Paying $1 to lose $2: Misperceptions of the value of information in predicting the performance of others

George Loewenstein, Department of Social and Decision Sciences, Carnegie Mellon University

Don Moore, Graduate School of Industrial Administration, Carnegie Mellon University

Roberto A. Weber, Department of Social and Decision Sciences, Carnegie Mellon University

April 30, 2003[*]

Abstract

Traditional economic and decision-making models allow for "free disposal" of information, meaning that more information will never make a decision maker worse off. This implies that those faced with making decisions should always place non-negative value on information. Building on previous research on the "curse of knowledge," we explore situations where this might not be so. In two experiments, we document situations in which participants place positive value on information, even when acquiring that information hurts their performance and earnings. In the first experiment, a significant number of participants pay for information – the solution to a puzzle – that hurts their ability to predict how many others will solve the puzzle. In the second experiment, a majority of participants choose to "hire" informed – rather than uninformed – agents, leading to lower earnings. We discuss implications of our results for the role of information and informed decision makers in economic situations.

Introduction

Information is typically assumed to be valuable for decision-making, and in most cases it is. Information helps resolve uncertainty concerning the likelihood and value of outcomes, which allows people to make informed choices between alternative courses of action. It can also shed light on the likely behavior and strategies of others, which helps people to behave more strategically themselves. Stigler's (1961) seminal analysis of the economics of information spawned a considerable literature exploring the extent to which people can derive rents from possessing information (e.g., Lewis and Sappington, 1997; Osband, 1989; Porter, 1995).

One fundamental assumption underlying almost all economic discussions of information is that more information is (weakly) better for decision-making.[1] Information is rarely thought of as bad, in part because it is widely assumed that decision-makers can ignore information that is not valuable or that should not be used. This "free disposal" assumption implies that the value placed on additional information can never be less than zero. In most cases, using information will lead to better decisions, and in those in which it doesn't, the information will be ignored.

This free disposal assumption may be of questionable validity. Camerer, Loewenstein, and Weber (1989) conducted experiments demonstrating that participants were not able to ignore previously received information when subsequently making a decision and ended up making worse decisions as a result, even though the information they received was accurate. In their experiments, one group of participants made guesses about the earnings of a series of companies based only on information in a report. A second group of participants then traded "assets" (one

---

[1] There are some cases in which information has been shown to hurt decision-making. For instance, information about what others have done can lead to "herd behavior" and outcomes that are unfavorable *ex post*, even when decision makers are Bayesian and the information use is rational *ex ante* (Banerjee, 1992). Another situation in which more information might hurt decision makers is when they experience "information overload" Specifically, the simultaneous arrival of too many pieces of information can lead to a situation where sorting through the additional information leads to additional costs and errors (see, for instance, Earl, 1990).

for each company) with underlying value equal to the average of the first group of participants' predictions for that company. When these traders were given the actual earnings for the companies (in addition to the reports also received by the original guessers), their trades revealed a bias away from the guesses of the group they were trying to predict, and in the direction of the actual earnings. This phenomenon, which the authors dubbed "the curse of knowledge," indicates that individuals cannot always recover mental states in which they did not possess unhelpful information, even when such recovery would be beneficial. Participants trying to predict the guesses of other participants who did not know the actual earnings should have ignored the actual earnings when making their predictions, but did not do so, either because they were unable to or because they did not think it was necessary.[2]

Camerer et al. did not try to measure whether or not participants would have *preferred* to receive the actual estimates. It might be the case that participants were aware of the negative effect of information but could not ignore it, and hence would have been unwilling to pay for it (or even might have paid to avoid receiving it). Alternatively, whether or not they could have ignored the information, they may have not recognized that it was affecting their judgments adversely. Therefore, while participants in their experiments exhibited the curse of knowledge, the experiments did not address the question of whether or not participants placed positive value on the information that ultimately hurt them. This is especially important since in real economic environments – unlike in the experiments by Camerer et al. – the decision to acquire information is usually endogenously made by economic actors themselves.

In this paper, we report experiments that further explore the implications of the potentially harmful effects of more information. Building on the earlier work demonstrating the

---

[2] Camerer et al. also found that market forces reduced the bias: when the "assets" were traded in a market, participants on average were less susceptible to the curse of knowledge than when they simply tried to predict what other participants had guessed.

curse of knowledge, we explore individuals' beliefs concerning the usefulness of additional information and knowledge. Like Camerer et al., we document situations in which people are unable to ignore "harmful" information. As a result, those with more information are worse at performing tasks in which they have to predict the behavior and performance of other participants who are not informed.

We extend this result to show that participants are unaware of the harmful effects of more information, and instead place positive value on such information. We find that a significant number of participants are willing to pay for information that causes them to make less money, both because of the adverse effects of the information they receive and because of the cost they incur to receive that information. In a second experiment, we show that when participants have to decide between hiring "informed" (and biased) agents and "uninformed" agents, they opt for the former, resulting in lower earnings.

Our results are consistent with the notion that people's naïve theories about their use of information parallel economic theories in assuming that more information is good (or at least not bad). While this rule of thumb will often lead to better decision-making, our studies show that this is not always the case. We conclude the paper by exploring possible implications for economically consequential situations. In particular, our results are important for decisions regarding whether to acquire information (or pay for the help of an informed agent), especially when decision makers' goal is to predict the behavior of uninformed others. In these situations, which parallel our experiments, information can be harmful; but people may have the opposite intuition and may therefore acquire too much information or employ overly informed agents.

Experiments

Participants in our experiments are given the goal of predicting the performance of others in solving a problem. We show that knowing the solution leads people to make worse predictions about the behavior of those trying to solve the problem. In the first experiment, we also show that a significant number of people who are given the choice of obtaining the harmful information – even at a cost – choose to do so. In the second experiment, we find that if people are given the choice of "hiring" an agent that is either informed or uninformed, a majority of participants tie their earnings to the informed agent and end up making less money as a result.

Experiment 1: Paying for cursed knowledge

Experimental Design

Participants in two sections of an introductory business class at Carnegie Mellon (n = 66) viewed three video clips. In each clip, two nearly identical images alternated appearing on the screen, each one appearing for about one second. The two images alternated for about 20 seconds. In between each appearance of the images, there was a very brief flash in which the screen was completely white. The two images differed in one important aspect. For instance, one set of images is pictured in Figure 1. Before reading on, try to distinguish the difference between the two images.

Figure 1. Sample of images used in experiment 1

These video clips have been previously used to demonstrate "change blindness" – the difficulty most people have noticing changes or inconsistencies in visual perception, even when these are as substantial as in Figure 1 (Rensink, O'Regan, and Clark, 1997; see also, Simons and Levin, 1997). Therefore, we predicted that participants would have a difficult time noticing the differences.

While most people have a hard time noticing the differences between the paired images, they are quite obvious once they are highlighted. For instance, notice that the two images in Figure 1 are identical except that the one on the right has the shadow cast by the helicopter below the jeep, while the one on the left does not. As with Camerer et al.'s experiments on the curse of knowledge, we predicted that participants who were informed of the difference would find it very difficult not to notice it and would tend to overestimate the extent to which other participants would notice the difference.

For each video clip (each pair of images), participants were first told that their goal was to identify the difference between the two images. Specifically, they were instructed that, "There is one difference between the pictures you will see in each clip. Look to see if you can spot the difference." Participants were also asked to predict what percentage of their classmates who did

5

not know the difference between the two clips would be able to spot the difference. Participants were paid for the accuracy of their predictions. If a participant's guess was within 2 percentage points of the actual percentage, then he or she would receive $10. If the guess was 3, 4, or 5 percentage points away, the payment was $5. Guesses off by more than 5 percent earned nothing. Participants repeated this task three times (once for each video clip) and their earnings were summed across all three video clips. Participants were not given any feedback until after the experiment.

Across the three clips, participants experienced each of three following information conditions:

- In the *Uninformed* condition, participants were not informed of the difference between the two pictures. They simply watched the video clip and were then asked to make their predictions.

- In the *Informed* condition, participants' written instructions informed them, in bold type, of the difference. For instance, for the clip with pictures represented in Figure 1, participants in the Informed condition were told, "CLUE: The helicopter's shadow disappears."

- In the *Choice* condition, participants were given the option of finding out what differed between the two images. Each participant received an envelope that revealed inside what the difference was. However, participants were told that by opening the envelope they would sacrifice a $0.50 bonus.

	Clip 1 – **"Statue"**	Clip 2 – **"City"**	Clip 3 – **"Chopper"**	Number of participants
Sequence 1	Uninformed	Informed	Choice	25
Sequence 2	Informed	Choice	Uninformed	20
Sequence 3	Choice	Uninformed	Informed	21

Table 1. Number of participants by sequence of conditions

Each participant saw all three information conditions. Table 1 presents the three sequences in which participants experienced the information conditions and the corresponding sample sizes. To minimize any effect of curiosity, all participants were told that they would be shown all three clips again and informed about the difference between the images at the conclusion of the experiment.

Results

When participants were uninformed about the change, 20 percent of them correctly identified the change, and this did not differ by video clip ($F(2,63) < 1$, ns). Our experiments, therefore, replicated the finding that the changes are difficult to detect.

Information condition		Mean prediction	Standard deviation	N
Uninformed		30.1 %	25.6	66
Informed		58.2 %	32.7	66
Choice		40.6 %	29.5	66
Choice (unopened)	(71%)	34.6 %	29.0	47
Choice (opened)	(29%)	55.4 %	25.8	19

Table 2. Predictions pooled by information condition across sequences

As the results in Table 2 indicate, uninformed participants on average guessed that 30 percent (SD = 26 percent) of their uninformed peers would spot the change; they earned an average of $1.21 (SD = $2.49). When participants were informed about the difference in the two pictures, they guessed that 58 percent (SD = 33 percent) of their uninformed peers would spot the difference, and earned an average of $0.45 (SD = $1.69). The average within-subject difference between guesses in the Informed and Uninformed conditions is significantly different from zero for both guesses (t(65) = 6.28, p < 0.001) and payoffs (t(65) = 2.19, p < 0.05). These results are consistent with the curse of knowledge: Participants who are told the difference between the two pictures are worse at predicting how frequently other participants who do not know the difference will be able to find it.

Among uninformed participants, some figured out the difference on their own (13 of 66). Since they did so before making their guesses, we might expect them to be more likely to correctly infer how difficult it is to notice the difference. This was not the case. For participants in the Uninformed treatment who figured out the difference, the mean guess was 63.4 percent, which is slightly higher than the mean guess in the Informed condition. Therefore, participants who figured out the difference on their own were no less likely to fall victim to the curse of knowledge than those who are told of the difference. Interestingly, the mean guess by uninformed participants who did not figure out the difference was 21.9 percent, which is very close to the actual percentage (20 percent).

The important question for our main hypothesis, however, has to do with what participants will do when given the choice of being informed or uninformed. This is exactly the decision faced by participants in the Choice condition. When given the choice of whether to

learn the difference between the two pictures before seeing the clip and making their guess, 19 of 66 participants (29 percent) chose to open the envelope and become informed. These participants all sacrificed $0.50 for doing so.

The pattern of earnings among participants in the Choice condition similarly reflects the curse of knowledge. The 47 participants who chose not to open their envelopes guessed, on average, that 35 percent (SD = 29 percent) of their uninformed peers would see the difference, while the 19 participants who chose to pay $0.50 to become more informed guessed, on average, that 55 percent (SD = 26 percent) of their uninformed peers would see the difference. This difference is significant (t(64) = 2.71, p < 0.01). As a result, those who chose to remain uninformed earned an average of $1.49 (SD = $3.10), whereas none of those who chose to open their envelopes earned anything. This difference is also significant (t(64) = 2.08, p < 0.05).

Table 3 presents the results broken down by sequence and condition. Each of the rows in the table presents results from the same group of subjects, with the results for the condition they were first exposed to in the second column, and the results for the condition they were last exposed to in the last column. The percentage in each cell refers to the mean prediction made by participants in that information condition and sequence. For participants in the Choice condition, the mean prediction is also broken down by whether or not participants chose to open the envelope.

Note first that the predictions do not vary by sequence for the Uninformed condition (30, 33, and 28 percent in cells a, f and h, respectively; none of the differences is significant). They do, however, differ somewhat for the Informed condition (59, 69, and 47 percent in cells b, d and i, respectively; only the two most extreme are significantly different from each other (t(39) = 2.29, p < 0.05)). Participants who had previously experienced the Choice and Uninformed

9

conditions (Sequence 3) made the lowest predictions while participants who had no previous experience (Sequence 2) made the highest predictions. This is expected, as participants who have previously had to try to figure out the differences between two similar pictures are more likely to realize how difficult it is to notice the difference. However, participants in all three of the Informed cells still made predictions that are on average greater than those in all three of the Uninformed cells, meaning that the increased perspective that comes with experience is insufficient to overcome the curse of knowledge. Therefore, our result that participants in the Uninformed condition are better calibrated on average than those in the Informed condition is supported even when we look for possible sequence effects, providing strong support for the curse of knowledge in this task.

	Prediction/clip 1 **Statue**	Prediction/clip 2 **City**	Prediction/clip 3 **Chopper**
Sequence 1 Uninf-Inf-Choice n = 25	29.9 % (SD = 28.2)	59.3 % (SD = 34.0)	46.9 % (SD = 26.0) Not open: 49.8 (16) Open: 41.8 (9)
	(a: uninformed)	(b: informed)	(c: choice)
Sequence 2 Inf-Choice-Uninf n = 20	68.7 % (SD = 32.5)	48.2 % (SD = 29.7) Not open: 29.1 (11) Open: 71.4 (9)	33.1 % (SD = 26.2)
	(d: informed)	(e: choice)	(f: uninformed)
Sequence 3 Choice-Uninf-Inf n = 21	26.0 % (SD = 29.2) Not open: 25.6 (20) Open: 34.0 (1)	27.6 % (SD = 22.6)	46.7 % (SD = 28.8)
	(g: choice)	(h: uninformed)	(i: informed)

Table 3. Predictions by sequence and information condition

(means in Choice condition also presented by choice)

Examining participants in the Choice condition across pictures and sequences reveals that participants in Sequence 3, who experience Choice before any other conditions, are very unlikely to open the envelope (1 of 21). There are at least three possibilities for why this might be the case. First, it is possible that participants in this condition do not want to open the envelope because they want to see if they can spot the difference on their own, a tendency which should act against our hypothesized effect. A second possibility is that participants were overconfident in their ability to detect the change, which also works against our hypothesis. The last possibility is that participants without prior experience with this task do not believe that knowing the difference between the images will increase their earnings beyond the $0.50 cost, but, after observing similar images, many of them change this belief, suggesting that the effect could worsen with experience.

Overall, the results support our main hypotheses. Examining the aggregate data, participants are clearly better off if they are not informed. However, a significant number of participants choose to become informed and pay a $0.50 fee to do so. While there are some interesting interactions when sequence effects are examined, none of these directly works against the main result. Still, the fact that participants in one of the sequences (Sequence 1) are more likely in the Choice condition to make better predictions when they open the envelope (mean guess: 41.8 percent) than when they do not (49.8 percent) – although participants in **all** the other cells of Table 3 do better when they are uninformed (mean guesses: 29.9, 29.1, 33.1, 25.6, and 27.6 percent) rather than informed (mean guesses: 59.3, 68.7, 71.4, 34.0, and 46.7 percent) – suggests that we should further explore the robustness of our main result.

In Experiment 2, we test the robustness of the aggregate result using a different task and explore whether this result persists when participants are not choosing whether to become

informed themselves, but are instead deciding whether it is better to hire an informed agent. Having participants decide about whether to hire informed or uninformed agents eliminates the issue of curiosity – that participants might choose not to be informed because they are curious to see if they can solve the problem themselves.

Experiment 2: Hiring "cursed" agents

Experimental Design

One large session was conducted with 166 students from Carnegie Mellon University and the University of Pittsburgh. Participants showed up to a large auditorium and were told that they would be paid based on their decisions in the experiment.[3] Upon arriving, they were seated and received written instructions that they were told to read. The instructions differed based on the role a participant was randomly assigned to.

Each participant was assigned to one of four roles: *Solver*, *Informed Predictor, Uninformed Predictor*, or *Chooser*. Each Solver was given one puzzle to solve. The puzzle was a simple logic problem in which participants needed to generate an insight to figure out the solution. We used two different puzzles to rule out the possibility that the results are idiosyncratic to one specific puzzle. The two different puzzles were the "boxes" puzzle and the "chains" puzzle, both shown in Figure 2. Roughly half the participants in each role had the chains puzzle, and the other half had the boxes puzzle (see Table 4 below).

Solvers were paid based on how quickly they solved the puzzle. Specifically, they were told that if they solved the puzzle immediately they would receive $6. The amount they received

---

[3] The experiment was the first part of several tasks that the participants completed (which included another experiment and filling out questionnaires). Since this was the first task in which they participated, and since they were not told about the other tasks until after this experiment was completed, it is unlikely that any of the other tasks affected performance in this one.

went down by one cent for each second they spent solving the puzzle. If they did not solve the puzzle after 10 minutes (600 seconds) then the payment was equal to zero. Fourteen participants were in the role of Solver.

Chains Puzzle: You have four chains of three links each, shown below. Your challenge is to take the four chains and form them into one continuous ring while breaking and re-connecting no more than three links. Which three (or less) links do you break and re-connect? When you have the answer, draw arrows to each of the links, and have the experimenter verify your answer.
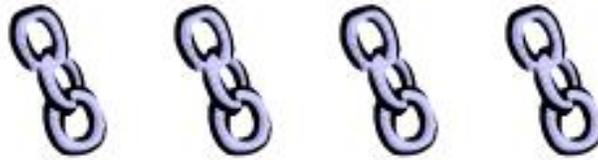
Boxes Puzzle: By repositioning only two of the matches in the following picture, how would you create four squares instead of five? Remember that the squares may be repositioned but the new squares will be the same size as the old ones. When you have the answer, draw the new arrangement of matchsticks, and have the experimenter verify your answer.
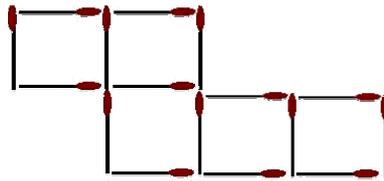

Figure 2. Puzzles used in experiment 2 ("chains" and "boxes")

Both types of Predictors were told that they would be paired with one randomly selected Solver. The Predictor's task was to predict how long the Solver would take to solve the puzzle. The Predictors were shown the puzzle. They were rewarded for predicting longer times (i.e., waiting longer), but were penalized for exceeding the actual time it took the Solver to solve the puzzle. Specifically, Predictors received one cent for every second they predicted the Solver would take, but that payment fell to zero if their prediction was longer than the actual time it took the Solver. In other words, Predictors maximized their payoffs when they predicted exactly how long it took the Solver to solve the puzzle, but not longer.

Ninety-nine participants were in the role of Predictor. Fifty of these predictors were told the solution to the puzzle (Informed Predictors) and Forty-nine were not (Uninformed Predictors). We designed the incentives faced by Predictors to very roughly mimic those that would be faced by someone trying to decide how long to wait to introduce a new invention to the market, where there is a threat that a competitor may also come up with the invention, and if so may introduce it to the market first. In such cases, it often pays to be the first mover, but delaying introduction of the invention to the market allows one to refine its design. If people in this situation suffer from the curse of knowledge, then one would expect them to exaggerate their competitors' progress and hence to introduce their own product too early. Thus, our first prediction for the results of the study was that Informed Predictors would underestimate solution times, and would make less money as a result than Uninformed Predictors, who were not expected to underestimate solution times.

Role	Boxes	Chains	Total
Solvers	8	6	14
Choosers	27	26	53
Uninformed predictors	23	26	49
Informed predictors	27	23	50
Total	85	81	166

Table 4. Number of participants by puzzle and condition

Another fifty-three of the participants were in the role of Chooser. Each Chooser's task was to decide whether to tie their payment to that of an Informed or Uninformed Predictor. The problem faced by Choosers is similar to that of a principal in the invention problem just described who must hire an agent to predict how long it will take competitors to come up with

the idea for the invention. Choosers were first asked to predict the average payoffs for the two different types of Predictors. Then they were told that their payment would be equal to that of one randomly chosen Predictor, but they could pick whether that Predictor would come from the set of the Informed or the Uninformed. Our second, and main, hypothesis is that Choosers will misjudge the benefit of information and will guess that Informed Predictors will make more money and will select an Informed Predictor as their "agent."

Table 4 presents the number of participants for each role and for each puzzle.

Results

Six of fourteen Solvers (43 percent) were able to solve the puzzle. The remaining eight Solvers worked on the puzzle the full ten minutes and did not solve it. The average time spent for all Solvers (including the ones who did not finish) was 7 minutes and 2 seconds (SD = 3:54), and did not significantly differ between the boxes (mean = 7:20, SD = 4:14) and chains (mean = 6:37, SD = 3:47) puzzles (t(12) = 0.32).

Predictors, on average, predicted that Solvers would require 4:08 (SD = 2:32) to solve the puzzle. Both Informed and Uninformed Predictors underestimated Solver solution times but, as in experiment 1, Informed Predictors did worse, and predicted that Solvers would solve the puzzle more quickly (mean = 3:36, SD = 2:28) than did Uninformed Predictors (mean = 4:41, SD = 2:30). This difference is significant at the $p < 0.05$ level (t(97) = 2.17). As a result, Informed Predictors earned less money on average (mean = \$1.45, SD = \$0.84) than Uninformed Predictors (mean = \$1.76, SD = \$0.80), and this difference is significant at the $p < 0.05$ level in a one-tailed test (t(97) = 1.88).

Much as in Experiment 1, the results of Experiment 2 demonstrate the curse of knowledge. Informed Predictors did a worse job predicting the performance of Solvers than did

15

Uninformed Predictors, and ended up making less money as a result. Given this, unbiased Choosers should correctly believe that Informed Predictors are likely to earn less money that Uninformed Predictors, and should select Uninformed Predictors as their "agents."

This is not the case. Choosers tended to believe that Informed Predictors would earn more money that Uninformed Predictors. Chooser's average estimates of earnings for Informed Predictors were \$3.43 (SD = 1.93) and for Uninformed Predictors they were \$2.77 (SD = 1.51). The average within-subject difference between these estimates (\$0.65) is significantly different from zero ($t(52) = 2.26$, $p < 0.05$). Of the 53 Choosers in the experiment, 28 (53 percent) gave an earnings prediction that was higher for the Informed Predictor than the Uninformed Predictor, 17 (32 percent) guessed that Uniformed Predictors would have greater earnings, and 8 (15 percent) guessed equal earnings for both types of predictors.

The misprediction by Choosers is even more dramatic when judged against the standard of accurately predicting earnings. Choosers on average guessed that both kinds of Predictors would make more money than they actually did. However, Choosers overestimated the earnings of Informed Predictors (mean overestimation = \$2.09, SD = \$1.94) by more than they did for Uninformed Predictors (mean overestimation = \$1.06, SD = \$1.51). The average within-subject difference between the degree of overestimation for Informed and Uninformed Predictors (\$1.03) is significantly different from zero ($t(52) = 3.54$, $p < 0.001$).

Finally, Choosers' expectations that Informed Predictors would earn more money are reflected in how they chose to have their earnings determined. The majority of Choosers (33 of 53, or 62 percent) chose to tie their payoffs to informed Predictors. This difference is significant at the $p < 0.05$ level in a one-tailed Binomial test using the normal approximation with adjustment for continuity ($z = 1.65$). As in Experiment 1, participants tend to believe that more

16

information is better, and end up making less money by choosing to have their earnings tied to those of an informed – rather than uninformed – agent.

The results of Experiment 2 provide stronger, and cleaner, support for our main hypothesis. Participants were clearly subject to the curse of knowledge: Informed Predictors did significantly worse than Uninformed ones in predicting the amount of time it would take solvers to complete the puzzle. Choosers' guesses, however, exhibited the opposite pattern – they tended to believe that Informed predictors would do better. In addition, when given the choice of selecting an agent to determine their earning, a majority of Choosers selected agents that were Informed rather than Uninformed, leading to lower payoffs.

Discussion

Taken together, our experiments demonstrate that participants exhibit the curse of knowledge when trying to predict how easy it will be for other people to obtain a necessary insight for solving a problem. In both experiments, participants who were given the solution to the problem or discovered it on their own tended to make biased predictions, leading to lower performance and earnings. This result is consistent with previous work demonstrating the curse of knowledge.

We also demonstrate that a significant number of people are unaware of this bias and believe that more information will be at least weakly better. In the first experiment, a significant number of participants (29 percent) were willing to pay \$0.50 to receive information that would hurt them. In the second experiment, a majority of participants (62 percent) opted to "hire" an informed agent and ended up making less money as a result.

Of course, we demonstrated this bias using decision tasks with very specific characteristics. In our experiments, the main problem to be solved consisted of a task in which a participant needed to discover an insight or solution that was not transparent at first. In both cases, the insight involved seeing the problem in a way that is different than the way that most people see it. Prior research has shown that outcome feedback (in this case the solution to the problems) biases people's predictions of others for insight problems, but not for all other types of problems (Hoch and Loewenstein, 1989). For example, being told the answers to trivia problems often helps one to predict whether others will be able to answer those problems correctly because, if the answer is surprising, one will recognize that few people will get it right. Therefore, one should be cautious of generalizing our results to too wide a domain of problems and tasks.

Our results do not address the question of whether more information will generally be better when decision makers are not trying to predict the performance of others, or when the underlying problem is not one in which insight plays a key role. Our main result should be viewed more as a demonstration of the combined facts that accurate information can, in some situations, be harmful and that a significant percentage of people are not aware of when this is true.

In fact, there are many consequential economic and organizational situations with these characteristics, such as the example mentioned earlier in which a firm is trying to figure out how quickly or easily a competitor will develop a product or innovation. Our results suggest that in these situations, knowing more about the key insight associated with the product or innovation may lead to worse predictions, but that key decision makers may delegate those decisions to those who know more. A similar problem surrounds the question of who should write product

documentation or instruction manuals. Our results suggest that the people who know the most about the product or the topic may over-estimate the ease with which others will be able to understand the necessary information. It has been shown, for example, that experts on the use of a telephone headset were worse than people with intermediate levels of experience when it came to predicting how long it would take novices to learn the basics of using the headset (Hinds, 1999). Therefore, the most informed or most knowledgeable individuals may be worst at writing such documentation than someone who is less informed. However, there may be a common bias to assume that those with more information will be the best at writing such documents. Both of the above examples would be similar to our experimental result that people tend to hire the wrong kind of agent to try to predict how much others know or how easily they will solve problems.

Final comments

Stigler's seminal paper on the economics of information initiated an extraordinarily productive line of research on the "new economics of information," which has encompassed phenomena such as signaling, adverse selection, asymmetric information in bargaining and "herd behavior." We hope that the work presented here will become part of a "new new" economics of information that draws on psychological research to revise some of the strong, unrealistic assumptions that economists typically make about the ways in which people use information.

Some of this new research calls into question conventional assumptions about information processing, such as the idea that information can be freely disposed of or that people update probabilities in a fashion consistent with Bayes' rule. For example, people exhibit "hindsight bias" (Fischhoff, 1975) that is, in a way, a within-person version of the curse of knowledge; people overestimate their own ability to have predicted events which they know have taken

place. They have a difficult time reverting back to their original beliefs after evidence on the basis of which they updated those beliefs is discredited (e.g., Ross, Lepper and Hubbard, 1975). And in some situations, they seem to underweight base-rates in forming expectations of future events (e.g., Bar-Hillel, 1990).

Another line of research challenges the conventional assumption that people process information in an impartial fashion. For example, research on the self-serving bias shows that people unconsciously and without deliberate intent interpret information in a fashion that is favorable to themselves (Babcock and Loewenstein, 1997). Research on the "confirmatory bias" shows that people behave in a "super-Bayesian" fashion, dismissing evidence that contradicts their preexisting beliefs and overweighing evidence that confirms them (e.g., Lord, Lepper and Ross, 1979; Rabin & Schrag, 1999).

Yet a third line or work focuses on the non-controversial idea that information can constitute a source of utility apart from its role in securing desired material outcomes. Several existing economic models incorporate utility from information – e.g., from anticipation of future outcomes (Loewenstein, 1987; Caplin and Leahy, 2001), beliefs about one's own self-worth (e.g., Koszegi, 2001; Loewenstein, 1999), perceptions of fairness (e.g., Rabin, 1993), and from feelings of identification with groups (Akerlof & Kranton, 2000).

Clearly, there is more to be learned about the economics of information.

References

Stigler, G. (1961). The economics of information. <u>Journal of Political Economy, 69</u>(3), 213-225.

Akerlof, G. A. & Kranton, R. E. (2000). Economics and identity. <u>Quarterly Journal of Economics</u>, <u>115</u>, 713-53.

Babcock, L., & Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving biases. <u>Journal of Economic Perspectives</u>. <u>11</u>, 109-126.

Banerjee, A. V. (1992). A simple model of herd behavior. <u>Quarterly Journal of Economics, 107</u>, 797-817.

Bar-Hillel, M. (1990). Back to base rates. In: Hogarth, Robin M.; Ed; Insights in decision making: A tribute to Hillel J. Einhorn; p. 200-216 Chicago, IL: The University of Chicago Press.

Camerer, C., Loewenstein, G. and Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. <u>The Journal of Political Economy, 97</u>, 1232-1254.

Caplin, A., Leahy, J. (2001). Psychological Expected Utility Theory and Anticipatory Feelings. <u>Quarterly Journal of Economics, 116</u>, 55-79.

Earl, P. E. (1990). Economics and psychology: A survey. <u>The Economic Journal, 100</u>(402): 718-755.

Fischhoff, B. (1975). Hindsight $\neq$ foresight: The effect of outcome knowledge on judgment under uncertainty. <u>Journal of Experimental Psychology: Human Perception and Performance, 104</u>, 288-299.

Hinds, P. J. (1999). The curse of expertise: The effects of expertise and debiasing methods on prediction of novice performance. Journal of Experimental Psychology: Applied 5(2): 205-221.

Hoch, S. & Loewenstein, G. (1989). Outcome feedback: Hindsight and information. <u>Journal of</u>
<u>Experimental Psychology: Learning, Memory and Cognition</u>, <u>15</u>, 605-619.

Koszegi, B. (2001). Who has anticipatory feelings. Working Paper. University of California,
Berkeley.

Lewis, T. R. & Sappington, D. E. M. (1997). Information management in incentive problems.
The Journal of Political Economy, 105(4): 796-821.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. <u>Economic</u>
<u>Journal</u>, <u>97</u>, 666-684

Loewenstein, G. (1999). Because it is there: The challenge of mountaineering...for utility theory.
<u>Kyklos</u>, <u>52</u>, 315-44.

Lord, C., Lepper, M. R. & Ross, L. (1979). Biased assimilation and attitude polarization: The
effects of prior theories on subsequently considered evidence. <u>Journal of Personality and</u>
<u>Social Psychology</u>, <u>37</u>, 2098-2110.

Osband, K. (1989). Optimal forecasting incentives. <u>The Journal of Political Economy, 97</u>(5):
1091-1112.

Porter, R. H. (1995). The role of information in U.S. offshore oil and gas lease auctions.
<u>Econometrica, 63</u>(1): 1-27.

Rabin, M. (1993). Incorporating fairness into game theory and economics. <u>American Economic</u>
<u>Review</u>, <u>83</u>(5), 1281-1302.

Rabin, M. & Schrag, J. (1999). First Impressions Matter: A Model of Confirmatory Bias.
<u>Quarterly Journal of Economics</u>, <u>114</u>(1), 37-82.

Rensink, R. A., O'Regan, J. K., and Clark, J. J. (1997). To See or Not to See: The Need for
Attention to Perceive Changes in Scenes. <u>Psychological Science, 8</u>: 368-373.

Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self perception and social

    perception: Biased attributional processes in the debriefing paradigm. <u>Journal of</u>

    <u>Personality and Social Psychology</u>, <u>32</u>, 880-892.

Simons, D. J., & Levin, D. T. (1997). Change Blindness. <u>Trends in Cognitive Science, 1</u>: 261-

    267.