

Computation in a Distributed Information Market¹

Joan Feigenbaum^{a,2}

^a*Yale University, Department of Computer Science, New Haven, CT 06520*

Lance Fortnow^{b,3}

^b*Department of Computer Science, University of Chicago, 1100 E 58th Street, Chicago, IL 60637*

David M. Pennock^{c,3}

^c*Yahoo! Research Labs, 74 N. Pasadena Ave, 3rd floor, Pasadena, CA 91103*

Rahul Sami^{d,4}

^d*MIT CS & AI Laboratory, 200 Technology Square, Cambridge, MA 02139*

Abstract

According to economic theory—supported by empirical and laboratory evidence—the equilibrium price of a financial security reflects all of the information regarding the security’s value. We investigate the computational process on the path toward equilibrium, where information distributed among traders is revealed step-by-step over time and incorporated into the market price. We develop a simplified model of an information market, along with trading strategies, in order to formalize the computational properties of the process. We show that securities whose payoffs cannot be expressed as weighted threshold functions of distributed input bits are not guaranteed to converge to the proper equilibrium predicted by economic theory. On the other hand, securities whose payoffs are threshold functions *are* guaranteed to converge, for all prior probability distributions. Moreover, these *threshold securities* converge in at most n rounds, where n is the number of bits of distributed information. We also prove a lower bound, showing a type of threshold security that requires at least $n/2$ rounds to converge in the worst case.

Key words: Distributed information market, market computation, information aggregation, convergence to equilibrium, rational expectations, efficient markets hypothesis

1 Introduction

The strong form of the *efficient markets hypothesis* states that market prices nearly instantly incorporate all information available to all traders. As a result, market prices encode the best forecasts of future outcomes given all information, even if that information is distributed across many sources. Supporting evidence can be found in empirical studies of options markets [15], political stock markets [8,9,23], sports betting markets [3,10,28], horse-racing markets [31], market games [24,25], and laboratory investigations of experimental markets [7,26,27].

The process of information incorporation is, at its essence, a distributed computation. Each trader begins with his or her own information. As trades are made, summary information is revealed through market prices. Traders *learn* or infer what information others are likely to have by observing prices, then update their own beliefs based on their observations. Over time, if the process works as advertised, all information is revealed, and all traders converge to the same information state. At this point, the market is in what is called a *rational expectations equilibrium* [12,17,20]. All information available to all traders is now reflected in the going prices, and no further trades are desirable until some new information becomes available.

While most markets are not designed with information aggregation as a primary motivation—for example, derivatives markets are intended mainly for risk management and sports betting markets for entertainment—recently, some markets have been created solely for the purpose of aggregating information on a topic of interest. The Iowa Electronic Market⁵ is a prime example, operated by the University of Iowa Tippie College of Business for the purpose of investigating how information about political elections distributed among traders gets reflected in securities prices whose payoffs are tied to actual election outcomes [8,9].

Email addresses: feigenbaum@cs.yale.edu (Joan Feigenbaum),
fortnow@cs.uchicago.edu (Lance Fortnow), david.pennock@overture.com
(David M. Pennock), sami@theory.lcs.mit.edu (Rahul Sami).

¹ This work was supported by the DoD University Research Initiative (URI) administered by the Office of Naval Research under Grant N00014-01-1-0795. The results reported here were reported in preliminary form at the Fourth Annual ACM Conference on Electronic Commerce (EC '03) [6].

² Supported in part by ONR and NSF.

³ This work was done while the author was at NEC Research Institute, Princeton, NJ.

⁴ Supported by ONR and NSF. Most of this work was done while the author was at Yale University.

⁵ <http://www.biz.uiowa.edu/iem/>

In this paper, we investigate the nature of the computational process whereby distributed information is revealed and combined over time into the prices in information markets. To do so, in Section 3, we propose a model of an information market that is tractable for theoretical analysis and, we believe, captures much of the important essence of real information markets. In Section 4, we present our main theoretical results concerning this model. We prove that only Boolean securities whose payoffs can be expressed as *threshold functions* of the distributed input bits of information are guaranteed to converge as predicted by rational expectations theory. Boolean securities with more complex payoffs may not converge under some prior distributions. We also provide upper and lower bounds on the convergence time for these *threshold securities*. We show that, for all prior distributions, the price of a threshold security converges to its rational expectations equilibrium price in at most n rounds, where n is the number of bits of distributed information. We show that this worst-case bound is tight within a factor of two by illustrating a situation in which a threshold security requires $n/2$ rounds to converge.

2 Relationship to related work

As mentioned, there is a great deal of documented evidence supporting the notion that markets are able to aggregate information in a number of scenarios using a variety of market mechanisms. The theoretically ideal mechanism requires what is called a *complete market*. A complete market contains enough linearly independent securities to span the entire state space of interest [1,32]. That is, the dimensionality of the available securities equals the dimensionality of the event space over which information is to be aggregated.⁶ In this ideal case, all private information becomes common knowledge in equilibrium, and thus any function of the private information can be directly evaluated by any agent or observer. However, this theoretical ideal is almost never achievable in practice, because it generally requires a number of securities exponential in the number of random variables of interest.

When available securities form an *incomplete market* [18] in relation to the desired information space—as is usually the case—aggregation may be partial. Not all private information is revealed in equilibrium, and prices may not convey enough information to recover the complete joint probability distribution over all events. Still, it is generally assumed that aggregation does occur along the dimensions represented in the market; that is, prices do reflect a consistent projection of the entire joint distribution onto the smaller-dimensional space

⁶ When we refer to independence or dimensionality of securities, we mean the independence or dimensionality of the random variables on which the security payoffs are based.

spanned by securities. In this paper, we investigate cases in which even this partial aggregation fails. For example, even though there is enough private information to determine completely the price of a security in the market, the equilibrium price may in fact reveal no information at all! So characterizations of when a rational expectations equilibrium is fully revealing do not immediately apply to our problem. We are not asking whether all possible functions of private information can be evaluated, but whether a particular target function can be evaluated. We show that properties of the function itself play a major role, not just the relative dimensionalities of the information and security spaces.

Our second main contribution is examining the *dynamics* of information aggregation before equilibrium, in particular proving upper and lower bounds on the time to convergence in those cases in which aggregation succeeds.

Shoham and Tennenholtz [30] define a *rationaly computable function* as a function of agents' valuations (types) that can be computed by a market, assuming agents follow rational equilibrium strategies. The authors mainly consider auctions of goods as their basic mechanistic unit and examine the communication complexity involved in computing various functions of agents' valuations of goods. For example, they give auction mechanisms that can compute the maximum, minimum, and k th-highest of the agents' valuations of a single good using 1 , 1 , and $n - k + 1$ bits of communication, respectively. They also examine the potential tradeoff between communication complexity and revenue.

3 Model of an information market

To investigate the properties and limitations of the process whereby an information market converges toward its rational-expectations equilibrium, we formulate a representative model of the market. In designing the model, our goals were two-fold: (1) to make the model rich enough to be realistic and (2) to make the model simple enough to admit meaningful analysis. Any modeling decisions must trade off these two generally conflicting goals, and the decision process is as much an art as a science. Nonetheless, we believe that our model captures enough of the essence of real information markets to lend credence to the results that follow. In this section, we present our modeling assumptions and justifications in detail. Section 3.1 describes the initial information state of the system, Section 3.2 covers the market mechanism, and Section 3.3 presents the agents' strategies.

3.1 Initial information state

There are n agents (traders) in the system, each of whom is privy to one bit of information, denoted x_i . The vector of all n bits is denoted $\mathbf{x} = (x_1, x_2, \dots, x_n)$. In the initial state, each agent is aware only of her own bit of information. All agents have a common prior regarding the joint distribution of bits among agents, but none has any specific information about the actual value of bits held by others. Note that this *common-prior assumption*—typical in the economics literature—does *not* imply that all agents agree. To the contrary, because each agent has different information, the initial state of the system is in general a state of disagreement.

Example 1: Suppose there are two agents, each with a single bit of information. Then, the information vector is $\mathbf{x} = (x_1, x_2)$. Further, if the prior probability distribution is uniform, then each of the four possible values for \mathbf{x} will have a prior probability of $\frac{1}{4}$. Now, agent 1 observes x_1 ; say she observes $x_1 = 0$. Then, conditioned on this information, agent 1 believes $P((x_1, x_2) = (0, 0)) = P((x_1, x_2) = (0, 1)) = \frac{1}{2}$. Similarly, if agent 2 observes $x_2 = 1$, he will condition on this information to get posterior probabilities $P((x_1, x_2) = (0, 1)) = P((x_1, x_2) = (1, 1)) = \frac{1}{2}$. Thus, the agents share a common prior but have different posterior beliefs because of their different information. Also, neither agent knows the exact value of (x_1, x_2) , despite the fact that their combined information is sufficient to determine it. \square

3.2 Market mechanism

The security being traded by the agents is a financial instrument whose payoff is a function $f(\mathbf{x})$ of the agents' bits. The form of f (the description of the security) is common knowledge⁷ among agents. We sometimes refer to the x_i as the *input bits*. At some time in the future after trading is completed, the true value of $f(\mathbf{x})$ is revealed,⁸ and every owner of the security is paid an amount $f(\mathbf{x})$ in cash per unit owned. If an agent ends with a negative quantity of the security (by selling short), then the agent must pay the amount $f(\mathbf{x})$ in cash per unit. Note that if someone were to have complete knowledge of all input bits \mathbf{x} , then that person would know the true value $f(\mathbf{x})$ of the security with certainty, and so would be willing to buy it at any price lower than $f(\mathbf{x})$ and (short) sell it at any price higher than $f(\mathbf{x})$.⁹

⁷ “Common knowledge” is information that all agents know, that all agents know that all agents know, and so on *ad infinitum* [5].

⁸ The values of the input bits themselves may or may not be publicly revealed.

⁹ Throughout this paper we ignore the time value of money.

Following Dubey, Geanakoplos, and Shubik [4], and Jackson and Peck [14], we model the market-price formation process as a multiperiod *Shapley-Shubik market game* [29]. The Shapley-Shubik process operates as follows: The market proceeds in synchronous rounds. In each round, each agent i submits a *bid* b_i and a *quantity* q_i . The semantics are that agent i is supplying a quantity q_i of the security and an amount b_i of money to be traded in the market. For simplicity, we assume that there are no restrictions on credit or short sales, and so an agent's trade is not constrained by her possessions. The market clears in each round by settling at a single price that balances the trade in that round: The clearing price is $p = \sum_i b_i / \sum_i q_i$. At the end of the round, agent i holds a quantity q'_i proportional to the money she bid: $q'_i = b_i/p$. In addition, she is left with an amount of money b'_i that reflects her net trade at price p : $b'_i = b_i - p(q'_i - q_i) = pq_i$. Note that agent i 's net trade in the security is a purchase if $p < b_i/q_i$ and a sale if $p > b_i/q_i$.

After each round, the clearing price p is publicly revealed. Agents then revise their beliefs according to any information garnered from the new price. The next round proceeds as the previous. The process continues until an equilibrium is reached, meaning that prices and bids do not change from one round to the next.

In this paper, we make a further simplifying restriction on the trading in each round: We assume that $q_i = 1$ for each agent i . This modeling assumption serves two analytical purposes. First, it ensures that there is *forced trade* in every round. Classic results in economics show that perfectly rational and risk-neutral agents will never trade with each other for purely speculative reasons (even if they have differing information) [21]. There are many factors that can induce rational agents to trade, such as differing degrees of risk aversion, the presence of other traders who are trading for liquidity reasons rather than speculative gain, or a market maker who is pumping money into the market through a subsidy. We sidestep this issue by simply assuming that the informed agents will trade (for unspecified reasons). Second, forcing $q_i = 1$ for all i means that the total volume of trade and the impact of any one trader on the clearing price are common knowledge; the clearing price p is a simple function of the agents' bids, $p = \sum_i b_i/n$. We will discuss the implications of alternative market models in Section 5.

3.3 Agent strategies

In order to draw formal conclusions about the price evolution process, we need to make some assumptions about how agents behave. Essentially we assume

that agents are risk-neutral, myopic,¹⁰ and bid truthfully: Each agent in each round bids his or her current valuation of the security, which is that agent’s estimation of the expected payoff of the security. Expectations are computed according to each agent’s probability distribution, which is updated via Bayes’ rule. When new information (revealed via the clearing prices) becomes available. We also assume that it is common knowledge that all the agents behave in the specified manner.

Example 2: Consider the initial information set-up in Example 1, and introduce a security F based on the OR function $f(\mathbf{x}) = x_1 \vee x_2$, that is F eventually pays off \$1 if $f(\mathbf{x})$ is 1. Suppose agent 1 observed $x_1 = 0$ and agent 2 observed $x_2 = 1$. Then agent 1’s initial expectation of the value of F is 0.5; hence, in our model, she would bid $b_1 = 0.5$ in the first round of trading. On the other hand, agent 2 would know for certain that f is 1, and would bid $b_2 = 1$. The clearing price of the market after the first round would thus be 0.75. \square

Would rational agents actually behave according to this strategy? It’s hard to say. Certainly, we do not claim that this is an equilibrium strategy in the game-theoretic sense. Furthermore, it is clear that we are ignoring some legitimate tactics, for example bidding falsely in one round in order to effect other agents’ judgments in the following rounds (non-myopic reasoning). However, we believe that the strategy outlined is a reasonable starting point for analysis. Solving for a true game-theoretic equilibrium strategy in this setting seems extremely difficult. Our assumptions seem reasonable when there are enough agents in the system such that extremely complex meta-reasoning is not likely to improve upon simply bidding one’s true expected value. In this case, according to the Shapley-Shubik mechanism, if the clearing price is below an agent’s expected value that agent will end up buying (increasing expected profit); otherwise, if the clearing price is above the agent’s expected value, the agent will end up selling (also increasing expected profit).

4 Computational properties

In this section, we study the computational power of information markets for a class of aggregation functions: Boolean functions of n variables. We characterize the set of Boolean functions that can be computed in our market

¹⁰ Risk neutrality implies that each agent’s utility for the security is linearly related to his or her subjective estimation of the expected payoff of the security. Myopic behavior means that agents treat each round as if it were the final round: They do *not* reason about how their bids may affect the bids of other agents in future rounds.

model for all prior distributions and then prove upper and lower bounds on the worst-case convergence time for these markets.

The information structure we assume is as follows: There are n agents, and each agent i has a single bit of private information x_i . All the agents also have a common prior probability distribution $\mathcal{P} : \{0, 1\}^n \rightarrow [0, 1]$ over the values of \mathbf{x} . We define a Boolean aggregate function $f(\mathbf{x}) : \{0, 1\}^n \rightarrow \{0, 1\}$ that we would like the market to compute. Note that \mathbf{x} , and hence $f(\mathbf{x})$, is completely determined by the combination of all the agents' information, but it is not known to any one agent. The agents trade in a Boolean security F , which pays off \$1 if $f(\mathbf{x}) = 1$ and \$0 if $f(\mathbf{x}) = 0$. So an omniscient agent with access to all the agents' bits would know the true value of security F —either exactly \$1 or exactly \$0. In reality, risk-neutral agents with limited information will value F according to their expectation of its payoff, or $E_i[f(\mathbf{x})]$, where E_i is the expectation operator applied according to agent i 's probability distribution.

For any function f , trading in F may happen to converge to the true value of $f(\mathbf{x})$ by coincidence if the prior probability distribution is sufficiently degenerate. More interestingly, we would like to know for which functions f does the price of the security F *always* converge to $f(\mathbf{x})$ *for all* prior probability distributions \mathcal{P} .¹¹ In Section 4.2, we prove a necessary and sufficient condition that guarantees convergence. In Section 4.3, we address the natural follow-up question, by deriving upper and lower bounds on the worst-case number of rounds of trading required for the value of $f(\mathbf{x})$ to be revealed.

4.1 Equilibrium price characterization

Our analysis builds on a characterization of the equilibrium price of F that follows from a powerful result on common knowledge of aggregates due to McKelvey and Page [20], later extended by Nielsen et al. [22].

Information markets aim to aggregate the knowledge of all the agents. Procedurally, this occurs because the agents *learn* from the markets: The price of the security conveys information to each agent about the knowledge of other agents. We can model the flow of information through prices as follows.

Let $\Omega = \{0, 1\}^n$ be the set of possible values of \mathbf{x} ; we say that Ω denotes the set of possible “states of the world.” The prior \mathcal{P} defines everyone's initial belief about the likelihood of each state. As trading proceeds, some possible states can be logically ruled out, but the relative likelihoods among the remaining states are fully determined by the prior \mathcal{P} . So the common knowledge after any

¹¹ We assume that the common prior is *consistent* with \mathbf{x} in the sense that it assigns a non-zero probability to the actual value of \mathbf{x} .

stage is completely described by the set of states that an external observer—with no information beyond the sequence of prices observed—considers possible (along with the prior). Similarly, the knowledge of agent i at any point is also completely described by the set of states she considers possible. We use the notation S^r to denote the common-knowledge possibility set after round r , and S_i^r to denote the set of states that agent i considers possible after round r .

Initially, the only common knowledge is that the input vector \mathbf{x} is in Ω , where Ω is the set of states assigned a non-zero probability by the prior, generically all possible states. In other words, the set of states considered possible by an external observer before trading has occurred is the set $S^0 = \Omega$. However, each agent i also knows the value of her bit x_i ; thus, her knowledge set S_i^0 is the set $\{\mathbf{y} \in \Omega | y_i = x_i\}$. Agent i 's first-round bid is her conditional expectation of the event $f(\mathbf{x}) = 1$ given that $\mathbf{x} \in S_i^0$. All the agents' bids are processed, and the clearing price p^1 is announced. An external observer could predict agent i 's bid if he knew the value of x_i . Thus, if he knew the value of \mathbf{x} , he could predict the value of p^1 . In other words, the external observer knows the function $price^1(\mathbf{x})$ that relates the first round price to the true state \mathbf{x} . Of course, he does not know the value of \mathbf{x} ; however, he can *rule out any vector \mathbf{x} that would have resulted in a different clearing price from the observed price p^1* .

Thus, the common knowledge after round 1 is the set $S^1 = \{\mathbf{y} \in S^0 | price^1(\mathbf{y}) = p^1\}$. Agent i knows the common knowledge and, in addition, knows the value of bit x_i . Hence, after every round r , the knowledge of agent i is given by $S_i^r = \{\mathbf{y} \in S^r | y_i = x_i\}$. Note that, because knowledge can only improve over time, we must always have $S_i^r \subseteq S_i^{r-1}$ and $S^r \subseteq S^{r-1}$. Thus, only a finite number of changes in each agent's knowledge are possible, and so eventually we must converge to an equilibrium after which no player learns any further information. We use S^∞ to denote the common knowledge at this point, and S_i^∞ to denote agent i 's knowledge at this point. Let p^∞ denote the clearing price at equilibrium.

Informally, McKelvey and Page [20] show that, if n people with common priors but different information about the likelihood of some event A agree about a “suitable” aggregate of their individual conditional probabilities, then their individual conditional probabilities of event A 's occurring must be identical. (The precise definition of “suitable” is described below.) There is a strong connection to rational expectation equilibria in markets, which was noted in the original McKelvey-Page paper: The market price of a security is common knowledge at the point of equilibrium. Thus, if the price is a “suitable” aggregate of the conditional expectations of all the agents, then in equilibrium they must have identical conditional expectations of the event that the security will pay off. (Note that their information may still be different.)

Definition 1 (Bergin and Brandenburger [2]) A function $g : \mathfrak{R}^n \rightarrow \mathfrak{R}$ is called **stochastically monotone** if it can be written in the form $g(\mathbf{x}) = \sum_i g_i(x_i)$, where each function $g_i : \mathfrak{R} \rightarrow \mathfrak{R}$ is strictly increasing.

Bergin and Brandenburger [2] proved that this simple definition of stochastically monotone functions is equivalent to the original definition in McKelvey-Page [20].

Definition 2 (McKelvey and Page [20]) A function $g : \mathfrak{R}^n \rightarrow \mathfrak{R}$ is called **stochastically regular** if it can be written in the form $g = h \circ g'$, where g' is stochastically monotone and h is invertible on the range of g' .

We can now state the McKelvey-Page result, as generalized by Nielsen et al. [22]. In our context, the following simple theorem statement suffices; more general versions of this theorem can be found in [20,22].

Theorem 3 (Nielsen et al. [22]) Suppose that, at equilibrium, the n agents have a common prior, but possibly different information, about the value of a random variable F , as described above. For all i , let $p_i^\infty = E(F|\mathbf{x} \in S_i^\infty)$. If g is a stochastically regular function and $g(p_1^\infty, p_2^\infty, \dots, p_n^\infty)$ is common knowledge, then it must be the case that

$$p_1^\infty = p_2^\infty = \dots = p_n^\infty = E(F|\mathbf{x} \in S^\infty) = p^\infty$$

In one round of our simplified Shapley-Shubik trading model, the announced price is the mean of the conditional expectations of the n agents. The mean is a stochastically regular function; hence, Theorem 3 shows that, at equilibrium, all agents have identical conditional expectations of the payoff of the security. It follows that the equilibrium price p^∞ must be exactly the conditional expectations of all agents at equilibrium.

Theorem 3 does not in itself say *how* the equilibrium is reached. McKelvey and Page, extending an argument due to Geanakoplos and Polemarchakis [11], show that repeated announcement of the aggregate will eventually result in common knowledge of the aggregate. In our context, this is achieved by announcing the current price at the end of each round; this will ultimately converge to a state in which all agents bid the same price p^∞ .

However, reaching an equilibrium price is not sufficient for the purposes of information aggregation. We also want the price to reveal the actual value of $f(\mathbf{x})$. It is possible that the equilibrium price p^∞ of the security F will not be either 0 or 1, and so we cannot infer the value of $f(\mathbf{x})$ from it.

Example 3: Consider two agents 1 and 2 with private input bits x_1 and x_2 respectively. Suppose the prior probability distribution is uniform, *i.e.*, $\mathbf{x} =$

(x_1, x_2) takes the values $(0, 0)$, $(0, 1)$, $(1, 0)$, and $(1, 1)$ each with probability $\frac{1}{4}$. Now, suppose the aggregate function we want to compute is the *XOR* function, $f(\mathbf{x}) = x_1 \oplus x_2$. To this end, we design a market to trade in a Boolean security F , which will eventually payoff \$1 iff $x_1 \oplus x_2 = 1$.

If agent 1 observes $x_1 = 1$, she estimates the expected value of F to be the probability that $x_2 = 0$ (given $x_1 = 1$), which is $\frac{1}{2}$. If she observes $x_1 = 0$, her expectation of the value of F is the conditional probability that $x_2 = 1$, which is also $\frac{1}{2}$. Thus, in either case, agent 1 will bid 0.5 for F in the first round. Similarly, agent 2 will also always bid 0.5 in the first round. Hence, the first round of trading ends with a clearing price of 0.5. From this, agent 2 can infer that agent 1 bid 0.5, but this gives her no information about the value of x_1 —it is still equally likely to be 0 or 1. Agent 1 also gains no information from the first round of trading, and hence neither agent changes her bid in the following rounds. Thus, the market reaches equilibrium at this point. As predicted by Theorem 3, both agents have the same conditional expectation (0.5) at equilibrium. However, the equilibrium price of the security F does not reveal the value of $f(x_1, x_2)$, even though the combination of agents' information is enough to determine it precisely. \square

4.2 Characterizing computable aggregates

We now give a necessary and sufficient characterization of the class of functions f such that, for any prior distribution on \mathbf{x} , the equilibrium price of F will reveal the true value of f . We show that this is exactly the class of *weighted threshold functions*:

Definition 4 A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a **weighted threshold function** iff there are real constants w_1, w_2, \dots, w_n such that

$$f(\mathbf{x}) = 1 \quad \text{iff} \quad \sum_{i=1}^n w_i x_i \geq 1$$

Theorem 5 If f is a weighted threshold function, then, for any prior probability distribution \mathcal{P} , the equilibrium price of F is equal to $f(\mathbf{x})$.

Proof:

Let S_i^∞ denote the possibility set of agent i at equilibrium. As before, we use p^∞ to denote the final trading price at this point. Note that, by Theorem 3, p^∞ is exactly agent i 's conditional expectation of the value of $f(\mathbf{x})$, given her final possibility set S_i^∞ .

First, observe that if p^∞ is 0 or 1, then we must have $f(\mathbf{x}) = p^\infty$, regardless of the form of f . For instance, if $p^\infty = 1$, this means that $E(f(\mathbf{y})|\mathbf{y} \in S^\infty) = 1$. As $f(\cdot)$ can only take the values 0 or 1, it follows that $P(f(\mathbf{y}) = 1|\mathbf{y} \in S^\infty) = 1$. The actual value \mathbf{x} is always in the final possibility set S^∞ , and, furthermore, it must have non-zero prior probability, because it actually occurred. Hence, it follows that $f(\mathbf{x}) = 1$ in this case. An identical argument shows that if $p^\infty = 0$, $f(\mathbf{x}) = 0$.

Hence, it is enough to show that, if f is a weighted threshold function, then p^∞ is either 0 or 1. We prove this by contradiction. Let $f(\cdot)$ be a weighted threshold function corresponding to weights $\{w_i\}$, and assume that $0 < p^\infty < 1$. By Theorem 3, we must have:

$$P(f(\mathbf{y}) = 1|\mathbf{y} \in S^\infty) = p^\infty \quad (1)$$

$$\forall i \ P(f(\mathbf{y}) = 1|\mathbf{y} \in S_i^\infty) = p^\infty \quad (2)$$

Recall that $S_i^\infty = \{\mathbf{y} \in S^\infty | y_i = x_i\}$. Thus, Equation (2) can be written as

$$\forall i \ P(f(\mathbf{y}) = 1|\mathbf{y} \in S^\infty, y_i = x_i) = p^\infty \quad (3)$$

Now define

$$J_i^+ = P(y_i = 1|\mathbf{y} \in S^\infty, f(\mathbf{y}) = 1)$$

$$J_i^- = P(y_i = 1|\mathbf{y} \in S^\infty, f(\mathbf{y}) = 0)$$

$$J^+ = \sum_{i=1}^n w_i J_i^+$$

$$J^- = \sum_{i=1}^n w_i J_i^-$$

Because by assumption $p^\infty \neq 0, 1$, both J_i^+ and J_i^- are well-defined (for all i): Neither is conditioned on a zero-probability event.

Claim: Eqs. 1 and 3 imply that $J_i^+ = J_i^-$, for all i .

Proof of claim: We consider the two cases $x_i = 1$ and $x_i = 0$ separately.

Case (i): $x_i = 1$. We can assume that J_i^- and J_i^+ are not both 0 (or else, the claim is trivially true). In this case, we have

$$\begin{aligned} & \frac{P(f(\mathbf{y}) = 1|\mathbf{y} \in S^\infty) \cdot J_i^+}{P(f(\mathbf{y}) = 1|\mathbf{y} \in S^\infty) \cdot J_i^+ + P(f(\mathbf{y}) = 0|\mathbf{y} \in S^\infty) \cdot J_i^-} \\ & = P(f(\mathbf{y}) = 1|y_i = 1, \mathbf{y} \in S^\infty) \quad (\text{Bayes' law}) \end{aligned}$$

$$\begin{aligned} \frac{p^\infty J_i^+}{p^\infty J_i^+ + (1 - p^\infty) J_i^-} &= p^\infty \quad (\text{by Eqs. 1 and 3}) \\ J_i^+ &= p^\infty J_i^+ + (1 - p^\infty) J_i^- \\ \Rightarrow J_i^+ &= J_i^- \quad (\text{as } p^\infty \neq 1) \end{aligned}$$

Case (ii): $x_i = 0$. When $x_i = 0$, observe that the argument of Case (i) can be used to prove that $(1 - J_i^+) = (1 - J_i^-)$. It immediately follows that $J_i^+ = J_i^-$ as well. \square

Hence, we must also have $J^+ = J^-$. But using linearity of expectation, we can also write J^+ as

$$J^+ = E \left(\left[\sum_{i=1}^n w_i y_i \right] \middle| \mathbf{y} \in S^\infty, f(\mathbf{y}) = 1 \right),$$

and, because $f(\mathbf{y}) = 1$ only when $\sum_i w_i y_i \geq 1$, this gives us $J^+ \geq 1$. Similarly,

$$J^- = E \left(\left[\sum_{i=1}^n w_i y_i \right] \middle| \mathbf{y} \in S^\infty, f(\mathbf{y}) = 0 \right),$$

and thus $J^- < 1$. This implies $J^- \neq J^+$, which leads to a contradiction. \square

Perhaps surprisingly, the converse of Theorem 5 also holds:

Theorem 6 *Suppose $f : \{0, 1\}^n \rightarrow \{0, 1\}$ cannot be expressed as a weighted threshold function. Then there exists a prior distribution \mathcal{P} for which the price of the security F does not converge to the value of $f(\mathbf{x})$.*

Proof: We start from a geometric characterization of weighted threshold functions. Consider the Boolean hypercube $\{0, 1\}^n$ as a set of points in \mathfrak{R}^n . It is well known that f is expressible as a weighted threshold function iff there is a hyperplane in \mathfrak{R}^n that separates all the points at which f has value 0 from all the points at which f has value 1.

Now, consider the sets

$$H^+ = \text{Conv}(f^{-1}(1))$$

and

$$H^- = \text{Conv}(f^{-1}(0)),$$

where $\text{Conv}(S)$ denotes the convex hull of S in \mathfrak{R}^n . H^+ and H^- are convex sets in \mathfrak{R}^n , and so, if they do not intersect, we can find a separating hyperplane

between them. This means that, if f is *not* expressible as a weighted threshold function, H^+ and H^- must intersect. In this case, we show how to construct a prior \mathcal{P} for which $f(\mathbf{x})$ is not computed by the market.

Let $\mathbf{x}^* \in \mathfrak{R}^n$ be a point in $H^+ \cap H^-$. Because \mathbf{x}^* is in H^+ , there exists some points $\mathbf{z}^1, \mathbf{z}^2, \dots, \mathbf{z}^m$ and constants $\lambda_1, \lambda_2, \dots, \lambda_m$, such that the following constraints are satisfied:

$$\begin{aligned} \forall k \quad \mathbf{z}^k &\in \{0, 1\}^n, \text{ and } f(\mathbf{z}^k) = 1 \\ \forall k \quad 0 &< \lambda_k \leq 1 \\ \sum_{k=1}^m \lambda_k &= 1 \\ \sum_{k=1}^m \lambda_k \mathbf{z}^k &= \mathbf{x}^* \end{aligned}$$

Similarly, because $\mathbf{x}^* \in H^-$, there are points $\mathbf{y}^1, \mathbf{y}^2, \dots, \mathbf{y}^l$ and constants $\mu_1, \mu_2, \dots, \mu_l$, such that

$$\begin{aligned} \forall j \quad \mathbf{y}^j &\in \{0, 1\}^n, \text{ and } f(\mathbf{y}^j) = 0 \\ \forall j \quad 0 &< \mu_j \leq 1 \\ \sum_{j=1}^l \mu_j &= 1 \\ \sum_{j=1}^l \mu_j \mathbf{y}^j &= \mathbf{x}^* \end{aligned}$$

We now define our prior distribution \mathcal{P} as follows:

$$\begin{aligned} P(\mathbf{z}^k) &= \frac{\lambda_k}{2} \quad \text{for } k = 1, 2, \dots, m \\ P(\mathbf{y}^j) &= \frac{\mu_j}{2} \quad \text{for } j = 1, 2, \dots, l, \end{aligned}$$

and all other points are assigned probability 0. It is easy to see that this is a valid probability distribution. Under this distribution \mathcal{P} , first observe that $P(f(\mathbf{x}) = 1) = \frac{1}{2}$. Further, for any i such that $0 < x_i^* < 1$, we have

$$\begin{aligned} P(f(\mathbf{x}) = 1 | x_i = 1) &= \frac{P(f(\mathbf{x}) = 1 \wedge x_i = 1)}{P(x_i = 1)} \\ &= \frac{\frac{x_i^*}{2}}{x_i^*} \end{aligned}$$

$$= \frac{1}{2}$$

and

$$\begin{aligned} P(f(\mathbf{x}) = 1 | x_i = 0) &= \frac{P(f(\mathbf{x}) = 1 \wedge x_i = 0)}{P(x_i = 0)} \\ &= \frac{\frac{(1-x_i^*)}{2}}{(1-x_i^*)} \\ &= \frac{1}{2} \end{aligned}$$

For indices i such that x_i^* is 0 or 1 exactly, i 's private information reveals no additional information under prior \mathcal{P} , and so here too we have $P(f(\mathbf{x}) = 1 | x_i = 0) = P(f(\mathbf{x}) = 1 | x_i = 1) = \frac{1}{2}$.

Hence, *regardless of her private bit x_i* , each agent i will bid 0.5 for security F in the first round. The clearing price of 0.5 also reveals no additional information, and so this is an equilibrium with price $p^\infty = 0.5$ that does not reveal the value of $f(\mathbf{x})$. \square

The *XOR* function is one example of a function that cannot be expressed as weighted threshold function; Example 3 illustrates Theorem 6 for this function.

4.3 Convergence time bounds

We have shown that the class of Boolean functions computable in our model is the class of weighted threshold functions. The next natural question to ask is: How many rounds of trading are necessary before the equilibrium is reached? We analyze this problem using the same simplified Shapley-Shubik model of market clearing in each round. We first prove that, in the worst case, at most n rounds are required.

The idea of the proof is to consider the sequence of common knowledge sets $\Omega = S^0, S^1, \dots$, and show that, until the market reaches equilibrium, each set has a strictly lower dimension than the previous set.

Definition 7 *For a set $S \subseteq \{0, 1\}^n$, the **dimension** of set S is the dimension of the smallest affine subspace of \mathbb{R}^n that contains all the points in S ; we use the notation $\dim(S)$ to denote it.*

Lemma 1 *If $S^r \neq S^{r-1}$, then $\dim(S^r) < \dim(S^{r-1})$.*

Proof: Let $k = \dim(S^{r-1})$. Consider the bids in round r . In our model, agent

i will bid her current expectation for the value of F ,

$$b_i^r = E(f(\mathbf{y}) = 1 | \mathbf{y} \in S^{r-1}, y_i = x_i).$$

Thus, depending on the value of x_i , b_i^r will take on one of two values $h_i^{(0)}$ or $h_i^{(1)}$. Note that $h_i^{(0)}$ and $h_i^{(1)}$ depend only on the set S^{r-1} , which is common knowledge before round r . Setting $d_i = h_i^{(1)} - h_i^{(0)}$, we can write $b_i^r = h_i^{(0)} + d_i x_i$. It follows that the clearing price in round r is given by

$$p^r = \frac{1}{n} \sum_{i=1}^n (h_i^{(0)} + d_i x_i) \quad (4)$$

All the agents already know all the $h_i^{(0)}$ and d_i values, and they observe the price p^r at the end of the r th round. Thus, they effectively have a linear equation in x_1, x_2, \dots, x_n that they use to improve their knowledge by ruling out any possibility that would not have resulted in price p^r . In other words, after r rounds, the common knowledge set S^r is the intersection of S^{r-1} with the hyperplane defined by Equation (4).

It follows that S^r is contained in the intersection of this hyperplane with the k -dimension linear space containing S^{r-1} . If S^r is not equal to S^{r-1} , this intersection defines a linear subspace of dimension $(k - 1)$ that contains S^r , and hence S^r has dimension at most $(k - 1)$. \square

Theorem 8 *Let f be a weighted threshold function, and let \mathcal{P} be an arbitrary prior probability distribution. Then, after at most n rounds of trading, the price reaches its equilibrium value $p^\infty = f(\mathbf{x})$.*

Proof: Consider the sequence of common knowledge sets S^0, S^1, \dots , and let r be the minimum index such that $S^r = S^{r-1}$. Then, the r th round of trading does not improve any agent's knowledge, and thus we must have $S^\infty = S^{r-1}$ and $p^\infty = p^{r-1}$. Observing that $\dim(S^0) = n$, and applying Lemma 1 to the first $r - 1$ rounds, we must have $(r - 1) \leq n$. Thus, the price reaches its equilibrium value within n rounds. \square

Theorem 8 provides an upper bound of $O(n)$ on the number of rounds required for convergence. We now show that this bound is tight to within a factor of 2 by constructing a threshold function with $2n$ inputs and a prior distribution for which it takes n rounds to determine the value of $f(\mathbf{x})$ in the worst case.

The functions we use are the *carry-bit functions*. The function \mathcal{C}_n takes $2n$ inputs; for convenience, we write the inputs as $x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_n$ or as a pair (\mathbf{x}, \mathbf{y}) . The function value is the value of the high-order carry bit when the binary numbers $x_n x_{n-1} \dots x_1$ and $y_n y_{n-1} \dots y_1$ are added together.

In weighted threshold form, this can be written as

$$\mathcal{C}_n(\mathbf{x}, \mathbf{y}) = 1 \quad \text{iff} \quad \sum_{i=1}^n \frac{x_i + y_i}{2^{n+1-i}} \geq 1.$$

For this proof, let us call the agents $A_1, A_2, \dots, A_n, B_1, B_2, \dots, B_n$, where A_i holds input bit x_i , and B_i holds input bit y_i .

We first illustrate our technique by proving that computing \mathcal{C}_2 requires 2 rounds in the worst case. To do this, we construct a common prior \mathcal{P}_2 as follows:

- The pair (x_1, y_1) takes on the values $(0, 0), (0, 1), (1, 0), (1, 1)$ uniformly (*i.e.*, with probability $\frac{1}{4}$ each).
- We extend this to a distribution on (x_1, x_2, y_1, y_2) by specifying the conditional distribution of (x_2, y_2) given (x_1, y_1) : If $(x_1, y_1) = (1, 1)$, then (x_2, y_2) takes the values $(0, 0), (0, 1), (1, 0), (1, 1)$ with probabilities $\frac{1}{2}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}$ respectively. Otherwise, (x_2, y_2) takes the values $(0, 0), (0, 1), (1, 0), (1, 1)$ with probabilities $\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{2}$ respectively.

Now, suppose x_1 turns out to be 1, and consider agent A_1 's bid in the first round. It is given by

$$\begin{aligned} b_{A_1}^1 &= P(\mathcal{C}_2(x_1, x_2, y_1, y_2) = 1 | x_1 = 1) \\ &= P(y_1 = 1 | x_1 = 1) \cdot P((x_2, y_2) \neq (0, 0) | x_1 = 1, y_1 = 1) \\ &\quad + P(y_1 = 0 | x_1 = 1) \cdot P((x_2, y_2) = (1, 1) | x_1 = 1, y_1 = 0) \\ &= \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} \\ &= \frac{1}{2} \end{aligned}$$

On the other hand, if x_1 turns out to be 0, agent A_1 's bid would be given by

$$\begin{aligned} b_{A_1}^1 &= P(\mathcal{C}_2(x_1, x_2, y_1, y_2) = 1 | x_1 = 0) \\ &= P((x_2, y_2) = (1, 1) | x_1 = 0) \\ &= \frac{1}{2} \end{aligned}$$

Thus, irrespective of her bit, A_1 will bid 0.5 in the first round. Note that the function and distribution are symmetric between x and y , and so the same argument shows that B_1 will also bid 0.5 in the first round. Thus, the price p^1 announced at the end of the first round reveals no information about x_1 or

y_1 . The reason this occurs is that, under this distribution, the second carry bit \mathcal{C}_2 is statistically independent of the first carry bit ($x_1 \wedge y_1$); we will use this trick again in the general construction.

Now, suppose that (x_2, y_2) is either $(0, 1)$ or $(1, 0)$. Then, even if x_2 and y_2 are completely revealed by the first-round price, the value of $\mathcal{C}_2(x_1, x_2, y_1, y_2)$ is not revealed: It will be 1 if $x_1 = y_1 = 1$ and 0 otherwise. Thus, we have shown that at least 2 rounds of trading will be required to reveal the function value in this case.

We now extend this construction to show by induction that the function \mathcal{C}_n takes n rounds to reach an equilibrium in the worst case.

Theorem 9 *There is a function \mathcal{C}_n with $2n$ inputs and a prior distribution \mathcal{P}_n such that, in the worst case, the market takes n rounds to reveal the value of $\mathcal{C}_n(\cdot)$.*

Proof: We prove the theorem by induction on n . The base case for $n = 2$ has already been shown to be true. Starting from the distribution \mathcal{P}_2 described above, we construct the distributions $\mathcal{P}_3, \mathcal{P}_4, \dots, \mathcal{P}_n$ by inductively applying the following rule:

- Let \mathbf{x}^{-n} denote the vector $(x_1, x_2, \dots, x_{n-1})$, and define \mathbf{y}^{-n} similarly. We extend the distribution \mathcal{P}_{n-1} on $(\mathbf{x}^{-n}, \mathbf{y}^{-n})$ to a distribution \mathcal{P}_n on (\mathbf{x}, \mathbf{y}) by specifying the conditional distribution of (x_n, y_n) given $(\mathbf{x}^{-n}, \mathbf{y}^{-n})$: If $\mathcal{C}_{n-1}(\mathbf{x}^{-n}, \mathbf{y}^{-n}) = 1$, then (x_n, y_n) takes the values $(0, 0), (0, 1), (1, 0), (1, 1)$ with probabilities $\frac{1}{2}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}$ respectively. Otherwise, (x_n, y_n) takes the values $(0, 0), (0, 1), (1, 0), (1, 1)$ with probabilities $\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{2}$ respectively.

Claim: Under distribution \mathcal{P}_n , for all $i < n$,

$$P(\mathcal{C}_n(\mathbf{x}, \mathbf{y}) = 1 | x_i = 1) = P(\mathcal{C}_n(\mathbf{x}, \mathbf{y}) = 1 | x_i = 0).$$

Proof of claim: A similar calculation to that used for \mathcal{C}_2 above shows that the value of $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$ under this distribution is statistically independent of $\mathcal{C}_{n-1}(\mathbf{x}^{-n}, \mathbf{y}^{-n})$. For $i < n$, x_i can affect the value of \mathcal{C}_n only through \mathcal{C}_{n-1} . Also, by construction of \mathcal{P}_n , given the value of \mathcal{C}_{n-1} , the distribution of \mathcal{C}_n is independent of x_i . It follows that $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$ is statistically independent of x_i as well. Of course, a similar result holds for y_i by symmetry.

Thus, in the first round, for all $i = 1, 2, \dots, n-1$, the bids of agents A_i and B_i do not reveal anything about their private information. Thus, the first-round price does not reveal any information about the value of $(\mathbf{x}^{-n}, \mathbf{y}^{-n})$.

On the other hand, agents A_n and B_n do have different expectations of $\mathcal{C}_n(\mathbf{x})$ depending on whether their input bit is a 0 or a 1; thus, the first-round price

does reveal whether neither, one, or both of x_n and y_n are 1. Now, consider a situation in which (x_n, y_n) takes on the value $(1, 0)$ or $(0, 1)$. We show that, in this case, after one round we are left with the residual problem of computing the value of $\mathcal{C}_{n-1}(\mathbf{x}^{-n}, \mathbf{y}^{-n})$ under the prior \mathcal{P}_{n-1} .

Clearly, when $x_n + y_n = 1$, $\mathcal{C}_n(\mathbf{x}, \mathbf{y}) = \mathcal{C}_{n-1}(\mathbf{x}^{-n}, \mathbf{y}^{-n})$. Further, according to the construction of \mathcal{P}_n , the event $(x_n + y_n = 1)$ has the same probability $(1/3)$ for all values of $(\mathbf{x}^{-n}, \mathbf{y}^{-n})$. Thus, conditioning on this fact does not alter the probability distribution over $(\mathbf{x}^{-n}, \mathbf{y}^{-n})$; it must still be \mathcal{P}_{n-1} .

Finally, the inductive assumption tells us that solving this residual problem will take at least $n - 1$ more rounds in the worst case and hence that finding the value of $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$ takes at least n rounds in the worst case. \square

5 Discussion

Our results have been derived in a simplified model of an information market. In this section, we discuss the applicability of these results to more general trading models.

Assuming that agents bid truthfully, Theorem 5 holds in any model in which the price is a *known* stochastically monotone aggregate of agents' bids. While it seems reasonable that the market price satisfies monotonicity properties, the exact form of the aggregate function may not be known if the volume of each user's trades is not observable; this depends on the details of the market process. Theorem 6 and Theorem 9 hold more generally; they only require that an agent's strategy depends only on her conditional expectation of the security's value. Perhaps the most fragile result is Theorem 8, which relies on the linear form of the Shapley-Shubik clearing price (in addition to the conditions for Theorem 5); however, it seems plausible that a similar dimension-based bound will hold for other families of non-linear clearing prices.

Up to this point, we have described the model with the same number of agents as bits of information. However, all the results hold even if there is competition in the form of a known number of agents who know each bit of information. Indeed, modeling such competition may help alleviate the strategic problems in our current model.

Another interesting approach to addressing the strategic issue is to consider alternative markets that are at least myopically incentive compatible. One example is a market mechanism called a *market scoring rule*, suggested by Hanson [13]. These markets have the property that a risk-neutral agent's best myopic strategy is to truthfully bid her current expected value of the secu-

rity. Additionally, the number of securities involved in each trade is fixed and publicly known. If the market structure is such that, for example, the current scoring rule is posted publicly after each agent's trade, then in equilibrium there is common knowledge of all agents' expectation, and hence Theorem 5 holds. Theorem 6 also applies in this case, and hence we have the same characterization for the set of computable Boolean functions. This suggests that the problem of eliciting truthful responses may be orthogonal to the problem of computing the desired aggregate, reminiscent of the revelation principle [19].

In this paper, we have restricted our attention to one class of aggregation functions: Boolean functions of Boolean inputs. The proofs of Theorems 6 and 9 also hold if we consider Boolean functions of real inputs, where each agent's private information is a real number. Further, Theorem 5 also holds *provided the market reaches equilibrium*. With real inputs and arbitrary prior distributions, however, it is not clear that the market will reach an equilibrium in a finite number of steps.

6 Conclusion

6.1 Summary

We have framed the process of information aggregation in markets as a computation on distributed information. We have developed a simplified model of an information market that we believe captures many of the important aspects of real agent interaction in an information market. Within this model, we prove several results characterizing precisely what the market can compute and how quickly. Specifically, we show that the market is guaranteed to converge to the true rational expectations equilibrium if and only if the security payoff function is a weighted threshold function. We prove that the process whereby agents reveal their information over time and learn from the resulting announced prices takes at most n rounds to converge to the correct full-information price in the worst case. We show that this bound is tight within a factor of two.

Some interesting and important next steps include gaining a better understanding of the following:

- *The effect of price accuracy and precision:* We have assumed that the clearing price is known with unlimited precision; in practice, this will not be true. Further, we have neglected influences on the market price other than from rational traders; the market price may also be influenced by other factors such as misinformed or irrational traders. It is interesting to ask what aggregates can be computed even in the presence of noisy prices.
- *Incremental updates:* If the agents have computed the value of the function and a small number of input bits are switched, can the new value of the function be computed incrementally and quickly?
- *Distributed computation:* In our model, distributed information is aggregated through a centralized market computation. In a sense, some of the computation itself is distributed among the participating agents, but can the market computation also be distributed? For example, can we find a good distributed-computational model of a decentralized market?
- *Agents' computation:* We have not accounted for the complexity of the computations that agents must do to accurately update their beliefs after each round.
- *Strategic market models:* For reasons of simplicity and tractability, we have directly assumed that agents bid truthfully. A more satisfying approach would be to assume only rationality and solve for the resulting game-theoretic solution strategy, either in our current computational model or another model of an information market.
- *The common-prior assumption:* Can we say anything about the market behavior when agents' priors are only approximately the same or when they differ greatly?
- *Average-case analysis:* Our negative results (Theorems 6 and 9) examine worst-case scenarios, and thus involve very specific prior probability distributions and initial information states. It is interesting to ask whether we would get very different results for generic prior distributions or expected information states.
- *Information market design:* Non-threshold functions can be implemented by layering two or more threshold functions together. What is the minimum number of threshold securities required to implement a given function? This is exactly the problem of minimizing the size of a neural network, a well-studied problem known to be NP-hard [16]. What configuration of securities can best approximate a given function? Are there ways to define and configure securities to speed up convergence to equilibrium? What is the relationship between machine learning (*e.g.*, neural-network learning) and information-market design?

We thank Joe Kilian for many helpful discussions. We thank Robin Hanson, Varsha Dani, the anonymous reviewers for the Fourth ACM Conference on Electronic Commerce, and the anonymous referees of this journal for useful insights, pointers, and suggestions.

References

- [1] K. J. Arrow. The role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies*, 31(2):91–96, 1964.
- [2] J. Bergin and A. Brandenburger. A simple characterization of stochastically monotone functions. *Econometrica*, 58(5):1241–1243, Sept. 1990.
- [3] S. Debnath, D. M. Pennock, S. Lawrence, E. J. Glover, and C. L. Giles. Information incorporation in online in-game sports betting markets. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce (EC’03)*, pages 258–259, June 2003.
- [4] P. Dubey, J. Geanakoplos, and M. Shubik. The revelation of information in strategic market games: A critique of rational expectations equilibrium. *Journal of Mathematical Economics*, 16:105–137, 1987.
- [5] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, Cambridge, MA, 1996.
- [6] J. Feigenbaum, L. Fortnow, D. M. Pennock, and R. Sami. Computation in a distributed information market. In *Proceedings of the Fourth Annual ACM Conference on Electronic Commerce (EC ’03)*, pages 156–165, June 2003.
- [7] R. Forsythe and R. Lundholm. Information aggregation in an experimental market. *Econometrica*, 58(2):309–347, 1990.
- [8] R. Forsythe, F. Nelson, G. R. Neumann, and J. Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.
- [9] R. Forsythe, T. A. Rietz, and T. W. Ross. Wishes, expectations, and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39:83–110, 1999.
- [10] J. M. Gandar, W. H. Dare, C. R. Brown, and R. A. Zuber. Informed traders and price variations in the betting market for professional basketball games. *Journal of Finance*, LIII(1):385–401, 1998.
- [11] J. Geanakoplos and H. Polemarchakis. We can’t disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.

- [12] S. J. Grossman. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies*, 48(4):541–559, 1981.
- [13] R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1), 2002.
- [14] M. Jackson and J. Peck. Asymmetric information in a strategic market game: Reexamining the implications of rational expectations. *Economic Theory*, 13:603–628, 1999.
- [15] J. C. Jackwerth and M. Rubinstein. Recovering probability distributions from options prices. *Journal of Finance*, 51(5):1611–1631, Dec. 1996.
- [16] J.-H. Lin and J. S. Vitter. Complexity results on learning by neural nets. *Machine Learning*, 6:211–230, 1991.
- [17] R. E. Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–24, 1972.
- [18] M. Magill and M. Quinzii. *Theory of Incomplete Markets, Vol. 1*. MIT Press, 1996.
- [19] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [20] R. D. McKelvey and T. Page. Common knowledge, consensus, and aggregate information. *Econometrica*, 54(1):109–127, 1986.
- [21] P. Milgrom and N. Stokey. Information, trade, and common knowledge. *Journal of Economic Theory*, 26:17–27, 1982.
- [22] L. T. Nielsen, A. Brandenburger, J. Geanakoplos, R. McKelvey, and T. Page. Common knowledge of an aggregate of expectations. *Econometrica*, 58(5):1235–1238, 1990.
- [23] D. M. Pennock, S. Debnath, E. J. Glover, and C. L. Giles. Modeling information incorporation in markets, with application to detecting and explaining events. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, pages 405–413, 2002.
- [24] D. M. Pennock, S. Lawrence, C. L. Giles, and F. Å. Nielsen. The real power of artificial markets. *Science*, 291:987–988, February 9 2001.
- [25] D. M. Pennock, S. Lawrence, F. Å. Nielsen, and C. L. Giles. Extracting collective probabilistic forecasts from web games. In *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 174–183, 2001.
- [26] C. R. Plott and S. Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988.

- [27] C. R. Plott, J. Wit, and W. C. Yang. Parimutuel betting markets as information aggregation devices: Experimental results. Technical Report Social Science Working Paper 986, California Institute of Technology, Apr. 1997.
- [28] C. Schmidt and A. Werwatz. How accurate do markets predict the outcome of an event? the Euro 2000 soccer championships experiment. Technical Report 09-2002, Max Planck Institute for Research into Economic Systems, 2002.
- [29] L. Shapley and M. Shubik. Trade using one commodity as a means of payment. *Journal of Political Economy*, 85:937–968, 1977.
- [30] Y. Shoham and M. Tennenholtz. Rational computation and the communication complexity of auctions. *Games and Economic Behavior*, 35(1–2):197–211, 2001.
- [31] R. H. Thaler and W. T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.
- [32] H. R. Varian. The arbitrage principle in financial economics. *J. Economic Perspectives*, 1(2):55–72, 1987.