



The Curse of Knowledge in Economic Settings: An Experimental Analysis

Author(s): Colin Camerer, George Loewenstein, Martin Weber

Source: *The Journal of Political Economy*, Vol. 97, No. 5 (Oct., 1989), pp. 1232-1254

Published by: [The University of Chicago Press](http://www.uchicago.edu)

Stable URL: <http://www.jstor.org/stable/1831894>

Accessed: 14/02/2011 14:13

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=ucpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

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.



The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Political Economy*.

The Curse of Knowledge in Economic Settings: An Experimental Analysis

Colin Camerer

University of Pennsylvania

George Loewenstein

University of Chicago

Martin Weber

Institut für Wirtschaftswissenschaften

In economic analyses of asymmetric information, better-informed agents are assumed capable of reproducing the judgments of less-informed agents. We discuss a systematic violation of this assumption that we call the “curse of knowledge.” Better-informed agents are unable to ignore private information even when it is in their interest to do so; more information is not always better. Comparing judgments made in individual-level and market experiments, we find that market forces reduce the curse by approximately 50 percent but do not eliminate it. Implications for bargaining, strategic behavior by firms, principal-agent problems, and choice under uncertainty are discussed.

We thank Andy Daughety, Marc Knez, Pete Kyle, Howard Kunreuther, Amy McCreedy, and Charles Plott for useful discussions and assistance. The financial support of the Risk and Decision Processes Center at the Wharton School is gratefully acknowledged. We also acknowledge support from the National Science Foundation to Camerer, from the Alfred P. Sloan Foundation, the Russell Sage Foundation, and the IBM Faculty Research Fund at the University of Chicago to Loewenstein, and from Deutsche Forschungsgemeinschaft to Weber during his visit at the Wharton School.

[*Journal of Political Economy*, 1989, vol. 97, no. 5]

© 1989 by The University of Chicago. All rights reserved. 0022-3808/89/9705-0004\$01.50

I. Introduction

In many economic transactions, some agents know more than others. For example, sellers are better informed about the true value of their products than buyers; workers know more about their ability and motivation than prospective employers. Attempts by informed agents to exploit the asymmetry in information can result in market-level consequences. Beneficial bargains go unstruck because of adverse selection (Akerlof 1970), bid-ask spreads in financial markets increase when insiders are present (Glosten and Milgrom 1985), and wages can diverge from productivity when workers' characteristics are unobservable (Salop and Salop 1976). The conventional assumption in such analyses of asymmetric information is that better-informed agents can accurately anticipate the judgments of less-informed agents. We discuss a systematic judgmental bias that challenges this apparently noncontroversial assumption.

In predicting the judgments of others, agents are unable to ignore the additional information they possess. This "curse of knowledge"¹ has two consequences: First, better-informed agents may suffer losses. More information can actually hurt. Second, the curse of knowledge can mitigate market consequences resulting from information asymmetry. For example, the seller of a "lemon" may lower its price to reflect unobservable defects, reducing the degree of market failure. Paradoxically, individual irrationality can enhance collective rationality.²

All the previous evidence of the curse of knowledge has been gathered in psychological studies of individual judgments (see, e.g., Fischhoff 1975). But the important question for economics is whether the curse harms the allocation of resources in economic settings. There are several reasons why economists might be skeptical of the psychologists' findings: (1) The curse may result from careless thinking by subjects who have no financial incentive to respond accurately. Given the right incentives, it could be conjectured, individuals would exert sufficient cognitive effort to overcome the curse. (2) In natural settings, people often receive feedback about the accuracy of their predictions. Over time, such feedback might reduce or eliminate judgmental bias. Subjects in earlier studies of the curse did not receive feedback and, hence, did not have an opportunity to learn from their mistakes. (3) Disciplining forces in markets are not present in psychology experiments. In markets, agents who are less subject to the curse of knowledge might exert disproportionate influence on prices and

¹ This term was suggested by Robin Hogarth.

² Analogously, contributions to public goods may be individually unprofitable but collectively profitable (e.g., Dawes and Thaler 1988).

allocations, effectively reducing or eliminating market-level effects even if many agents are biased.

We test arguments 1–3 by using market experiments to see whether financial incentives, learning from feedback, and market forces make the curse of knowledge disappear. We find that feedback alone has little effect, while market forces reduce the magnitude of the curse by approximately 50 percent. After describing the experiments and results, we return to the central question of how these judgment biases might affect economic settings.

Our experiments are one example of empirical efforts to determine whether violations of normative theories of judgment and choice, typically found in psychological studies, tend to persist in economic settings (e.g., Camerer 1987). Such tests may help answer longstanding theoretical questions about the behavioral foundations of economic theory (e.g., Akerlof and Yellen 1985; Russell and Thaler 1985; Hogarth and Reder 1986).

II. Formal Representation of the Curse of Knowledge

By expressing the curse of knowledge formally, we shall see that it violates a normative rule—the “law of iterated expectations”—much as choices violate the normative model of expected utility theory (Machina 1987; Weber and Camerer 1987) and probability judgments violate normative principles such as Bayes’s rule and the conjunction rule (e.g., Kahneman, Slovic, and Tversky 1982).

Call the random variable being forecasted X . If X is a discrete event, then it has the value zero or one. Forecasts of X depend on the information set available to the forecaster. Assume that there are two information sets I_0 and I_1 , where I_0 is a subset of I_1 . A forecaster with information set I_1 knows everything that the forecaster with information set I_0 knows, and more.³ Denote the optimal forecast of X given the information set I_0 by $E(X|I_0)$. We are interested in forecasts of forecasts, which are useful when agents need to forecast behavior of other agents. An agent with information set I_1 who forecasts the forecast of an agent with information I_0 is estimating $E[E(X|I_0)|I_1]$.

The law of iterated expectations states that if I_1 includes I_0 , then $E[E(X|I_0)|I_1]$ must equal $E(X|I_0)$ (Chow and Teicher 1978, p. 204). Better-informed agents should ignore their additional information when forecasting the forecasts of less-informed agents. When the

³ In technical treatments, I_1 is a finer partition of a probability space than I_0 ; the additional knowledge contained in I_1 , like the additional information given to subjects in our experiments, presumably enables one to make a finer partition.

curse of knowledge occurs, the forecaster with information I_1 *overestimates* the scope of I_0 . Formally, the curse of knowledge means that $E[E(X|I_0)|I_1]$ is *not* equal to $E(X|I_0)$, but is somewhere between $E(X|I_0)$ and $E(X|I_1)$. A simple model we test in our experiments is

$$E[E(X|I_0)|I_1] = wE(X|I_1) + (1 - w)E(X|I_0). \quad (1)$$

If $w = 0$, an agent is applying the law of iterated expectations correctly. If $w = 1$, agents who know I_1 think that all other agents know I_1 too. The parameter w thus measures the *degree* of curse of knowledge.

III. Experimental Design

The market experiments consisted of two stages. In the first stage we collected predictions of eight actual companies' earnings from 51 Wharton students. In the second stage, conducted 2 months after the first, we informed a second group of subjects of the actual earnings and had them trade assets that paid a liquidating dividend equal to the predictions of the first group. This second group knew that the asset dividend was determined by the first group's predictions, but they did not know the exact amount of the dividend.

In the first stage, subjects were given a Value Line⁴ report about each company's prospects in 1980, along with a brief summary of the company's business activity, annual earnings per share from 1970 through 1979, and quarterly earnings per share from 1977 through 1979.⁵ A sample report is shown in the Appendix (fig. A2). Each report had a blank box in the upper right-hand corner, labeled "1980." Subjects were instructed to write their estimate of the company's actual 1980 earnings in the box.

For the sake of credibility, the companies' names were not disguised. However, no well-known companies were selected, and there was no evidence that any subject knew any company well enough to recall any details of its history other than those given in the report. These data were collected as a classroom exercise in a quantitative methods course. As an incentive for accuracy, subjects were paid \$1.00 for each estimate that was within 10 percent of actual 1980 earnings.

⁴ Value Line issues reports used by investors to assess future earnings potential of companies.

⁵ Most market experiments involving uncertainty rely on random probability devices such as bingo cages or dice. These devices are used because experimenters wish to control subjective probabilities as tightly as possible. Since our study is specifically focused on differences in subjective probability, we used natural stimuli that permit such differences instead of using bingo cages or dice.

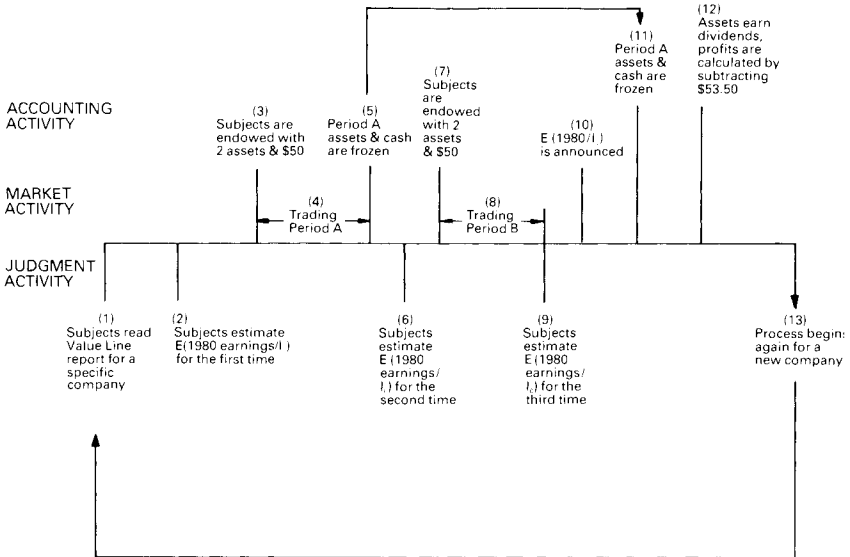


FIG. 1.—Sequence of steps in the market experiments

In the second stage, market experiments 1–2, the subjects were groups of nine Wharton students who had participated in other market experiments.⁶ Instructions were read aloud. The experiments consisted of eight identical markets, one for each of the eight companies. Markets were split into two periods, the A and B periods.

The chronological sequence of steps in the market experiment is summarized in figure 1. At the beginning of each period, these “informed subjects” were endowed with two assets and \$50 in cash, which had to be repaid at the end of the period (along with a \$3.50 “tax”). They could trade these assets to others⁷ in a double-oral auction⁸ or keep them. At the end of the B period for a particular company, subjects received a dividend for each asset they held at the end

⁶ These subjects knew from their earlier experience that our experiments do not involve any form of deception and that all rewards would be delivered as claimed. Their experience with trading also reduced the noise in market behavior created by inexperienced subjects.

⁷ Subjects could also sell assets short by “creating” assets and paying dividends for assets created. Short selling increases market discipline because if the curse of knowledge causes prices to be too high, unbiased traders can discipline biased traders by selling more units than they are endowed with.

⁸ Buyers and sellers shouted out bids or offers at which they were willing to buy or sell. When a bid and offer matched, a trade took place; previous bids and offers disappeared. All bids, offers, and trades in a current period were recorded on a transparency shown on an overhead projector for subjects to see. No record of previous periods was posted. The A and B trading periods were each 4 minutes long.

of each of the A and B periods. After the dividends were paid, the assets expired worthless.

The dividend amount was equal to the mean earnings prediction of subjects in the first stage of the experiment. In the notation used earlier, the value of assets was $E(1980 \text{ earnings} | I_0)$, where I_0 was the information given to the first group of Wharton students. Subjects in the market stage of the experiment saw exactly the same Value Line reports that subjects had seen in the first stage except that *the blank box in the upper right-hand corner was filled in with the actual 1980 earnings*. Therefore, the subjects in the market experiment had information I_1 , which included I_0 ⁹ and the actual 1980 earnings. Because the assets paid a dividend of $E(1980 \text{ earnings} | I_0)$, the subjects in the market experiment had to estimate $E[E(1980 \text{ earnings} | I_0) | I_1]$. If these subjects suffer from the curse of knowledge and market forces do not erase it, market prices will be biased away from $E(1980 \text{ earnings} | I_0)$ in the direction of $E(1980 \text{ earnings} | I_1)$ (which is simply the 1980 earnings).¹⁰

In addition to trading assets, market subjects also estimated the first group's predictions at three points: before the A period, between the A and B periods, and after the B period. They earned 25 cents for each response that was within 10 percent of the first group's mean prediction.

The traders in these markets are like investment banks who underwrite issues of corporate securities. Underwriters compete for securities they sell to investors who know less about the true value of the issues than they know (because underwriters typically have access to detailed information beyond that included in the public prospectus). To price the issue correctly, underwriters must ignore the extra information they have. To price the assets in our experiments correctly, traders must anticipate the judgments of a "public" (the uninformed subjects) that knows less than they do. Instead of selling assets directly

⁹ Of course, we cannot be certain that I_0 is a subset of I_1 because different subjects participated in the two parts of the experiment and the experiment was conducted after 1980. However, subjects were drawn from a common pool of students with homogeneous education, and there was no evidence that the two groups had different knowledge about company earnings.

¹⁰ The Value Line report for Diamond Shamrock, shown in the Appendix (fig. A2), illustrates how the curse of knowledge might operate. Diamond's earnings per share for 1979 were \$3.37 and 1980 earnings were \$3.66, up about 30 cents. The report states that Diamond's "acquisition of Falcon Seaboard . . . reduced 1979 earnings per share by about 30¢. The impact should be smaller in 1980." A subject who knew that the 1980 earnings were \$3.66 could think that uninformed subjects who read the report would know enough to correct the 1979 earnings of \$3.37 for the temporary 30-cent reduction and expect 1980 earnings to be around \$3.67. But the 30-cent reduction looms large (it is psychologically "available") only because it explains the difference between 1979 earnings and 1980 earnings.

to the public, as underwriters do, they exchange the assets for a dividend determined by the public's judgment.

IV. Experimental Results

First we describe the results of the market experiments. Then we contrast judgments that subjects made during the market experiment with judgments made by individuals outside of a market setting.

Market Prices

Figures 2 and 3 summarize the mean prices and judgments in the two market experiments. For each company, arrowheads show the mean estimate by market subjects of the first group's prediction. Vertical lines extend two standard errors around each arrowhead. Two plus signs indicate the mean trade price in the A and B periods. (Means are plotted because prices did not vary much within periods.) The numbers of trades in the A and B periods are shown underneath the figures. Notice also that the y-axis scales are different for each

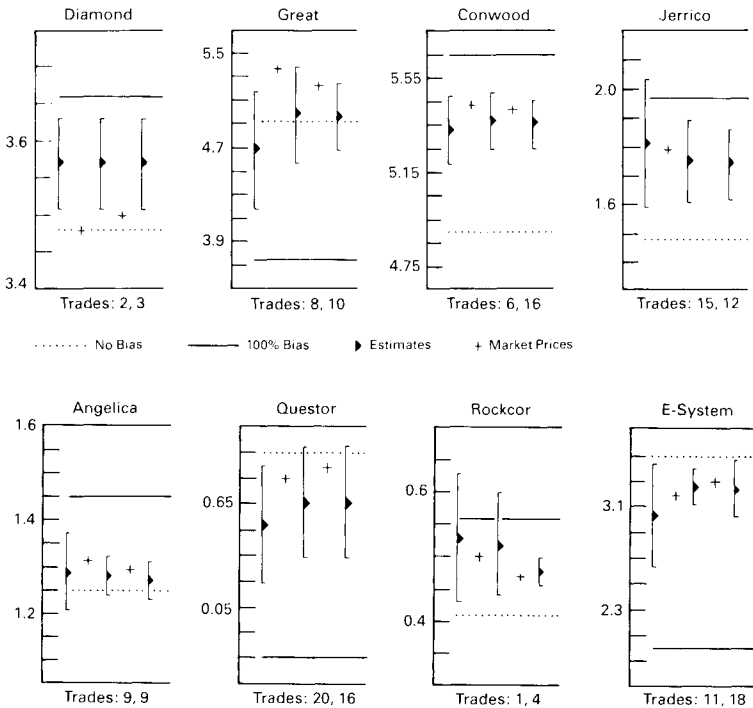


FIG. 2.—Mean prices and judgments by company (experiment 1)

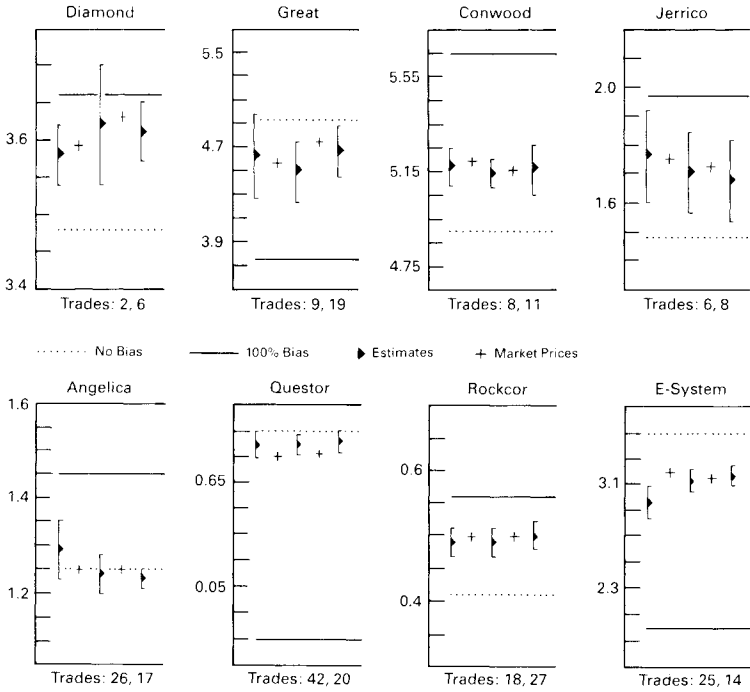


FIG. 3.—Mean prices and judgments by company (experiment 2)

company. Horizontal dotted lines show the prediction of the uninformed group about actual earnings, $E(1980 \text{ earnings}|I_0)$; this is the no-bias prediction ($w = 0$ in the model [1]). Solid lines show the actual 1980 earnings, $E(1980 \text{ earnings}|I_1)$, the pure-bias ($w = 1$) prediction. Notice that the no-bias prediction is sometimes above the pure-bias prediction and sometimes below it, so risk aversion that lowers prices will not favor one prediction over another.

Several effects are apparent in the data. Judgments and prices are close together, as one would expect. Prices generally begin between the pure-bias and no-bias predictions and move away from the pure-bias prediction slightly. There were an average of 12.8 trades per period, including 5.8 short sales.

Equilibrium Prices

The equilibrium price that a time series of trade prices is converging toward can be assessed by estimating a simple partial adjustment model (see Camerer 1987). Regressions show that equilibrium prices are between the pure-bias and no-bias predictions, as figures 2 and 3 suggest (details are available from the authors). There is little varia-

tion in the time series of prices, so the estimates are extremely precise; standard errors are typically a penny or two. We can strongly reject both the no-bias ($w = 0$) and pure-bias ($w = 1$) hypotheses.

Movements in Prices and Judgments

By including a dummy variable for B periods in equilibrium price regressions, we can tell whether the degree of bias shrinks between periods A and B. When three companies for which bias could not be measured were excluded, bias was smaller in period B than in period A for 11 of 13 comparisons (significant at the 10 percent level in four comparisons) and larger in two of 13 comparisons (both significant).

We can also ask whether the judgments that market subjects made about $E(1980 \text{ earnings}|I_0)$ —before the A period, between the A and B periods, and after the B periods—differed because of their market experience. In 54 percent of the cases, judgments did not change at all or moved toward or away from both predictions. In the remaining 131 cases, 83 (63 percent) indicated a reduction in bias and 48 (37 percent) indicated an increase in bias. (These fractions are different from 50 percent at the 1 percent level.) Market experience does reduce bias.

Judgments under Incentives, Feedback, and Markets

Markets occupy a special place in economics, but to a psychologist market forces are simply another “treatment variable,” akin to incentives, instructions, education, and so forth. We can test the impact of market experience as a treatment variable by comparing the judgments made by individuals in the market experiments (market condition) with judgments made by other subjects who were given incentives but no feedback (incentives condition, $n = 19$) or incentives and feedback (feedback condition, $n = 13$).

The subjects in the incentives and feedback conditions were University of Chicago MBA students. Their judgment task was the same as the judgment task of subjects in the market experiments: knowing actual 1980 earnings, they were asked to guess the mean prediction of the uninformed subjects who had predicted 1980 earnings previously. They earned \$1.00 if their judgment was within 10 percent of the actual mean prediction. In the feedback condition, subjects were told the actual mean prediction of the uninformed subjects after each guess.

To compare degrees of bias, it is useful to convert subjects' judgments $E[E(1980 \text{ earnings}|I_0)|I_1]$ into a curse of knowledge bias index:

$$w = \frac{E[E(1980 \text{ earnings}|I_0)|I_1] - E(1980 \text{ earnings}|I_0)}{E(1980 \text{ earnings}|I_1) - E(1980 \text{ earnings}|I_0)}. \quad (2)$$




FIG. 4.—Percentage bias: market vs. individual judgment

This index measures the parameter w in model (1). Index values below zero or above one are conceivable but rare. If $w = 0$, there is no bias;¹¹ if $w = 1$, market subjects act as if 1980 earnings were perfectly predictable. The index measures the distance between judgments and the dotted lines in figures 2 and 3, as a fraction of the distance between the solid and dotted lines.

There is no difference between the biases of subjects in the incentives and feedback conditions; feedback apparently did not reduce bias.¹² Therefore, we pool data from subjects in the incentives and feedback conditions.

Figure 4 shows the mean degree of bias among subjects in the incentives and feedback conditions (labeled “individual judgments”) and the mean degree of bias in judgments made after period B of the market experiments (labeled “market judgments”). The t -statistics below each company’s name test whether the two mean biases are equal.

Individuals in both conditions exhibit some positive degree of bias, but the market reduces bias by about half. (The reduction is signifi-

¹¹ To be sure that uninformed subjects could guess other uninformed subjects’ estimates without bias, two groups of Chicago subjects were asked to guess the uninformed group’s mean prediction without knowing the actual 1980 earnings. One group got feedback ($n = 14$) and one group did not ($n = 17$). Their judgments were randomly distributed around the judgments of the uninformed Wharton subjects.

¹² For the eight companies, the t -statistics comparing incentive group and feedback group means were 1.6, -0.4 , 2.9, 0.3, -1.9 , -1.3 , and 0.6 (a positive number means that feedback group subjects show less bias).

cant at 5 percent for three of eight companies.) With more than two trading periods per company, the market might reduce bias even further.

How Does the Market Reduce Bias?

Market experience clearly reduces bias more than individual judgment tempered by incentives and feedback. Why? A common argument is that markets correct irrationalities because more rational traders drive less rational traders into bankruptcy or somehow correct the errors of less rational traders. Since traders are unlikely to go bankrupt in these experiments, the latter explanation is more natural: perhaps less biased subjects trade earlier or more often, thus signaling their information to more biased subjects.

The process we have in mind is analogous to models of information aggregation in which uninformed traders can infer the information of insiders from their trading activity (e.g., Grossman 1981; Plott and Sunder 1982). In our experiments there are no true insiders because everyone receives the same information. However, less biased traders are like insiders because they are better informed about the true value of the assets.

For the information aggregation analogy to apply it is essential that subjects have enough self-insight to know whether they are more or less biased than others. (If all traders think that they are the least biased, the most biased will trade as often as the least biased.) To measure self-insight, after the experiment¹³ we asked subjects to predict their ranks, relative to the other traders, in (i) trading profits and (ii) earnings from predicting the uninformed group's estimates accurately. Subjects earned \$2.00 if their predicted rank exactly matched their actual rank.

Rank-order correlations between predicted and actual ranks in the two experiments were (i) .99 ($p < .005$) and .60 ($p < .04$) for trading profits and (ii) .58 ($p < .05$) and .78 ($p < .01$) for earnings from judgments.¹⁴ Since these ranks were measured after the experiment was over, the correlations probably represent upper bounds on self-insight, but they are rather large nonetheless.

The substantial correlations indicate that less biased traders seem to know who they are. Were they also more aggressive? To find out, we can rank traders according to the degree of bias evident in their initial

¹³ We asked for ranks after the experiment because doing so before trading began gives them an incentive to achieve their rank, possibly by trading unprofitably.

¹⁴ The rank-order correlations between ranks in judgment profits and ranks in trading profits were .39 and $-.03$ in the two experiments. Estimating the dividend accurately is apparently not a strong predictor of trading profit.

TABLE 1
DEGREE OF TRADER BIAS BY ACTION NUMBERS (Positive Number = Less than
Average Bias)

STATISTIC	ACTION NUMBERS				
	1-5	6-10	11-15	16-20	All
Mean rank	.43	.36	.03	.01	.30
Median rank	.90	.60	.65	.50	.25
Percentage of positive ranks	55.8	59.5	54.8	57.4	55.8
z-statistic*	1.03	1.09	.82	1.21	3.04

NOTE.—Ranks are from -4 (most biased) to +4 (least biased). If least biased traders account for a disproportionate fraction of transactions, mean and median ranks will be positive.

* The z-statistic has a unit normal distribution under the null hypothesis that the percentage of positive ranks is 50 percent, by a normally approximated binomial test.

judgment, made before trading began. The trader with the judgment farthest from the no-bias prediction was ranked -4 and the closest trader was ranked +4. Then we analyzed the ranks of the traders who took each market action—a bid, offer, or acceptance of a bid or offer—in the A trading periods.

The data are summarized in table 1. The table is read as follows: For the first five actions (1-5), the mean rank in bias of the traders taking those actions was .43, the median rank was .90, and 55.8 percent of those traders had positive ranks (excluding median ranks of zero), indicating less bias than the median trader. Since the mean rank of all traders is zero, by definition, the fact that the mean rank of traders taking actions is positive implies that less biased (positive-rank) traders were taking more than their share of actions, throughout the experiment.

Less biased traders know roughly who they are and act disproportionately often, as the information aggregation account suggests, but the effect is small in magnitude.

Summary of Results

The results answer several questions: *Q*: To what equilibrium are prices converging? *A*: To an equilibrium between the pure-bias and no-bias predictions. *Q*: Are prices and judgments moving away from the pure-bias prediction more often than chance? *A*: Yes. We know this from comparing prices in A and B periods and studying changes in individuals' judgments between periods. *Q*: Do judgments made in markets show less bias than judgments made by individuals not subject to market forces? *A*: Yes. Subjects in markets show about half as much bias. *Q*: Why do markets show less bias? *A*: Less biased traders

seem to know that they are less biased, and they take actions more often than more biased traders do.

Our objective has been to test whether incentives, experience, and market forces reduce or eliminate the curse of knowledge. The experimental data suggest that incentives and feedback do not reduce the bias, but market forces do. Of course, our subjects have less incentive and experience than professionals because they earn less money and train for hours, not years. But the available laboratory evidence suggests that very large incentives and extensive training do not improve learning much.¹⁵ Furthermore, our subjects' incentives and learning experience are intense and compacted; they may learn *better* in a simple laboratory task than professionals can in a noisy natural environment.

V. Implications for Economics

There are at least two major economic situations in which the curse of knowledge may be important. The first is the classic case of asymmetric information (as discussed in the Introduction); the second occurs when agents gain information over time and then try to recall what they or others previously believed (hindsight bias).

Asymmetric Information

Often agents have private information about their characteristics (causing adverse selection) or actions (causing moral hazard) that less-informed agents lack. Most theoretical analyses of such situations focus on less-informed agents' attempts to learn the private information or create contracts that minimize efficiency losses from the information asymmetry. Implicit in such analyses is the assumption that better-informed agents can optimally exploit their informational advantage. Our data and earlier studies (Ross, Greene, and House 1977; Nickerson, Baddeley, and Freeman 1987) suggest that this simulation is difficult; better-informed agents do not fully appreciate the informational advantage they possess. For example, in bargaining, one agent may know the size of the amount being divided while the other does not. However, to fully exploit their advantage, they should make the same offer regardless of the amount to be divided (Myerson 1986). But agents actually offer more when the amount to be divided is larger (Forsythe, Kennan, and Sopher 1987; cf. Banks, Camerer, and Porter 1988). Agents are unable to act "inscrutably"; they cannot ignore their better information when they should.

¹⁵ On incentives, see, e.g., Camerer (1987). On learning, see the calibration literature summarized in Lichtenstein, Fischhoff, and Phillips (1982).

The curse of knowledge may also influence strategic behavior by firms. Firms often do not know how many other firms compete with them. Knowing the information and reasoning that motivate their own ventures and exaggerating the degree to which the rationale is shared, businesses may tend to overestimate the amount of competition (i.e., firms act like price takers when they are not). During the silver bubble of the 1980s, one broker abandoned the business of buying coins and selling bullion and the bid-asked spread between coins and bullion rose dramatically. The futures-stock spread rose similarly when many index futures arbitrageurs left the business after the October 1987 stock market crash. In both of these examples, firms did not seem to realize how profitable spreads would be after they left, as if they underestimated their own influence in the market.

Perhaps the setting closest in structure to our market experiments is underwriting, in which well-informed experts price goods that are sold to a less-informed public. Investment bankers value securities, experts taste wines, store buyers watch clothes being modeled, and theater owners see motion pictures before they are released. They then sell those goods to a less-informed public. If they suffer from the curse of knowledge, high-quality goods will be overpriced and low-quality goods underpriced relative to profit-maximizing values; prices will reflect characteristics (e.g., quality) that are unobservable to uninformed buyers.

The curse of knowledge has a paradoxical effect in these settings. By making better-informed agents think that their knowledge is shared by others, the curse helps alleviate the inefficiencies that result from information asymmetries, bringing outcomes closer to complete information (first-best) outcomes. In such settings, the curse on individuals may actually improve social welfare.

Hindsight

The second economic situation in which the curse of knowledge may operate is when agents accumulate information over time and must attempt to reconstruct their earlier perspective. The analogue to the tendency to overestimate the knowledge of less-informed agents is a tendency to exaggerate what one knew before, when one was less informed. This has been called "hindsight bias" (Fischhoff 1975). When one looks backward, events seem to have been more predictable than they were. For instance, public discussions of medical cases make doctors think that diseases were easier to diagnose than they actually were (Dawson et al. 1986).

This exaggeration interferes with the evaluation of decision quality. Outcomes are an imperfect indicator of decision quality; good deci-

sions can lead to bad outcomes and vice versa. But principals must often judge decisions of agents on the basis of outcomes because actions or decision criteria are unobservable (Ross 1973; Jensen and Meckling 1976). The curse of knowledge suggests that outcome information will be overused; principals will tend to think that *ex ante* optimal decisions with unfavorable outcomes were nonoptimal and that nonoptimal decisions with favorable outcomes were optimal (Baron and Hershey 1988). As a result, agents will be excessively penalized for negative outcomes and insufficiently rewarded for favorable results.

Although the principal and agent can contract today to avoid tomorrow's hindsight bias,¹⁶ when contracts are implicit, hindsight bias will cause a principal to recall the terms of yesterday's contract incorrectly (see, e.g., Fischhoff and Beyth 1975). This problem is especially acute in public decision making, in which the principals are a diffuse group of voters and contracts are rarely explicit. Hindsight bias appears regularly in investigations of disasters (such as the Challenger explosion or the MOVE incident in Philadelphia). It may also lead judges and juries to exaggerate the culpability of defendants in liability litigation, leading to inflated awards.

The curse of knowledge may also influence individual decision making under uncertainty. Exaggerating the predictability of events intensifies the regret people feel when choices yield outcomes worse than those that would have resulted from forgone options. Choice theories incorporating regret have been axiomatized (Bell 1982; Loomes and Sugden 1982, 1987*a*; Fishburn 1988) and have received empirical support (Loomes and Sugden 1987*b*; Loomes 1988).

VI. Conclusion

There has been much debate about the impact of judgmental errors in economic settings. Economists usually contend that in natural settings people either learn from personal experience or are surrounded by institutions—such as advice of relatives or consultants—that provide advice in unfamiliar situations. How well people learn from personal experience, and from the experience of others, is therefore a central question in the debate about the behavioral foundations of economics.

Our research suggests grounds for pessimism about both kinds of learning. Hindsight bias narrows the gap between what occurred and what predictions are recalled, reducing valuable feedback and inhibit-

¹⁶ However, contracts tying compensation to *ex post* performance appear to be rare (Baker, Jensen, and Murphy 1988).

ing learning. The curse of knowledge makes personal expertise seem more widely shared than it is, making it difficult for people to convey their expertise to others and reducing the apparent need (from the perspective of the better-informed individual) for such a transfer of knowledge.

As with all judgment errors by individuals, the curse of knowledge may not persist in a competitive market. To examine the effect of the market we ran a series of experiments in which better-informed subjects predicted the judgments of less well informed subjects. The curse of knowledge suggests that informed subjects will be unable to ignore the information they have that the uninformed subjects lack, causing a bias in their predictions. We found that bias in markets was half as large as bias in individual judgments. Our data suggest that the error-correcting power of markets derives not from the feedback they provide, but from the disproportionate activity of more rational traders.

Appendix

Instructions for Market Experiment and Experimental Materials

General

This is an experiment about decision making in a market. Various research foundations have provided funds for this research. The instructions are simple, and if you follow them carefully and make good decisions, you can earn a considerable amount of money which will be paid to you in cash.

In this experiment you will buy and sell certificates in a series of market periods. Each period consists of two sessions, session A and session B.

Specific Instructions

Your market profits come from two sources—from collecting dividends on all certificates you hold at the end of each market session and from buying and selling certificates during the market session. During each session you are free to buy and sell as many certificates as you wish provided you follow the rules below.

For each certificate you hold at the end of a session you will earn a specific dollar amount (a “dividend”), which will be announced at the end of each market period. This amount will be recorded on row 11 of the Information and Record Sheet [see fig. A1] for each market session. The method by which the dividend is determined in each period is explained later in these instructions.

You will calculate your total certificate earnings for a session by multiplying the dividend per certificate by the number of certificates held. That is,

$$\begin{aligned} &\text{number of certificates held} \times \text{dividend per certificate} \\ &= \text{total certificate earnings.} \end{aligned}$$

Trader #

Period: DIAMOND SHAMROCK Session A

Please Estimate the Other Group's Prediction: _____ \$.25
 (before trading begins) (circle if within 10%)

INFORMATION AND RECORD SHEET

Transaction Number	Transaction Purchases	Transaction Price Sales	Certificates on hand	Dollars on hand
0	-----initial holdings-----		2	\$50
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				
11	Dividend (Other Group's Prediction) = \$ _____			
12	Total Certificate Earnings = Dividend x Number of Certificates			
13	Total Dollars = Dollars on Hand + Row 12			
14	NET PROFIT FOR SESSION A = Row 13 - \$53.50			

Record this number on your profit sheet

FIG. 11a.—Information and record sheet, session A

For example, suppose that you hold three certificates at the end of session A. If for that session your dividend is \$1.27 per certificate, then your total certificate earnings in the session would be $3 \times \$1.27$, or \$3.81. This number should be recorded on row 12 at the end of the market period.

At the beginning of each session you are provided with an initial holding of certificates, which are recorded on row 0 of your Information and Record Sheet. You may sell these certificates if you wish, or you may hold them. If you hold a certificate, then you receive the dividend for each certificate you hold at the end of the period.

Sales from your certificate holdings increase your dollar holdings by the amount of the sale price. Similarly, purchases reduce your holdings by the

Trader #

Period: DIAMOND SHAMROCK Session B

Please Estimate the Other Group's Prediction: _____ \$.25
 (before trading begins) (circle if within 10%)

 INFORMATION AND RECORD SHEET

Transaction Number	Transaction Price Purchases	Sales	Certificates on hand	Dollars on hand
0	-----initial holdings-----		2	\$50
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				
11	Dividend (Other Group's Prediction) = \$ _____			
12	Total Certificate Earnings = Dividend x Number of Certificates			
13	Total Dollars = Dollars on Hand + Row 12			
14	NET PROFIT FOR SESSION A = Row 13 - \$53.50			

Record this number on your profit sheet

Please Estimate the Other Group's Prediction: _____ \$.25
 (after trading ends) (circle if within 10%)

FIG. 11b.—Information and record sheet, session B

amount of the purchase price. Thus you can gain or lose money on the purchase and resale of certificates.

In addition to your certificates, in each session you are given an initial amount of dollars, which is recorded on row 0 of your Information and Record Sheet. You may keep these dollars if you wish or you may use them to purchase certificates.

Suppose you sell all the certificates you began the session with, so that your holding of certificates is zero. You may continue to sell certificates by "creat-

ing" additional certificates. If you create and sell a certificate, you receive the sale price of this certificate, but you will have the dividend amount of this certificate deducted from your earnings at the end of the period. For example, suppose you created two certificates and the dividend was \$1.27 per certificate. The sale prices of the two certificates increase your holdings of dollars on hand, and your total certificate dividends deducted will be $2 \times \$1.27$, or \$2.54.

Your dollars on hand at the end of a period are determined by your initial amount of dollars on hand, by dividends earned for certificates held (or deducted for certificates created), and by gains and losses from purchases and sales of certificates. All dollars on hand at the end of each session in excess of the "fixed cost" of \$53.50 are yours to keep. These are your profits for the session.

Information about Dividends

The dividends that the certificates earn are related to a financial report about an actual company, written by a financial reporting service called Value Line. In each market period, a different company report will be used to determine the dividend. We will refer to each period by the company report which determines the dividend for that period—for example, the Angelica Corp. period.

The excerpted Value Line reports [see fig. A2] you will see show four things about a company: (1) the annual earnings per share of stock, for each period from 1970 through 1979; (2) the earnings per share for each quarter (three-month period), from 1977 through 1979; (3) a verbal and statistical description of the basic nature of the company's business, which was written by Value Line reporters; (4) a verbal description, written by Value Line reporters in January 1980, of the company's prospects in 1980.

On the excerpted reports that you see, you will also be told a fifth piece of information—the company's actual earnings per share for 1980.

Earlier this week, we showed the excerpted Value Line report that you will see to a number of Wharton students, much like yourselves. The only difference is that the students were not told the fifth piece of information, the company's actual 1980 earnings per share. The other group of students were asked to predict the actual earnings per share in 1980. Students were paid \$1 if they were within 10% of the true earnings. We calculated an average of their predictions for each company by adding up all their predictions and dividing the total by the number of predictions. We will call this average prediction "the other group's prediction." The other group's prediction of 1980 earnings per share will determine the amount of dividend that certificates will earn. For example, suppose the average prediction of 1980 earnings per share—by students who did not know the actual earnings—was \$1.27. Then for each certificate you hold, you will earn a dividend of \$1.27. For each certificate created, \$1.27 will be deducted from your earnings.

Estimating the Other Group's Prediction

Before each trading session, you must make an estimate of the other group's prediction. Write these estimates on the top of your Information and Record Sheet. You should also make an estimate of the other group's prediction at the end of session B. After the market period is over, we will announce the

DIAMOND SHAMROCK

Earnings per Share											
1970	71	72	73	74	75	76	77	78	79	80	
\$.75	.57	.85	1.42	2.89	3.41	3.90	4.28	3.26	3.37		
		1st quarter		.72	.72	.56					
		2nd quarter		1.34	.93	1.01					
		3rd quarter		1.08	.62	.70					
		4th quarter		1.14	.99	1.10					

BUSINESS: Diamond Shamrock Corp. produces chemicals and plastics for the agricultural, automotive, furniture, metal, and textile industries. Also engages in the exploration, production and refining of crude oil and natural gas and the sale of refined petroleum products and natural gas. Acquired Falcon Seaboard (coal producer) in 1979.

Avg. '78 gas & oil production: 290 mill. cu. ft. and 8,900 bl. a day. R&D costs, 2.1% of sales; wages, 16% '78 deprec. rate: 5.8% Est'd plant age: 6 yrs. Has 11,644 empl., 40,204 stkhldrs. Insiders hold 1.6% of comm. Chrmn: C.A. Cash Pres. & C.E.O.: W.H. Brcker, Inc., Delaware. Address: 2300 Southland Center, Dallas, Texas 75201.

This is an energy stock: Most of the company's earnings emanate from oil and gas exploration, production, refining and marketing. We estimate that the contribution from oil and gas will grow significantly over the next several years. The refining operation, however, faces a significant challenge. Diamond Shamrock's wells don't provide nearly enough crude, particularly sweet crude, to profitably sustain operations. Previously, the company relied exclusively on other domestic suppliers. But in another oil crisis, Diamond Shamrock could get cut short. So management is looking abroad. Diamond Shamrock's reserves are an important factor. We have roughly calculated the company's oil and gas reserves at \$30-\$32 a share. And recently acquired Falcon Seaboard adds a new stitch--an estimated \$2.5 million (about \$4 a share) of untapped coal reserves. These assets make this stock look cheap. The recession dampens our enthusiasm. Thus far, demand for most of Diamond's chemical and plastics products has remained firm. Slippage has occurred only in goods directly affected by the depressed

automotive and construction industries. But we expect chemical and product sales to fall off soon. Reason: Consumer spending is on the decline, and the chemicals industry generally lags the economy by about 6 months. The acquisition of Falcon Seaboard, Diamond's coal mining subsidiary, is still diluting earnings. We estimate that Diamond's new addition reduced 1979 earnings per share by about 30c. The impact should be smaller in 1980 and succeeding years as Falcon's profits grow. Diamond Shamrock has excellent long-term prospects. The coal and oil and gas operations are destined to be winners in the energy hungry Eighties. The chemicals, too, should thrive as the economy recovers. The result may well be significant earnings, progress coupled with a healthy advance in Diamond's stock price.

A.H.S./F.B.N.

FIG. A2.—Value Line report on Diamond Shamrock Corp.

other group's prediction. For each of your estimates which are within 10% of the other group's prediction, you will earn \$0.25 (twenty-five cents) in addition to your earnings from the markets.

Sessions within Market Periods

In each market period there are two sessions, A and B. One company's excerpted Value Line report will be used for both the A and the B sessions in a single market period. We will only announce the other group's prediction

(i.e., the dividend which certificates will earn) at the end of session B, which also concludes the market period.

At the end of session A you can record your holdings of dollars and certificates, but you cannot calculate your profits because you will not be told the dividend until the end of session B. Therefore, after we announce the dividend at the end of session B, you can return to your session A Information and Record Sheet and record your dividends and calculate your profits from session A.

Trading and Recording Rules

1. All transactions are for one certificate at a time. After each of your transactions you must record the transaction price, in the "purchases" column if you bought a certificate or in the "sales" column if you sold a certificate. Your first transaction in a session is recorded on row 1 of that session's Information and Record Sheet, your second transaction is recorded on row 2, and so on. After each transaction you must calculate and record your new holdings of certificates and your new amount of dollars on hand. Your holding of dollars may never go below zero.

2. If your holding of certificates goes to zero and you create certificates, the number of created certificates will be recorded as a negative number of certificates held. For instance, if you sell all your certificates, then create three certificates, your certificate holding will be -3 . If you buy certificates after creating certificates, then your purchase reduces the number of certificates created. If you created three certificates, so that your certificate holding is -3 , and you buy two certificates, then your new holding of certificates is $-3 + 2$, or -1 . If you have created more certificates than you began with and bought, your amount of certificates will be negative at the end of the period, so your dividends (recorded in row 12) will be negative.

3. At the end of the period, record your total certificate dividends in row 12. At the end of the period, add dividends earned on row 12 to your dollars on hand, and write the total in row 13. Subtract the "fixed cost" of \$53.50 from this total, and write the difference in row 14. This is your profit for the market session and is yours to keep. At the end of each market period, record this number on your Profit Sheet.

4. At the end of each market period, after we announce the other group's prediction, you can evaluate whether your estimates were within 10% of the other group's prediction or not. For each estimate which is within 10%, you earn \$0.25. These earnings should also be recorded on your Profit Sheet.

5. At the end of the experiment add up your total profit from the market period, and enter this sum on row P1 of your Profit Sheet. Add up your earnings from estimating the other group's prediction, and write this total in row P2. Add the profits in rows P1 and P2, and record the total in row P3. The experimenter will pay this amount of money in cash at the end of the experiment.

Market Organization

The market will be conducted in a series of periods. Each trading session will last 3 minutes. Anyone wishing to purchase a certificate is free to raise his or her hand and make a verbal bid to buy one certificate at a specified price, and anyone who is willing to sell or create a certificate is free to accept or not

accept the bid. Likewise, anyone wishing to sell a certificate is free to raise his or her hand and make a verbal offer to sell one certificate at a specified price. If a bid or offer is accepted, a binding contract has been closed for a single certificate, and the contracting traders will record the transaction on their Information and Record Sheets. Bids and offers may not be removed after they are recorded. Any ties in bids or acceptances will be resolved by random choice. There are likely to be many bids and offers that are not accepted, but you are free to keep trying. You are free to make as much profit as you can. Except for the bids and offers and their acceptance, you are not to speak to any other subject about the experiment.

References

- Akerlof, George A. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Q.J.E.* 84 (August 1970): 488–500.
- Akerlof, George A., and Yellen, Janet L. "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?" *A.E.R.* 75 (September 1985): 708–20.
- Baker, George P.; Jensen, Michael C.; and Murphy, Kevin J. "Compensation and Incentives: Practice vs. Theory." *J. Finance* 43 (July 1988): 593–616.
- Banks, Jeff; Camerer, Colin F.; and Porter, David. "Experimental Tests of Nash Refinements in Signaling Games." Working paper. Philadelphia: Univ. Pennsylvania, Dept. Decision Sci., 1988.
- Baron, John, and Hershey, John C. "Outcome Bias in Decision Evaluation." *J. Personality and Soc. Psychology* 54 (April 1988): 569–79.
- Bell, David E. "Regret in Decision Making under Uncertainty." *Operations Res.* 30 (September–October 1982): 961–81.
- Camerer, Colin F. "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence." *A.E.R.* 77 (December 1987): 981–97.
- Chow, Yuan-Shih, and Teicher, Henry. *Probability Theory, Independence, Interchangeability, and Martingales*. New York: Springer-Verlag, 1978.
- Dawes, Robyn M., and Thaler, Richard H. "Anomalies: Cooperation." *J. Econ. Perspectives* 2 (Summer 1988): 187–97.
- Dawson, Neal V.; Arkes, Hal R.; Siciliano, Carl; Blinkhorn, Richard; Lakshmanan, Mark; and Petrelli, Mary. "Clinicopathological Conferences: What Does the Audience Learn?" Paper presented at the Society for Medical Decision Making meeting, Chicago, October 1986.
- Fischhoff, Baruch. "Hindsight \neq Foresight: The Effect of Outcome Knowledge on Judgment under Uncertainty." *J. Experimental Psychology: Human Perception and Performance* 1 (August 1975): 288–99.
- Fischhoff, Baruch, and Beyth, R. "I Knew It Would Happen—Remembered Probabilities of Once-Future Things." *Organizational Behavior and Human Performance* 13 (February 1975): 1–16.
- Fishburn, Peter C. *Nonlinear Preference and Utility Theory*. Baltimore: Johns Hopkins Univ. Press, 1988.
- Forsythe, Robert; Kennan, John; and Sopher, Barry. "An Experimental Analysis of Bargaining and Strikes with One Sided Private Information." Working Paper no. 87-4. Iowa City: Univ. Iowa, Dept. Econ., 1987.
- Glosten, Lawrence R., and Milgrom, Paul R. "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *J. Financial Econ.* 14 (March 1985): 71–100.
- Grossman, Sanford J. "An Introduction to the Theory of Rational Expecta-

- tions under Asymmetric Information." *Rev. Econ. Studies* 48 (October 1981): 541-59.
- Hogarth, Robin M., and Reder, Melvin W., eds. *The Behavioral Foundations of Economic Theory*. Suppl., *J. Bus.*, vol. 59 (October 1986).
- Jensen, Michael C., and Meckling, William H. "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure." *J. Financial Econ.* 3 (October 1976): 305-60.
- Kahneman, Daniel; Slovic, Paul; and Tversky, Amos, eds. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge Univ. Press, 1982.
- Lichtenstein, Sarah; Fischhoff, Baruch; and Phillips, Lawrence D. "Calibration of Probabilities: The State of the Art to 1980." In *Judgment under Uncertainty: Heuristics and Biases*, edited by Daniel Kahneman, Paul Slovic, and Amos Tversky. Cambridge: Cambridge Univ. Press, 1982.
- Loomes, Graham. "When Actions Speak Louder than Prospects." *A.E.R.* 78 (June 1988): 463-70.
- Loomes, Graham, and Sugden, Robert. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty." *Econ. J.* 92 (December 1982): 805-24.
- . "Some Implications of a More General Form of Regret Theory." *J. Econ. Theory* 41 (April 1987): 270-87. (a)
- . "Testing for Regret and Disappointment in Choice under Uncertainty." *Econ. J.* 97 (suppl., 1987): 118-29. (b)
- Machina, Mark J. "Choice under Uncertainty: Problems Solved and Unsolved." *J. Econ. Perspectives* 1 (Summer 1987): 121-54.
- Myerson, Roger B. "Negotiation in Games: A Theoretical Overview." In *Uncertainty, Information, and Communication: Essays in Honor of Kenneth J. Arrow*, vol. 3, edited by Walter P. Heller, Ross M. Starr, and David A. Starrett. New York: Cambridge Univ. Press, 1986.
- Nickerson, Raymond S.; Baddeley, Alan; and Freeman, Barbara. "Are People's Estimates of What Other People Know Influenced by What They Themselves Know?" *Acta Psychologica* 64 (March 1987): 245-59.
- Plott, Charles R., and Sunder, Shyam. "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models." *J.P.E.* 90 (August 1982): 663-98.
- Ross, Lee; Greene, David; and House, Pamela. "The 'False Consensus Effect': An Egocentric Bias in Social Perception and Attribution Processes." *J. Experimental Soc. Psychology* 13 (May 1977): 279-301.
- Ross, Stephen A. "The Economic Theory of Agency: The Principal's Problem." *A.E.R. Papers and Proc.* 63 (May 1973): 134-39.
- Russell, Thomas, and Thaler, Richard. "The Relevance of Quasi Rationality in Competitive Markets." *A.E.R.* 75 (December 1985): 1071-82.
- Salop, Joanne, and Salop, Steven. "Self-Selection and Turnover in the Labor Market." *Q.J.E.* 90 (November 1976): 619-27.
- Weber, Martin, and Camerer, Colin F. "Recent Developments in Modelling Preferences under Risk." *OR Spektrum* 9 (September 1987): 129-51.