

Eliminating Public Knowledge Biases in Information Aggregation Mechanisms

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

HP Laboratories

Palo Alto, CA 94304

March 21, 2003

Abstract

We present a novel methodology for identifying public knowledge and eliminating the biases it creates when aggregating information in small group settings. A two-stage mechanism consisting of an information market and a coordination game is used to reveal and adjust for individuals' public information. A nonlinear aggregation of their decisions then allows for the calculation of the probability of the future outcome of an uncertain event, which can then be compared to both the objective probability of its occurrence and the performance of the market as a whole. Experiments show that this nonlinear aggregation mechanism outperforms both the imperfect market and the best of the participants.

Introduction

The prediction of the future outcomes of uncertain situations is an important problem for individuals and organizations. As a result, large resources are devoted to producing reliable forecasts of technology trends, revenues, growth, and other valuable insights. To complicate matters, in the case of organizations 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 (Gordon et al 1964), 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 et al 1982, O'Brien et al 1991, Plott et al 1982, Plott et al 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 et al 1991, Nöth et al 1999), illiquidity (Sunder 1992), manipulation (Forsythe et al 1990, Nöth et al 1998), and lack of meaningful equilibria (Anderson et al 1997, Scharfstein et al 1990)ⁱ. 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 earlier work, (Chen et al 2001), we demonstrated the comparative efficacy of a nonlinear aggregation mechanism with behavioral components to that of a market. Specifically, we showed that one could take past predictive performance of participants in information markets and to create weighting schemes that help predict future events, even if they are not the same event on which the performance was measured. Furthermore, our two-stage approach successfully harnessed distributed knowledge in a manner that alleviated the problems that arise from low levels of participation. Therefore, this mechanism is most useful in environments in which markets fail due to low participation and the resulting lack of liquidity.

However, these results were not immune to the presence of public information, that is, information that is commonly known to multiple individuals in the group. This is because public information is bound to introduce strong correlations in the knowledge possessed by members of the group, correlations that were not explicitly taken into account by our aggregation algorithm.

Nevertheless, the success of our two-stage forecasting mechanism with private information led 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. 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 created 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 both the market and the best player in the group.

In what follows, we first outline the original two-stage mechanism for information aggregation and then explain the modified second stage that allows for public information to be extracted. Next, we present laboratory experiments that quantitatively measure the performance of this new mechanism and established its

superiority with respect to both the information market and the participating members. An appendix provides the mathematical details of the coordination game.

Extracting Private Information

We start by reviewing the original nonlinear aggregation scheme, presented in Chen, Fine and Huberman (Chen et al 2001). This aggregation scheme applies to a group of individuals that hold private information.

Consider, for a moment, an environment in which a set of N people has 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 probability of an event by collecting individual beliefs and 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 . 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 above restrictions on information and risk neutrality.

In Chen, Fine and Huberman (Chen et al 2001), we show that in order to account for both the diverse levels of risk aversion and information strengths, we can add a stage to the mechanism. Before individuals are asked to report their beliefs, they are asked to 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. It is assumed that there exists a sequence of similar events for us to conduct information markets on. We use information markets on earlier events 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 be perfectly efficient, it is a powerful enough mechanism to help us elicit the needed information.

In the second stage, each player i is asked to report a vector of perceived state-probabilities, $\hat{p}_i = \{p_{1i}, p_{2i}, \dots, p_{Ni}\}$, with p_{si} the perceived probability that a given state s will be realized, and with the constraint that the vector sums to one. When the true state x is revealed, each player is paid an amount equal to $c_1 + c_2 \cdot \log(p_{xi})$, where c_1 and c_2 are positive numbers. This payoff function ensures that risk-neutral expected utility maximizers would report their true beliefs. So, each player should report his perceived probability distribution over the N possible states.

In order to compute the probability distribution, we aggregate the individual reports by using the following nonlinear aggregation function, which is a modification of Bayes' rule:

$$P(s | I) = \frac{p_{s_1}^{b_1} p_{s_2}^{b_2} \dots p_{s_N}^{b_N}}{\sum_{\forall s} p_{s_1}^{b_1} p_{s_2}^{b_2} \dots p_{s_N}^{b_N}} \quad (2)$$

where s is a given possible state, I is the available information, and β_i is the exponent assigned to individual i .

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

$$b_i = r(V_i/S_i)c \quad (3)$$

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 s_i is the variance of his holdings over time. We use c as a normalization factor so that if $r=1$, $\sum b_i$ equals the number of players; it is chosen to make the average b equal to one.

To derive r , notice that when 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 the case. 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. This derivation is not standard in the finance literature, this is because typical market analysis assume a no arbitrage condition. However, in our case with a small market, it is clear from experimental evidence (Chen et al 2001) that the no arbitrage condition was violated. Furthermore, the amount and direction (whether on the buy or the sell side) of arbitrage opportunities in the market becomes an important indication of the risk attitude of the market as a whole. If people are driving the price up and beyond reasonable levels suggested by the condition of no arbitrage, we conclude that people are willing to pay more for risks (since all the assets are risky).

The ratio of value to risk, (V_i/s_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, his amount of risk can be measured by the variance of his values using normalized market prices as probabilities of the possible outcomes. There are many ways of measuring risk attitudes. Equation 3 is one rule of the thumb that delivered good performance in multiple experiments (Chen et al 2001).

Identifying Public Information

Although this mechanism works well with private independent information (see Table 1 and Table 2, Experiments 1 through 5), its performance can be significantly degraded by the introduction of public information. The introduction of public information implies that the probabilities that enter into Equation (2) are no longer independent of each other, and therefore they are no longer aggregated correctly. Equation (2) over counts information that is observed by more than one individual since it adds (in the probability space) probabilities disregarding 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 in the second stage a coordination game. In this game, we ask players to announce what they believe to be common information. That is, we ask them to reveal what they think others will reveal. Since the private information a player holds is, by definition, unique to that player, a rational player should reveal only what he believes to be public.

So, 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) $P(s|O_i^{\text{priv}} \text{ and } O_j^{\text{priv}}) = P(s|O_i^{\text{priv}}) P(s|O_j^{\text{priv}})$ 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 their $P(s|O_i^{\text{priv}})$ and $P(s|O_i^{\text{pub}})$

Notice that (A3) does not place any restrictions on the subset of individuals who observe public information and does not require public information to be common knowledge. The only requirement is that any public information is observed by more than one individual. Furthermore, (A4) is also not as restrictive as it might seem. We do not require the structure of public information (who knows what) to be common knowledge. All we assume is that each individual knows which part of his information is private and which is public without having to know who has the same public information.

So, in the second stage, each player i is asked to report two probability distributions, $\hat{p}_i = \{p_{1i}, p_{2i}, \dots, p_{Ni}\}$ (from MYBB) and $\hat{q}_i = \{q_{1i}, q_{2i}, \dots, q_{Ni}\}$ (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:

$$P = c_1 + c_2 * \log(p_{xi}) + f(\hat{q}_i, \hat{q}_j) * (c_3 + c_4 * \log(q_{xj})) \quad (4)$$

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

$$f(\hat{x}, \hat{y}) = \left(1 - \left[\sum_s |x_s - y_s| \right] / 2 \right)^2 \quad (5)$$

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 Equation (4) and is detailed in Equation (5) above.

As shown earlier (Chen et al 2001), the first part of the payoff function in Equation (4), $c_1 + c_2 * \log(p_{xi})$, 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 Equation (4), we first note that player i can only affect it through his matching level, which is given by the function $f(\hat{q}_i, \hat{q}_j)$. Since $f(\hat{x}, \hat{x}) \geq f(\hat{y}, \hat{x})$ for all \hat{y} , player i 's best response is to report $\hat{q}_i = \hat{q}_j$. Further,

since j is chosen such that $f(\hat{q}_i, \hat{q}_j) \geq f(\hat{q}_i, \hat{q}_k)$ for all k , player i only needs to coordinate his \hat{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 \hat{q} reported by both players i and j is a potential Nash equilibrium. Therefore, we designed the payoff function given by in Equations (4) and (5) 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 , $c_3 + c_4 \cdot \log(q_{j|x})$ induces a different payoff for each Nash equilibrium \hat{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 important to preserve the revelatory qualities of the game. Since player i does not know to whom he will be matched, only that it will be a player with revelation similar to his, his best strategy is simply to reveal what he thinks others will.

We thus designed the payoff such that the more information revealed in the reports \hat{q} , the higher the potential payoff to the subjects involved, which will yield an information-rich equilibrium. Additionally, since private information is independent across individuals (it is truly private), the most lucrative 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 (\hat{p}_i) and their public information (\hat{q}_i). Once these vectors are reported, we still need to aggregate them, which we discuss in the next section.

Aggregating Information

Once we have a mechanism for extracting public beliefs from private ones, it is straightforward to add a public information generalization to Equation (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

$$P(s | I) = \frac{\left(\frac{p_{s1}}{q_{s1}}\right)^{b_1} \left(\frac{p_{s2}}{q_{s2}}\right)^{b_2} \dots \left(\frac{p_{sN}}{q_{sN}}\right)^{b_N}}{\sum_{\forall s} \left(\frac{p_{s1}}{q_{s1}}\right)^{b_1} \left(\frac{p_{s2}}{q_{s2}}\right)^{b_2} \dots \left(\frac{p_{sN}}{q_{sN}}\right)^{b_N}} \quad (6)$$

where the $\overset{i}{q}$ s are extracted from individuals' reports before they are aggregated. This correction allows us to isolate the private information from the individual reports. Note that Equation (6) yields a report on only private information; it entirely eliminates anything public.

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 $\overset{i}{p}_i = \{p_{1i}, p_{2i}, \dots, p_{Ni}\}$ and $\overset{i}{q}_i = \{q_{1i}, q_{2i}, \dots, q_{Ni}\}$, 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 add in public information more efficiently, we aggregate the individual reports of public information $\overset{i}{q}_i = \{q_{1i}, q_{2i}, \dots, q_{Ni}\}$ into a single vector $\overset{i}{q} = \{q_1, q_2, \dots, q_N\}$. 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 , thusly:

$$q_s = \frac{\sum_{i=1}^N \beta_i q_{si}}{\sum_{i=1}^N \beta_i} \quad (7)$$

Once we have completed this aggregation process, we can use the new vector $\overset{i}{q}$ in place of $\overset{i}{q}_i$ in the original function in Equation (6). If $\overset{i}{q}$ 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 Equation (6), we use a limited set J , a subset of the whole group:

$$q_s = \frac{\sum_{i \in J} \beta_i q_{si}}{\sum_{i=1}^N \beta_i} \quad (7a)$$

The β_i is determined by equation 3. The resultant forecast is then given by a modification of the GPIC in Equation (6). 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.

$$P(s | I) = \frac{q_s \left(\frac{p_{s1}}{q_s} \right)^{\beta_1} \left(\frac{p_{s2}}{q_s} \right)^{\beta_2} \dots \left(\frac{p_{sN}}{q_s} \right)^{\beta_N}}{\sum_{\forall s} q_s \left(\frac{p_{s1}}{q_s} \right)^{\beta_1} \left(\frac{p_{s2}}{q_s} \right)^{\beta_2} \dots \left(\frac{p_{sN}}{q_s} \right)^{\beta_N}} \quad (8)$$

Note that this mechanism yields a probability distribution that reflects both private and public information. Table 3 walks the reader through an example in which Equations (2), (6), and (8) are all used to aggregate information, with varying levels of success.

Experimental Design

In order to test this mechanism we conducted a number of experiments at Hewlett-Packard Laboratories in Palo Alto, California. The subjects were a mix of undergraduate and graduate students in economics and computer science at Stanford University and knew the experimental parameters discussed below, as they were part of the instructions and training for the sessions. The instructions and quizzes can be found online at: <http://www.hpl.hp.com/econexperiment/InfoExpte.htm>. The five sessions were run with nine to eleven subjects in each. Most students had participated in similar experiments some months earlier, which involved private information only. By the

final experiment, 90% of subjects were experienced. Table 4 details the percent of players with experience in each aggregation experiment.

We implemented the two-stage mechanism in a laboratory setting. Possible outcomes were referred to as states in the experiments. There were ten possible states, A through J, in all the experiments. Each had an Arrow-Debreu state security associated with it. The information available to the subjects consisted of observed sets of random draws from a computerized urn, with replacement. After privately drawing the state for the ensuing period, we created the virtual urn. This urn had twelve entries (or balls), one for each of the ten possible states, plus an additional two for the true (preselected) 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. A private draw is a draw from the urn which is shown to a single subject. A public draw was shown to multiple subjects. We used four variants on the information structure to ensure that the results obtained were robust. In the first structure, we provide two private and two public draws to all participants. In this case, since a subject would observe 3 draws, he has relatively little information about the outcome. However, the a group of ten would have observed 22 ($10 \times \text{number of private draws} + \text{number of public draws}$) draws and knows quite precisely what the outcome should be.

In the remaining experiments, all subjects received three private draws. In one, they also received one public draw, in another, only half of the cohort received the public draw, and in the final treatment all players received a public draw, but there were two different public draws available. Further details of the treatments can be found in Table 1.

The information market we constructed consists of an artificial call market in which the securities are traded (described below). The states are equally likely and randomly drawn. If a state occurred, the associated state security paid off at a value of 1,000 francs. Hence, the expected value of any given security, a priori, is 100 francs. Subjects were provided with some securities and francs at the beginning of each period.

Each period consisted of six rounds, lasting 90 seconds each. At the end of each round, the bids and asks were gathered and a market price and volume was determined. We employed a market clearing condition, such that supply equals demand in each round. No one sold an item for less than they asked for it, and no one paid more for an item than the offered to pay. The price at which all items transacted was the average of the marginal bids. The transactions were then completed and another call round began. At the end of six trading rounds the period was over, the true state security was revealed, and subjects were paid according to the holdings of that security. This procedure was then repeated in the next period, with no correlation between the states drawn in each period. The market phase of the experiments consisted of six to nine periods.

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.

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 Equation (3). The chart the subjects received showed the payoff for every possible ticket expenditure, and an excerpt from the chart is shown in

Table 5. The MYBB game is identical to the second stage played in Chen, Fine and Huberman (Chen et al 2001).

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 5.

Figure 1 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 has 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 $.49 \times 178 = 87$ francs for the AK game, for a total of 749 francs.

Figure 1: Sample Page from Stage Two of the Experiment

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			

Analysis

In order to analyze these results we first calculate an omniscient probability distribution for each period using every observation that was available to the individuals. This distribution is used as a limit-case benchmark. That is, only a perfect information aggregation mechanism should be able to achieve this distribution. We compare the resultant probabilities from information aggregation mechanisms to this benchmark by using the Kullback-Leibler measure (Kullback et al

1952). The Kullback-Leibler measure of two probability distributions p and q is given by:

$$KL(p, q) = E_p \left(\log \left(\frac{p}{q} \right) \right) \quad (9)$$

where p is the "true" distribution (in our case, the omniscient probability distribution). In the case of finite number of discrete states, the above Equation (9) can be rewritten as:

$$KL(p, q) = \sum_s p_s \log \left(\frac{p_s}{q_s} \right) \quad (10)$$

It can be shown that $KL(p, q) = 0$ if and only if the two distributions being compared, p and q , are identical, and that $KL(p, q) \geq 0$ for all probability distributions p and q . Therefore, the smaller the Kullback-Leibler number, the closer that two probabilities are 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 statistic for an entire experiment.

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 no information prediction serves as the first baseline to determine if any information is contained in our mechanism's predictions. Further, if a mechanism is really aggregating information, then it should be doing at least as well as the best individual. Therefore, the predictions of the best individual in the experiment serve as the second baseline, which helps us to determine if information aggregation indeed occurred in the experiments.

The first two information aggregation mechanisms we evaluate are the market prediction and the Chen, Fine, and Huberman (Chen et al 2001) mechanism in Equation (2). We calculate the market prediction by using the last traded prices of the assets. We use 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 aggregation mechanism is the original IAM, found in Equation (2). Recall that this mechanism was designed on the assumption of no public information. The purpose of its inclusion is to measure the performance degradation due to the double-counting issue of inherent to the presence of public information.

The third mechanism is our proposed improvement, referred to as the General Public Information Correction (GPIC) mechanism, given by Equation (6). 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.

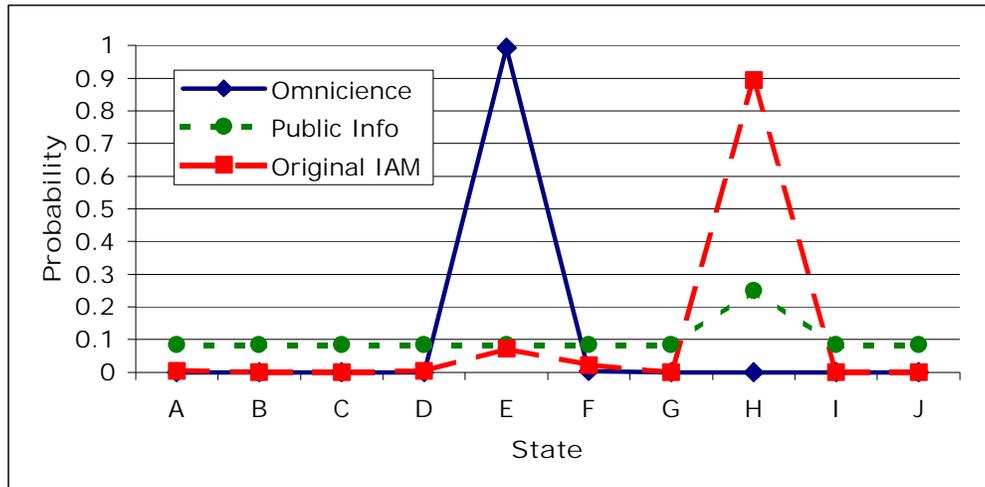
Lastly, we address the special case in which the experimenter knows that every individual receives the same public information. This fifth mechanism, referred to as the Special Public Info Correction mechanism (SPIC), recovers the public information by using the reports of only the best two individuals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

Results

We start by reporting not only the result of our public information experiments, but those of the original Chen, Fine, and Huberman (Chen et al 2001) paper as well (Experiments 1 through 5 on Tables 1 and 2). Recall that in these experiments, all information was independent and private. As is shown in Table 1 and Table 2, once even a small amount of public information is introduced into the system

(Experiments 6 through 10), the performance of the original IAM decreases dramatically. In Figure 2 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 8, 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 counting issue is so severe that the results are worse than that of the no information measure (see, for example, Experiments 6, 7 and 8). Thus, this verifies the necessity to derive a method correcting for the biases introduced by public information.

Figure 2: Illustration of the Double Counting Issue



Before we evaluate the public information mechanisms as a whole, we evaluate the accuracy of the AK game. If this game is working as planned, the \hat{q}_i s reported by the players should mirror the true public information distributed to them. Empirically, we have found that reporting varies a great deal, with some players reporting almost exactly the public information they have been provided, and some reporting far flatter distributions. In the experiments where there was more public information (two draws, as opposed to one), the KL between the true public information and the average report was 0.125. Furthermore, if we use the SPIC, and only take the best pair for each experiment, the KL is 0.027, as opposed to 0.15 were we to use a flat distribution instead.

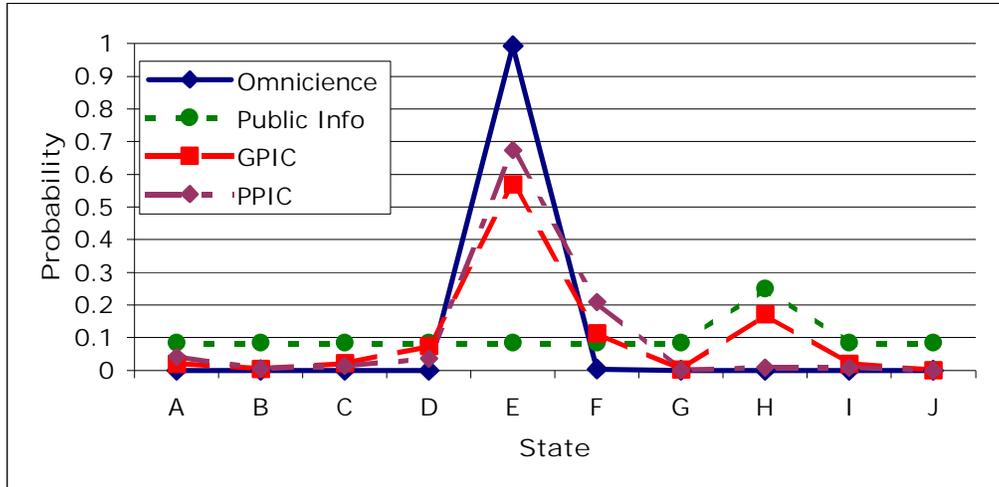
In Table 1 we summarize the relative performance, in terms Kullback-Leibler measures, of all of the benchmarks mechanisms enumerated above. Table 2 reports the same results in terms of the percentage relative to the no information Kullback-Leibler measure (indicating the level of improvement over this benchmark). Note that the amount of aggregate information available in an experiment varied across the treatments. Because the pure KL measure reported in is affected by the amount of underlying information, the percentage measurement in Table 2 are more useful when comparing results across experiments.

The GPIC mechanism (Equation 6) 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 algorithm in Equation (8). 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 (Experiments 1 through 5). 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 information is correct. Therefore, there is validity in our approach to teasing out this public information in the GPIC.

Figure 3 illustrates the efficacy of the GPIC. In this figure, once again, the results from Experiment 8, period 9 are plotted. The GPIC mechanism eliminates the false peak shown in Figure 2. However, the correction is not perfect. There is still some residual positive probability being placed on state H, the site of the false peak. When the PPIC is used to perform the correction, the false peak is completely eliminated.

Figure 3: Information Aggregation with Public Information Correction



In addition to using KL numbers, we can judge an information aggregation process by evaluating how often the correct outcome is predicted. For these purposes, we assume that the state with the highest probability weight is the mechanisms' prediction.

Table 6 looks at the public information experiments (experiments 6 to 10) and evaluates the frequency with which each mechanism correctly predicts the outcome.

Notice that, because of the manner in which information is generated in the experiment, even using all the information will not result in perfect predictions every time. The omniscient probability, which is calculated as if we have all the information in the experiments, predicts correctly an average of 74% of the time. Even if the mechanism were doing a perfect job, this would be the maximal accuracy level. As one can see, the general method (GPIC) is accurate 47% of the time. The SPIC predicts correctly 58% of the time.

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 Table 1 and Table 2 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.

Conclusion and Extensions

Accurate predictions are essential to individuals and organizations, and the ability to quickly gather, aggregate, and act on information is a key asset in today's economy. For large communities, information relevant to forecasts is often dispersed across people, frequently in different geographical areas. Our methodology addresses the needs for an implementable mechanism 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. It also mitigates the issues of redundant, public signals in a 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.

References

1. Anderson, L. and C. Holt (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87, 847-862.
2. Camerer, C. and K. Weigelt (1991): "Information Mirages in Experimental Asset Markets," *Journal of Business*, 64, 463-493.
3. Chen, K., L. Fine and B. Huberman (2001): "Forecasting Uncertain Events with Small Groups," *Proceedings of the ACM EC '01 Conference*.
4. Forsythe, R. and R. Lundholm (1990): "Information Aggregation in an Experimental Market," *Econometrica*, 58, 309-347.
5. Forsythe, R., T. Palfrey, and C. Plott (1982): "Asset Valuation in an Experimental Market," *Econometrica*, 50, 537-567.
6. Gordon T.J. and O. Helmer (1964): "Report on a Long Range Forecasting Study," The RAND Corporation-P-2982, (DDC AD 607777).
7. Hayek, F. (1945): "The Use of Knowledge in Society," *American Economic Review*, XXXV, No. 4, 519-30.
8. Iowa Electronic Markets: <http://www.biz.uiowa.edu/iem/>
9. Kullback, S. and R.A. Leibler, (1952): "On Information and Sufficiency," *Annals of Mathematical Statistics*, 22, 79-86.
10. Nöth, M., C. Camerer, C. Plott, and M. Weber (1999): "Information Traps in Experimental Asset Markets," submitted to *Review of Financial Studies*.
11. Nöth, Markus, and Martin Weber (1998): "Information Aggregation with Random Ordering: Cascades and Overconfidence," University of Mannheim, Discussion Paper, presented at the Summer 1998 ESA Meetings.
12. O'Brien, J. and S. Srivastava (1991): "Dynamic Stock Markets with Multiple Assets," *Journal of Finance*, 46, 1811-1838.
13. Pennock, D., S. Lawrence, C. Giles, and F. Nielsen. (2000): "The Power of Play: Efficient and Forecast Accuracy in Web Market Games," NEC Research Institute Technical Report 2000-168.
14. Plott, C. and S. Sunder (1982): "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational Expectations Models," *Journal of Political Economy*, 90, 663-698.
15. Plott, C. and S. Sunder (1988): "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 56, 1085-1118.
16. Scharfstein, D. and J. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80, 465-479.

17. Sunder, S. (1992): "Markets for Information: Experimental Evidence," *Econometrica*, 60, 667-695.

Table 1: Kullback-Leibler Numbers, by Experiment

Expt	Experimental Structure			Kullback-Liebler Values (Standard Deviation)						
	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
1	13	3 draws for all	None	1.977 (0.312)	1.222 (0.650)	0.844 (0.599)	0.553 (1.057)	N/A	N/A	N/A
2	9	3 draws for all	None	1.501 (0.618)	1.112 (0.594)	1.128 (0.389)	0.214 (0.195)	N/A	N/A	N/A
3	11	½: 5 draws ½: 1 draw	None	1.689 (0.576)	1.053 (1.083)	0.876 (0.646)	0.414 (0.404)	N/A	N/A	N/A
4	8	½: 5 draws ½: 1 draw	None	1.635 (0.570)	1.136 (0.193)	1.074 (0.462)	0.413 (0.260)	N/A	N/A	N/A
5	10	½: 3 draws ½: varied draws	None	1.640 (0.598)	1.371 (0.661)	1.164 (0.944)	0.395 (0.407)	N/A	N/A	N/A
6	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)
7	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)
8	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)
9	10	3 draws for all	½: 1 draw	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)
10	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 2: Percentage of No-Info Kullback-Leibler Numbers, by Experiment

Expt	Experimental Structure			Kullback-Liebler Values, as a Percent of the No Info Case						
	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
1	13	3 draws for all	None	100%	61.8%	42.7%	28.0%	N/A	N/A	N/A
2	9	3 draws for all	None	100%	74.1%	75.2%	14.3%	N/A	N/A	N/A
3	11	½: 5 draws ½: 1 draw	None	100%	62.3%	51.9%	24.5%	N/A	N/A	N/A
4	8	½: 5 draws ½: 1 draw	None	100%	69.5%	65.7%	25.3%	N/A	N/A	N/A
5	10	½: 3 draws ½: varied draws	None	100%	83.6%	71.0%	24.1%	N/A	N/A	N/A
6	10	2 draws for all	2 draws for all	100%	63.6%	70.0%	157.3%	61.94%	20.94%	24.53%
7	9	2 draws for all	2 draws for all	100%	69.0%	64.7%	205.0%	56.2%	18.2%	32.6%
8	11	3 draws for all	1 draws for all	100%	80.9%	61.9%	151.7%	43.0%	22.0%	40.1%
9	10	3 draws for all	½: 1 draw	100%	53.3%	67.1%	59.6%	50.0%	44.1%	49.7%
10	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%

Table 3: Example demonstrating the Aggregation Mechanisms

Private Signals			
	State A	State B	State C
Player 1	-	2	1
Player 2	-	1	2

Probability Predictions, by Player			
	State A	State B	State C
Player 1 p	40%	40%	20%
Player 2 p	40%	20%	40%
Player 1 q	70%	15%	15%
Player 2 q	50%	25%	25%

Public Signals			
	State A	State B	State C
Player 1	2	-	-
Player 2	2	-	-

Total Signals			
	State A	State B	State C
Omniscient signal	2	3	3
Omniscient Prob.	25%	38%	38%

Player 1's •	Private Info. Mechanism (Eq. 1)			General Public Info. Mechanism (Eq. 5)			Special Public Info. Mechanism (Eq. 7)		
	P(A)	P (B)	P (C)	P(A)	P (B)	P (C)	P(A)	P (B)	P (C)
0.2	46%	13%	40%	15%	21%	64%	29%	18%	54%
0.4	48%	16%	36%	14%	26%	60%	28%	22%	50%
0.6	49%	19%	32%	13%	32%	55%	27%	26%	46%
0.8	50%	22%	29%	11%	38%	51%	26%	32%	42%
1.0	50%	25%	25%	10%	45%	45%	25%	38%	38%
1.2	50%	29%	22%	8%	52%	40%	23%	44%	33%
1.4	49%	32%	19%	7%	59%	34%	21%	50%	29%
1.6	48%	36%	16%	6%	66%	29%	19%	56%	25%
1.8	46%	40%	13%	4%	72%	24%	17%	62%	21%
2.0	44%	44%	11%	4%	77%	19%	15%	68%	17%

Table 4: Percent of Participants with Experience, by Experiment

Expt	% Inexperienced	% Experienced in Private Info. Treatment	% Experienced in Public Info. Treatment	% Experienced in Either Treatment
1	100%	0%	0%	0%
2	62%	38%	0%	38%
3	78%	22%	0%	22%
4	45%	55%	0%	55%
5	50%	50%	0%	50%
6	70%	30%	0%	30%
7	56%	0%	44%	44%
8	27%	9%	64%	73%
9	40%	10%	50%	60%
10	10%	20%	70%	90%

Table 5: Excerpt from Payoff Chart used in the MYBB Game

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 6: Percent Accuracy of Information Mechanisms, by Experiment

Expt.	Omniscience	GPIC	PPIC	SPIC
6	86%	45%	68%	77%
7	67%	38%	52%	52%
8	78%	67%	83%	61%
9	70%	40%	55%	45%
10	70%	43%	61%	57%

Appendix I

Consider the following game.

There are M players facing an uncertain world. There are N possible states of the world. Each player receives information about the likelihood of possible states. There are two types of information: private and public. Each player is asked to report two probability distributions and their payoff in the game depends on these reports.

Definitions and notations:

- A set of players indexed by $i = \{1, \dots, M\}$
- A set of possible outcomes (states) $S = \{s_1, \dots, s_N\}$
- O_i^{priv} private information of player i
- O_i^{pub} public information of player i . Public information is defined as information observed by more than one player.
- $P(s|O)$ denotes the conditional probabilities of outcome s conditioned on information O .

The Game

- Stage one: Each player observes his or her information (private & public)
- Stage two: Each player is asked to report two probability distributions

$\mathbf{p}_i = \{p_{1i}, p_{2i}, \dots, p_{Ni}\}$ and $\mathbf{q}_i = \{q_{1i}, q_{2i}, \dots, q_{Ni}\}$ with the constraints $\sum_{s=1}^N p_{si} = 1$ and

$\sum_{s=1}^N q_{si} = 1$, and $q_{si} > \epsilon$ for all i (where ϵ is a small number chosen to bound the payoff away from negative infinity)

- Stage three: The true state, $x \in S$ is revealed and each player is paid

$$P = c_1 + c_2 * \log(p_{xi}) + f(\mathbf{q}_i, \mathbf{q}_j) * (c_3 + c_4 * \log(q_{xj}))$$

With the following conditions:

Condition F1: $f(\mathbf{x}, \mathbf{y})$ is maximal when $x_s = y_s$ for all s

Condition F2: c_2 and c_3 are positive real numbers.

Condition F3: c_1 and c_4 are non-negative real numbers.

Condition F4: j is chosen such that $j = \arg \max_k f(\mathbf{q}_i, \mathbf{q}_k)$

Condition F5: c_3 and c_4 are chosen such that $[c_3 + c_4 * \log(\epsilon)] > 0$

Note: The specific functional form $f(\overset{r}{x}, \overset{r}{y}) = \left(1 - \left[\sum_s |x_s - y_s|\right]/2\right)^2$ is used in the experiments. However, any continuously differentiable function that satisfies condition (F1) will be sufficient for the results derived below.

Assumptions about information:

Assumption A1: $P(s|O_i) = P(s|O_i^{\text{priv}}) P(s|O_i^{\text{pub}})$ for all i, s

Assumption A2: $P(s|O_i^{\text{priv}} \text{ and } O_j^{\text{priv}}) = P(s|O_i^{\text{priv}}) P(s|O_j^{\text{priv}})$ for all i, j, s

Assumption A3: There exists a j for every i such that $O_i^{\text{pub}} = O_j^{\text{pub}}$

Assumption A4: All individuals know their $P(s|O_i^{\text{priv}})$ and $P(s|O_i^{\text{pub}})$

Lemma 1: $\{\overset{i}{p}_i = P(s|O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}), \overset{i}{q}_i = P(s|O_i^{\text{pub}})$ for all $s, i\}$ is a Bayesian Nash equilibrium. That is, each player will report as $\overset{i}{p}_i$ his true conditional probability beliefs and as $\overset{i}{q}_i$ the beliefs conditioned solely on his public information.

Proof:

Assume all players but i are playing an equilibrium strategy. Player i 's maximization problem is

$$\begin{aligned} \underset{\{p_{is}, q_{is}\}_{s=1}^N}{\text{Max}} \sum_{s=1}^N P(s|O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}) \{c_1 + c_2 \log(p_{is}) + f(\overset{r}{q}_i, \overset{r}{q}_j) * (c_3 + c_4 * \log(q_{js}))\} \\ \text{s.t. } \sum_{s=1}^N p_{is} = 1 \text{ and } \sum_{s=1}^N q_{is} = 1. \end{aligned}$$

There will be at least one other player j that plays $q_j = P(s|O_j^{\text{pub}})$, since at least one player other than i observes the same public information and arrives at the same distribution Q_i .

The resulting Lagrangian is

$$\begin{aligned} L = \sum_{s=1}^N P(s|O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}) \{c_1 + c_2 \log(p_{si}) + f(\overset{r}{q}_i, \overset{r}{q}_j) * (c_3 + c_4 * \log(q_{sj}))\} \\ - l \left(\sum_{i=1}^N p_{si} - 1 \right) - m \left(\sum_{i=1}^N q_{si} - 1 \right) \end{aligned}$$

The first order condition is $\frac{P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{pub}})}{p_{si}} = 1$ for all i , which implies

$$P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}) = \lambda p_{si}$$

Summing over both sides, we get $1 = \lambda$. Thus $p_{si} = P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{pub}})$ for all i .

Rewriting the terms in the payoff functions that depends on q_{si} , we have:

$$f(q_i, q_j) = \sum_{s=1}^N P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}) (c_3 + c_4 * \log(q_{sj}))$$

Recalling condition (F1), $q_i = P(s | O_i^{\text{pub}})$ maximizes $f(q_i, P(s | O_i^{\text{pub}}))$. So all we

need to show is $\sum_{s=1}^N P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}) (c_3 + c_4 * \log(q_{js})) \geq 0$.

Using condition (F5):

$c_3 + c_4 * \log(\epsilon) > 0$, implies

$c_3 + c_4 * \log(q_{sj}) > 0$ for all j, s since $q_{js} \geq \epsilon$, therefore

$$\sum_{s=1}^N P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}) (c_3 + c_4 * \log(q_{sj})) \geq 0 \text{ since } P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{pub}}) \geq 0$$

QED.

Lemma 2: There are multiple equilibria to this game.

The same proof applies to : $\{ \hat{p}_i = P(s | O_i^{\text{priv}} \text{ and } O_i^{\text{priv}}) , \hat{q}_i = P(s | O_i^{\text{pub}}) = \frac{1}{N} \text{ for all } s, i \}$ or for that matter, any set of \hat{q}_i on which players coordinate.

Footnotes

ⁱ 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, Lawrence, Giles and Nielsen (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).