



Predicting the Future

Kay-Yut Chen, Leslie R. Fine and Bernardo A. Huberman

HP Laboratories, Palo Alto, CA 94304, USA

Abstract. We present a novel methodology for predicting future outcomes that uses small numbers of individuals participating in an imperfect information market. By determining their risk attitudes and performing a nonlinear aggregation of their predictions, we are able to assess the probability of the future outcome of an uncertain event and compare it to both the objective probability of its occurrence and the performance of the market as a whole. Experiments show that this nonlinear aggregation mechanism vastly outperforms both the imperfect market and the best of the participants. We then extend the mechanism to prove robust in the presence of public information.

Key Words. information aggregation, information markets, public information, experimental economics, mechanism design

1. Introduction

The prediction of the future outcomes of uncertain situations is both an important problem and a guiding force behind the search for the regularities that underlie natural and social phenomena. While in the physical and biological sciences the discovery of strong laws has enabled the prediction of future scenarios with uncanny accuracy, in the social sphere no such accurate laws are known. To complicate matters further, in social groups the information relevant to predictions is often dispersed across people, making it hard to identify and aggregate it. Thus, while several methods are presently used in forecasting, ranging from committees and expert consultants to aggregation techniques such as the Delphi method (Anderson and Holt, 1997), the results obtained suffer in terms of accuracy and ease of implementation.

In this paper, we propose and experimentally verify a market-based method to aggregate scattered information so as to produce reliable forecasts of uncertain events. This method is based on the belief shared by most economists that markets efficiently collect and disseminate information (Hayek, 1945). In particular,

rational expectations theory tells us that markets have the capacity not only to aggregate information held by individuals, but also to convey it via the price and volume of assets associated with that information. Therefore, a possible methodology for the prediction of future outcomes is the construction of markets where the asset is information rather than a physical good. Laboratory experiments have determined that these markets do indeed have the capacity to aggregate information in this type of setting (Forsythe, Palfrey, and Plott, 1982; O'Brien and Srivastava, 1991; Plott and Sunder, 1982, 1988).

Information markets generally involve the trading of state-contingent securities. If these markets are large enough and properly designed, they can be more accurate than other techniques for extracting diffuse information, such as surveys and opinions polls. There are problems however, with information markets, as they tend to suffer from information traps (Camerer and Weigelt, 1991; Nöth, et al., 1999), illiquidity (Sunder, 1992), manipulation (Forsythe and Lundholm, 1990; Nöth and Weber, 1998), and lack of equilibrium (Anderson and Holt, 1997; Scharfstein and Stein, 1990).¹ These problems are exacerbated when the groups involved are small and not very experienced at playing in these markets. Even when possible, proper market design is very expensive, fragile, and context-specific.

In spite of these obstacles, it is worth noting that certain participants in information markets can have either superior knowledge of the information being sought, or are better processors of the knowledge harnessed by the information market itself. By keeping track of the profits and final holdings of the members, one can determine which participants have these talents, along with their risk attitudes.

In this paper, we propose a method of harnessing the distributed knowledge of a group of individuals by

using a two-stage mechanism. In the first stage, an information market is run among members of the group in order to extract risk attitudes from the participants, as well as their ability at predicting a given outcome. This information is used to construct a nonlinear aggregation function that allows for collective predictions of uncertain events. In the second stage, individuals are simply asked to provide forecasts about an uncertain event, and they are rewarded according to the accuracy of their forecasts. These individual forecasts are aggregated using the nonlinear function and used to predict the outcome. As we show empirically, this nonlinear aggregation mechanism vastly outperforms both the imperfect market and the best of the participants.

However, these results are achieved in a very particular environment, that of no public information. Public information is bound to introduce strong correlations in the knowledge possessed by members of the group, correlations that are not explicitly taken into account by the above-described aggregation algorithm. So, we propose a set of suitable modifications that would allow the detection of the amount of public information present in a group so as to subtract it. Assuming that subjects can differentiate between the public and private information they hold, that the private aspect of their information is truly private (held only by one individual), and that the public information is truly public (held by at least two individuals), we create a coordination variant of the mechanism which allows for the identification of public information within a group and its subtraction when aggregating individual predictions about uncertain outcomes. Experiments in the laboratory show that this aggregation mechanism outperforms the market, the best player in the group, and the initially proposed aggregation mechanism.

2. *Aggregation Mechanism Design*

We consider first an environment in which a set of N people have purely private information about a future event. If all players had the same amount of information about the event and were perfectly risk-neutral, then it would be easy to compute the true posterior probabilities using Bayes' rule. If individuals receive independent information conditioned on the true outcome, their prior beliefs are uniform (no other information is available other than the event sequence), and they each report the true posterior probabilities given

their information, then the probability of an outcome s , conditioned on all of their observed information I , is given by:

$$P(s | I) = \frac{p_{s_1} p_{s_2} \cdots p_{s_N}}{\sum_{\forall s} p_{s_1} p_{s_2} \cdots p_{s_N}} \quad (1)$$

where p_{s_i} is the probability that individual i ($i = 1, \dots, N$) assigns to outcome s (please see Appendix 1 for a discussion). This result allows us simply to take the individual predictions, multiply them together, and normalize them in order to get an aggregate probability distribution. However, this will only work under the extremely restrictive constraints enumerated above. The first of these issues we will consider is how to design a mechanism that elicits truthful reporting from individuals. We demonstrate in Appendix 2 that the following mechanism will induce risk neutral utility maximizing individuals to report their prior probabilities truthfully. We ask each player to report a vector of perceived state-probabilities, $\{q_1, q_2, \dots, q_N\}$ with the constraint that the vector sums to one. Then the true state x is revealed and each player paid $c_1 + c_2 \times \log(q_x)$, where c_1 and c_2 are positive numbers.

While this very simple method might seem to aggregate dispersed information well, it suffers from the fact that, due to their risk attitude, most individuals do not necessarily report their true posterior probabilities conditioned on their information. In most realistic situations, a risk averse person will report a probability distribution that is flatter than her true beliefs as she tends to spread her bets among all possible outcomes. In the extreme case of risk aversion, an individual will report a uniform probability distribution regardless of her information. In this case, no predictive information is revealed by her report. Conversely, a risk-loving individual will tend to report a probability distribution that is more sharply peaked around a particular prediction, and in the extreme case of risk loving behavior a subject's optimal response will be to put all his weight on the most probable state according to his observations. In this case, his report will contain some, but not all the information contained in his observations.

In order to account for both the diverse levels of risk aversion and information strengths, we add a stage to the mechanism. Before individuals are asked to report their beliefs, they participate in an information market designed to elicit their risk attitudes and other relevant behavioral information. This information market

is driven by the same information structure in the reporting game. We use information markets to capture the behavioral information that is needed to derive the correct aggregation function. Note that, although the participant pool is too small for the market to act perfectly efficiently, it is a powerful enough mechanism to help us illicit the needed information.

The nonlinear aggregation function that we constructed is of the form:

$$P(s | I) = \frac{p_{s_1}^{\beta_1} p_{s_2}^{\beta_2} \cdots p_{s_N}^{\beta_N}}{\sum_{\forall s} p_{s_1}^{\beta_1} p_{s_2}^{\beta_2} \cdots p_{s_N}^{\beta_N}} \quad (2)$$

where β_i is the exponent assigned to individual i . The role of β_i is to help recover the true posterior probabilities from individual i 's report. The value of β for a risk neutral individual is one, as he should report the true probabilities coming out of his information. For a risk averse individual, β_i is greater than one so as to compensate for the flat distribution that he reports. The reverse, namely β_i smaller than one, applies to risk loving individuals. In terms of both the market performance and the individual holdings and risk behavior, a simple functional form for β_i is given by

$$\beta_i = r \left(\frac{V_i}{\sigma_i} \right)^c \quad (3)$$

where r is a parameter that captures the risk attitude of the whole market and is reflected in the market prices of the assets, V_i is the utility of individual i , and σ_i is the variance of his holdings over time. We use c as a normalization factor so that if $r = 1$, $\sum \beta_i$ equals the number of individuals. Thus the problem lies in the actual determination of both the risk attitudes of the market as a whole and of the individual players.

To do so, notice that if the market is perfectly efficient then the sum of the prices of the securities should be exactly equal to the payoff of the winning security. However, in the thin markets characterized here, this efficiency condition was rarely met. Moreover, although prices that do not sum to the winning payoff indicate an arbitrage opportunity, it was rarely possible to realize this opportunity with a portfolio purchase (once again, due to the thinness of the market). However, we can use these facts to our advantage. If the sum of the prices is below the winning payoff, then we can infer that the market is risk-averse, while if the price is above this payoff then the market exhibits risk-loving behavior. Thus, the ratio of the winning payoff to the sum of the prices provides a proxy for the risk attitude of the market as a whole.

The ratio of value to risk, (V_i/σ_i) , captures individual risk attitudes and predictive power. An individual's value V_i is given by the market prices multiplied by his holdings, summed over all the securities. As in portfolio theory (Markowitz, 1959), his amount of risk can be measured by the variance of his values using normalized market prices as probabilities of the possible outcomes.

3. Experimental Design for Private Information Experiments

In order to test this mechanism we conducted a number of experiments at Hewlett-Packard Laboratories, in Palo Alto, California. The subjects were undergraduate and graduate students at Stanford University and knew the experimental parameters discussed below, as they were part of the instructions and training for the sessions. The five sessions were run with eight to thirteen subjects in each.

We implemented the two-stage mechanism in a laboratory setting. Possible outcomes were referred to as "states" in the experiments. There were 10 possible states, A through J, in all the experiments. Each had an Arrow-Debreu² state security associated with it. The information available to the subjects consisted of observed sets of random draws from an urn with replacement. After privately drawing the state for the ensuing period, we filled the urn with one ball for each state, plus an additional two balls for the just-drawn true state security. Thus it is slightly more likely to observe a ball for the true state than others.

We allowed subjects to observe different number of draws from the urn in order to control the amount of information given to the subjects. Three types of information structures were used to ensure that the results obtained were robust. In the first treatment, each subject received three draws from the urn, with replacement. In the second treatment, half of the subjects received five draws with replacement, and the other half received one. In a third treatment, half of the subjects received a random number of draws (averaging three, and also set such that the total number of draws in the community was $3N$) and the other half received three, again with replacement.

The information market we constructed consists of an artificial call market in which the securities are traded. The states were equally likely and randomly

drawn. If a state occurred, the associated state security paid off at a value of 1,000 francs.³ Hence, the expected value of any given security, *a priori*, was 100 francs. Subjects were provided with some securities and francs at the beginning of each period.

Each period consists of six rounds, lasting 90 seconds each. At the end of each round, our system gathers the bids and asks and determines market price and volume. The transactions are then completed and another call round began. At the end of six trading rounds the period is over, the true state security is revealed, and subjects are paid according to the holdings of that security. This procedure is then repeated in the next period, with no correlation between the states drawn in each period.

In the second-stage, every subject played under the same information *structure* as in the first stage, although the draws and the true states were independent from those in the first. Each period they received their draws from the urn and 100 tickets. They were asked to distribute these tickets across the 10 states with the constraint that all 100 tickets must be spent each period and that at least one ticket is spent on each state. Since the fraction of tickets spent determines p_{si} , this

implies that p_{si} is never zero. The subjects were given a chart that told them how many francs they would earn upon the realization of the true state as a function of the number of tickets spent on the true state security. The payoff is a linear function of the log of the percentage of tickets placed in the winning state (Please see Appendix 2 for a discussion of the payoff function). The chart the subjects received showed the payoff for every possible ticket expenditure, and an excerpt from the chart is shown in Table 1.

We conducted a total of five experiments. The number of subjects in the experiments ranged from eight to thirteen. The speed of the experiments depended on how fast the subjects were making their decisions, the length of the training sessions and a number of other variables. Therefore, we have completed different number of periods in different experiments. Table 2 provides a summary.

4. Analysis

Notice that if the aggregation mechanism were perfect, the probability distribution of the states would be as

Table 1. Excerpt from payoff chart used in the second-stage games

Number of tickets	Possible payoff in MYBB game	Possible payoff in AK game	Number of tickets	Possible payoff in MYBB game	Possible payoff in AK game
1	33	-1244	50	854	1515
10	516	388	60	893	1642
20	662	873	70	925	1750
30	747	1157	80	953	1844
40	808	1359	90	978	1926

Table 2. Summary of private information experiments

Expt	Experimental structure		Kullback-Liebler values (Standard deviation)				
	Number of players	Private info	No info	Market prediction	Best player	Simple IAM	Original IAM
1	13	3 draws for all	1.977 (0.312)	1.222 (0.650)	0.844 (0.599)	1.105 (2.331)	0.553 (1.057)
2	9	3 draws for all	1.501 (0.618)	1.112 (0.594)	1.128 (0.389)	0.207 (0.215)	0.214 (0.195)
3	11	1/2: 5 draws 1/2: 1 draw	1.689 (0.576)	1.053 (1.083)	0.876 (0.646)	0.489 (0.754)	0.414 (0.404)
4	8	1/2: 5 draws 1/2: 1 draw	1.635 (0.570)	1.136 (0.193)	1.074 (0.462)	0.253 (0.325)	0.413 (0.260)
5	10	1/2: 3 draws 1/2: varied draws	1.640 (0.598)	1.371 (0.661)	1.164 (0.944)	0.478 (0.568)	0.395 (0.407)

if one person had seen all of the information available to the community. Therefore, we can use the probability distribution conditioned on all the information as a benchmark to which we can compare alternative aggregation mechanisms. In order to compute this omniscient probability distribution, recall that there are twelve balls in the information urn, three for the true state and one for each of the other nine states. Using Bayes' rule one obtains the omniscient probability distribution, i.e.

$$P(s | O) = \frac{\left(\frac{3}{12}\right)^{\#(s)} \left(\frac{1}{12}\right)^{\#(\bar{s})}}{\sum_{\forall s} \left(\frac{3}{12}\right)^{\#(s)} \left(\frac{1}{12}\right)^{\#(\bar{s})}} \quad (4)$$

where s denotes the states, O is a string of observations, $\#(s)$ is the number of draws of the state s in the string, and $\#(\bar{s})$ is the number of draws of all other states.

Once we have this benchmark, the next step is to find a measure by which we can compare it to the probabilities provided by different aggregation mechanisms against this benchmark. The obvious measure to use is the Kullback-Leibler measure, also known as the relative entropy (Kullback and Leibler, 1952). The Kullback-Leibler measure of two probability distributions p and q is given by:

$$\text{KL}(p, q) = E_p \left(\log \left(\frac{p}{q} \right) \right) \quad (5)$$

where p is the "true" distribution (in our case, the omniscient probability distribution). In the case of finite number of discrete states, the above Eq. (5) can be rewritten as:

$$\text{KL}(p, q) = \sum_s p_s \log \left(\frac{p_s}{q_s} \right) \quad (6)$$

It can be shown that $\text{KL}(p, q) = 0$ if and only if the distribution p and q are identical, and that $\text{KL}(p, q) \geq 0$ for all probability distributions p and q . Therefore, a smaller Kullback-Leibler number indicates that two probabilities are closer to each other. Furthermore, the Kullback-Leibler measure of the joint distribution of multiple independent events is the sum of the Kullback-Leibler measures of the individual events. Since periods within an experiment were independent events, the sum or average (across periods) of Kullback-Leibler measures is a good summary statistics of the whole experiment.

5. Results

We compare three information aggregation mechanisms to the benchmark distribution given by Eq. (4) by using the Kullback-Leibler measure. The first of the three information aggregation mechanisms is the market prediction. The market prediction is calculated using the last traded prices of the assets. We used the last traded prices rather than the current round's price because sometimes there was no trade in a given asset in a given round. From these prices, we infer a probability distribution on the states. The second and the third mechanisms are the simple aggregation function given by the risk neutral formula in Eq. (1), and the market-based nonlinear aggregation function in Eq. (2).

In addition, we also report the Kullback-Leibler measures of the no information prediction (uniform distribution over all the possible states) and the predictions from the best individual. The no information prediction serves as the first baseline to determine if any information is contained in the predictions of the mechanisms. Further, if a mechanism is really aggregating information, then it should be doing at least as well as the best individual. So, the predictions of the best individual serve as the second baseline, which helps us to determine if information aggregation indeed occurred in the experiments.

The results are shown in Table 2. The entries are the average values and standard deviations (in parentheses) of the Kullback-Leibler number, which was used to characterize the difference between the probability distributions resulting from a given mechanism and that of the omniscient probability.

As can easily be seen, the nonlinear aggregation function worked extremely well in all the experiments. It resulted in significantly lower Kullback-Leibler numbers than the no information case, the market prediction, and the best a single player could do. In fact, it performed almost three times as well as the information market. Furthermore, the nonlinear aggregation function exhibited a smaller standard deviation than the market prediction, which indicates that the quality of its predictions is more consistent than that of the market. In three of five cases, it also offered substantial improvements over the simple aggregation function.

The results displayed in the second column show that the market was not sufficiently liquid to aggregate information properly, and it was only marginally better than the *a priori* no information case. In almost

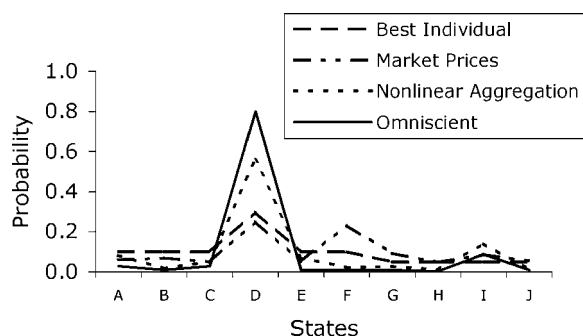


Fig. 1. Accuracy of prediction, by mechanism.

all cases, the best player in the reporting game conveyed more information about the probability distribution than the market did. However, even in situations where the market performs quite poorly, it does provide some information, enough to help us construct an aggregation function with appropriate exponents.

These results are illustrated in Fig. 1, where we show the probability distributions generated by the market mechanisms, the best individual in a typical experiment, the nonlinear aggregation function, as well as the omniscient probability distribution generated by Eq. (4).⁴ Notice that the nonlinear aggregation function exhibits a functional form very similar to the omniscient probability, and with low variance compared to the other mechanisms. This is to be contrasted with the market prediction, which exhibits an information trap at state F and a much larger variance.

These experiments confirm the utility of our nonlinear aggregation mechanism for making good forecasts of uncertain outcomes. This nonlinear function applies the predictions of a group of people whose individual risk attitudes can be extracted by making them participate in an information market. Equally important, our results show that many of the shortcomings associated with information markets can be bypassed by this two-stage method, without having to resort to designing complicated market games. In this context it is worth pointing out that even with such small groups we were able to obtain information whose accuracy, measured by Kullback-Leibler, surpasses by a factor of seven even more complicated institutions such as pari-mutuel games (Plott, Wit, and Yang, 1997).

Lastly, unlike the standard information aggregation implied by the Condorcet theorem, our mechanism allows us to extract probability distributions rather than the validity of a discrete choice obtained via a majority

vote. Moreover, our mechanism provides a signal even in situations when an overall system itself does not contain accurate information as to the outcome. Equally important, unlike Condorcet our two-stage mechanism does not demand risk neutrality and access to the same information by all participants in the system.

6. Distilling Public Information

As we discussed earlier, the above results are not robust to the presence of public information. Nevertheless, the success of our two-stage forecasting mechanism with private information leads us to search for suitable modifications that would allow the detection of the amount of public information present in a group so as to subtract it. In this section, we will propose a method for the identification of public information within a group and its subtraction when aggregating individual predictions about uncertain outcomes. In the following sections, we will present our methodology for aggregating the information resulting from our mechanism, and then the results of experiments in the laboratory that show that this aggregation mechanism outperforms the market, the best player in the group, and the above-described aggregated mechanism.

The introduction of public information implies that the probabilities that enter into Eqs. (1) and (2) are no longer independent of each other, and therefore they are no longer aggregated correctly. Equations (1) and (2) over count information that is observed by more than one individual since they add (in the probability space) probabilities without regard to whether the reports are coming from the same information source. Thus the mechanism has to incorporate a feature that distinguishes the public information from the private, so that it can be suitably subtracted when aggregating the individual predictions. We achieve this by using a coordination game in the second stage, which incentivizes players to reveal what they believe others will reveal. This coordination game is similar to the Battle of the Sexes game.

In the Battle of the Sexes, a couple enjoys spending time together, but each member would rather do so while engaged in his or her preferred activity. As an example, a payoff matrix is shown for an instance in which he'd like them both to go to the baseball game (upper-left), and she'd prefer they went to the opera together (lower-right). If they disagree, no one goes anywhere and no one is happy (off-diagonals).

		SHE	
		B	O
HE	B	3,1	0,0
	O	0,0	1,3

This game has multiple mixed-strategy Nash equilibria, in which both players mix with the goal of landing in the upper-left and lower-right quadrants of the payoff matrix. Notice that with these payoffs each member of the couple is incented to reveal the information that they believe the other will.

In much the same way, our matching game asks players that, in addition to making their best bet (MYBB), they reveal what they believe they all know (AK). The first half, MYBB, works as in the original experiments. That is, players report a vector of bets on the possible states, and are paid according to a log function of these bets. In the AK game however, the subjects try to guess the bets placed by *someone else* in the room, and these bets are then matched to another player whose bets are most similar to theirs. The payout from this part of the game is a function of both their matching level and the possible payout from the number of tickets allocated by the other member of the pair. The payoffs are constructed such that participants have the incentive to match their peers in their public reports. The design of this game is discussed further in the Experimental Design section.

In order to design a payoff function that induces both truthful revelation and maximal matching, we assume that: (A1) the public and private information held by an individual are independent of one another, (A2) that private information is independent across individuals, (A3) that public information is truly public (observed by more than one individual), and (A4) that an given individual can distinguish between the public and the private information he holds. In other words:

For each individual i with observed information O_i , there exists information O_i^{priv} and O_i^{pub} such that:

- (A1) $P(s|O_i) = P(s|O_i^{\text{priv}})P(s|O_i^{\text{pub}})$ for all i, s
- (A2) $P(s|O_i^{\text{priv}})$ and $P(s|O_j^{\text{priv}}) = P(s|O_i^{\text{priv}})P(s|O_j^{\text{priv}})$ for all i, j, s
- (A3) There exists a j for every i such that $O_i^{\text{pub}} = O_j^{\text{pub}}$
- (A4) All individuals know $P(s|O_i^{\text{priv}})$ and $P(s|O_i^{\text{pub}})$

So, in the second stage, each player i is asked to report *two* probability distributions, $\vec{p}_i = \{p_{i1}, p_{i2}, \dots, p_{iN}\}$ (from MYBB) and $\vec{q}_i = \{q_{i1}, q_{i2}, \dots, q_{iN}\}$ (from AK), by allocating a set of tickets to each of the possible states. Let x be the true outcome. The payoff function for each player i is given by the following expression:

$$P = c_1 + c_2 \times \log(p_{ix}) + f(\vec{q}_i, \vec{q}_j) \times (c_4 + c_5 \times \log(q_{jx})) \quad (7)$$

where c_1, c_2, c_3 and c_4 are positive constants, j is chosen in such a way that $f(\vec{q}_i, \vec{q}_j) \geq f(\vec{q}_i, \vec{q}_k)$ for all k , and the function $f(\cdot)$ is given by:

$$f(\vec{x}, \vec{y}) = \left(\frac{1 - [\sum_s |x_s - y_s|]}{2} \right)^2 \vec{y} \quad (8)$$

In words, subjects are paid according to a log function of their reports in the MYBB game, plus a payment from the AK game. This payment is a function of the player with whom he has a maximal match, and is the product of the matching level and a scaled log function of the *matched* player's report in the AK game. This match level is given by the second term of Eq. (7) and is detailed in Eq. (8) above.

As shown earlier, the first part of the payoff function in Eq. (7), $c_1 + c_2 \times \log(p_{ix})$, will induce risk neutral subjects who maximize their expected utility to report their true belief, conditioned on *both* their private and public information. Concerning the last term of Eq. (7), we first note that player i can only affect it through his matching level, which is given by the function $f(\vec{q}_i, \vec{q}_j)$. Since $f(\vec{x}, \vec{x}) \geq f(\vec{y}, \vec{x})$ for all \vec{y} , player i 's best response is to report $\vec{q}_i = \vec{q}_j$. Further, since j is chosen such that $f(\vec{q}_i, \vec{q}_j) \geq f(\vec{q}_i, \vec{q}_k)$ for all k , player i only needs to co-ordinate his \vec{q}_i with only one other individual in the group to achieve an optimal payoff. Additionally, it is easy to show that this part of the game has multiple Nash equilibria, since any common report vector \vec{q} reported by both players i and j is a potential Nash equilibrium. Therefore, we designed the payoff function given by in Eqs. (7) and (8) to encourage individuals to coordinate on the probability distribution induced by the public information. Lastly, the third piece of the payoff function for player i , $c_4 + c_5 \times \log(q_{jx})$ induces a different payoff for each Nash equilibrium \vec{q} on which the two individuals coordinate. Since this factor depends on the strategy of player i 's partner j , no one player can directly affect it. This is crucial to preserve the equilibrium structure.

We thus designed the payoff such that the more information revealed in the reports \vec{q} , the higher the potential payoff to the subjects involved, which implies an information-rich equilibrium. Additionally, since private information is independent across individuals (it is truly private), the best equilibrium on which individuals can coordinate on is the probability distribution induced by using the public information only. Therefore, this mechanism will induce individuals to report both their true beliefs (\vec{p}_i) and their public information (\vec{q}_i). Once these vectors are reported, we still need to aggregate them, which we discuss in the next section.

7. Aggregating Public and Private Information

Once we have a mechanism for extracting public beliefs from private ones, it is straightforward to add a public information generalization to Eq. (2). By dividing the perceived probability distributions of the players by the distributions induced by the public information only, we develop what we call a *General Public Information Mechanism* (GPIC), which is given by

$$P(s | I) = \frac{\left(\frac{p_{s1}}{q_{s1}}\right)^{\beta_1} \left(\frac{p_{s2}}{q_{s2}}\right)^{\beta_2} \cdots \left(\frac{p_{sN}}{q_{sN}}\right)^{\beta_N}}{\sum_{\forall s} \left(\frac{p_{s1}}{q_{s1}}\right)^{\beta_1} \left(\frac{p_{s2}}{q_{s2}}\right)^{\beta_2} \cdots \left(\frac{p_{sN}}{q_{sN}}\right)^{\beta_N}} \quad (9)$$

where the q_s are extracted from individuals' reports before they are aggregated. This correction allows us to isolate the private information from the individual reports.

While this mechanism is quite general, and outperforms both the market prediction and that of our original IAM, there are potential improvements to it that can be implemented. Thus, we developed modifications to the aggregation function to address issues of uncertain information structures and multiple equilibria. In theory, knowledge of the individuals' reports $\vec{p}_i = \{p_{i1}, p_{i2}, \dots, p_{iN}\}$ and $\vec{q}_i = \{q_{i1}, q_{i2}, \dots, q_{iN}\}$, should make information aggregation straightforward since for a given individual i , his probability assignment to state s , with respect to private information, should be proportional to p_{si}/q_{si} . To more efficiently add in public information, we aggregate the individual reports of public information $\vec{q}_i = \{q_{i1}, q_{i2}, \dots, q_{iN}\}$ into a single vector $\vec{q} = \{q_1, q_2, \dots, q_N\}$. In order to do this, we employ one additional assumption, that every

individual observes the *same* public information, O^{pub} . We then aggregate by averaging the reports, weighted by each individual's β , thus:

$$q_s = \frac{\sum_{i=1}^N \beta_i q_{si}}{\sum_{i=1}^N \beta_i} \quad (10)$$

Once we have completed this aggregation process, we can use the new vector \vec{q} in place of \vec{q}_i in the original function in Eq. (9). If \vec{q} is derived correctly, it will resolve the matter of parsing the private information from the public. Furthermore, in much the same way that some people process their private signals better than others, there are some individuals that report public information more accurately than others. If one can identify these individuals, one can recover public information more efficiently than by taking a weighted average of everyone's report. Thus, instead of using the whole group to recover public information, as in Eq. (10), we use a limited set J , a subset of the whole group:

$$q_s = \frac{\sum_{i \in J} \beta_i q_{si}}{\sum_{i=1}^N \beta_i} \quad (10a)$$

The resultant forecast is then determined by a modification of the GPIC in Eq. (9). It uses a small subset of players to determine the public information so as to parse it from the private. While this mechanism is quite efficient, it only applies to the special case where the public information is completely public and identical. Therefore, we refer to it as the *Special Public Information Correction Mechanism*, or SPIC.

$$P(s | I) = \frac{q_s \left(\frac{p_{s1}}{q_s}\right)^{\beta_1} \left(\frac{p_{s2}}{q_s}\right)^{\beta_2} \cdots \left(\frac{p_{sN}}{q_s}\right)^{\beta_N}}{\sum_{\forall s} q_s \left(\frac{p_{s1}}{q_s}\right)^{\beta_1} \left(\frac{p_{s2}}{q_s}\right)^{\beta_2} \cdots \left(\frac{p_{sN}}{q_s}\right)^{\beta_N}} \quad (11)$$

8. Experimental Design for Public Information Experiments

As in the private information experiments, all sessions were conducted at Hewlett-Packard Laboratories in Palo Alto, California, with a similar cohort to those in the private information sessions. The information structure was also identical to that of the private information experiments, as was the first-stage call market. The only difference was in the second stage, which we will now describe.

In the second-stage, every subject played under the same information *structure* as in the first stage, although the draws and the true states were independent from those in the first. There are two parts to this game, described in the Identifying Public Information section above, which were referred to as the "What Do We All Know" (AK) and the "Make Your Best Bet" (MYBB) games. Each period, the subjects received their draws of information, as in the market game. They also received two sets of 100 tickets each, one set for AK, and one for MYBB. We will discuss these two games in turn.

The MYBB game is identical to the second stage played in the private information game. That is, in MYBB, the subjects were asked to distribute their tickets across the ten states with the constraint that all 100 tickets must be spent each period and that at least one ticket is spent on each state. Since the fraction of tickets spent determines p_{si} , this implies that p_{si} is never zero. The subjects were given a chart that told them how many francs they would earn upon the realization of the true state as a function of the number of tickets spent on the true state security. The payoff was a linear function of the log of the percentage of tickets placed in the winning state as given by the first half of Eq. (7). The chart the subjects received showed the payoff for every possible ticket expenditure, and was identical to that shown in Table 1.

We also played the matching game in this stage, known as AK. In this stage, subjects received 100 tickets, but with a different goal. They tried to guess the bets placed by someone else in the room. After they placed the bets, they were matched to another player, one whose bets were most similar to theirs. The more similar the bets were to their nearest match, the higher the reported "Percent Match with Partner." The payoffs for any given ticket expenditure were higher in the AK game than the MYBB game, and are detailed in Table 1.

Fig. 2 shows a screenshot from the second stage of the game, which displays the bets placed in a sample Period 1. As shown on the upper right, the true state was F. Following down the items reported in the upper right of the screen, we see that this player bet 20 tickets on F in the MYBB game, which corresponds to a Possible Payout of 662 francs. He was matched with a partner whose AK distribution of tickets matched his at a 49% level. This partner bet enough tickets to have a Possible Payout of 178 francs. Our sample player thus earned 662 francs for the 20 tickets bet in the MYBB game, plus $0.49 \times 178 = 87$ francs for the AK game, for a total of 749 francs.

9. Experimental Results

We compare five information aggregation mechanisms to the benchmark distributions. In addition, we also report the Kullback-Leibler measures of the no information prediction (uniform distribution over all the possible states) and the best (most accurate) individual's predictions. The first two information aggregation mechanisms we evaluate are the market prediction and the original private information mechanism in Eq. (2).

The third mechanism is our proposed improvement, referred to as the *General Public Information Correction (GPIC) mechanism*, given by Eq. (9). It uses both individuals' reports of public information regarding outcomes as well as the individuals' perceived probabilities of these outcomes. If this mechanism is working as predicted by the theory, it should provide a superior outcome to that of the original IAM.

As an additional benchmark the fourth mechanism, referred to as the *Perfect Public Info Correction (PPIC)*, replaces individuals' reports of public information with the true public information that they have observed. Obviously, this is not possible in a realistic environment, since we do not know the true public information (or, this exercise would be pointless). However, it allows us to validate the behavioral assumptions we make in the design of the mechanism. Our model implicitly assumes that individuals aggregate their public and private information by a modified version of Bayes' rule to arrive at their reports, and we can use this benchmark to validate this assumption.

Lastly, we address the special case in which the experimenter knows that every individual receives the same public information. This fifth mechanism, referred to as the *Special Public Info Correction mechanism (SPIC)*, recovers the public information by using the reports of only the best two individuals to correct the public information bias in all participants' reports.

As is shown in Tables 3 and 4, once even a small amount of public information is introduced into the system (Experiments 1a through 5a), the performance of the original IAM decreases dramatically. In Fig. 3 we illustrate the double counting issue before the GPIC modification. In this figure, we plot the probability distributions generated by omniscience, the prediction from the original IAM and the available public information from a sample period (Experiment 3a, period 9). As one can see, using the original IAM results in a false peak at state H, which is the state on which public information was available. In some cases, the double

Possible States:	A	To	J	Period	1
Drawing from a urn with replacement containing:				State	F
3	balls for the true state		Total	Tickets on this State	20
1	ball for each false state		Payoff	Percent Match with Partner	49
	Private Information		749	Maximal Payoff from Partner	178
>>>	Public Information			Payoff	749
	MAKE YOUR BEST BET		WHAT DO WE ALL KNOW?	My Information	
State	Number of Tickets	Possible Payout	Number of Tickets	State	Count
A	20	662	10	A	1
B	3	264	5	B	0
C	3	264	5	C	0
D	3	264	5	D	0
E	20	662	25	E	1 >>>
F	20	662	10	F	1
G	3	264	5	G	0
H	3	264	5	H	0
I	20	662	25	I	1 >>>
J	5	371	5	J	0
	Total Tickets Spent	100	Total Tickets Spent	100	

Fig. 2. Sample page from stage two of the experiment.

counting issue is so severe that the results are worse than that of the no information measure (see, for example, Experiments 1a, 2a and 3a). Thus, this verifies the necessity to derive a method correcting for the biases introduced by public information.

In Table 3 we summarize the relative performance, in terms Kullback-Leibler measures, of all of the benchmarks mechanisms enumerated above. Table 4 reports the same results in terms of the percentage relative to the no information Kullback-Leibler measure (indicating the level of improvement over this benchmark). Note that the amount of aggregate information available in an experiment varied across the treatments. Because the pure KL measure reported in Table 3 is affected by the amount of underlying information, the

percentage measurement in Table 4 are more useful when comparing results across experiments.

The GPIC mechanism (Eq. (9)) outperforms the best single individual's guesses reports in all five experiments. It also outperforms the market prediction in four out of five experiments. The GPIC mechanism uses the reports of public information of individuals to perform the correction. As expected, this mechanism recovers enough public information to perform well compared to an information market. However, there is room for improvement compared to the case where the true public information is used.

To understand this inefficiency, let us assume that the information aggregator knows the true public information seen by every individual and applies the

Table 3. Summary of public information experiments

Experimental structure				Kullback-Liebler values (Standard deviation)						
Expt	Number of players	Private info	Public info	No info	Market prediction	Best player	Original IAM	General public info correction	Perfect public info correction	Special public info correction
1a	10	2 draws for all	2 draws for all	1.332 (0.595)	0.847 (0.312)	0.932 (0.566)	2.095 (.196)	0.825 (0.549)	0.279 (0.254)	0.327 (0.247)
2a	9	2 draws for all	2 draws for all	1.420 (0.424)	0.979 (0.573)	0.919 (0.481)	2.911 (2.776)	0.798 (0.532)	0.258 (0.212)	0.463 (0.492)
3a	11	3 draws for all	1 draws for all	1.668 (0.554)	1.349 (0.348)	1.033 (0.612)	2.531 (1.920)	0.718 (0.817)	0.366 (0.455)	0.669 (0.682)
4a	10	3 draws for all	1/2 : 1draw	1.596 (0.603)	0.851 (0.324)	1.072 (0.604)	0.951 (1.049)	0.798 (0.580)	0.704 (0.691)	0.793 (0.706)
5a	10	3 draws for all	1 draws for all 2 sets of public info	1.528 (0.600)	0.798 (0.451)	1.174 (0.652)	0.886 (0.763)	1.015 (0.751)	0.472 (0.397)	0.770 (0.638)

Table 4. Percentage of no-info Kullback-Leibler numbers (Public information)

Experimental structure				Kullback-Liebler values, as a percent of the no info case						
Expt	Number of players	Private info	Public info	No info	Market prediction	Best player	Original IAM	General public info correction	Perfect public info correction	Special public info correction
1a	10	2 draws for all	2 draws for all	100%	63.6%	70.0%	157.3%	61.94%	20.94%	24.53%
2a	9	2 draws for all	2 draws for all	100%	69.0%	64.7%	205.0%	56.2%	18.2%	32.6%
3a	11	3 draws for all	1 draws for all	100%	80.9%	61.9%	151.7%	43.0%	22.0%	40.1%
4a	10	3 draws for all	1/2 : 1draw	100%	53.3%	67.1%	59.6%	50.0%	44.1%	49.7%
5a	10	3 draws for all	1 draws for all 2 sets of public info	100%	52.2%	76.9%	57.9%	66.4%	30.9%	50.4%

algorithm in Eq. (11). The accuracy of the results obtained (Perfect Public Info Correction, or PPIC) are almost as good as the performance of the original IAM mechanism in the private information case. Furthermore, this method outperforms any other method by a large margin. Although this is not an implementable mechanism, since no one knows the true public information, it does show the correctness of our behavioral model as to how people mix private and public infor-

mation is correct. Therefore, there is validity in our approach to teasing out this public information in the GPIC.

Fig. 4 illustrates the efficacy of the GPIC. In this figure, once again, the results from Experiment 3a, period 9 are plotted. The GPIC mechanism eliminates the false peak shown in Fig. 3. However, the correction is not perfect. There is still some residual positive probability being placed on state H, the site of the false peak.




Fig. 3. Illustration of the double counting issue.




Fig. 4. Information aggregation with public information correction.

When the PPIC is used to perform the correction, the false peak is completely eliminated.

It is important to realize that while algorithms that explicitly aggregate private and public information are sensitive to the underlying information structures, markets are not. In all the experiments, including the ones with only private information, the performance of the market, measured as a percentage of the no information KL, is fairly consistent, albeit somewhat inaccurate.

It is interesting to note that if we assume that every individual receives the same public information, we may not need to use everyone's report to recover public information, as described in the SPIC mechanism. By searching for pairs with the best performance, we can achieve improvements over our GPIC. However, these pairs were found *ex post*. That is, we calculate the performance for every pair and then choose the best. So, this extension shows merely the possibility of using pairs (or larger subgroups) to recover public information. Simple intuitive ad hoc rules, such as choosing the pairs that are closest together in the KL sense, can find good pairs in some experiments. We include the results from such an attempt in Tables 3 and 4 as the *Special Public Information Correction*, or SPIC. The issue of identifying subgroups to recover either public, or for that matter, private information is subject of future research.

10. Conclusions

Accurate predictions are essential to individuals and organizations. For large communities, information relevant to forecasts is often dispersed across people, frequently in different geographical areas. Examples include forecasting sales of a product, aggregating the financial predictions of the venture capital community, and public opinion polls. The methodology described in this paper addresses many of the needs to aggregate this information accurately and with the correct incentives. One can take past predictive performance of participants in information markets and create weighting schemes that will help predict future events, even if they are not the same event on which the performance was measured. Furthermore, our two-stage approach can improve upon predictions by harnessing distributed knowledge in a manner that alleviates problems with low levels of participation. The typical business forecast cycle also lends itself to this approach. Since forecasts cycles in organizations typically in-

volve the prediction of similar events on a periodic basis, it is possible to set up an initial market to obtain consistent measures of skills and risk attitudes and then use the reporting mechanism to extract and aggregate information in the future.

Obviously, this approach can also be extended to work across organizations. One possible use is to aggregate and create consensus estimates in the financial analyst community. Another one is to provide the venture capital community a way of forming predictions about the viability of new ventures. The Hollywood Stock Exchange has shown that information markets can be used to predict movie ticket sales, which are tremendously important to studio executives. In the same vein, our methodology can be used with smaller groups of movie screen test subjects to create forecasts before a movie is released. One can imagine a world in which focus groups are no longer run solely on survey questions and discussions, but where each member has a financial stake in the information coming out of the focus group.

The rapid advances of information technologies and the understanding of information economics have opened up many new possibilities for applying mechanism design to gather and analyze information. This paper discusses one such design and provides empirical evidence about its validity. Although the results we presented are particular to events with finite number of outcomes, they can be generalized to a continuous state space. We are currently pursuing these extensions. Equally intriguing is the possibility of having this mechanism in the context of the Web, thus enabling information aggregation over large geographical areas, perhaps asynchronously. This leads to issues of information cascades and the optimal time to keep an aggregation market open, which we will explore in turn.

Appendices

Appendix 1: Conditional probabilities and products of reports

Lemma. *If O_1 through O_n are independent observations conditioned on a given state and the a priori beliefs of the probabilities of the states are uniform*

$$\text{Then } P(s | O_1, O_2, \dots, O_n) = \frac{\prod_{i=1}^n P(s | O_i)}{\sum_{s'} \prod_{i=1}^n P(s' | O_i)}.$$

In other words, if N people observe independent information about the likelihood of a given state and they

report those probabilities, one can find the probability conditioned on all of their observations by multiplying their reported probabilities and then normalizing the results.

Proof: For any observation O ,

$$P(s | O) = \frac{P(s)P(O | s)}{\sum_{s'} P(s')P(O | s')} \quad (\text{Bayes' Rule})$$

Because O_1 through O_n are independent observations conditioned on a given state, $P(O_1, \dots, O_n | s) = P(O_1 | s), \dots, P(O_n | s)$. So, applying Bayes' Rule again,

$$\begin{aligned} P(s | O_1, \dots, O_n) &= \frac{P(s)P(O_1, \dots, O_n | s)}{\sum_{s'} P(s')P(O_1, \dots, O_n | s')} \\ &= \frac{P(s)P(O_1 | s), \dots, P(O_n | s)}{\sum_{s'} P(s')P(O_1 | s'), \dots, P(O_n | s')} \end{aligned}$$

Because the *a priori* beliefs of the probabilities of the states are uniform, $P(s)$ is a constant independent of s . That is, if there are K possible states of the world, $P(s) = 1/K$ for all s in K . So,

$$\begin{aligned} \frac{P(s | O_1, \dots, O_n)}{\prod_{i=1}^n P(s | O_i)} &= \frac{\frac{P(s)P(O_1, \dots, O_n | s)}{\sum_{s'} P(s')P(O_1 | s'), \dots, P(O_n | s')}}{\prod_{i=1}^n \frac{P(s)P(O_i | s)}{\sum_{s'} P(s')P(O_i | s')}} \\ &= \frac{P(s)P(O_1, \dots, O_n | s) \prod_{i=1}^n \sum_{s'} P(s')P(O_i | s')}{\prod_{i=1}^n P(s)P(O_i | s) \sum_{s'} P(s')P(O_i | s'), \dots, P(O_n | s')} \\ &= \frac{P(O_1, \dots, O_n | s) \prod_{i=1}^n \sum_{s'} P(s')P(O_i | s')}{P(s)^{n-1} \prod_{i=1}^n P(O_i | s) \sum_{s'} P(s')P(O_i | s'), \dots, P(O_n | s')} \\ &= \frac{\prod_{i=1}^n \sum_{s'} P(s')P(O_i | s')}{P(s)^{n-1} \sum_{s'} P(s')P(O_i | s'), \dots, P(O_n | s')} \end{aligned}$$

We can now call this right hand side λ , which is only related to observations O_1 to O_n and is therefore independent of the state s . Now, we simply need to rearrange and solve.

Rearranging, we have

$$P(s | O_1, \dots, O_n) = \lambda \prod_{i=1}^n P(s | O_i), \quad \text{with}$$

$$\lambda = \frac{\prod_{i=1}^n \sum_{s'} P(s')P(O_i | s')}{P(s)^{n-1} \sum_{s'} P(s')P(O_1 | s'), \dots, P(O_n | s')}$$

We know that $\sum_{s'} P(s' | O_1, \dots, O_n) = 1$, so that

$$\begin{aligned} \sum_{s'} P(s' | O_1, \dots, O_n) &= \lambda \sum_{s'} \prod_{i=1}^n P(s' | O_i) = 1 \\ \Rightarrow \lambda &= \frac{1}{\sum_{s'} \prod_{i=1}^n P(s' | O_i)} \end{aligned}$$

Therefore,

$$\begin{aligned} P(s | O_1, \dots, O_n) &= \lambda \prod_{i=1}^n P(s | O_i) \\ &= \frac{1}{\sum_{s'} \prod_{i=1}^n P(s' | O_i)} \\ &\quad \times \prod_{i=1}^n P(s | O_i) \\ &= \frac{\prod_{i=1}^n P(s | O_i)}{\sum_{s'} \prod_{i=1}^n P(s' | O_i)} \end{aligned}$$

■

Appendix 2: Risk neutrality and log payoff functions in the reporting game

Consider the following game:

- There are N possible states of the world.
- A player is given information about the state of the world $x \in \{1, 2, \dots, N\}$. His belief on the probabilities of these states of the world, conditioned on his information, are $P_i, i \in \{1, 2, \dots, N\}$.
- The player is asked to report a vector $\{q_1, q_2, \dots, q_N\}$ with the constraint $\sum_{i=1}^N q_i = 1$. Then the true state x is revealed and he is paid $f(q_x)$.

Lemma. *If the player is risk neutral and $f(y) = c_1 + c_2 \log(y)$, then $q_i = P_i$ for all i . That is, players will report their true beliefs on the probabilities.*

Proof: The player's maximization problem is:

$$\begin{aligned} \text{Max}_{\{q_i\}} \sum_{i=1}^N P_i (c_1 + c_2 \log(q_i)) \\ \text{s.t.} \quad \sum_{i=1}^N q_i = 1 \end{aligned}$$

This is equivalent to $\text{Max}_{\{q_i\}} \sum_{i=1}^N P_i \log(q_i)$ s.t. $\sum_{i=1}^N q_i = 1$ for all positive c_1 and c_2 .

The Lagrangian for this problem is

$$L = \sum_{i=1}^N P_i \log(q_i) - \lambda \left(\sum_{i=1}^N q_i - 1 \right)$$

The first order condition is: $\frac{P_i}{q_i} = \lambda$ for all $i \Rightarrow P_i = \lambda q_i$. Summing over all $i, 1 = \lambda$. Thus $q_i = P_i$ for all i . ■

Appendix 3: The stage-two coordination game

Consider the following scenario:

- N possible states of the world.
- M players indexed by $i = 1 \dots M$.
- Player i is given information about the state of the world $x \in \{1, 2, \dots, N\}$.
 - His beliefs as to the probabilities of the states of the world conditioned on his information are P_{ix} .
 - Some of player i 's information is observed by at least one other player j . Let Q_{ix} be the probability conditioned on i 's public information only (does not consider his private information).
- Each player i is asked to report two probability distributions $p_i = \{p_{i1}, p_{i2}, \dots, p_{iN}\}$ and $q_i = \{q_{i1}, q_{i2}, \dots, q_{iN}\}$ with the constraints $\sum_{s=1}^N p_{is} = 1$ and $\sum_{s=1}^N q_{is} = 1$.
- The true state x is revealed and he is paid $f(p_i, q_i, q_{-i} | x)$.

Assumptions:

- (a) Players are risk neutral utility maximizers, and
- (b) $f(p_i, q_i, q_{-i} | x) = c_1 + c_2 \log(p_{ix}) + H(g(q_i, q_j), \log(q_{ix}))$.
- (c) $g : q \times q \rightarrow \Re$ is any real function of two probability distribution such that $y = \text{Max}_x g(x, y)$.
- (d) $H(x, y)$ is increasing both in x and y .
- (e) j is determined by: $j = \arg \max_{k \in \{1 \dots M\}} g(q_i, q_k)$.

Lemma 1. $\{p_i = P_i, q_i = Q_i \text{ for all } i\}^5$ is a Bayesian Nash equilibrium. That is, each player will report his true conditional probability beliefs and the beliefs conditioned solely on his public information.

Proof: Assuming all players but i are playing an equilibrium strategy, player i 's maximization problem is

$$\begin{aligned} & \text{Max}_{\{p_i, q_i\}} \sum_{s=1}^N P_{is} \{c_1 + c_2 \log(p_{is}) + H(g(q_i, Q_i), \log(Q_{is}))\} \\ & \text{s.t. } \sum_{s=1}^N p_{is} = 1 \text{ and } \sum_{s=1}^N q_{is} = 1. \end{aligned}$$

There will be at least one other player j that plays $q_j = Q_{ix}$ since at least one player other than i observes the same public information and arrives at the same distribution Q_i .

The resultant Langrangian is

$$\begin{aligned} L = & \sum_{s=1}^N P_{is} \{c_1 + c_2 \log(p_{is}) + H(g(q_i, Q_i), \log(Q_{is}))\} \\ & - \lambda \left(\sum_{i=1}^N p_{is} - 1 \right) - \mu \left(\sum_{i=1}^N q_{is} - 1 \right) \end{aligned}$$

The first order condition is $\frac{P_{is}}{p_{is}} = \lambda$ for all $i \Rightarrow P_{is} = \lambda p_{is}$.

Summing over both sides, we get $1 = \lambda$. Thus $p_{is} = P_{is}$ for all i .

Recalling assumption (c), $q_i = Q_i$ maximizes $g(q_i, Q_i)$. Since H is increasing in g , it also maximizes $H(g(q_i, q_j), \log(q_{ix}))$. \square

Lemma 2. There are multiple equilibria to this game.

The same proof applies to $\{p_i = P_i \text{ for all } i; q_{ix} = \frac{1}{N} \text{ for all } i, x\}$ or for that matter, any set of q_i on which players coordinate.

Notes

1. Notable exceptions: The Iowa Electronic Market (<http://www.biz.uiowa.edu/iem>) has shown that political events can be accurately predicted using markets when they are large enough. Their predictions have consistently been more accurate than those resulting from major news polls. Additionally, recent work by Pennock, et al. (2000) show that the Hollywood Stock Exchange (HSX) does a remarkable job of predicting box office revenues and Oscar winners. However, both of these institutions have many traders, while we focus on systems with small number of participants (fewer than 15).
2. These securities have lottery-like properties, and they pay off one unit contingent on the positive outcome of an event linked to that security, and zero otherwise.
3. An experimental currency, exchanged for dollars at the end of the experiment according to an announced exchange rate.
4. While different, independent events are used for the market stage and the reporting stage, we found one period in both stages that contained the exact same information. Thus, we can compare results from these two periods in this figure.
5. Notice that p_i and q_i are probability distributions. Therefore, the statement is equivalent to $\{p_{ix} = P_{ix} \text{ for all } i, x; q_{ix} = Q_{ix} \text{ for all } i, x\}$.

References

- Anderson L., Holt C. Information cascades in the laboratory. *American Economic Review* 1997;87:847–862.
- Camerer C., Weigelt K. Information mirages in experimental asset Markets. *Journal of Business* 1991;64:463–493.

- Forsythe R., Lundholm R. Information aggregation in an experimental market. *Econometrica* 1990;58:309–347.
- Forsythe R., Palfrey T., Plott C. Asset valuation in an experimental market. *Econometrica*, 1982;50:537–567.
- Hayek F. The use of knowledge in society. *American Economic Review*, 1945;35(4):519–530.
- Iowa Electronic Markets. <http://www.biz.uiowa.edu/iem>.
- Kullback S., Leibler R.A. On information and sufficiency. *Annals of Mathematical Statistics*, 1952;22:79–86.
- Markowitz H. *Portfolio Selection*, NY: Wiley, 1959.
- Nöth M., Camerer C., Plott C., Weber M. Information traps in experimental asset markets, Submitted to *Review of Financial Studies*, 1999.
- Nöth M., Weber M. Information aggregation with random ordering: Cascades and overconfidence. University of Mannheim, Discussion Paper, presented at the Summer 1998 ESA Meetings, 1998.
- O'Brien J., Srivastava S. Dynamic stock markets with multiple assets. *Journal of Finance* 1991;46:1811–1838.
- Pennock D., Lawrence S., Giles C., Nielsen F. The power of play: Efficient and forecast accuracy in web market games. NEC Research Institute Technical Report 2000–168, 2000.
- Plott C., Sunder S. Efficiency of experimental security markets with insider information: An application of rational expectations models. *Journal of Political Economy* 1982;90:663–698.
- Plott C., Sunder S. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 1988;56:1085–1118.
- Plott C., Wit J., Yang W. Pari-mutuel betting markets as information aggregation devices: Experimental results. Social Science Working Paper 986, California Institute of Technology, 1997.
- Scharfstein D., Stein J. Herd behavior and investment. *American Economic Review* 1990;80:465–479.
- Sunder S. Markets for information: Experimental evidence. *Econometrica* 1992;60:667–695.

Kay-Yut Chen came to HP Labs in 1994 after graduating from Caltech and started the HP Labs experimental economics program. His work in applied experimental economics has helped HP business to develop better channel strategies. He employs experimental methodologies to develop a business “wind tunnel” to test business policies, process and strategies. Chen’s work also includes research in the area of information aggregation, reputation mechanisms and quantum game theory.

Leslie Fine is a Scientist in the Information Dynamics Lab at Hewlett Packard Laboratories. She works on new methods and technologies for the practical application of information economics, including incentive design, market design, and information aggregation mechanisms.

Bernardo Huberman is a HP Fellow and Director of the Systems Research Center at Hewlett Packard Laboratories, where he also heads the research effort in Information Dynamics. He received his Ph.D. in Physics from the University of Pennsylvania, and is currently a Consulting Professor in the Department of Applied Physics at Stanford University. He recently published the book: “The Laws of the Web: Patterns in the Ecology of Information” with MIT Press.