

Information aggregation in smooth markets

Krishnamurthy Iyer
Department of Management
Science and Engineering
Stanford University
kriyer@stanford.edu

Ramesh Johari
Department of Management
Science and Engineering
Stanford University
ramesh.johari@stanford.edu

Ciamac C. Moallemi
Graduate School of Business
Columbia University
ciamac@gsb.columbia.edu

ABSTRACT

Recent years have seen extensive investigation of the information aggregation properties of prediction 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. We consider a market model involving finitely many informed risk-averse traders interacting with a market maker. Our main result identifies a basic *smoothness* condition on the price in the market that ensures information will be aggregated. We give conditions under which cost function market makers (or, equivalently, market makers based on market scoring rules) satisfy the smoothness requirement. We further 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.

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 events 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 in the hands of traders [13].

This basic goal has led to extensive investigation of the information aggregation properties of prediction markets. More broadly, however, the notion that information is aggregated in market prices has been a longstanding topic of investigation in economics; for example, see [6, 9, 10, 14]. Informally, the idea that market prices should reflect the information possessed by market participants is a form of the *efficient markets hypothesis* [8], and underlies much of the appeal of decentralized markets in settings where traders

possess asymmetric information.

Despite the central connection between information aggregation 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. The central goal of this paper is to investigate conditions under which information is aggregated in such scenarios.

Our main contributions are as follows:

1. *Smoothness and information aggregation.* Our main result identifies a basic *smoothness* condition on prices in the market that ensures information will be aggregated. Formally, we show that 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 the market will aggregate information in any perfect Bayesian equilibrium. 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. In particular, if a rational market maker is sufficiently concerned that traders possess significantly superior information regarding the value of a traded asset, then the bid-ask spread may become so large as to completely preclude trading and thus the dissemination of information [6]. Our main contribution is to demonstrate that with a smoothness assumption on the market, such behavior cannot arise. 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.

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 [3, 4, 12], 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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

markets as insurance to hedge against risks inherent in other elements of their portfolio. Second, if all traders are risk neutral, then in general, the *no-trade theorem* applies and precludes trade among rational traders without additional assumptions on the market [11]. 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 averse traders trading can proceed even if all traders are rational: traders may trade purely on the motive of hedging.

3. *The information content of prices.* When traders are risk averse, prices may not reflect the true posterior of an event occurring; this is because prices must also reflect the marginal expected utility of the traders. We show that if even one trader in the market is risk neutral, however, then prices are accurate posterior probabilities. We further 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.

The most closely related paper to our own work is by Ostrovsky [12]. 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 both the Kyle market maker model [10], as well as in markets based on market scoring rules. The markets we consider are complete (i.e., there exists one contingent contract for each possible event that can occur); and with the partition model for signal structure that Ostrovsky considers (a special case of our signal structure), a complete market is always separable. 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 Section 3, we formally define information aggregation. In Section 4, we prove that if the portfolios of traders converge and the market is asymptotically smooth, then information is aggregated. In Section 5, we show that the portfolio convergence assumption can be relaxed in markets where the signal space of traders is finite, and the loss to the market maker is bounded; the latter condition is satisfied by a wide range of algorithmic market makers based on cost functions, as we observe in Section 6. Finally, in Section 7, we provide insight into the information content of the limiting prices.

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 known *a priori*

to all the traders and the market maker. We further assume that each trader visits the market infinitely often.

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 *state of the world*, and assume that all traders have a common prior distribution for $\tilde{\omega}$. We assume that the state is only revealed after all trades are completed. Further, we assume a *complete market* over the possible states of the world. 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 0 otherwise.

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 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, we have

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

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 let $H_f = \cup_{t \geq 1} H_t$ denote the set of all possible finite histories. Finally, let $h_\infty = (y_1, y_2, \dots)$ 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 $\mathbf{1} = (1, \dots, 1)$, 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*, 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.

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*. 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 $\mathbf{0} = (0, \dots, 0)$. Furthermore, we assume that the pricing function satisfies the following assumption:

$$K(h_t, y + a\mathbf{1}) = K(h_t, y) + a, \text{ for all } a \in \mathbb{R} . \quad (1)$$

The above condition ensures that the traders do not profit via an exchange of money with the market maker. Another implication is that the market maker prices the trades considering only the risk involved in the trade.

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. In our model, uncertainty arises because the state of the world is unknown; however, we assume traders are informed, and receive signals regarding the true state. In this section we define signals, beliefs, strategies, utility, and, ultimately, 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 B we provide a formal measure-theoretic description of each of the elements introduced here.

Signals. Before trading begins, each trader i receives a private signal $\tilde{s}_i \in \Sigma_i$. We assume that the private signals are drawn independently conditional on the state $\tilde{\omega}$, and that the joint distribution of $\tilde{\omega}$ and the signals is common knowledge among the traders. Let \mathbf{P} denote the joint prior distribution of $\tilde{\omega}$ and $\tilde{s}_1, \dots, \tilde{s}_n$. We refer to $\tilde{s} = (\tilde{s}_1, \dots, \tilde{s}_n)$ as the joint signal and $\Sigma = \times_i \Sigma_i$ as the joint signal space.

Beliefs. Let $\mathcal{S} = H_\infty \times \Omega \times \Sigma$. Observe that an element of \mathcal{S} captures all uncertainty in the model: the path of the market, the state of the world, 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_i , trader i 's belief $\nu_i(h_t, s_i)$ is a probability distribution over \mathcal{S} . This represents trader i 's forecast of the future actions by traders, the state $\tilde{\omega}$, and the signal vector \tilde{s} .

Strategies. The 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 $\delta_i(h, s_i)$, where s_i denotes the private signal received by the trader. (Of course, a 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 on the realized history h_t and the private signal \tilde{s}_i . When the context is clear, we will omit the explicit history dependence. Let $\delta = (\delta_1, \dots, \delta_n)$ denote the 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 player except i .

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

Utility and the value function. We assume that if trader i is holding a portfolio w_i , and state $\tilde{\omega} = \omega$ is realized, then trader i obtains a utility $u_i(w_i(\omega))$. The utility function $u_i : \mathbb{R} \rightarrow \mathbb{R}$ 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 to hold more quantity of any security ω . These two characteristics of the traders are captured by assuming that each trader's utility function is strictly increasing and concave. We further make the assumption in our analysis that the utility functions are differentiable. Note that all these assumptions hold trivially for risk neutral traders, whose utility functions are given by $u_i(x) = x$ for all $x \in \mathbb{R}$.

While we assume that traders act as expected utility maximizers, as trading proceeds over an infinite horizon, defining their 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:

$$J_{i,t}(\delta) = \mathbf{E}_{\nu_i(h_t, s_i), \delta} \left[\liminf_{T \rightarrow \infty} u_i(w_{i,T}(\tilde{\omega})) \right]. \quad (2)$$

We assume that traders act at each time period to maximize their value function. One may think of the quantity $\liminf_{T \rightarrow \infty} u_i(w_{i,T}(\tilde{\omega}))$ as the final utility of the trader i at the conclusion of trading in the market. Thus the value function captures the utility of the eventual portfolio of the trader. It must be noted that the value function is not time separable over the trades at different time periods, and there is no discounting involved for trades occurring at later periods. We assume that for all cases of interest, the value functions are well defined at all times for all traders.

Perfect Bayesian equilibrium. Finally, we describe the notion of PBE for the market defined above. 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 almost all $s_i \in \Sigma_i$, we have

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

2. For each i , we have $\nu_i(h, s_i)(\cdot) = Q^\delta(\cdot | h, \tilde{s}_i = s_i)$ for all histories h where the latter conditional distribution is well defined.

The first condition requires that after any history, the actions specified by the strategy δ_i for trader i are 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.

Whenever we consider a particular PBE with 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 will define the notion of information aggregation. Informally, this is when the private signals observed by individual traders are aggregated into a common market belief. In order 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 state and 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 as the *common belief of the market* over the state and signals at time t , by

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

Similarly, 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 state and signals at time t , and in a PBE this will be given by

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

Note that φ_t and $\varphi_{i,t}$ have implicit history dependence. Moreover, trivially, $\varphi_{i,t}$ can be obtained from φ_t by conditioning on trader i 's signal: $\varphi_{i,t}(\cdot) = \varphi_t(\cdot | \tilde{s}_i = s)$.

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 (resp., $\varphi_{i,t}$) converges weakly to a probability distribution φ_∞ (resp., $\varphi_{i,\infty}$), 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.

DEFINITION 1. (Information Aggregation) *The market aggregates the information of the traders 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 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 above 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 private joint signal. Subsequently, we will show that an uninformed outsider can infer the information in the private signal under much milder requirements. Further, we will demonstrate the relationship between prices and the utility functions of the traders.

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 so that almost surely under the distribution Q^δ , $K(h_t, \cdot)$ converges uniformly over \mathcal{N} as $t \rightarrow \infty$. Further, defining the limit*

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

for $y \in \mathcal{N}$, we require that $K(h_\infty, \cdot)$ be almost surely continuously differentiable over \mathcal{N} .

The above condition ensures 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 [5]. Further, 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 will make the following assumption on the equilibria under consideration:

ASSUMPTION 2. (Portfolio Convergence) *Under the strategy profile δ , the portfolio of each trader converges as trading proceeds in the market. In other words, for each trader i , the limit*

$$w_i^\delta = \lim_{t \rightarrow \infty} w_{i,t}^\delta$$

exists (and is finite) almost surely.

Portfolio convergence captures the idea that, asymptotically, traders will stop trading despite the fact that infinitely many trading opportunities are available. Note that in our model, the private information about security payoffs is revealed to individual traders at the start of trading. Over time, no new information is revealed to the collective of market participants, existing information is merely diffused. Hence, it seems reasonable in this setting that if the market is well-functioning, the portfolios will eventually converge. In this way, one can view Assumption 2 as excluding pathological equilibria. In what follows, however, we will develop conditions on the underlying model primitives which allow the verification of portfolio convergence of the equilibrium and demonstrate that this holds in many common markets.

The following theorem, which is the main result of this section, shows information is aggregated if the previous two assumptions hold.

THEOREM 1. *In any PBE where asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2) hold, the market aggregates the information of the traders.*

In order to prove Theorem 1, we need the following lemma; the proof can be found in the appendix. It makes the intuitive observation that under Q^δ , even after observing past history, the signals are independent conditional on the state. (Note that this result holds even after infinite history, i.e., $t = \infty$.)

LEMMA 2. *Under Q^δ , conditional on the history h_t and the state $\tilde{\omega}$, the private signals $\tilde{s}_1, \dots, \tilde{s}_n$ are independent, for any $1 \leq t \leq \infty$.*

The proof of Theorem 1 uses the fact that in a PBE, any unilateral deviation by a trader should result in lower utility for the 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 . We then obtain the result via Lemma 5 in Appendix A, a basic lemma on independent random variables.

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

$$J_{i,t}(\delta) = \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)} \left[u_i(w_i^\delta(\tilde{\omega})) \right]. \quad (3)$$

Now, consider the deviation where trader i follows the strategy δ_i up to time t , then trades a fixed quantity $z \in \mathbb{R}^m$ at time t and never trades again thereafter. Let $w_{i,t}^z$ denote the payoff vector of the trader i after such a trade, i.e.,

$$w_{i,t}^z = w_{i,t-1}^\delta + z - K(h_t, z)\mathbf{1}. \quad (4)$$

As δ is a PBE strategy profile, any deviation after history h_t should lead to lower utility for the trader. Thus, almost surely,

$$\mathbf{E}_{\nu_i(h_t, \tilde{s}_i)} [u_i(w_{i,t}^z(\tilde{\omega}))] \leq \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)} [u_i(w_i^\delta(\tilde{\omega}))], \quad (5)$$

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

$$\begin{aligned} & \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)} [u_i(w_{i,t}^z(\tilde{\omega}))] \\ &= \sum_{\omega \in \Omega} \varphi_{i,t}(\tilde{\omega} = \omega) u_i(w_{i,t}^\delta(\omega) + z(\omega) - K(h_t, z)). \end{aligned} \quad (6)$$

By Assumption 2, the portfolio of the trader converges. Moreover, as the pricing function satisfies Assumption 1 (asymptotic smoothness), for all $z \in \mathcal{N}$, we know $K(h_t, z) \rightarrow K(h_\infty, z)$. Using these facts and Lemma 1, we obtain for all $z \in \mathcal{N}$, almost surely,

$$\begin{aligned} & \lim_{t \rightarrow \infty} \mathbf{E}_{\nu_i(h_t, \tilde{s}_i)} [u_i(w_{i,t}^z(\tilde{\omega}))] \\ &= \sum_{\omega \in \Omega} \varphi_{i,\infty}(\tilde{\omega} = \omega) u_i(w_i^\delta(\omega) + z(\omega) - K(h_\infty, z)) \\ &= \mathbf{E}_{\varphi_{i,\infty}} [u_i(w_{i,t}^z(\tilde{\omega}))], \end{aligned} \quad (7)$$

where w_i^z is defined as

$$w_i^z = w_i^\delta + z - K(h_\infty, z)\mathbf{1}. \quad (8)$$

On the other hand, almost surely,

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

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

Taking the limit as $t \rightarrow \infty$ in (5), we obtain almost surely for all $z \in \mathcal{N}$ that

$$\mathbf{E}_{\varphi_{i,\infty}} [u_i(w_{i,t}^z(\tilde{\omega}))] \leq \mathbf{E}_{\varphi_{i,\infty}} [u_i(w_i^\delta(\tilde{\omega}))], \quad (9)$$

Since $w_i^0 = w_i^\delta$, we conclude that, almost surely,

$$\mathbf{0} \in \arg \max_{z \in \mathcal{N}} \mathbf{E}_{\varphi_{i,\infty}} (u_i(w_{i,t}^z(\tilde{\omega}))). \quad (10)$$

By asymptotic smoothness, the pricing function is $K(h_\infty, z)$ is differentiable at $z = \mathbf{0}$. Moreover, the traders have differentiable utility functions. This implies the following first order necessary condition holds almost surely, for each $\omega \in \Omega$:

$$\begin{aligned} & \varphi_{i,\infty}(\tilde{\omega} = \omega) u_i'(w_i^\delta(\omega)) \\ &= \nabla_\omega K(h_\infty, \mathbf{0}) \mathbf{E}_{\varphi_{i,\infty}} (u_i'(w_i^\delta(\tilde{\omega}))). \end{aligned}$$

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

$$\frac{\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)}{\mathbf{E}_{\varphi_\infty} [u_i'(w_i^\delta(\tilde{\omega})) | \tilde{s}_i]} = \frac{\nabla_\omega K(h_\infty, \mathbf{0})}{u_i'(w_i^\delta(\omega))}. \quad (11)$$

Now, observe that if h_∞ is known, then the right hand side of the above equation is fixed. On the other hand, the left hand side depends on the realization of \tilde{s}_i . For the equality to hold, we infer that the left hand side is also independent of the realization of \tilde{s}_i . On summing the left hand side over $\omega \in \Omega$, we obtain that $\mathbf{E}_{\varphi_\infty} [u_i'(w_i^\delta(\tilde{\omega})) | \tilde{s}_i]$ is independent of \tilde{s}_i , since the numerator sums to 1. This in turn implies that for all ω , $\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)$ is independent of the realization of \tilde{s}_i . From this argument we conclude that under the asymptotic common belief φ_∞ , the state $\tilde{\omega}$ is independent of the private signal \tilde{s}_i of each trader i .

By Lemma 2, under measure φ_∞ , the private signals $\tilde{s}_1, \dots, \tilde{s}_n$ are independent given the state $\tilde{\omega}$. Together with Lemma 5 from Appendix A, this implies that under measure φ_∞ , the state $\tilde{\omega}$ and the joint signal of the traders $\tilde{s} = (\tilde{s}_1, \dots, \tilde{s}_n)$ are independent. In other words, we have, almost surely, that

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

The right hand side is equal to $Q^\delta(\tilde{\omega} = \omega | h_\infty)$. For the left hand side, note that at any time t , we have $\varphi_t(\tilde{\omega} = \omega | \tilde{s}) = \mathbf{P}(\tilde{\omega} = \omega | \tilde{s})$, since the history cannot contain any more information about $\tilde{\omega}$ than the pooled signal. Taking limits, we obtain the same equality for φ_∞ , as required. \square

5. INFORMATION AGGREGATION WITH A FINITE SIGNAL SPACE

Although the assumption of portfolio convergence is reasonable, it is a condition that depends on the equilibrium itself, rather than model primitives. In particular, one goal of our paper is to develop an information aggregation result that applies to a wide range of algorithmic market makers. For many such market makers, especially those based on cost functions (or equivalently, market scoring rules), the market maker employs a pricing function that ensures the total loss by the market maker over any trajectory of the market remains bounded. In this section, we specialize to a model where the signal space Σ_i of each trader i is finite, and show that if the loss to the market maker is bounded, then information is aggregated. In the next section, we discuss the application of this result to algorithmic market makers.

The loss incurred by the market maker is given by the net change in the portfolio of all the traders. Thus the loss of the market maker at time t , if the state realized is $\tilde{\omega} = \omega$, is given by $\sum_{i=1}^n (w_{i,t}(\omega) - w_{i,0}(\omega))$. The assumption we make is that in the PBE, the market maker's loss is bounded; this is equivalent to the following statement.

ASSUMPTION 3. (Bounded Loss) For the PBE with strategy profile δ , there exists $B^\delta \in \mathbb{R}$ such that, almost surely, the payoff vector w_i^δ satisfies

$$\limsup_{t \rightarrow \infty} \sum_{i=1}^n w_{i,t}^\delta(\omega) < B^\delta, \quad (12)$$

for all $\omega \in \Omega$.

The following theorem shows that under the preceding assumption, information is aggregated in any *pure strategy* PBE, i.e., a PBE where traders play deterministic actions after each history.

THEOREM 2. *If the signal space is finite, then in any pure strategy PBE satisfying asymptotic smoothness (Assumption 1) and bounded loss (Assumption 3), the market aggregates the information of the traders.*

PROOF. The proof of the theorem involves showing that portfolio convergence holds in the present setting. Suppose δ is the strategy profile in the given PBE. First note that since δ is pure, the history h_t is a deterministic function of the joint signal \tilde{s} . Since the set Σ of all possible joint signals σ is finite, it follows that the set of possible infinite histories under the PBE with strategy profile δ is a finite set—this is the main simplification gained by assuming the signal space is finite.

Let $H_\infty^\delta = \{h^{(1)}, \dots, h^{(K)}\}$ denote the set of all possible infinite histories under the PBE strategy profile δ . Thus, we have $Q^\delta(H_\infty^\delta) = 1$ and for all $h \in H_\infty^\delta$, we have $Q^\delta(h) > 0$. Moreover, we see that the joint signal space Σ can be partitioned into $\{\Sigma^1, \dots, \Sigma^K\}$, such that all realizations of \tilde{s} in the partition Σ^l lead to the same infinite history $h^{(l)}$.

Let $T_m = \inf\{t : \text{for all } k \neq l, h_t^{(k)} \neq h_t^{(l)}\}$. In other words, T_m is the earliest time by which all the possible histories in H_∞^δ differ. Thus, at time T_m , each trader, by looking at the previous actions of the traders in the market, can infer the remainder of the equilibrium path, and thus the set Σ^l containing the true joint signal \tilde{s} .

Since the future equilibrium path is known at T_m , the belief of the traders will no longer be updated subsequent to time T_m . Thus, for all i , we have $\varphi_{i,t} = \varphi_{i,T_m}$ for all $t \geq T_m$, and we have a similar relation for φ_t . For convenience of notation, denote φ_{T_m} and φ_{i,T_m} by φ and φ_i respectively. Also, without loss of generality, assume that for all $\omega \in \Omega$, we have $\varphi(\tilde{\omega} = \omega) > 0$ (since otherwise we can restrict attention to only those states for which the condition holds). Define

$$\underline{w}_i^\delta(\omega) = \liminf_{T \rightarrow \infty} w_{i,T}^\delta(\omega), \quad \text{for } \omega \in \Omega. \quad (13)$$

Lemmas 3 and 4 below establish that in the setting of the theorem, for each trader i , $w_{i,t}^\delta \rightarrow \underline{w}_i^\delta$. As portfolio convergence holds, this implies through Theorem 1 that the market aggregates the information of the traders. \square

To complete the proof, we present the following two lemmas.

LEMMA 3. *In the setting of Theorem 2, for each trader i , and each $\omega \in \Omega$, we have almost surely under Q^δ :*

$$\begin{aligned} \liminf_{t \rightarrow \infty} w_{i,t}^\delta(\omega) &> -\infty; \text{ and} \\ \limsup_{t \rightarrow \infty} w_{i,t}^\delta(\omega) &< \infty. \end{aligned} \quad (14)$$

PROOF. Recall that we assume the utility of the trader is well defined in the PBE under consideration. Since not participating in the market from time T_m onwards is always an option for each trader, we must have:

$$\mathbf{E}_\varphi \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \mid \tilde{s}_i \right] \geq \mathbf{E}_\varphi \left[u_i(w_{i,T_m}^\delta(\tilde{\omega})) \mid \tilde{s}_i \right]. \quad (15)$$

Since this is true for all values of \tilde{s}_i possible after history h_{T_m} , taking expectations yields:

$$\mathbf{E}_\varphi \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \right] \geq \mathbf{E}_\varphi \left[u_i(w_{i,T_m}^\delta(\tilde{\omega})) \right]. \quad (16)$$

Since the right-hand side is finite and the utility functions are concave and strictly increasing, we must have $\underline{w}_i^\delta(\omega) > -\infty$, almost surely, for all $\omega \in \Omega$ and for all traders i .

Note that:

$$w_{i,t}^\delta(\omega) = \sum_{k=1}^n w_{k,t}^\delta(\omega) - \sum_{k \neq i} w_{k,t}^\delta(\omega), \quad (17)$$

for all $\omega \in \Omega$. Taking limits, we obtain that almost surely:

$$\begin{aligned} \limsup_{t \rightarrow \infty} w_{i,t}^\delta &\leq \limsup_{t \rightarrow \infty} \sum_{k=1}^n w_{k,t}^\delta(\omega) - \liminf_{t \rightarrow \infty} \sum_{k \neq i} w_{k,t}^\delta(\omega) \\ &\leq B^\delta - \sum_{k \neq i} \liminf_{t \rightarrow \infty} w_{k,t}^\delta(\omega) \\ &= B^\delta - \sum_{k \neq i} \underline{w}_k^\delta(\omega) \\ &< \infty. \end{aligned} \quad (18)$$

\square

LEMMA 4. *In the setting of Theorem 2, we have almost surely that $w_{i,t}^\delta(\omega) \rightarrow \underline{w}_i^\delta(\omega)$ for each trader i and for all $\omega \in \Omega$.*

PROOF. From Lemma 3, we observe that there exist $L, U \in \mathbb{R}$ with $L \leq U$, such that for all times t , for any trader i , and for any $\omega \in \Omega$, $w_{i,t}^\delta(\omega) \in [L, U]$ almost surely. Let $w_i^\delta(\omega) \in [L, U]$ be some accumulation point of $w_{i,t}^\delta(\omega)$ for $\omega \in \Omega$. By the definition of $\underline{w}_i^\delta(\omega)$, we observe that for all ω , $\underline{w}_i^\delta(\omega) \leq w_i^\delta(\omega)$. Since u_i is strictly increasing, we conclude:

$$\mathbf{E}_{Q^\delta} \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \mid h_{T_m}, \tilde{s}_i \right] \leq \mathbf{E}_{Q^\delta} \left[u_i(w_i^\delta(\tilde{\omega})) \mid h_{T_m}, \tilde{s}_i \right]. \quad (19)$$

The above inequality holds for all values of \tilde{s}_i that are possible after the history h_{T_m} . After taking expectations with respect to \tilde{s}_i , we obtain

$$\mathbf{E}_\varphi \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \right] \leq \mathbf{E}_\varphi \left[u_i(w_i^\delta(\tilde{\omega})) \right], \quad (20)$$

where the expectation is written in terms of $\varphi(\cdot) = \varphi_{T_m}(\cdot) = Q^\delta(\cdot | h_{T_m})$.

Let $\{t_{i,k} : k \geq 1\}$ be a subsequence of time periods where, for all $\omega \in \Omega$, $w_{i,t_{i,k}}^\delta(\omega) \rightarrow w_i^\delta(\omega)$ as $k \rightarrow \infty$. Moreover, as δ is a PBE strategy profile, if the trader i trades according to δ until time $t_{i,k}$ and subsequently stops trading in the market, then the utility obtained by the trader will be less than that corresponding to trading according to δ forever. This implies that

$$\mathbf{E}_{Q^\delta} \left[u_i(w_{i,t_{i,k}}^\delta(\tilde{\omega})) \mid h_{t_{i,k}}, \tilde{s}_i \right] \leq \mathbf{E}_{Q^\delta} \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \mid h_{t_{i,k}}, \tilde{s}_i \right]. \quad (21)$$

As δ is a pure strategy profile, from the discussion above we have $Q^\delta(\cdot | h_{t_{i,k}}, \tilde{s}_i) = Q^\delta(\cdot | h_{T_m}, \tilde{s}_i)$ for large enough k . Moreover, as the preceding inequality is true for all possible values of \tilde{s}_i , we obtain by taking expectations:

$$\mathbf{E}_\varphi \left[u_i(w_{i,t_{i,k}}^\delta(\tilde{\omega})) \right] \leq \mathbf{E}_\varphi \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \right]. \quad (22)$$

Observing that $\mathbf{E}_\varphi \left[u_i(w_{i,t_i,k}^\delta(\tilde{\omega})) \right] = \sum_{\omega \in \Omega} \varphi(\tilde{\omega} = \omega) u_i(w_{i,t_i,k}^\delta(\omega))$ security sold by the market maker, and not on the actual sequence of trades leading to the final position. by taking the limit as $k \rightarrow \infty$ we obtain:

$$\mathbf{E}_\varphi \left[u_i(w_i^\delta(\tilde{\omega})) \right] \leq \mathbf{E}_\varphi \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \right]. \quad (23)$$

Thus, from (20) and (23), we have:

$$\mathbf{E}_\varphi \left[u_i(w_i^\delta(\tilde{\omega})) \right] = \mathbf{E}_\varphi \left[u_i(\underline{w}_i^\delta(\tilde{\omega})) \right]. \quad (24)$$

Now recall that $\underline{w}_i^\delta(\omega) \leq w_i^\delta(\omega)$ and $\varphi(\tilde{\omega} = \omega) > 0$ for all $\omega \in \Omega$. Thus we must have $w_i^\delta(\omega) = \underline{w}_i^\delta(\omega)$ for all $\omega \in \Omega$. As this is true for all accumulation points w_i^δ , we conclude

$$w_{i,t}^\delta(\omega) \rightarrow \underline{w}_i^\delta(\omega) \quad \text{for all } \omega \in \Omega. \quad (25)$$

□

6. COST FUNCTION MARKET MAKER

In this section we focus our attention on a class of algorithmic market makers defined via *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, and of particular significance is the fact that market scoring rules can be reformulated as cost functions [1, 2]. 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. Our main observation is that such markets immediately satisfy the smoothness condition developed earlier in this paper for information aggregation. Further, in many of these markets, the loss to the market maker is bounded on any path of the market, and thus the information aggregation result of the preceding section applies if signal spaces are finite.

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 :

$$q_t = \sum_{\tau=1}^t y_\tau. \quad (26)$$

A cost function market maker is defined by a continuously differentiable function $C : \mathbb{R}^m \rightarrow \mathbb{R}$. After any history h_t , the market maker prices the trades at time t according to

$$K(h_t, z) = C(q_{t-1} + z) - C(q_{t-1}). \quad (27)$$

Note that the condition imposed by (1) on the pricing function K requires the cost function C to satisfy $C(q + a\mathbf{1}) = C(q) + a$, for all $a \in \mathbb{R}$.

Markets based on cost functions are a special case of the market studied in this paper. In particular, the market maker's pricing function depends on the history only through the total quantity of securities sold up to that time. Also, observe that the total revenue obtained by the market maker at time t is given by:

$$\begin{aligned} \sum_{\tau=1}^t K(h_\tau, y_\tau) &= \sum_{\tau=1}^t C(q_\tau) - C(q_{\tau-1}) \\ &= C(q_t) - C(\mathbf{0}), \end{aligned} \quad (28)$$

This shows that at any instant the total revenue of the market maker is dependent only on the total quantity of the

security sold by the market maker, and not on the actual sequence of trades leading to the final position.

Differentiability of the cost function then yields the following information aggregation result.

THEOREM 3. *Given a cost function market maker and a PBE with strategy profile δ satisfying portfolio convergence (Assumption 2), the market aggregates the information of the traders.*

PROOF. The proof involves showing that any PBE of a cost function market satisfies the asymptotic smoothness condition.

For a PBE with strategy profile δ , we first define $\Gamma(q_t^\delta) \triangleq q_t^\delta - (\sum_{\omega} q_t^\delta(\omega)) \mathbf{1}$, and note that, as $C(q + a\mathbf{1}) = C(q) + a$,

$$\begin{aligned} K(h_t, z) &= C(q_{t-1}^\delta + z) - \left(\sum_{\omega} q_t^\delta(\omega) \right) \mathbf{1} \\ &\quad - \left(C(q_{t-1}^\delta) - \left(\sum_{\omega} q_{t-1}^\delta(\omega) \right) \mathbf{1} \right) \\ &= C(\Gamma(q_{t-1}^\delta) + z) - C(\Gamma(q_{t-1}^\delta)). \end{aligned} \quad (29)$$

Also, 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 get

$$q_t = \sum_{i=1}^n (w_{i,t} - w_{i,0}) + \left(\sum_{\tau=1}^t K(h_\tau, y_\tau) \right) \mathbf{1}. \quad (30)$$

This implies that,

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

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 1 that information aggregation is obtained. □

The preceding result implies that for any algorithmic market maker using a cost function, if the traders' portfolios converge, then information is aggregated. It is important to note that in almost all cost function markets of interest, the loss of the market maker is bounded; and using the results in the previous section, we can conclude that if the PBE is pure and the signal space is finite, then information will be aggregated. Formally, we consider a class of cost function market makers satisfying the following condition.

ASSUMPTION 4.

$$\sup_{q \in \mathbb{R}^m} \left\{ \max_{1 \leq j \leq m} q(j) - C(q) \right\} < \infty.$$

This assumption is satisfied by a number of well-studied cost functions considered in the literature, and can be readily

verified (see, e.g., [1] for examples). For example, the assumption holds for the cost function C corresponding to the *logarithmic market scoring rule* [7], given by

$$C(q) = b \log \left(\sum_{\omega \in \Omega} \exp(q_j/b) \right), \quad (31)$$

for some $b > 0$. Similarly, the assumption holds true for the *quadratic market scoring rule* [2], given by

$$C(q) = \frac{1}{m} \sum_{\omega \in \Omega} q(\omega) + \frac{1}{4b} \sum_{\omega \in \Omega} q(\omega)^2 - \frac{1}{4mb} \left(\sum_{\omega \in \Omega} q(\omega) \right)^2 - \frac{b}{m}, \quad (32)$$

for some $b > 0$. Note that the assumption imposes a restriction on the model primitives, in particular the cost function, rather than on the equilibrium.

To see that Assumption 4 implies bounded loss (Assumption 3), note that, from (30), the loss of the market maker at time t is given by:

$$\begin{aligned} \sum_{i=1}^n (w_{i,t} - w_{i,0}) &= q_t - \left(\sum_{\tau=1}^t K(h_\tau, y_\tau) \right) \mathbf{1} \\ &= q_t - (C(q_t) - C(\mathbf{0})) \mathbf{1}. \end{aligned} \quad (33)$$

It follows from Assumption 4 that the loss of the market maker is bounded above uniformly for all t . This in turn implies that bounded loss holds in this case in any PBE. Thus, in a market where, additionally, the traders' signal space is finite, Theorem 3 applies, and thus *if signal spaces are finite, then in any pure strategy PBE of a cost function market satisfying Assumption 4, the market aggregates the information of the traders.*

7. INFORMATION CONTENT OF PRICES

We now study the information content of the limiting prices. Although the market aggregates information, it is not true in general that an observer seeing only the limiting prices has the same information as the traders. Under what conditions does such an observer have the same quality of information as the pooled signals of the traders?

We let $\phi \triangleq \nabla K(h_\infty, \mathbf{0})$ denote the eventual price vector in a PBE with strategy profile δ satisfying asymptotic smoothness (Assumption 1). Thus $\phi(\omega)$ for $\omega \in \Omega$ denotes *the asymptotic price charged by the market maker for selling an infinitesimal quantity of security ω* . It can be seen that the prices in general will not be equal to the posterior state distribution $\mathbf{P}(\tilde{\omega} = \omega | \tilde{s})$, simply because prices must also reflect the risk aversion of the traders. Thus just looking at the limiting prices will not inform an outside observer of the information contained in the pooled private signals of the traders.

However, in the special case when at least one of the traders in the market is *risk neutral*, the prices are equal to the posterior probability. A trader i is risk neutral if their utility function is linear, $u_i(x) = x$: in this case, the trader's marginal utility is independent of her wealth level. Note that in this case the first order conditions (11) for the risk-neutral trader i reduce to $\phi(\omega) = \varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)$, since $u'_i(x) = 1$. As the state $\tilde{\omega}$ and the private signal \tilde{s}_i are independent under φ_∞ , we obtain $\phi(\omega) = \varphi_\infty(\tilde{\omega} = \omega) = \mathbf{P}(\tilde{\omega} = \omega | \tilde{s})$. This

shows that in this case all the information of the traders is reflected in the price vector.

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, and in general if all traders are strictly risk averse, then an outside observer who knows only the final price cannot share the same posteriors as the traders. As a remedy, we identify a milder set of requirements under which an uninformed observer could identify the posterior distribution of $\tilde{\omega}$.

THEOREM 4. *Consider a PBE with strategy profile δ satisfying asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2). An uninformed observer with access to the final price $\nabla K(h_\infty, \mathbf{0})$ and the final portfolio of a single trader i with their utility function u_i can infer the posterior distribution φ_∞ of $\tilde{\omega}$.*

PROOF. Note that since \tilde{s}_i is independent of $\tilde{\omega}$ under the measure φ_∞ , the first order condition (11) can be rewritten as:

$$\frac{\varphi_\infty(\tilde{\omega} = \omega)}{\mathbf{E}_{\varphi_\infty} [u'_i(w_i^\delta(\tilde{\omega}))]} = \frac{\nabla_\omega K(h_\infty, \mathbf{0})}{u'_i(w_i^\delta(\omega))}. \quad (34)$$

For ease of notation, let $R_i(\omega) = \frac{\nabla_\omega K(h_\infty, \mathbf{0})}{u'_i(w_i^\delta(\omega))}$ for $\omega \in \Omega$. Then, we have for all $\omega \in \Omega$,

$$\varphi_\infty(\tilde{\omega} = \omega) = R_i(\omega) \cdot \mathbf{E}_{\varphi_\infty} [u'_i(w_i^\delta(\tilde{\omega}))]. \quad (35)$$

Let $U_i = (u'_i(w_i^\delta(1)), \dots, u'_i(w_i^\delta(m)))$. Then the preceding set of equations can be rewritten in vector notation as:

$$(I_m - R_i U_i) \varphi_\infty = \mathbf{0}, \quad (36)$$

where I_m is the $m \times m$ identity matrix, and we view R_i as a column vector, and U_i as a row vector.

The result then follows from observing that the matrix $I_m - R_i U_i$ has rank $m-1$, which implies that φ_∞ is uniquely determined as the probability vector in the null space of the aforementioned matrix. Thus, an uninformed observer with access to R_i and U_i can infer the common belief φ_∞ . \square

The fact that prices may not reflect the posterior is easiest to see when we observe that the limiting allocation and prices in the market constitute a competitive equilibrium of the economy where traders maximize their expected utility conditioned on \tilde{s} , with endowment w_i^δ . Consider an exchange economy among the traders where each trader shares the pooled information $\mathbf{P}(\cdot | \tilde{s})$, and trader i has an initial portfolio given by w_i^0 . A *competitive equilibrium* for this economy is specified by an allocation $\{w_i : i \leq n\}$ 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 $\{w_i : i \leq n\}$ should be the same as that in the initial allocation. Formally, we require

$$\sum_{i=1}^n w_i(\omega) = \sum_{i=1}^n w_i^0(\omega), \text{ for all } \omega \in \Omega. \quad (37)$$

2. *Utility maximization.* For each trader i , the allocation w_i of the security to trader should i maximize her expected utility (with respect to the pooled belief $\mathbf{P}(\cdot | \tilde{s})$),

given the price vector p . Formally, for each trader i , the allocation w_i is a solution to the following maximization problem.

$$\begin{aligned} & \max \mathbf{E}_{\mathbf{P}} [u_i(w_i(\tilde{\omega})) | \tilde{s}] \\ \text{subject to, } & \sum_{\omega \in \Omega} p(\omega) w_i(\omega) = \sum_{\omega \in \Omega} p(\omega) w_i^0(\omega). \end{aligned} \quad (38)$$

We connect competitive equilibrium and the limiting prices formally in the following theorem.

THEOREM 5. *Consider a PBE with strategy profile δ satisfying asymptotic smoothness (Assumption 1) and portfolio convergence (Assumption 2). Then the limiting portfolio $\{w_i^\delta : i \leq n\}$ and the price vector $\phi \triangleq \nabla K(h_\infty, \mathbf{0})$ together constitute a competitive equilibrium of an exchange economy with pooled information, where trader i 's initial portfolio is specified by w_i^δ .*

PROOF. Observe that the market clearance condition (37) holds trivially in this case, since the final and initial allocation are the same. So, in order for the allocation $\{w_i^\delta : i \leq n\}$ 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 (38) is a convex program, the sufficient conditions of the problem can be written as

$$\mathbf{P}(\tilde{\omega} = \omega | \tilde{s}) u'_i(w_i(\omega)) = \lambda p(\omega), \quad (39)$$

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

Now we know from Theorem 1 that the market aggregates the information of the traders. Thus we obtain $\mathbf{P}(\tilde{\omega} = \omega | \tilde{s}) = \varphi_\infty(\tilde{\omega} = \omega) = \varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i)$. Thus the sufficiency conditions become

$$\varphi_\infty(\tilde{\omega} = \omega | \tilde{s}_i) u'_i(w_i(\omega)) = \lambda p(\omega). \quad (40)$$

It is now easily seen from (11) that the sufficiency conditions are satisfied by the allocation $w_i = w_i^\delta$ for the price vector $p = \nabla K(h_\infty, \mathbf{0})$ by setting the Lagrange multiplier $\lambda = \mathbf{E}_{\varphi_\infty} [u'_i(w_i(\tilde{\omega}))]$. This implies that the allocation w_i^δ is utility maximizing for trader i with the price vector ϕ .

Thus we obtain that ϕ and $\{w_i^\delta : i \leq n\}$ together constitute a competitive equilibrium for the exchange economy described. \square

The first fundamental theorem of welfare economics states that the allocation in any competitive equilibrium is Pareto optimal. In particular, the above argument proves that the final portfolio of each trader is *ex post* Pareto optimal, in the sense that if the traders are eventually allowed to trade among themselves circumventing the market maker, they would not have any incentive to do so. In this respect, the market described above is efficient.

8. REFERENCES

- [1] Shipra Agrawal, Erick Delage, Mark Peters, Zizhuo Wang, and Yinyu Ye. A Unified Framework for Dynamic Pari-Mutuel Information Market Design. In *Proceedings of the Tenth ACM Conference on Electronic Commerce*, pages 255–264, 2009.
- [2] Yiling Chen and David M. Pennock. A Utility Framework for Bounded-Loss Market Makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 49–56, July 2007.

- [3] Yiling Chen, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. Bluffing and Strategic Reticence in Prediction Markets. In *International Workshop on Internet and Network Economics (WINE)*, pages 70–81, 2007.
- [4] Stanko Dimitrov and Rahul Sami. Non-myopic Strategies in Prediction Markets. In *Proceedings of the Ninth ACM Conference on Electronic Commerce*, pages 200–209, 2008.
- [5] Lawrence R. Glosten. Is the Electronic Open Limit Order Book Inevitable? *The Journal of Finance*, 49(4):1127–1161, September 1994.
- [6] Lawrence R. Glosten and Paul R. Milgrom. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics*, 14(1):71–100, March 1985.
- [7] Robin Hanson. Combinatorial Information Market Design. *Information Systems Frontiers*, 5(1):107–119, January 2003.
- [8] Larry Harris. *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press, USA, 2003.
- [9] F. A. Hayek. The Use of Knowledge in Society. *The American Economic Review*, 35(4):519–530, 1945.
- [10] Albert S. Kyle. Continuous Auctions and Insider Trading. *Econometrica*, 53(6):1315–1335, 1985.
- [11] Paul Milgrom and Nancy Stokey. Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26(1):17–27, February 1982.
- [12] Michael Ostrovsky. Information Aggregation in Dynamic Markets with Strategic Traders. In *Proceedings of the Tenth ACM Conference on Electronic Commerce*, pages 253–254, 2009.
- [13] David M. Pennock and Rahul Sami. Computational Aspects of Prediction Markets. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, pages 651–675. Cambridge University Press, Cambridge, UK, 2007.
- [14] Roy Radner. Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices. *Econometrica*, 47(3):655–678, 1979.

APPENDIX

A. PROOFS

PROOF OF LEMMA 2. The proof is by induction. For each i , let $f_i : \Sigma_i \rightarrow \mathbb{R}$ be any continuous bounded function. Since under the prior distribution Q^δ , conditional on the state $\tilde{\omega}$, the private signals $\tilde{s}_1, \dots, \tilde{s}_n$ are independent, we have:

$$\mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_1, \tilde{\omega} \right] = \prod_i \mathbf{E}_{Q^\delta} [f_i(\tilde{s}_i) | h_1, \tilde{\omega}].$$

Thus, we have the base case for $t = 1$. Assume that the claim holds for $t = T$. In particular, we have

$$\mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_T, \tilde{\omega} \right] = \prod_i \mathbf{E}_{Q^\delta} [f_i(\tilde{s}_i) | h_T, \tilde{\omega}].$$

Let trader j arrive at the market at time $T+1$, i.e. $i_{T+1} =$

j. Using tower property of conditional expectation, we have:

$$\begin{aligned}
& \mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_{T+1}, \tilde{\omega} \right] \\
&= \mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_T, y_T, \tilde{\omega} \right] \\
&= \mathbf{E}_{Q^\delta} \left[\mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_T, y_T, \tilde{\omega}, \tilde{s}_j \right] \middle| h_T, y_T, \tilde{\omega} \right] \\
&= \mathbf{E}_{Q^\delta} \left[\mathbf{E}_{Q^\delta} \left[\prod_{i \neq j} f_i(\tilde{s}_i) | h_T, y_T, \tilde{\omega}, \tilde{s}_j \right] f_j(\tilde{s}_j) \middle| h_T, y_T, \tilde{\omega} \right].
\end{aligned}$$

Note that as y_T depends only on \tilde{s}_j and h_T , we can write:

$$\begin{aligned}
& \mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_{T+1}, \tilde{\omega} \right] \\
&= \mathbf{E}_{Q^\delta} \left[\mathbf{E}_{Q^\delta} \left[\prod_{i \neq j} f_i(\tilde{s}_i) | h_T, \tilde{\omega}, \tilde{s}_j \right] f_j(\tilde{s}_j) \middle| h_T, y_T, \tilde{\omega} \right] \\
&= \mathbf{E}_{Q^\delta} \left[\prod_{i \neq j} \mathbf{E}_{Q^\delta} [f_i(\tilde{s}_i) | h_T, \tilde{\omega}] f_j(\tilde{s}_j) \middle| h_T, y_T, \tilde{\omega} \right],
\end{aligned}$$

where we have used the inductive hypothesis that conditional on h_T and $\tilde{\omega}$, the private signals are independent. As each expectation involved in the product on the right hand side are determined once h_T and $\tilde{\omega}$ are known, we obtain

$$\begin{aligned}
& \mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_{T+1}, \tilde{\omega} \right] \\
&= \prod_{i \neq j} \mathbf{E}_{Q^\delta} [f_i(\tilde{s}_i) | h_T, \tilde{\omega}] \mathbf{E}_{Q^\delta} [f_j(\tilde{s}_j) | h_{T+1}, \tilde{\omega}].
\end{aligned}$$

Finally, choosing f_k as the constant function $\mathbf{1}$ for all $k \neq i$, we see that $\mathbf{E}_{Q^\delta} [f_i(\tilde{s}_i) | h_T, \tilde{\omega}] = \mathbf{E}_{Q^\delta} [f_i(\tilde{s}_i) | h_{T+1}, \tilde{\omega}]$ for all $i \neq j$. Thus we have:

$$\mathbf{E}_{Q^\delta} \left[\prod_i f_i(\tilde{s}_i) | h_{T+1}, \tilde{\omega} \right] = \prod_i \mathbf{E}_{Q^\delta} [f_i(\tilde{s}_i) | h_{T+1}, \tilde{\omega}]. \quad (41)$$

As this is true for all continuous bounded functions f_i , we obtain that conditional on the history h_{T+1} and the state $\tilde{\omega}$, the private signals are independent. This completes the induction step.

To prove the result for $t = \infty$, we take the limit as $T \rightarrow \infty$ in (41), and apply Lemma 1. \square

We use the following Lemma in the proof of Theorem 1.

LEMMA 5. *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 continuous functions. As Y_1, \dots, Y_n are independent conditional on X , we have

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

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

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

We then reason as follows:

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

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

B. 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 which takes place in the market, given by $h_\infty = (y_1, \dots, y_t, \dots)$. 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 point-wise convergence, and let \mathcal{G}_∞ denote the Borel σ -algebra on H_∞ .

Let $(\Sigma_i, \mathcal{T}_i)$ 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 $\mathcal{S} = H_\infty \times \Omega \times \Sigma$, 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 w.r.t 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 a 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 a 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.