ signals are hierarchically conditionally independent. We have the following extension of Theorem 2, proving information aggregation with hierarchically conditionally independent signals in a complete market.

THEOREM 6. *In any PBE of a complete market with hierarchically conditionally independent signals, if asymptotic smoothness*

(Assumption 1) and portfolio convergence (Assumption 2) hold, then the market aggregates the information of the traders.

We prove the theorem using two lemmas, analogous to the proof of Theorem 2. The first lemma shows that the information revealed by the trades in the market preserves the hierarchical conditional independent structure of the signals.

LEMMA 5. *Suppose that the traders' signals are hierarchically conditionally independent with respect to a graph G before the start of trading. Then in any PBE with strategy profile δ , under Q^δ , conditional on the history h_t , the traders' signals $\tilde{s}_1, \dots, \tilde{s}_n$ are hierarchically conditionally independent with respect to the same graph G for any $1 \leq t \leq \infty$.*

PROOF. Let G be the graph with no simple cycles involving node 0, under which the traders' signals are hierarchically conditionally independent under Q^δ (without conditioning on h_t). Recall that $\tilde{s}_0 = \tilde{\omega}$ by definition.

The proof of the lemma is by induction. The base case for $t = 1$ is trivial. Assume the claim holds for $t = T$; i.e., conditioning on the history h_T , the set of random variables $\{\tilde{s}_0, \dots, \tilde{s}_n\}$ forms a Markov random field with respect to G . This implies that for every set $S, A, B \subset \{0, \dots, n\}$ with A and B disconnected in $G \setminus S$, $\{\tilde{s}_i: i \in A\}$ is conditionally independent of $\{\tilde{s}_i: i \in B\}$, given $\{\tilde{s}_i: i \in S\}$ and h_T . Fix one such S, A , and B .

Let trader k arrive at the market at time T , i.e., $i_T = k$, and let y_T denote the trade conducted by the trader. We have the following three cases, depending on whether $k \in S$, $k \in A$, or $k \in B$.

1. $k \in S$: Since y_T depends only on h_T and \tilde{s}_k and $k \in S$, conditioning on $h_{T+1} = (h_T, y_T)$ and $\{\tilde{s}_i: i \in S\}$ is same as conditioning on h_T and $\{\tilde{s}_i: i \in S\}$. Thus it follows trivially that the induction step holds.

2. $k \in A$: Again, since y_T depends only on \tilde{s}_k and h_T , we see that conditional on h_T and $\{\tilde{s}_i: i \in S\}$, the trade y_T depends only on the set $\{\tilde{s}_i: i \in A\}$. Thus, further conditioning on y_T cannot change the independence of $\{\tilde{s}_i: i \in A\}$ and $\{\tilde{s}_i: i \in B\}$. Thus, again the induction step holds.

3. $k \in B$: The argument is similar as that for $k \in A$.

Thus, regardless of which trader arrives at the market, the induction step carries through. As S, A , and B were arbitrary, this proves the result for all $t < \infty$. The result for $t = \infty$ follows through the use of Lemma 1. \square

The following lemma plays an analogous role to Lemma 4 in the proof of Theorem 2.

LEMMA 6. *Let $\{X_v: v \in V\}$ be a Markov random field with respect to a connected graph $G = (V, E)$. Let $u \in V$ be such that there exists no simple cycle involving u in G . Moreover, let X_u be independent of X_v for each $v \in \partial u$. Then, X_u is independent of $\{X_v: v \in V, v \neq u\}$.*

PROOF. Let A_v denote an event for each $v \in V$. As $\{X_v: v \in V\}$ is a Markov random field, we get

$$\begin{aligned} & \mathbf{P}(X_v \in A_v \forall v \notin \bar{\partial}u, X_u \in A_u | X_i \forall i \in \partial u) \\ &= \mathbf{P}(X_v \in A_v \forall v \notin \bar{\partial}u | X_i \forall i \in \partial u) \cdot \mathbf{P}(X_u \in A_u | X_i \forall i \in \partial u). \end{aligned}$$

Next, note that conditional on X_u , the random variables in the set $\{X_v: v \in \partial u\}$ are independent of each other. Furthermore, we have that X_u is independent of X_v for each $v \in \partial u$. Thus,

from Lemma 4, we obtain that X_u is independent of $\{X_v: v \in \partial u\}$. This implies that $\mathbf{P}(X_u \in A_u | X_i \forall i \in \partial u) = \mathbf{P}(X_u \in A_u)$. Thus, we get

$$\begin{aligned} & \mathbf{P}(X_v \in A_v \forall v \notin \bar{\partial}u, X_u \in A_u | X_i \forall i \in \partial u) \\ &= \mathbf{P}(X_v \in A_v \forall v \notin \bar{\partial}u | X_i \forall i \in \partial u) \cdot \mathbf{P}(X_u \in A_u). \end{aligned}$$

This implies by definition that

$$\mathbf{P}(X_v \in A_v \forall v \in V) = \mathbf{P}(X_v \in A_v \forall v \neq u) \cdot \mathbf{P}(X_u \in A_u).$$

As this is true for all events A_v , we obtain that X_u is independent of $\{X_v: v \in V, v \neq u\}$. \square

The proof of Theorem 6 follows directly from the preceding two lemmas.

PROOF OF THEOREM 6. From Lemma 5, we obtain that, under φ_∞ , the state $\tilde{\omega}$ and the signals $\tilde{s}_1, \dots, \tilde{s}_n$ satisfy hierarchical conditional independence. Furthermore, recall from Theorem 1 that, under φ_∞ , the payoff-relevant state $\tilde{\omega}$ is independent of any trader i 's private signal \tilde{s}_i . Thus, using Lemma 6, we obtain that the state $\tilde{\omega}$ is independent of the joint signal \tilde{s} . The result then follows from Lemma 2. \square

Can the signal structure be generalized further and still yield information aggregation? To partly investigate this question, we end this section with an example of a market where the traders' signals have a more general dependence: we show that no trade is a PBE in this market, thus proving that information aggregation may not occur with stronger correlation in the signals of traders.

EXAMPLE 1. Let $\Omega = \{0, 1\}$, and let the state $\tilde{\omega}$ be distributed uniformly at random. The signal space for each trader is given by $\Sigma_i = \{0, 1\}$. The signals of the traders are distributed according to the following distribution:

$$\begin{aligned} & \mathbf{P}(\tilde{s} = (0, 0) | \tilde{\omega} = 0) = \mathbf{P}(\tilde{s} = (1, 1) | \tilde{\omega} = 0) = 0.5, \\ & \mathbf{P}(\tilde{s} = (0, 1) | \tilde{\omega} = 1) = \mathbf{P}(\tilde{s} = (1, 0) | \tilde{\omega} = 1) = 0.5. \end{aligned}$$

That is, if the payoff-relevant state $\tilde{\omega} = 0$, then the traders' signals are the same and equal to 0 or 1 with equal probability. On the other hand, if the state $\tilde{\omega} = 1$, then the traders' signals are again 0 or 1 with equal probability, but now the signal are complementary to each other. Note that by pooling their information, the two traders can always exactly determine the state $\tilde{\omega}$.

Let the market maker use the cost function corresponding to the logarithmic market scoring rule to price the trades. We now present a system of beliefs and a strategy for each trader, which together constitutes a PBE. After any history h , let each trader's belief about the state $\tilde{\omega}$ and the other trader's signal be given by the distribution \mathbf{P} conditioned on their private signal. Note that this belief structure does not condition on the observed history, and thus this prevents any signaling between the traders. Consider a strategy δ^i for the trader i , such that at any time t after any history h , the trade specified by δ^i is given by $y_t = -q_t$, where q_t denotes the total quantity of securities sold by the market maker until time t . Given the belief of a trader about the state $\tilde{\omega}$ and the other trader's future actions, it can be shown that the optimal action for a trader to take is to always bring the total quantity of securities sold by the market maker to zero. Since the actions

of the traders do not depend on their signals, the beliefs satisfy the Bayesian rule. Thus, it follows that the strategy profile δ and the belief structure form a PBE.

Looking at the equilibrium path for the above PBE, one obtains that there is no trade occurring in the market if the market maker starts with $q_0 = \mathbf{0}$. Thus at all time periods, the belief of the traders regarding the payoff-relevant state $\tilde{\omega}$ assigns equal probability to both possible alternatives. However, if the signals were to be pooled, then the true state would be known with complete certainty. Thus we observe that, with general signal structure, there exists a PBE where information aggregation does not occur. \square

The preceding example is closely related to an example presented by Ostrovsky (see Ostrovsky 2012, Example 1), where he shows the importance of separability of the security for information aggregation in his market model. The main distinction is that in our model, we have two securities; in Ostrovsky's example, trade is over one security (his security is a linear combination of those in the preceding example).

Appendix C. Measure-Theoretic Formulation

Recall that the history at time t is given by the trades until time t :

$$h_t = (y_1, \dots, y_{t-1}).$$

Moreover, the infinite history is defined as the entire sequence of trade that takes place in the market, given by $h_\infty \triangleq (y_1, \dots, y_t, \dots)$. Let H_t denote the set of all possible histories till time t , and let H_∞ denote the set of all possible infinite history. We endow H_∞ with the topology of pointwise convergence, and we let \mathcal{G}_∞ denote the Borel σ -algebra on H_∞ .

Let $(\Sigma_i, \mathcal{F}_i)$ be a measure space, and let \mathcal{T} denote the product σ -algebra on Σ . Moreover, let \mathcal{C} denote the complete σ -algebra on Ω . Finally, let $\mathcal{S} \triangleq H_\infty \times \Omega \times \Sigma$, and let \mathcal{F}_∞ denote the product σ -algebra on \mathcal{S} .

Note that $(\mathcal{S}, \mathcal{F}_\infty)$ captures all the uncertainty in our model. In particular, the random variables $\tilde{\omega}$ and $\tilde{s}_1, \dots, \tilde{s}_n$ are all measurable with respect to the measure space $(\mathcal{S}, \mathcal{F}_\infty)$.

The joint prior distribution \mathbf{P} over the state $\tilde{\omega}$ and the signals $\tilde{s}_1, \dots, \tilde{s}_n$ shared by all the traders and the market maker can now be represented as a measure over the restricted probability space $(\Omega \times \Sigma, \sigma(\mathcal{C} \times \mathcal{T}))$.

Beliefs. The belief $\nu_i(h_t, s_i)$ of trader i after each history h_t , and having observed signal s_i , is represented as a probability measure on the set $(\mathcal{S}, \mathcal{F}_\infty)$.

Strategies. The strategy δ_i of trader i can be represented as a set of random variables $\delta_i(h_t, s_i)$ such that $\delta_i(h_t, \tilde{s}_i)$ is measurable with respect to the sub- σ -algebra generated by h_t and \tilde{s}_i . Note that if δ_i is a mixed strategy, then we have to consider a larger measure space to include the randomization of the trader.

Given the joint prior \mathbf{P} on $\tilde{\omega}$ and the signals $\tilde{s}_1, \dots, \tilde{s}_n$, a PBE with strategy profile δ and the belief ν_i for each trader i induces a probability measure Q^δ on $(\mathcal{S}, \mathcal{F}_\infty)$ through the Kolmogorov extension theorem. The measures $\varphi_{i,t}$ and φ_t can then be defined as in the main text.

Appendix D. Information Aggregation and Separable Orientedness

For the sake of completeness, in this section, we provide the formal statement and the proof of the results claimed in the

discussion at the end of §4. We follow the same notation as introduced in that section.

The following theorem states that the limiting exchange economy satisfies the quasicompleteness condition, as specified in Definition 2 of DeMarzo and Skiadas (1998, p. 9; hereafter DS).

THEOREM 7. *The limiting exchange economy is quasicomplete, with the measure Q given by*

$$Q(\tilde{\omega} = \omega, \tilde{s}_i \in A_i \text{ for all } i) \\ = \nabla_{y(\omega)} K(h_\infty, \mathbf{0}) \varphi_\infty(\tilde{s}_i \in A_i \text{ for all } i \mid \tilde{\omega} = \omega). \quad (D1)$$

(Note that Q is well defined, as $\sum_\omega \nabla_{y(\omega)} K(h_\infty, \mathbf{0}) = 1$ from (6).)

PROOF. First note that since the traders share a common prior belief, we have (as per the notation in DS)

$$P_i = \varphi_\infty$$

for all i . This can be rewritten as

$$P_i(\tilde{\omega} = \omega, \tilde{s}_i \in A_i \text{ for all } i) \\ = \varphi_\infty(\tilde{\omega} = \omega) \varphi_\infty(\tilde{s}_i \in A_i \text{ for all } i \mid \tilde{\omega} = \omega).$$

With Q as in the theorem statement, this implies that

$$\frac{dQ}{dP_i}(\tilde{\omega} = \omega, \tilde{s}_i) = \frac{\nabla_{y(\omega)} K(h_\infty, \mathbf{0})}{\varphi_\infty(\tilde{\omega} = \omega)} \\ = \frac{\varphi_\infty(\tilde{\omega} = \omega \mid \tilde{s}_i)}{\varphi_\infty(\tilde{\omega} = \omega) \mathbf{E}_{\varphi_{i,\infty}}[u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))]} \\ \cdot u'_i(\omega, w_{i,\infty}^\delta(\omega)). \quad (D2)$$

For any admissible price vector p , define the following quantities:

$$\lambda_i^p(\tilde{\omega} = \omega, \tilde{s}_i) = \frac{\varphi_\infty(\tilde{\omega} = \omega \mid \tilde{s}_i)}{\varphi_\infty(\tilde{\omega} = \omega) \mathbf{E}_{\varphi_{i,\infty}}[u'_i(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}))]} \quad (D3) \\ \theta_i^p = \mathbf{0}.$$

Note that we have

$$\lambda_i^p = \frac{\varphi_\infty(\tilde{\omega} = \omega \mid \tilde{s}_i)}{\varphi_\infty(\tilde{\omega} = \omega) \sum_{\hat{\omega} \in \Omega} \varphi_\infty(\tilde{\omega} = \hat{\omega} \mid \tilde{s}_i) u'_i(\hat{\omega}, w_{i,\infty}^\delta(\hat{\omega}))} \\ \leq \frac{1}{\varphi_\infty(\tilde{\omega} = \omega) u'_i(\omega, w_{i,\infty}^\delta(\omega))} \\ \leq \frac{1}{\min_{\hat{\omega} \in \Omega} \varphi_\infty(\tilde{\omega} = \hat{\omega}) u'_i(\hat{\omega}, w_{i,\infty}^\delta(\hat{\omega}))} \\ < \infty.$$

Here, the first inequality follows from the fact that u_i is strictly increasing, and the final bound follows from the finiteness of Ω .

With these substitutions, we see that for any admissible price vector p , we have, almost surely,

$$\lambda_i^p u_i \left(\tilde{\omega}, w_{i,\infty}^\delta(\tilde{\omega}) + \sum_{\hat{\omega}} \theta_i^p(\hat{\omega}) (\mathbf{I}\{\tilde{\omega} = \hat{\omega}\} - p(\hat{\omega})) \right) = \frac{dQ}{dP_i},$$

with λ_i^p measurable with respect to the information in \tilde{s}_i and the price vector p . This implies that the exchange economy is quasicomplete. \square

Define $f: \Sigma \rightarrow [0, 1]^{|I|}$ as

$$f(s, \omega) = \mathbb{E}_Q[\mathbf{I}\{\tilde{\omega} = \omega\} \mid \tilde{s} = s]. \quad (\text{D4})$$

From the results in DS, we obtain that if, almost surely, the mapping f in the limiting exchange economy is separably oriented, then the aforementioned REE is unique and *fully revealing*. Since the market is complete, this implies that information aggregation holds in the PBE.

The next theorem states that the converse result is also true.

THEOREM 8. *If information aggregation holds in the dynamic market, then, almost surely, the function f , as defined in (D4) in the limiting exchange economy, is separably oriented.*

PROOF. For f to be separably oriented, we require the existence of product measurable and bounded functions $g_i: \Sigma_i \times \mathbb{R}^{|I|} \rightarrow \mathbb{R}^{|I|}$ such that the following condition holds for all $s \in \Sigma$ and $r \in \mathbb{R}^{|I|}$:

$$f(s) \neq r \Rightarrow (f(s) - r) \cdot \sum_{i=1}^n g_i(s_i, r) > 0. \quad (\text{D5})$$

Note that if information aggregation holds, then by Lemma 2, under φ_∞ , the state $\tilde{\omega}$ is independent of the joint signal \tilde{s} . This implies that $\varphi_\infty(\tilde{s}_i \in A_i \text{ for all } i \mid \tilde{\omega} = \omega) = \varphi_\infty(\tilde{s}_i \in A_i \text{ for all } i)$, and hence, from (D1), we obtain

$$Q(\tilde{\omega} = \omega, \tilde{s}_i \in A_i \text{ for all } i) = \nabla_\omega K(h_\infty, \mathbf{0}) \varphi_\infty(\tilde{s}_i \in A_i \text{ for all } i).$$

Thus, even under measure Q , the state $\tilde{\omega}$ is independent of the joint signal \tilde{s} . This implies that for all $s \in \Sigma$,

$$f(s) = \nabla K(h_\infty, \mathbf{0}).$$

As $f(s)$ is independent of s , we define the functions $g_i(s_i, r)$ as

$$g_i(s_i, r) = \nabla K(h_\infty, \mathbf{0}) - r$$

and observe that (D5) holds. Thus, f is separably oriented. \square

The preceding theorems imply that, in complete markets, information aggregation is equivalent to the condition that, almost surely, the function f in the limiting economy is separably oriented.

References

- Agrawal S, Delage E, Peters M, Wang Z, Ye Y (2011) A unified framework for dynamic prediction market design. *Oper. Res.* 59(3):550–568.
- Banerjee AV (1992) A simple model of herd behavior. *Quart. J. Econom.* 107(3):797–817.
- Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom and cultural change as information cascades. *J. Political Econom.* 100(5):992–1026.
- Blume LE, Easley D (1998) Rational expectations and rational learning. Majumdar M, ed. *Organizations with Incomplete Information: Essays in Economic Analysis: A Tribute to Roy Radner* (Cambridge University Press, Cambridge, UK), 61–109.
- Chen Y, Pennock DM (2007) A utility framework for bounded-loss market makers. Paar R, van der Gaag L, eds. *Proc. 23rd Conf. Uncertainty in Artificial Intelligence (UAI 2007)* (AUAI Press, Corvallis, OR), 49–56.
- Chen Y, Dimitrov S, Sami R, Reeves DM, Pennock DM, Hanson RD, Fortnow L, Gonen R (2010) Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica* 58(4):930–969.
- Cochrane JH (2005) *Asset Pricing*, revised ed. (Princeton University Press, Princeton, NJ).
- Cowgill B, Wolfers J, Zitzewitz E (2009) Using prediction markets to track information flows: Evidence from Google. Das S, Ostrovsky M, Pennock D, Szymanski B, eds. *Auctions, Market Mechanisms and Their Applications*, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, Vol. 14 (Springer, Berlin), 3.
- DeMarzo P, Skiadas C (1998) Aggregation, determinacy, and informational efficiency for a class of economies with asymmetric information. *J. Econom. Theory* 80(1):123–152.
- DeMarzo P, Skiadas C (1999) On the uniqueness of fully informative rational expectations equilibria. *Econom. Theory* 13(1):1–24.
- Fudenberg D, Levine D (1983) Subgame-perfect equilibria of finite and infinite-horizon games. *J. Econom. Theory* 31(2):251–268.
- Fudenberg D, Tirole J (1991) Perfect Bayesian equilibrium and sequential equilibrium. *J. Econom. Theory* 53(2):236–260.
- Glosten LR (1994) Is the electronic open limit order book inevitable? *J. Finance* 49(4):1127–1161.
- Glosten LR, Milgrom PR (1985) Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financial Econom.* 14(1):71–100.
- Hayek FA (1945) The use of knowledge in society. *Amer. Econom. Rev.* 35(4):519–530.
- Healy PJ, Linardi S, Lowery JR, Ledyard JO (2010) Prediction markets: Alternative mechanisms for complex environments with few traders. *Management Sci.* 56(11):1977–1996.
- Jian L, Sami R (2012) Aggregation and manipulation in prediction markets: Effects of trading mechanism and information distribution. *Management Sci.* 58(1):123–140.
- Kyle AS (1985) Continuous auctions and insider trading. *Econometrica* 53(6):1315–1335.
- Lo AW (2008) Efficient markets hypothesis. Durlauf SN, Blume LE, eds. *The New Palgrave Dictionary of Economics*, 2nd ed. (Palgrave Macmillan, Basingstoke, UK).
- Mailath GJ, Samuelson L (2006) *Repeated Games and Reputations: Long-Run Relationships* (Oxford University Press, New York).
- Mas-Colell A, Whinston MD, Green JR (1995) *Microeconomic Theory* (Oxford University Press, Oxford, UK).
- Milgrom P, Stokey N (1982) Information, trade and common knowledge. *J. Econom. Theory* 26(1):17–27.
- Ostrovsky M (2012) Information aggregation in dynamic markets with strategic traders. *Econometrica* 80(6):2595–2647.
- Othman A, Sandholm T (2013) The Gates Hillman prediction market. *Rev. Econom. Design* 17(2):95–128.
- Pennock DM, Sami R (2007) Computational aspects of prediction markets. Nisan N, Roughgarden T, Tardos É, Vazirani VV, eds. *Algorithmic Game Theory* (Cambridge University Press, Cambridge, UK), 651–675.
- Radner R (1979) Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica* 47(3):655–678.
- Smith L, Sorensen PN (2006) Informational herding and optimal experimentation. Cowles Foundation Discussion Paper 1552, Cowles Foundation, Yale University, New Haven, CT.