

Public and Private Information: An Experimental Study of Information Pooling

Author(s): Richard D. McKelvey and Talbot Page

Source: *Econometrica*, Vol. 58, No. 6 (Nov., 1990), pp. 1321-1339

Published by: The Econometric Society

Stable URL: <http://www.jstor.org/stable/2938318>

Accessed: 05-09-2017 21:19 UTC

REFERENCES

Linked references are available on JSTOR for this article:

http://www.jstor.org/stable/2938318?seq=1&cid=pdf-reference#references_tab_contents

You may need to log in to JSTOR to access the linked references.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://about.jstor.org/terms>



The Econometric Society is collaborating with JSTOR to digitize, preserve and extend access to *Econometrica*

PUBLIC AND PRIVATE INFORMATION: AN EXPERIMENTAL STUDY OF INFORMATION POOLING¹

BY RICHARD D. MCKELVEY AND TALBOT PAGE

This paper reports on an experimental study of the way in which individuals make inferences from publicly available information. We compare the predictions of a theoretical model of a common knowledge inference process with actual behavior. In the theoretical model, “perfect Bayesians,” starting with private information, take actions; an aggregate statistic is made publicly available; the individuals do optimal Bayesian updating and take new actions; and the process continues until there is a common knowledge equilibrium with complete information pooling. We find that the theoretical model roughly predicts the observed behavior, but the actual inference process is clearly less efficient than the standard of the theoretical model, and while there is some pooling, it is incomplete.

KEYWORDS: Information pooling, common knowledge, Bayesian learning.

1. INTRODUCTION

THIS PAPER REPORTS ON AN EXPERIMENTAL STUDY OF A COMMON KNOWLEDGE INFERENCE PROCESS. We consider a situation in which individuals, starting with differing private information, make decisions. A summary statistic of these decisions is aggregated into a publicly available signal, and individuals can then use this public data to augment their original private information. Then the process is iterated. The experiment is designed so that, under certain behavioral assumptions, truthful reporting of current private information is a Bayes-Nash equilibrium, and the equilibrium path leads to a common knowledge equilibrium, characterized by complete pooling of information.

The assumptions leading to a Bayes-Nash equilibrium are strong but they are typical of the assumptions in game theoretic analysis: Individuals are assumed to have common knowledge of the underlying information structure of the game and to have common knowledge of the rationality of the other players. They are assumed to make complicated mathematical calculations and apply Bayes theorem without error to update their beliefs.

The purpose of the experimental study reported here is to test how closely the predictions of a theoretical model of common knowledge inference approximate actual behavior. We investigate this question for a model which incorporates Bayesian learning from a public signal (it is one of the simplest models of a common knowledge inference process which does not pool all the information in the first iteration). The inference process of the model is the same as in Kobayashi (1977), Jordan (1982), and McKelvey and Page (1984). Thus the experiment provides, in a simple setting, a test of these theoretical models.

¹We acknowledge support of NSF Grant No. IST 85-13679 to the California Institute of Technology. We thank Jim Snyder for help with the experiments and data analysis, and Colin Camerer and a referee for useful comments on an earlier draft. The data from the experiments reported here are available on request from the authors.

To explore how well the theoretical model predicts behavior, we looked at three types of evidence. First we looked at individual use of private information. Second we looked at use of public information. And third we constructed an empirical measure of the closeness of fit between the prediction and the actual behavior. In focusing on the predictive performance of the model, we did not attempt to evaluate directly the realism of the assumptions. (Our own view is that the assumptions are unrealistically demanding, and this provides a motivation to investigate the predictive performance of this and similar models.)

The literature relating to our experiment can be briefly traced. Aumann (1976) initially defined the notion of common knowledge and showed that when two individuals start with a common prior over a probability space, if their posterior probabilities for some event become common knowledge, then the two posterior probabilities must be the same, with a consensus in beliefs. Geanakoplos and Polemarchakis (1982) specified an iterated process of information exchange and showed that for the finite case the process would terminate in a state where the posterior probabilities were the same and (generically) would fully reveal the original, privately held information.

In a previous paper (McKelvey and Page (1985)), we generalized these results to the case where there are n individuals, and only an aggregate statistic (or signal) of the individuals' posterior probabilities is made public. In the generalization, the iterated process is a sequential reaction to the public signal (identical to Kobayashi's and Jordan's). We showed in the finite case that, under a monotonicity condition on the signal, the iterated process leads to consensus, and in many cases to complete pooling of the original, privately held information. Brandenberger and Geanakoplos (1987) and Nielsen (1987) generalized the common knowledge theorem from an event to a random variable, and several researchers independently discovered a simple proof of the more general result (see Nielsen et al. (1989)). The experiment consists of a simple example of the general learning process and uses a signal which satisfies the monotonicity condition.

In previous experimental work, Forsythe, Palfrey, and Plott (1982), Plott and Sunder (1988), and McKelvey and Ordeshook (1985) also investigated how individuals use public information to augment their original private information, and whether in doing so, a rational expectations equilibrium is attained. In a laboratory securities market, Forsythe, Palfrey, and Plott found that through an iterated process there was rapid convergence to a rational expectations equilibrium. In another securities market experiment, Plott and Sunder found that when individuals have the same preferences, and when there is a complete set of contingent claims markets, there is rapid pooling of information, consistent with a rational expectations equilibrium; however, they also found that when preferences differ and when there is a single asset without contingent claims, there is relatively little pooling of information. In a voting model, where voters have access to poll data, McKelvey and Ordeshook (1985) found information pooling in which initially uninformed voters gain information from informed voters through the aggregate signals of a poll.

In Forsythe, Palfrey, and Plott and in Plott and Sunder the inference processes are complicated because of the enormous number of potential interactions among the individuals, and the optimal inference processes are not analyzed. In McKelvey and Ordeshook the inference process is analyzed but the working assumption is not altogether satisfactory (that during the adjustment process each of the uninformed voters believes each of the others is informed in spite of the changing poll results which imply that others are uninformed as well). Because of the simplicity of design used in this experiment we are able to analyze the optimal inference process in closed form.

2. EXPERIMENTAL DESIGN

The study consists of several experimental sessions, each consisting of four individuals participating in four experiments. At the beginning of a session the instructions (Appendix A) are handed out and read aloud. So the instructions are common knowledge. In each experiment of a session there are three subjects and one assistant, with a different assistant for each experiment. Thus, each of the four individuals participates as a subject in three of the four experiments and as an assistant in the fourth.

Each experiment in a session consists of an initial seeding of information (described below) followed by six rounds. At the end of every round each subject reports a number between zero and one. Each report is made privately to the experimenter and is unknown to the other subjects. The subjects know that at the end of a session they will be rewarded on the basis of their reports, the true state (A or B), and the other experiments of the session. The payoff rule, discussed in more detail below, provides an incentive, for each round, for each subject j to report j 's current posterior probability for A . (We will call a report "truthful" if it equals the current posterior probability for A .) The subjects in the experiment are not informed of these theoretical expectations, and in fact, most are not acquainted with probability theory. The experiments are conducted in a way so that the subject need not understand that their reports are posterior probabilities.

We conducted six sessions with inexperienced subjects who were participating in the experiment for the first time, and four sessions with experienced subjects who had participated previously. In all we conducted 10 sessions with 40 experiments. The subjects were undergraduates from the California Institute of Technology.

Each experiment has the following structure:

Seeding of Initial Information: At the beginning of each experiment the assistant chooses, by chance, one of two true states of the world, which we call state A and state B . Each state has probability of one-half of being chosen. The subjects do not know which state is chosen, but each is given the following private information on the event: if the state is A , each subject observes the number of successes in ten trials from a Bernoulli process with parameter 0.6. If the state is B , each subject observes the number of successes in ten trials of a

Bernoulli process with parameter 0.4. So, among the three subjects, a total of 30 independent trials are observed, and each subject's private, initial information is based on a different ten trials from the same process.

Round One: After receiving their private information, each subject is asked to report a number between zero and one. The three reports are then collected, and the simple average of the three is made publicly available, by writing it on the blackboard. The subjects can then make inferences from it to augment their private information.

Round Two: Next, a second report is elicited from each of the subjects. It is important to note that the subjects observe no new private information prior to making their second and subsequent reports. The second round reports of all the subjects are collected and the average written on the blackboard, as the second public signal.

Rounds Three through Five: The same procedure is repeated: a third report is elicited and the average made public; a fourth report elicited and the average made public; and a fifth report elicited (but not made public).

Round Six: In the sixth round of each experiment we make all the original private information public, by writing on the blackboard the total number of successes in the 30 Bernoulli trials. Then the sixth and final report is elicited. The last round provides evidence on whether or not the subjects believe that there is a pooled information equilibrium by the fifth round. If they do, there should be no change from the fifth to the sixth reports.

The Incentive Mechanism: The incentive mechanism used in the experiment is based upon a Proper Scoring Rule (PSR). PSRs provide incentives for an individual j to reveal j 's current posterior probability of an event. More specifically, if j is asked to report a number and is rewarded by a PSR for an event E , j will maximize j 's expected payoff by reporting j 's current posterior probability for E (see Savage (1971)).

The PSR used in this study is a quadratic loss rule (or the Brier rule). This rule has been used extensively in the evaluation of probabilistic weather forecasts (see Murphy and Winkler (1970)), and is the same in expectation as the rule used by Grether (1981), except for an extra gift in Grether's mechanism. The quadratic loss rule has the property of providing incentives for research which are invariant with respect to the current posterior probability (see Page (1988) for a discussion of research incentives from PSR).

The quadratic loss rule is defined as follows. If the state is A , write $s = 1$; if the state is B , write $s = 0$. Let individual j 's report be r (and remember that j makes j 's report before j observes the true state). Then define individual j 's payoff as

$$(1) \quad 1 - (r - s)^2.$$

Writing j 's current posterior probability for A as q , it is easily checked that j 's expected payoff is

$$(2) \quad (2r - r^2)q + (1 - r^2)(1 - q)$$

and that this is maximized at $r = q$.

In an experimental context, there is a well-known problem with PSRs. They are “flat”—even though they provide an incentive for (risk neutral) individuals to report their current posterior probability, they do not provide a strong incentive, relative to the experimenter’s expected cost of running the experiment, for the subjects to make inferences and calculate probabilities. The quadratic loss rule illustrates this problem. In our normalization of the rule (1), the individual’s payoff is bounded between zero and one. Thus the rule exhibits individual rationality in the sense that individual j cannot lose by participating in the experiment. However, n can be guaranteed 0.75 by neglecting j ’s private information and reporting 0.5. In the extreme case where j has perfect information j could only earn an additional 0.25 by reporting j ’s current posterior probability (by reporting $r = 1$ when the state is A , and $r = 0$ when the state is B). Usually, as in the case of this experiment, an individual’s information is imperfect, so that only a fraction of the remaining 0.25 is obtainable through optimal inference and reporting.

To sharpen the incentive of the scoring rule, and to relax the assumption of risk neutrality, we modify the payoff rule in two ways in our experiments. First, like Grether, we use a lottery version of a PSR to relax the assumption of risk neutrality. For each report of each subject, the following lottery is conducted at the end of the session. If subject j reports r for a particular round and the true state is A , j ’s probability of winning the lottery is $2r - r^2$; if j reports r and the true state is B , j ’s probability of winning the lottery is $1 - r^2$. A table of probabilities, for each r in increments of 0.02, is provided in the instructions to the subjects (See Table A1 in Appendix A). By the construction of the lottery, an individual’s probability of winning a lottery is given by expression (2) and that probability is maximized by reporting $r = q$.

Second, instead of paying a fixed amount for each lottery won, we pay according to a sliding scale. After all the lotteries in the session are conducted, we determine the total number of wins (6 reports per subject in each experiment, times 3 experiments per session, equals 18 reports and 18 lotteries for each subject in a session). Then an individual’s payment is made according to Table I.

The second modification helps increase the research incentive of the scoring rule by increasing the possible variance in payoffs in the range of likely

TABLE I

Number of Wins	Payment in Dollars
18	22
17	19
16	16
15	13
14	10
13	7
12	4
11 or fewer	3

outcomes. However, this modification introduces the potential of “bankruptcy” midway through the experiment: i.e., if the lotteries are computed as the experiment proceeds, then subjects might find themselves in a position where they have already lost more than six lotteries, and hence any reports they make from then on are irrelevant for their payoff. To avoid this problem, the lotteries are conducted at the end of the session. Thus, although there is an ex-post possibility of bankruptcy, this does not affect the incentive properties of the rule in an ex-ante sense, since the subject always has some positive expectation of not being bankrupt when making the decision.

Under a fairly weak assumption on preferences, truthful reporting of current posterior probabilities is a Bayes-Nash equilibrium to the game defined by a session. The assumption is that each individual prefers one lottery to a second if the first stochastically dominates the second. See Appendix B for details.

3. THE COMMON KNOWLEDGE INFERENCE PROCESS

In this section, we describe the common knowledge inference process, and show how it applies to the experiment we have designed.

We start with a set of n individuals, $N = \{1, 2, \dots, n\}$, who each have a common prior probability distribution over a probability space. Thus, we have the following structure:

- $(\Omega, \mathcal{F}, \rho)$, a probability space,
- $A \subseteq \Omega$, an event of interest,
- $\pi^1 = (\pi_1^1, \pi_2^1, \dots, \pi_n^1)$, a collection of initial information partitions,
- $h: [0, 1]^n \rightarrow \mathbb{R}$, an aggregation function.

We assume Ω is finite, with $\rho(\omega) > 0$ for all $\omega \in \Omega$. We let \mathcal{F} be the set of all subsets of Ω , and $\rho: \mathcal{F} \rightarrow \mathbb{R}$ be a probability measure representing the common prior. The partition, π_j^1 , is individual j 's initial information partition, with $\pi_j^1(\omega)$ representing the element of π_j^1 containing ω . So if ω occurs, then individual j is informed that $\omega \in \pi_j^1(\omega)$. Since Ω is finite, we can enumerate $\pi_j^1 = \{\pi_{ji}^1: 1 < i < I_j\}$, where π_{ji}^1 is the i th element of j 's initial partition. Finally, $h: [0, 1]^n \rightarrow \mathbb{R}$ is an aggregation function, where $[0, 1]$ is the closed unit interval.

We now consider the following iterative process, which is defined inductively on the period t . For each $j \in N$, define, for any $\omega \in \Omega$,

$$q_j^t(\omega) = \rho(A | \pi_j^t(\omega)),$$

$$q^t(\omega) = (q_1^t(\omega), \dots, q_n^t(\omega)),$$

and, for any $\omega \in \Omega$, define $H^t(\omega) = \Omega$, and

$$H^{t+1}(\omega) = \{\omega' \in H^t(\omega) | h(q^t(\omega')) = h(q^t(\omega))\},$$

$$H^{t+1} = \{H^{t+1}(\omega) | \omega \in \Omega\},$$

$$\pi_j^{t+1} = H^{t+1} \vee \pi_j^t = H^{t+1} \vee \pi_j^1.$$

Here $H^{t+1} \vee \pi_j^t$ is the refinement of the two partitions, H^{t+1} and π_j^t .

This process represents the sequence of events that occurs if individuals iteratively learn from public information. Thus, $q_j^t(\omega)$ represents j 's posterior probability for A given j 's current private information. We assume that at time t , each individual computes $q_j^t(\omega)$, and then $h(q^t(\omega))$ is publicly announced. Then H^{t+1} is a partition representing the publicly available information at time $t + 1$, while π_j^{t+1} is a partition representing individual j 's private information (as refined by the public information) at time $t + 1$.

Using results of Bergin and Brandenburger (1990) on the equivalence of additive separability and stochastic regularity, the theorem of McKelvey and Page (1986) can be stated:

THEOREM: *If h is additively separable into strictly monotonic components, then there is a T such that, for all $t \geq T$, and all $\omega \in \Omega$, $q_j^t(\omega) = \rho(A|H^t(\omega))$.*

Thus, eventually, the iterative process of public announcement will achieve consensus.

Our experiment parallels exactly the above theoretical model: In our experiment, $n = 3$, so $N = \{1, 2, 3\}$. Then, we set

$$\Omega = S \times Y^3, \quad \text{where} \quad S = \{0, 1\} \quad \text{and} \quad Y = \{0, 1, 2, \dots, 10\},$$

$$\rho(\omega) = \begin{cases} (.5)(.6^{y_1+y_2+y_3})(.4^{30-y_1-y_2-y_3}) & \text{if } s = 1, \\ (.5)(.4^{y_1+y_2+y_3})(.6^{30-y_1-y_2-y_3}) & \text{if } s = 0, \end{cases}$$

where $\omega = (s, y_1, y_2, y_3)$,

$$A = \{\omega \in \Omega | s = 1\},$$

$$\pi_j^1 = \{\pi_{ji}^1 | 0 \leq i \leq 10\},$$

where

$$\pi_{ji}^1 = \{\omega \in \Omega | y_j = i\},$$

and

$$h(r) = (r_1 + r_2 + r_3)/3,$$

where h is additively separable into the strictly monotonic components, $r_j/3$.

By direct calculation we find that for the experiment the refinement process terminates in three periods ($T = 3$) and that at the beginning of the third period if all the players are fully rational Bayesians, each has obtained the pooled information. Moreover, there is complete pooling by the third period if the subjects are constrained to report in even hundredths as in the table of Appendix A.

To illustrate the process of inference, we give an example. We suppose that the first subject observes 4 successes, the second subject observes 8, and the third observes 5, so $y = (y_1, y_2, y_3) = (4, 8, 5)$. Before the private information is observed, each subject has a prior of $\rho(A) = .5$ that the state is A . With the private information, each player's posterior probability is calculated by Bayes

TABLE II
FIRST PERIOD POSTERIOR PROBABILITIES OF A

Number of Successes y_j	Posterior Probability $\rho(A y_j)$
0	.02
1	.04
2	.08
3	.16
4	.30
5	.50
6	.70
7	.84
8	.92
9	.96
10	.98

theorem, as shown in Table II. Thus, the first player's posterior probability is .30, the second player's is .92, and the third is .50. So $q^1(\omega) = (.30, .92, .50)$. At this point, the public information, $h(q^1(\omega)) = (.30 + .92 + .50)/3 = .57$ is announced, and each subject can now try to learn from the public information.

Now consider subject 1's inference problem. Using the probabilities of Table II, subject 1 can construct Table III, which associates a public signal with each possible combination of private information for subjects 2 and 3. The (i, j) th element of the table is the public signal which results if $y = (4, i, j)$. From this calculation, subject 1 learns that the actual distribution of information must be one of two cases: (i) both subjects 2 and 3 observed 6 successes, or (ii) one subject observed 5 and the other observed 8. In this case, by using the public signal to augment 1's private information, subject 1 comes close to inferring (a sufficient statistic for) the pooled information. Subject 1 learns that the total number of successes for the 30 Bernoulli trials is either $4 + 6 + 6 = 16$ or

TABLE III
THE FIRST PUBLIC SIGNAL AS A FUNCTION OF PRIVATE
INFORMATION FOR SUBJECTS 2 AND 3^a

		Subject 3's Private Information (y_3)									
	0	1	2	3	4	5	6	7	8	9	10
0	.11	.12	.13	.16	.21	.27	.34	.39	.41	.43	.43
1	.12	.13	.14	.17	.21	.28	.35	.40	.42	.43	.44
2	.13	.14	.15	.18	.23	.29	.36	.41	.43	.45	.45
3	.16	.17	.18	.21	.25	.32	.39	.44	.46	.47	.48
4	.21	.21	.23	.25	.30	.37	.43	.49	.51	.52	.53
5	.27	.28	.29	.32	.37	.43	.50	.55	.57	.59	.59
6	.34	.35	.36	.39	.43	.50	.57	.62	.64	.65	.66
7	.39	.40	.41	.44	.49	.55	.62	.67	.69	.71	.71
8	.41	.42	.43	.46	.51	.57	.64	.69	.71	.73	.73
9	.43	.43	.45	.47	.52	.59	.65	.71	.73	.74	.75
10	.43	.44	.45	.48	.53	.59	.66	.71	.73	.75	.75

^aSubject 1's private information is $y_1 = 4$. The (i, j) th element of the table is the public signal resulting if $y_2 = i$ and $y_3 = j$.

$4 + 5 + 8 = 17$. With this new information, subject 1 calculates an updated posterior probability of \mathcal{A} , namely $q_1^2(\omega) = .74$, and submits this number as 1's second period report. Similarly, subjects 2 and 3 make their own inferences and submit their second period reports. The second aggregate signal is made public. From a similar line of reasoning, subject 1 can now make further inferences from the new public data and infer that the only way for this public signal to come about is for one of the others to have observed 5 successes and the other subject to have observed 8. Thus in the third period, subject 1 can infer that the total number of successes in the pooled sample is 17. The other subjects are also able to infer the pooled information after the second period signal, hence in the third and successive rounds, all subjects will have arrived at consensus, and their estimates of the posterior probability will all be based on the total pooled information available to the subjects.

4. EXPERIMENTAL RESULTS

Figure 1 illustrates the sequence of reports, as well as the optimal reports, for four typical experiments. It is apparent from the figure that, at best, the match of the actual to the optimal path is inexact. Other plots of actual and optimal paths show similar divergences. Given the demanding assumptions and complicated nature of the model, we do not find the inexact match surprising, and we go on to ask more specifically how closely the model predicts actual behavior.

Predictions for an Individual Bayesian Rationalist: In this subsection we consider predictions about individual choice that do not depend on that individual's ability to analyze the common knowledge inference process.

First, because of the rounding off in the table of probabilities in Appendix A, it is clear that it is never in the interest of a rational subject to report either 0 or 1.² By the table, if j 's posterior probability is q and j reports 0, j 's lottery probability is $(1 - q)$; if j reports .02, j 's lottery probability is $(1 - q) + (.04)(q)$. Since information is always incomplete in the experiment, $0 < q < 1$ and by monotonicity the report .02 is always preferable to the report 0. By a similar argument the report of .98 is always preferred to a report of 1.00. Of the total of the 720 reports in all sessions, 47 were either 0 or 1. The inexperienced subjects reported 0 or 1 about 7% of the time, the experienced about 6%.

Second, in the first and sixth rounds, we can analyze how well the subjects update based on an application of Bayes rule. In the first round, if all subjects are rational and adopt their Bayes-Nash equilibrium strategies for the entire experiment then they should each report truthfully their posterior probabilities. In the sixth round, all the original private information is made public, and each subject observes the total number of successes in the 30 Bernoulli trials. It is now a dominant strategy for each j to truthfully report j 's posterior probability. In Figure 2a and Figure 2b we plot the predicted optimal reports against the actual reports for rounds 1 and 6. If the predictions were perfect, all the points

² In the sixth round, it is a dominated strategy to report 0 or 1. In the earlier rounds, if all other subjects are rational and adopt their Bayes-Nash equilibrium strategies for the entire experiment, then a similar argument can be made that for these rounds there should be no reports of 0 or 1.

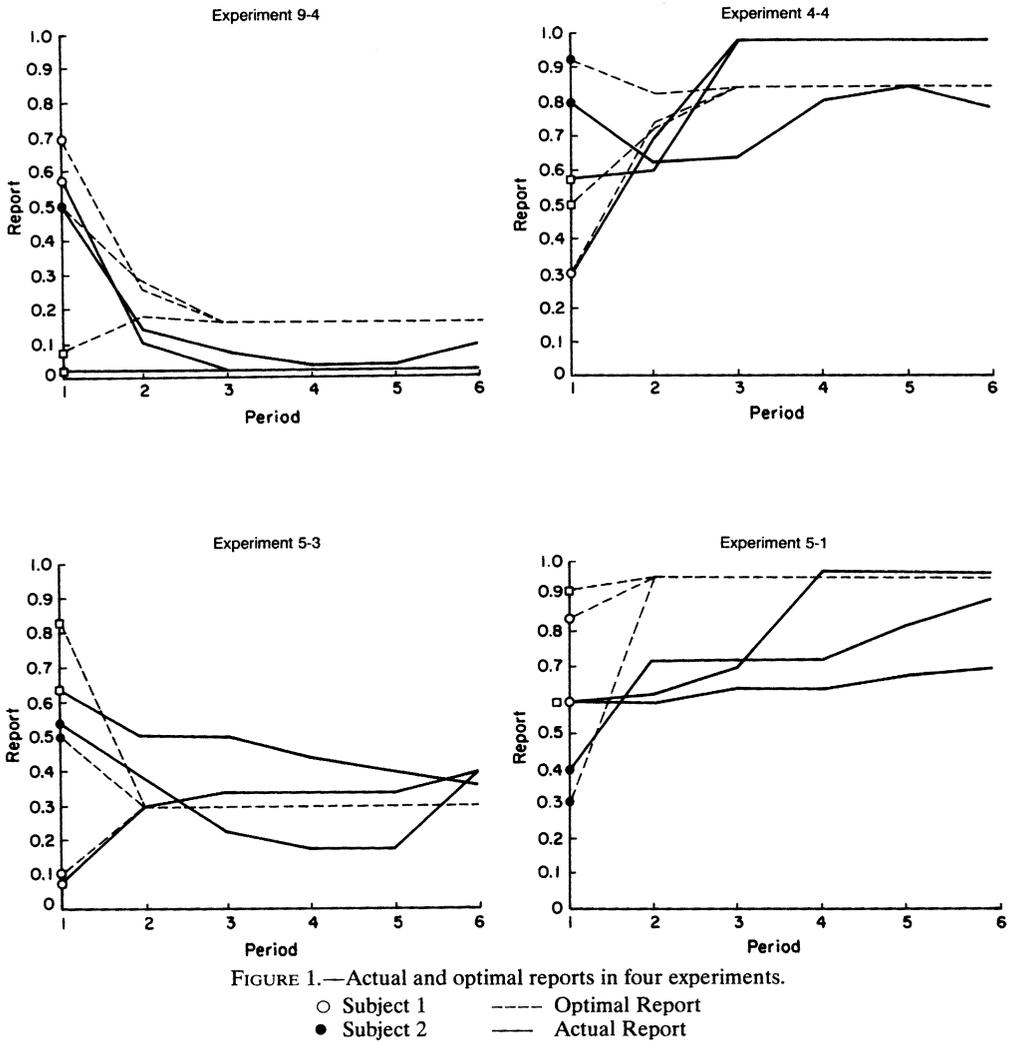


FIGURE 1.—Actual and optimal reports in four experiments.

- Subject 1
- Subject 2
- Optimal Report
- Actual Report

would be on the diagonal. Instead, what we see is substantial variance, with subjects tending to report posteriors that are a convex combination of the prior and the correct posterior. For example, in round 1, a subject who observes four successes reports $r = .4$ instead of the correct posterior probability of 0.3; a subject who observes 9 successes reports $r = .8$ instead of the correct .96. This pattern, referred to as underconfidence by Edwards (1982), suggests that when faced with a single application of Bayes theorem, subjects underestimate the power of the data to modify the prior (see also Grether (1980)).³ We see the same pattern in the sixth round reports.

³ An alternate interpretation is that the subjects maintain a skeptical attitude towards the veracity of the experimental procedures. While we did not attempt to deceive the subjects and went out of our way to insure that the subjects could verify the experimental procedures we used were those we described in the instructions, we cannot eliminate this possibility. If the subjects distrusted any aspect of the experimental procedures, underconfidence could be a rational response.

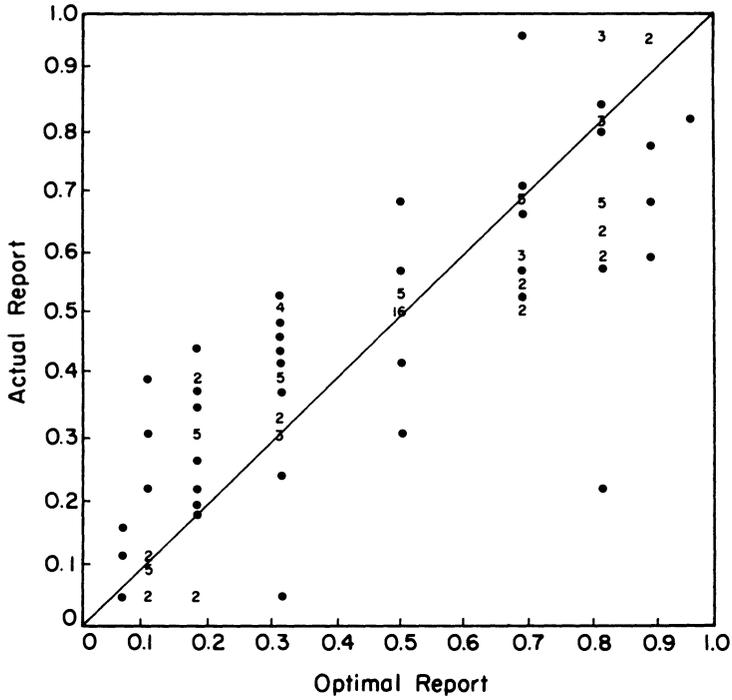


FIGURE 2a.—Actual report vs. optimal report (round 1).

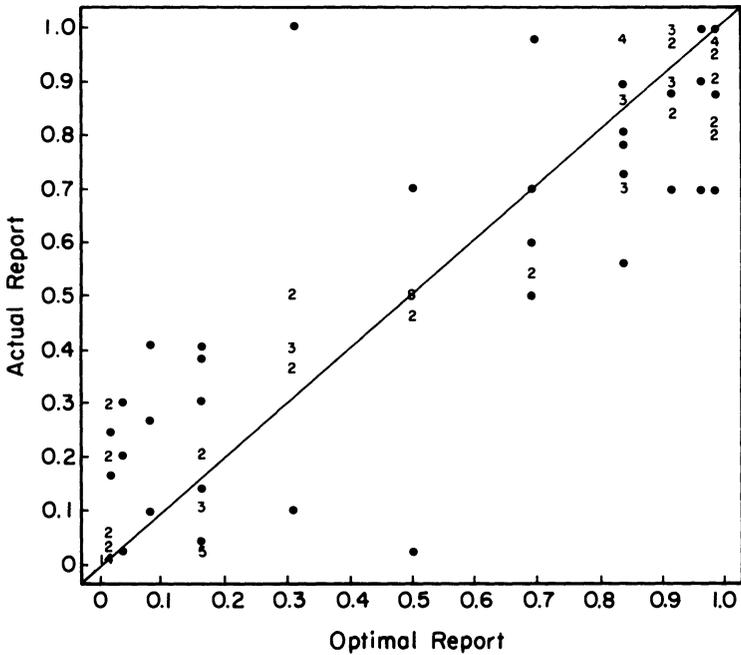


FIGURE 2b.—Actual report vs. optimal report (round 6).

TABLE IV
REGRESSIONS FOR ALL TEN SESSIONS
(Standard errors in parentheses, $n = 120$)

	C	$LPRI$	$LPUB_{t-1}$	$LPOOL$	R^2
$LREP_1$.009 (.088)	.805 (.059)			.61
$LREP_2$	-.013 (.087)	.642 (.069)	.966 (.110)		.68
$LREP_3$.193 (.104)	.340 (.084)	1.093 (.080)		.77
$LREP_4$.092 (.110)	.332 (.093)	.889 (.066)		.77
$LREP_5$.147 (.119)	.219 (.100)	.892 (.066)		.75
$LREP_6$.052 (.110)	.141 (.105)	.537 (.095)	.382 (.091)	.79

Qualitative Predictions from the Whole Model: In the experiment, the more complicated inference problem arises in rounds 1–5, where there is both private and public information to untangle. Here, roughly speaking, there could be two types of departure from the predicted, optimal path. The subjects could rely on the public information too little, in which case the public signal would not pool the private information. Or they could rely on the public signal too much, by discounting their private information, with the danger that the public signal becomes a random walk. If the subjects make the optimal inferences discussed in Section 2, they would begin by relying heavily on the private information and then gradually shift their reliance to the public signal as it came to incorporate the pooled information.

Table IV shows a series of regressions for all 10 sessions. $LREP_t$ is the log odds of the subjects' report in the t th round (reports of 0 were set to .02 and reports of 1.00 were set to 0.98). $LPRI$ is the log odds of the correct Bayesian posterior probability, given a subject's private information. $LPUB_{t-1}$ is the log odds of the public information in round $t - 1$, purged of the individuals' own immediate contribution. (From a public signal, each subject can infer the average of the others' reports for that round. $LPUB_{t-1}$ is the log odds of the average of the others' reported probabilities for the preceding round. When subject j is making a report, it is the public signal from the preceding round which is available to j .) $LPOOL$ is the log odds of the correct Bayesian posterior probability, given the pooled information. We use the log odds form to provide greater symmetry in the error term. Regressions with the untransformed variables, with the odds, and with the log forms yield similar results, as do regressions disaggregated to the sessions of inexperienced and experienced subjects.

We define the null hypothesis to be the case in which all the players are fully rational Bayesians. Under this hypothesis, in the first regression, if the subjects were making optimal inferences, the coefficient for the constant, C , would be 0,

and the coefficient for $LPRI$ would be 1. On the basis of the data, we can reject this hypothesis at level $\alpha = .01$. The underconfidence noted in Figure 2 shows up here as a coefficient of $LPRI$ less than 1 (it is .81).

Under the null hypotheses, in regressions 4 and 5, the coefficients for C would be 0 and the coefficient for $LPUB_{t-1}$ would be 1. These regressions also suggest rejecting the null hypothesis (although not so strongly). Under the null hypothesis, in regression 6, if the subjects were rational Bayesians (following the Bayes-Nash equilibrium path), $LPUB_{t-1}$ and $LPOOL$ would be equal, implying a singular moment matrix. If the subjects were “almost Bayes rational” or there were round off errors, the moment matrix would be nonsingular and the variables $LPUB_{t-1}$ and $LPOOL$ should explain most of the variance, but because of the high colinearity of $LPUB_{t-1}$ and $LPOOL$, these variables would have large standard errors, and perhaps be individually insignificant. The coefficients of C and $LPRI$ would be zero. If each subject were rational, but believed the others were not, the coefficient of $LPUB_{t-1}$ would be zero, and the coefficient of $LPOOL$ would be 1. The regression in period 6 does not fall into any of these categories. The coefficients of C and $LPRI$ are not significantly different from zero, but we see that both $LPUB_{t-1}$ and $LPOOL$ have significant effects.

Regressions 1, 4, 5, and 6 are correctly specified under the null hypothesis, but not 2 and 3. The correct specification for the latter two regressions would be a highly complicated nonlinear form. Nonetheless, these regressions are useful for qualitative comparison over the entire iterative process.

The regressions rule out complete Bayesian rationality. But they nonetheless suggest a partial rationality. Note that in the first round, subjects rely heavily on their private information, but in each succeeding round they rely less on their private information and more on the public information. Here we see clearly the pattern whereby the subjects gradually discard their private information ($LPRI$) as it gets incorporated into the public signal ($LPUB_{t-1}$). The coefficients of $LPRI$ begin at .81 in round one and decline monotonically through round five, with the standard errors monotonically increasing. This is the pattern we would expect if the subjects were making inferences somewhat close to the optimal inferences.

An Empirical Measure of Predictive Fit: To see how close the subjects come to optimal inference, we construct a measure of efficiency. First, we calculate ERO_t —the expected payoff in period t if all the subjects are rational. (To simplify the calculation we use a payoff structure in which winning a lottery is worth \$1 and losing a lottery \$3.) Next, we calculate ERA_t —the expected payoff in period t for the actual reports. These two expectations are shown in Figure 3, for sessions 1–6 (inexperienced subjects), and sessions 7–10 (experienced subjects). By construction of this measure of efficiency, if subject j neglected j 's initial private information and the subsequent public information and reported only on the basis of j 's original prior, j 's ERA_t would be zero for all rounds.

The plots of ERA_t show that for all three groups in each succeeding period from 1 to 4 the expected payoff increases for the actual reports. This means that

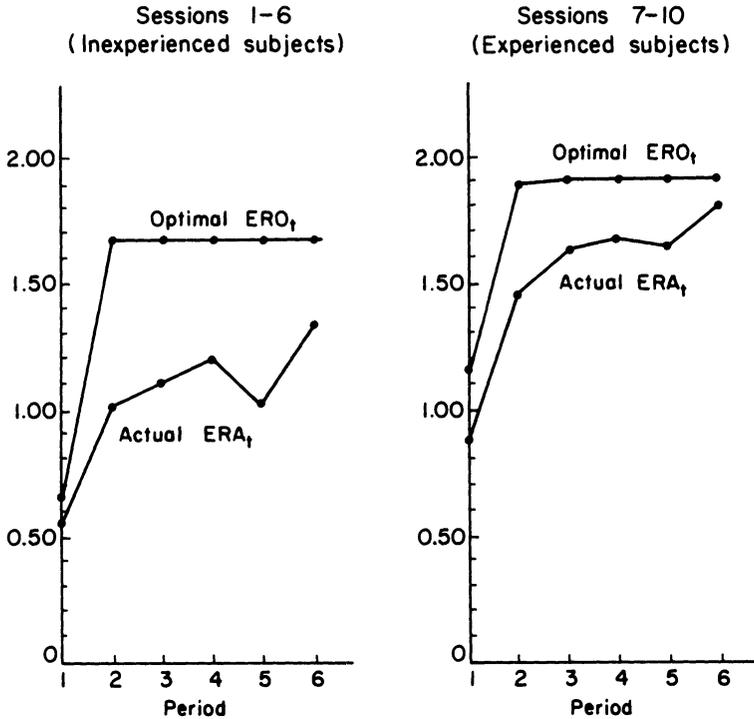


FIGURE 3.—Optimal and actual expected payoffs, by period.

the subjects are making use of the private information in the first period and both private and public information in succeeding periods up to period 5. Compared to the optimal standard, the subjects are a little slow in beginning their adjustment and a little slow in finishing it (they seem to have “inferential inertia”).

In period 5, ERA_t actually declines. By the 5th period, preceding errors in inference and computation are fed back into the public signal, some several times, with possible compounding of errors. Thus, while the public signal tends to aggregate the originally dispersed information, as can be seen by the increase of ERA_t in rounds 1-4 and by the regressions, it may also become contaminated as it compounds preceding errors. In sessions 1-6 and 7-10 the jump in ERA_t from round 5 to 6 suggests that the subjects do not believe that the fourth public signal is an error-free pooled information posterior probability (which it isn't).

We define the relative efficiency for a whole session to be ERA_t/ERO_t . The efficiency for the inexperienced subjects (sessions 1-6) was 69 percent; for the experienced subjects (sessions 7-10) was 85 percent. But while it appears that there is indeed improvement in inference through learning, not all of this difference is attributable to learning.

5. CONCLUSION

In this experiment, individuals make inferences on the basis of both private and public information. The experiment illustrates operational steps by which the behavior of rational individuals leads to a common knowledge equilibrium. The experiment is designed so that if the subjects are rational and each believes the others are rational, there should be complete pooling of information (in a utility relevant sense) by the third period. The optimal inference process provides a standard of reference to measure the efficiency of inference of actual individuals.

In the experiment, we find clear evidence that individuals make use of the public information to augment their original private information and increase their expected payoffs. The public signal comes to incorporate the original private information. However, there is also clear evidence that the actual process of information pooling is imperfect.

In the optimal inference process, rational individuals should begin by relying on their private information and then gradually shift to rely on the public signal as it comes to incorporate more and more of the original private information. We observe this shift in the regression analysis, but we also observe errors in the use of Bayes theorem, and other computational or inferential errors in reaction to the public signal.

We define a measure of the efficiency of inference to be the ratio of the aggregate expected payoff under optimal inference. In the experienced group, the relative efficiency of the actual inference process is 85%; for the inexperienced group 69%.

An important empirical question is how closely actual behavior approximates the rational behavior of theoretical models. As a general conclusion, we would say that a rough approximation of the common knowledge inference process (similar but simpler than those in Kobayashi (1977), Jordan (1982), and McKelvey and Page (1986)) is borne out in the experimental evidence. However, the actual inference process is clearly less efficient than the standard in the theoretical models, and while there is some pooling, it is incomplete.

*Division of Humanities and Social Sciences, California Institute of Technology,
Pasadena, CA 91125, U.S.A.,*

and

Department of Economics, Brown University, Providence, RI 02912, U.S.A.

Manuscript received November, 1985; final revision received October, 1989.

APPENDIX A

EXPERIMENT INSTRUCTIONS

This is an experiment on individual decision making. The purpose of the experiment is to study how individuals use two sources of information in making decisions about an unknown fact. You will

be paid for your participation in the experiment. Your payments will depend partly on your decisions and partly on chance. If you are careful, and make good decisions, you have a good chance of making a considerable amount of money.

The experiment is symmetric with respect to all the subjects. The instructions, the reward rule, and the lottery number table are the same for each of you. The experiment will be repeated four times. In each experiment one of you will be an assistant, and the rest will be subjects.

The Unknown Fact of the Experiment: Before you, you see two containers, each containing ten dice. All the dice are ten sided. In one container the dice are red and in the other the dice are white. Each red die has 6 marked faces and 4 unmarked faces, while each white die has 4 marked faces and 6 unmarked. At the beginning of each experiment, the assistant will select by chance one container of dice. The container of dice which is not selected will be set aside, and during the remaining part of the experiment, only the selected container of dice will be used. The unknown fact of the experiment, for you and the other subjects, is whether the dice selected are red or white.

Your Rewards: There are several rounds in each experiment. For each round, you will be required to choose a number, r , which we call your report. You will be paid for each round by means of a lottery corresponding to that round. Your probability of winning the lottery for a particular round depends upon the number r you report for that round, and upon the unknown fact. If the dice are red, and you report r , your probability of winning is read off from column 2 of the accompanying table. If the dice are white, your probability of winning is read off from column 3. For example, suppose you report $r = .32$ for some round of the experiment. If the dice are red your probability of winning the lottery for that round is .538. If the dice are white, your probability of winning is .898. The probability of winning a lottery is related to your reported r as follows. If the dice are red, your probability of winning is $2r - r^2$. If the dice are white, your probability of winning is $1 - r^2$.

At the beginning of each round of the experiment you will be given some information which you can use to assess how likely you think it is that the dice are red or white. At the end of each round you report on the worksheet: in column (1) your r ; in column (2) your probability of winning if the dice are red; in column (3) your probability of winning if the dice are white.

At the end of the experiment, you will be told whether the dice are red or white. If they are red, circle all the numbers in column (2), for the experiment. If they are white, circle all the numbers in column (3), for the experiment.

At the end of all 4 experiments your lotteries will be drawn and your earnings calculated. For each of your rounds, a three digit random number between .000 and .999 will be drawn. If it is less than the number you circled for that round, you win. If the random number is greater than or equal to your circled number, you lose. If you win, write a + in column (4) of the worksheet. Count the total number of wins for all three experiments, and read your payoff from the payoff chart on the blackboard.

Your Information: At the beginning of the first round, you will be given your own, private information concerning the unknown fact. After the container is selected the assistant will throw the dice from the container selected. You will be told how many of the dice have a marked face up. Record this information in the box, column (A) of the worksheet. This is your own, private information. Do not share it with the other subjects. Each of the others will also have his own private information in the form of an independent throw of the selected dice.

At the end of the first round, the assistant will record reported r 's from your worksheets. The average of all subjects' reported r 's (including yours) will be written on the blackboard. This is your public information, which you can use to augment your private information in round two. Record this average of the reports in the box, column (B) of the worksheet (corresponding to round 2).

At the end of round 2, the reports for that round will be collected and their average written on the blackboard. Write this public information in the box in column (B) (corresponding to round 3). In round 3, you can use this new public information to augment your original private information and the earlier public information.

The following rounds will proceed in the same fashion. At the beginning of a round, you will be given new, public information, which is the average of the reports for the previous round. At the end of the round, you will make your new report, based upon the new information and the old information.

The last round, round six, is a little different. The total number of marked faces appearing in all three samples (30 observations in all) will be written on the blackboard. At the end of this last round, with this additional public information, you will make your final report.

TABLE A
TABLE OF LOTTERIES FOR EXPERIMENTS

(1) Your Report <i>r</i>	(3) Probability of Winning		(1) Your Report <i>r</i>	(3) Probability of Winning	
	If Dice are Red	If Dice are White		If Dice are Red	If Dice are White
0.00	0.000	1.000	0.50	0.750	0.750
0.02	0.040	1.000	0.52	0.770	0.730
0.04	0.078	0.998	0.54	0.788	0.708
0.06	0.116	0.996	0.56	0.806	0.686
0.08	0.154	0.994	0.58	0.824	0.664
0.10	0.190	0.990	0.60	0.840	0.640
0.12	0.226	0.986	0.62	0.856	0.616
0.14	0.260	0.980	0.64	0.870	0.590
0.16	0.294	0.974	0.66	0.884	0.564
0.18	0.328	0.968	0.68	0.898	0.538
0.20	0.360	0.960	0.70	0.910	0.510
0.22	0.392	0.952	0.72	0.922	0.482
0.24	0.422	0.942	0.74	0.932	0.452
0.26	0.452	0.932	0.76	0.942	0.422
0.28	0.482	0.922	0.78	0.952	0.392
0.30	0.510	0.910	0.80	0.960	0.360
0.32	0.538	0.898	0.82	0.968	0.328
0.34	0.564	0.884	0.84	0.974	0.294
0.36	0.590	0.870	0.86	0.980	0.260
0.38	0.616	0.856	0.88	0.986	0.226
0.40	0.640	0.840	0.90	0.990	0.190
0.42	0.664	0.824	0.92	0.994	0.154
0.44	0.686	0.806	0.94	0.996	0.116
0.46	0.708	0.788	0.96	0.998	0.078
0.48	0.730	0.770	0.98	1.000	0.040
0.50	0.750	0.750	1.00	1.000	0.000

APPENDIX B

In this appendix we sketch the argument showing that truthful reporting is a Bayes-Nash equilibrium. (Recall that a truthful strategy for *j* is for *j* to report *j*'s posterior probability for *A*, conditioned on *j*'s current information set.)

First, consider *j*'s strategy for a given, fixed information set in round *t*; call it π_{jt}^t . Recalling (2) and the first modification of the Brier rule, we note that in expectation for $\omega \in \pi_{jt}^t$, *j*'s probability of winning the lottery in round *t* corresponding to *j*'s report is given by (2) where *q* is *j*'s posterior probability of *A* based on π_{jt}^t . As in Section 2, this probability is maximized by *j* reporting $r = q$. Setting $r = q$, we write *j*'s maximized probability as

$$\bar{p}(q) = (2q - q^2)q + (1 - q^2)(1 - q) = 1 - q + q^2.$$

Thus \bar{p} is a convex function of *q*.

Second, consider the effect of a possible refinement of an element of *j*'s partition, π , say from π to $\{\pi_1, \pi_2\}$, where $\{\pi_1, \pi_2\}$ partitions π . Then writing $q_1 = \rho(A|\pi_1)$ for *j*'s truthful strategy when $\omega \in \pi_1$ and $q_2 = \rho(A|\pi_2)$ when $\omega \in \pi_2$, we have

$$q = \rho(A|\pi) \frac{\rho(\pi_1)}{\rho(\pi_1) + \rho(\pi_2)} + \rho(A|\pi_2) \frac{\rho(\pi_2)}{\rho(\pi_1) + \rho(\pi_2)}$$

so q is a convex combination of q_1 and q_2 . Since $\bar{p}(\cdot)$ is a convex function,

$$\bar{p}(q) \leq \bar{p}(q_1) \frac{\rho(\pi_1)}{\rho(\pi_1) + \rho(\pi_2)} + \bar{p}(q_2) \frac{\rho(\pi_2)}{\rho(\pi_1) + \rho(\pi_2)}$$

where the right-hand side is j 's probability of winning the lottery when j reports truthfully under the refinement. More generally we can use Jensen's inequality to show that in expectation for $\omega \in \pi$ a refinement of π (with truthful reporting) sometimes increases and never decreases j 's probability of winning the lottery (compared with reporting $q = \rho(A|\pi)$).

Third, we consider what happens when every subject is reporting truthfully (and "myopically") round by round, and j considers a change from j 's truthful strategy. Under truthful reporting, every subject every round, there is complete refinement of j 's information partition in rounds 3, 4, 5, and 6. Thus truthful reporting for all 6 rounds leads to maximal probabilities of winning the lotteries corresponding to rounds 3, 4, 5, and 6. Truthful reporting in round 1 also maximizes j 's probability of winning the first round's lottery, since nothing j can do will affect j 's information set for the first round. Individual j cannot affect j 's information set for round 2 either, because this information set is determined by the other subjects' reports (along with j 's first period information set 1), but not j 's first period report (j can disentangle j 's report from the public signal so the signal is equivalent to the sum of the others' reports, as in the discussion of Table III). Thus truthful reporting in round 2 maximizes the probability of winning the second lottery. We conclude that truthful reporting in all six rounds leads to maximal probabilities of winning the round by round lotteries. Another way of putting the matter is that by departing from truthful reporting j might be able to change j 's information sets in rounds 3, 4, 5, or 6. But in doing so j could only make things worse for j , since j 's information sets for these rounds were already fully refined under truthful reporting.

Fourth, we are now ready to take into account the second modification, from wins and losses in individual lotteries to the composite lottery and the sliding scale of Table I. Consider two situations, situation 1 where everyone reports truthfully every round and situation 2 where j makes some departure from situation 1. We have just seen that in making the departure, none of the probabilities for the individual rounds will go up for j and some may go down. Now consider the cumulative distribution function (CDF) for the composite lottery over the payoff outcomes (3, 4, 7, 10, 13, 16, 19, 22). It is easy to check that the CDF for j in situation 1 stochastically dominates the CDF for j in situation 2 (or is equal to it).

But our behavioral assumption is that each subject prefers one composite lottery to a second if the first stochastically dominates the second. Thus "myopic" truthful reporting for every round by every subject is a Bayes-Nash equilibrium. (By a similar argument it is easy to see that if the first part of the experiment were only 3 instead of 5 rounds, truthful reporting would be a dominant strategy.)

REFERENCES

- AUMANN, R. J. (1976): "Agreeing to Disagree," *Annals of Statistics*, 4, 1236–1239.
- BERGIN, J., AND A. BRANDENBURGER (1990): "A Simple Characterization of Stochastically Monotone Functions," *Econometrica*, 58, 1241–1243.
- BRANDENBURGER, A., AND J. GEANAKOPOLOS (1987): "Common Knowledge of Summary Statistics," unpublished, Harvard Business School.
- EDWARDS, W. (1982): "Conservatism in Human Information Processing," in *Judgment under Uncertainty: Heuristics and Biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky. Cambridge: Cambridge University Press, pp. 359–369.
- FORSYTHE, R., T. PALFREY, AND C. PLOTT (1982): "Asset Valuation in an Experimental Market," *Econometrica*, 50, 537–582.
- GEANAKOPOLOS, J. D., AND H. M. POLEMARCHAKIS (1982): "We Can't Disagree Forever," *Journal of Economic Theory*, 28, 192–200.
- GREYER, D. M. (1980): "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, 95, 537–557.
- (1981): "Financial Incentives Effects and Individual Decisionmaking," Social Science Working Paper No. 401, California Institute of Technology.
- JORDAN, J. S. (1982): "A Dynamic Model of Expectations Equilibrium" *Journal of Economic Theory*, 28, 235–254.

- KOBAYASHI, T. (1977): "A Convergence Theorem on Rational Expectations Equilibrium with Price Information," Working Paper No. 79, Institute for Mathematical Studies in Social Sciences, Stanford University.
- McKELVEY, R. D., AND P. C. ORDESHOOK (1985): "Elections with Limited Information: A Fulfilled Expectations Model Using Contemporaneous Poll and Endorsement Data as Information Sources," *Journal of Economic Theory*, 35, 55–85.
- McKELVEY, R. D., AND T. PAGE (1986): "Common Knowledge, Consensus, and Aggregate Information," *Econometrica*, 54, 109–127.
- NIELSEN, L. T. (1987): "Aggregation of Expectations, Common Information, and Revealing Rational Expectations Equilibrium," unpublished, University of Texas at Austin.
- NIELSEN, L. T., A. BRANDENBURGER, J. D. GEANAKOPOLOS, R. D. McKELVEY, AND T. PAGE (1990): "Common Knowledge of an Aggregate of Expectations," *Econometrica*, 58, 1235–1239.
- MURPHY, A. H., AND R. L. WINKLER (1970): "Scoring Rules in Probability Assessment and Evaluation," *Acta Psychologica*, 34, 273–286.
- PAGE, T. (1988): "Pivot Mechanisms as a Link Between Probability and Preference Revelation," *Journal of Economic Theory*, 44, 43–62.
- PLOTT, C. R., AND S. SUNDER (1988): "Rational Expectations and the Aggregation of Diverse Information in Laboratory Securities Markets," *Econometrica*, 56, 1085–1118.
- SAVAGE, L. J. (1971): "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 66, 783–801.