

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 Weber
Department of Social and Decision Sciences
Carnegie Mellon University

June 6, 2002*

Traditional economic and decision-making models allow for “free disposal” of information, meaning that more information will always make a decision maker (weakly) better off. This implies that those faced with 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 three experiments, we document situations in which subjects place positive value on information, even when learning that information hurts their performance and earnings. In the first experiment, a significant number of subjects 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 subjects choose to “hire” informed – rather than uninformed – agents, leading to lower earnings. The third experiment reveals that the phenomenon is not reduced with experience, but that also that there are individual differences in the degree to which subjects fall victim to the bias. We discuss implications of our results for the role of information and informed decision makers in real economic situations.

* Preliminary and incomplete draft. The authors thank Colin Camerer, and participants at the 2001 Economic Science Association meetings in Tucson for helpful comments and suggestions. We also thank Daylian Cain, Wemi Peters, Erin Morgan, Jeff Crilley, and Sapna Shah for valuable research assistance. The authors also gratefully acknowledge the financial support of the National Science Foundation (SES-0095570 to Weber).

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

George Loewenstein, Don Moore, Roberto Weber

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. Beginning with Stigler's (1961) seminal analysis of the economics of information, there is a considerable literature dealing with the extent to which people can derive rents from information (e.g., Osband, 1989; Porter, 1995; Lewis and Sappington, 1997).

One fundamental assumption underlying almost all economic discussions of information is that more information is better for decision-making. 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.

While information is generally valuable, there are 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 unfavorable outcomes (Banerjee, 1992). This is true even when decision-makers are Bayesian – the fact that others moved first and were observed can lead to outcomes where all later movers follow the “leaders” even when these leaders received bad information. Thus, while using others' signals makes sense from an *ex ante* perspective, there are outcomes in which more people end up taking a

“worse” action than if they had simply ignored what others did – or had not received this information – and had instead followed their own private signals.

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 more errors that are not incurred when less information is received (see, for instance, Earl, 1990). In these cases, the benefits of the additional information can be outweighed by the costs of having to figure out which information to use.

The free disposal assumption may itself be of questionable validity. Camerer, Loewenstein, and Weber (1989) conducted experiments demonstrating that subjects 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 subjects made guesses about earnings of a series of companies based only on information in a report. A second group of subjects then traded “assets” (one for each company) with underlying value equal to the average of the first group of subjects’ 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, labeled “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. Subjects trying to predict the guesses of other subjects 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. Camerer et al. also found that market forces reduced the bias: when the “assets” were traded in a market, subjects on average were less susceptible to the curse of knowledge than when they simply tried to predict what other subjects had guessed.

Importantly, however, Camerer et al. did not try to measure whether or not subjects would have *preferred* to receive the actual estimates. It might be the case that

subjects 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 subjects in their experiments exhibited the curse of knowledge, their experiments do not address the question of whether or not subjects placed positive value on the information that ultimately hurt them.

In this paper, we report experiments that further explore the implications of the potentially harmful effects of more information. Building on the work demonstrating the curse of knowledge, we explore individuals' beliefs concerning the usefulness of additional information and knowledge. Like Camerer et al., we document situations in which subjects are unable to ignore "bad" 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 subjects who are not informed. In addition, we extend this result to show that subjects are unaware of the harmful effects of more information, and instead place positive value on such information. We find that a significant number of subjects are willing to pay for information, which 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. Moreover, we show that when subjects are given the choice between decisions made with and without harmful information, a large number choose to receive the harmful information.

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 since many situations in economics involve deciding whether or not to acquire information (or pay for the help of an informed agent) when making decisions. Our experiments indicate that people are likely to perform worse with more information when, as in our experiment, their goal is to predict the behavior of uninformed others.

But, people may have the opposite intuition and may therefore acquire too much information or employ overly informed agents.

Experiments

All our experiments use situations in which the goal of most subjects is to predict the behavior or performance of others. The subjects whose behavior is being predicted perform a task in which they either have to solve a problem or find a “flaw.” In the experiments, the problems all involve obtaining some insight or seeing the problem in a particular way in order to obtain the solution. We show that subjects who possess the insight or solution make worse predictions about those trying to solve the problem than those who do not. In the first experiment, we also show that a significant number of subjects who are given the choice of obtaining the harmful information – even at a cost – choose to do so. In the second and third experiments, we find that if subjects are given the choice of “hiring” an agent that is either informed or uninformed, a majority of subjects tie their earnings to the informed agent and end up making less money as a result.

Experiment 1: Paying for the curse of knowledge

Experimental Design

Subjects 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 1.¹ Before reading on, try to distinguish the difference between the two images.

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

¹ The other two pairs of images we used appear in the Appendix.

also, Simons and Levin, 1997). Therefore, we predicted that subjects would have a difficult time noticing the differences.



Figure 1. Sample of images used in experiment 1

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. Once people are made aware of the difference between the images, it becomes difficult not to see. Therefore, as with Camerer et al.’s experiments on the curse of knowledge, we predicted that subjects 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 subjects would notice the difference.

For each video clip (each pair of images), subjects 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.”

In addition, subjects were also asked to predict how many of their classmates who did 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 subject’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, each subject participated in each of three 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 prediction.
- In the *Informed* condition, participants’ written instructions informed them of the difference in bold type. For instance, for the clip with pictures represented in Figure 1, subjects 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 he or she could open to learn what difference there was. However, subjects were told that by opening the envelope they would sacrifice a \$0.50 bonus.

Each subject participated in all three information conditions. Table 1 presents the three sequences in which subjects experienced the information conditions and the corresponding sample sizes. To minimize any effect of curiosity, all subjects 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.

	Clip 1 – “Statue”	Clip 2 – “City”	Clip 3 – “Chopper”	Number of subjects
Sequence 1	Uninformed	Informed	Choice	25
Sequence 2	Informed	Choice	Uninformed	20
Sequence 3	Choice	Uninformed	Informed	21

Table 1. Number of subjects by sequence of conditions

Results

When participants were uninformed about the change, 20 percent of them correctly identified the change, and this did not differ by video clip ($F(2,63) < 1$, not significant). 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 guessed that 30 percent (standard deviation = 26 percent) of their uninformed peers would spot the change, and earned an average of \$1.21 (std. dev. = \$2.49). When participants were informed about the difference in the two pictures, they guessed that 58 percent (std. dev. = 33 percent) of their uninformed peers would spot the difference, and earned an average of \$0.45 (std. dev. = \$1.69). The average 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 – subjects who are told the difference between the two pictures are worse at predicting how frequently other subjects who do not know the difference will be able to find it.

Among uninformed subjects, some figured out the difference on their own (13 of 66). Since they did so before making their guesses, we might expect them to suffer from

a similar curse of knowledge as subjects in the Informed condition. Alternatively, subjects who figured out the difference on their own – while viewing the clip – might be more likely to correctly infer how difficult it is not notice the difference, and might therefore be better calibrated than those in the Informed condition. The former is clearly the case. For subjects in the Uninformed treatment who figured out the difference, the mean guess was 63.4%, which is slightly higher than the mean guess in the Informed condition. Therefore, subjects who figure out the difference on their own are no less likely to fall subject to the curse of knowledge than those who are told of the difference. Interestingly, the mean guess by uninformed subjects who did not figure out the difference was 21.9%, which is close to the actual percentage (20%).

The important question for our main hypothesis, however, has to do with what subjects will do when given the choice of being informed or uninformed. This is exactly the decision faced by subjects 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 out of 66 participants (29 percent) chose to open the envelope and become informed. These subjects all paid \$0.50 for doing so.

The pattern of earnings among subjects 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 (std. dev. = 29) 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 (std. dev. = 26) 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 (std. dev. = \$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. The percentage in each cell refers to the mean prediction made by subjects in that information condition and sequence. For subjects in the Choice condition, the mean prediction is also broken down by whether or not subjects 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 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 ($t(39) = 2.29, p < 0.05$), indicating that subjects who have previously experienced the Choice and Uninformed conditions (Sequence 3) made the lowest predictions while subjects who have had no previous experience (Sequence 2) made the highest predictions. This is expected, as subjects 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, subjects in all three of the Informed cells still make 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 subjects 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
<u>Sequence 1</u> Uninf-Inf-Choice n = 25	29.9 % (std. dev. = 28.2) (a: uninformed)	59.3 % (std. dev. = 34.0) (b: informed)	46.9 % (std. dev. = 26.0) Not open: 49.8 (16) Open: 41.8 (9) (c: choice)
<u>Sequence 2</u> Inf-Choice-Uninf n = 20	68.7 % (std. dev. = 32.5) (d: informed)	48.2 % (std. dev. = 29.7) Not open: 29.1 (11) Open: 71.4 (9) (e: choice)	33.1 % (std. dev. = 26.2) (f: uninformed)
<u>Sequence 3</u> Choice-Uninf-Inf n = 21	26.0 % (std. dev. = 29.2) Not open: 25.6 (20) Open: 34.0 (1) (g: choice)	27.6 % (std. dev. = 22.6) (h: uninformed)	46.7 % (std. dev. = 28.8) (i: informed)

Table 3. Predictions by sequence and information condition
(means in Choice condition also presented by choice)

Examining subjects in the Choice condition across pictures and sequences leads to some interesting observations. First, subjects in Sequence 3, who experience Choice before any other conditions, are very unlikely to open the envelope (1 of 21). There at least three possibilities for why this might be the case. First, it is possible that subjects in this condition do not want to open the envelope because they want to see if they can spot the difference on their own. This would be true as long as the consumption value from attempting to figure it out exceeded the value they placed on the information minus the 50-cent cost. This would be less likely the case in the other sequences, where they have already experienced the task at least once with a similar video clip. While this possibility would likely influence the results, it would do so in the direction of decreasing the

number of subjects willing to pay for the information (4.8% in Sequence 3 vs. 45% in Sequence 1 and 36% in Sequence 2). This means it goes against our hypothesis (that subjects will be willing to pay for the information) and therefore makes our findings more conservative.

A second possibility is that subjects were overconfident in their ability to detect the change. They were told that “something fairly major” would change between the two photographs they would see. Naïve theories of perception do not include the phenomenon of change blindness. If most subjects believed that they would be able to spot the change on their own, then it would not have made sense for them to pay additional money to be told what changed. However, in the other sequences, after having seen how difficult it was to spot the change in the first video clip, more participants were willing to spend the \$.50 to open the envelope and learn what changed for subsequent clips.

The last possibility is that subjects 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. This is consistent with the result that almost no subjects are willing to pay for the information for Clip 1. This would imply that subjects are initially better calibrated about the value of the information they will get from being told what changes, but they develop a bias in the direction of valuing the harmful information after gaining experience with the task. Since more subjects are willing to pay for information for the second and third predictions that they make, this explanation would mean that experience with this kind of task makes it more likely that subjects will place positive value on harmful information. At the very least, this would imply that the number of subjects willing to pay for the curse of knowledge increases – rather than decreases – with experience and again would not go against our results.

A final interesting observation in Table 3 results from comparing the predictions of subjects in the Choice condition in Sequences 1 and 2. Note that while roughly comparable numbers of subjects (S1: 9 of 25; S2: 9 of 20) choose to open the envelope, subjects who have previously experienced both Informed and Uninformed conditions are slightly better off by opening the envelope, though this difference is not significant. This

is not true of subjects who have previously experienced only the Informed condition. They are the most likely to exhibit the curse of knowledge (mean prediction = 71 percent) if they open the envelope and are much better calibrated if they do not, and this difference (between the mean estimates of 71 and 29 percent) is significant ($t(18) = 4.54$, $p < 0.001$).

Overall, the results support our main hypotheses. Examining the aggregate data, subjects are clearly better off if they are not informed. However, a significant number of subjects choose to become informed and pay a 50-cent 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 subjects 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 subjects 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 subjects decide about whether to hire informed or uninformed agents eliminated the issue of curiosity – that subjects might choose not to be informed because they are curious to see if they can solve the problem.

Experiment 2: Hiring “cursed” agents

Experimental Design

One large session was conducted with 166 students from Carnegie Mellon University and the University of Pittsburgh. Subjects were recruited to show up to a large auditorium and were told that they would be paid based on their decisions in the experiment.² Upon arriving, they were seated and received written instructions that they

² The experiment was the first part of several tasks that the subjects completed (which included another experiment and filling out questionnaires). Since this was the first task in which they participated, and

were told to read. The instructions differed based on the role a subject 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 subjects 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 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.

Both types of Predictors were told that they would be paired with one randomly chosen 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.

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.

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.



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”)

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 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 have 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.

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 subjects for each role and for each puzzle.

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 subjects by puzzle and condition

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 (std. dev. = 3:54), and was did not significantly differ between the boxes (mean = 7:20, std. dev. = 4:14) and chains (mean = 6:37, std. dev. = 3:47) puzzles ($t(12) = 0.32$).

Predictors, on average, predicted that Solvers would require 4:08 (std. dev. = 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, std. dev. = 2:28) than did Uninformed Predictors (mean = 4:41, std. dev. = 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, std. dev. = \$0.84) than Uninformed Predictors (mean = \$1.76, std. dev. = \$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 do a worse job predicting the performance of Solvers than do Uninformed Predictors, and end up making less money as a result. Given this, unbiased Choosers should correctly believe that Informed Predictors are likely to earn less money than 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 than Uninformed Predictors. Chooser’s average estimates of earnings for Informed Predictors were \$3.43 (std. dev. = 1.93) and for Uninformed Predictors they were \$2.77 (std. dev. = 1.51). The average difference between these estimates (\$0.65) is significantly different from zero ($t(52) = 2.26$, $p < 0.05$). Moreover, of the 53 Choosers in the experiment, 28 gave an earnings prediction that was higher for the Informed Predictor than the Uninformed Predictor, 17 guessed that Uninformed Predictors would have greater earnings, and 8 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, std. dev. = \$1.94) by more than they did for Uninformed Predictors (mean overestimation = \$1.06, std. dev. = \$1.51). The average difference between the degree of overestimation for the two conditions (\$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, subjects tend to believe that more 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 additional, and cleaner, support for our main hypothesis. Subjects 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, Choosers selected agents that were Informed rather than Uninformed, leading to lower payoffs.

While the above two experiments clearly show that a significant percentage of subjects value information that is ultimately harmful, all subjects in these experiments only revealed this value one time. Therefore, it is possible that this effect is eliminated with experience and feedback. The following experiments explore whether or not this is the case.

Experiment 3: Learning not to pay for the curse

Experimental Design

In this experiment, we combined elements from the first two and added repetition to explore the effect of feedback on the tendency to value information that results in the curse of knowledge. We used the task and video stimuli from the first experiment and, as in experiment 2, put subjects in a situation where they were hiring “agents” to make predictions for them.

In the experiment, subjects in two sections of a large introductory business class at Carnegie Mellon ($n = 59$) were told about the task posed to subjects in Experiment 1: guessing what percentage of people would spot the change between two pictures. To help them understand the task, they were shown the first of the three clips from Experiment 1. They were not told what changed or how many people saw the change.

Instead, they were told that there had been two types of guessers in Experiment 1: Group I and Group U. Subjects were told that members of Group U did not know what changed before they watched the clip and made their guess and that members of Group I were told what changed before they watched the clip and they made their guesses.

Subjects in this experiment were then faced with a decision task, which they performed for six rounds. In each round, they had to select an “agent” from experiment 1 and would receive the same payoff as that guesser received. They made the choice by selecting from a stack of sheets containing the guesses (and monetary earnings) of subjects in experiment 1. Their choice was whether to select a guesser from a stack containing only guesses from Group I subjects or from one containing guesses only from Group U subjects. There was a \$0.10 fee for choosing a guesser from Group I. After selecting, subjects were shown the guess made by the randomly selected subject about what percentage of people would see the change, and saw their payoff from that choice.

There was a slight difference between rounds 1 and 2 and rounds 3 through 6. After making their choices in rounds 3 through 6, participants were given feedback with one additional piece of information: the true percentage of those seeing the change. This was information was provided to better explore the effect of the feedback and to see if an even stronger form of feedback could correct the bias if the initial feedback did not.

At the end of the experiment, two out of the six rounds were randomly selected to determine payoffs for the experiment. Therefore, even though the monetary payment a subject received only depended on two out of the six choices, subjects did not know which ones these would be until the end.

After they had completed all six rounds, participants were asked several questions. First, they were asked whether they were able to identify the change in the clip they saw at the beginning of the experiment. They were also asked whether they had chosen more people from Group U or I, and why they did so. Finally, they were asked the maximum price they would have been willing to pay to obtain the payoff from a member of Group I rather than a member of Group U.

Results

Table 5 presents the results across six rounds. Each entry in the table represents the percentage of subjects in that particular round, who selected an agent from Group I. As the results indicate, subjects chose to draw their earnings from the informed group 29 percent of the time, and this percentage did not change significantly over the course of the six rounds.

Round	1	2	3	4	5	6	Total
Percentage choosing I	22%	37%	32%	34%	27%	19%	29%

Table 5. Percentage of subjects choosing the informed guesser

Naturally, because uninformed guesses were more accurate, those who chose Group I on average earned less than those who chose Group U. The mean earnings for rounds in which subjects chose from Group U were \$1.17, whereas mean earnings from Group I were \$0.40, less the \$.10 cost of choosing I, or \$0.30.

Table 6 presents the distribution of subjects' behavior across rounds. Roughly one quarter of the subjects correctly inferred that the uninformed agents would, on average, yield more money. Of the 59 subjects in the experiment, 15 subjects (25%) chose U on all six rounds. On the other hand, only one subject (2%) chose I on all six rounds. The results indicate that a substantial number of subjects do not fall subject to the bias.

Number of rounds choosing I	0	1	2	3	4	5	6
Number of subjects	15 (25%)	16 (27%)	12 (20%)	6 (10%)	8 (14%)	1 (2%)	1 (2%)

Table 6. Distribution of subjects by number of Group I choices

The presence of such subjects is further confirmed by the responses concerning why they made the choices that they did. Several subjects' responses indicate that they were precisely aware of the bias. These subjects indicated that knowing the answer

would make informed guessers more likely to think the problem was easy and overestimate the percentage. When we classify responses that clearly indicate this reasoning, we see that 22 of 59 subjects made this type of statement about why they chose as they did. These subjects selected from Group I only 13 percent of the time. At the same time, 6 of 59 subjects made statements that clearly indicated that they believed that the information was likely to lead to better predictions. These subjects selected from Group I 75 percent of the time.

The results of this experiment indicate that a significant number of subjects exhibit the bias even after six rounds of experience, and even after four rounds in which they are aware of the true percentage. These results indicate that the tendency to favor the harmful information is present, though not to the extent it was in experiment 2. Also, there is clear heterogeneity in the extent to which subjects are aware of the bias and respond optimally. Still the presence of a significant percentage of subjects who continue to select the informed agent, even after six rounds, provides further support for the existence of the bias.

Conclusion

Taken together, our experiments demonstrate convincingly that subjects 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. Across the experiments, subjects who are given the solution to the problem or discover 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 demonstrated that a significant number of subjects exhibit a tendency to be unaware of this bias and believe that more information will be at least weakly better. Specifically, our main result has to do with the choices that subjects make based on their perception of the value of information. In the first experiment, a significant number of participants were willing to pay \$0.50 to receive information that would hurt them. In the second experiment, a majority of participants opted to “hire” an informed agent and

ended up making less money as a result. In the third experiment, a significant proportion of subjects selected the informed agent, at a \$0.10 cost, even after several rounds of experience. However, while clearly demonstrating the bias among a large proportion of subjects, we have also found that a significant proportion rarely succumb to it and are, in fact, aware of the negative value of information in the kinds of situations we study.

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 subject 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 – in the first task subjects needed to see a change that was difficult for most people to see and in the second task solvers needed to approach the problem in a slightly different way than most people initially approach 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 a wide 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. Based on the prior research just discussed, however, we believe that this is especially likely to be true when the underlying problem to be solved is one in which insight plays a key role.

In fact, there are many real-world economic 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 indicate that in these types of situations, knowing more about the key insight associated with the product

or innovation may lead to worse predictions, but that decision makers will often delegate those decisions to those who know more. Another example is in trying to decide who should write product documentation or instructions. Our results indicate 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 more biased 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 or knowledgeable. However, there may be a common bias to assume that those with more information will be the best at writing such documents. Both of these 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 a "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 of 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 a lot more to be learned about the economics of information.

References

- Akerlof, G. & Kranton, R. (2000). Economics and identity. Quarterly Journal of Economics.
- Babcock, L., & Loewenstein, G. (1997). Explaining bargaining impasse: the role of self-serving biases. Journal of Economic Perspectives, 11, 109-126.
- Banerjee, Abhijit V. 1992. A simple model of herd behavior. Quarterly Journal of Economics 107(3): 797-817.
- Bar-Hillel, Maya 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, Colin, George Loewenstein and Martin Weber. 1989. The curse of knowledge in economic settings: An experimental analysis. The Journal of Political Economy 97(5): 1232-1254.
- Caplin, Andrew; Leahy, John 2001 Psychological Expected Utility Theory and Anticipatory Feelings. Quarterly Journal of Economics 116(1): 55-79.
- Earl, Peter E. 1990. Economics and psychology: A survey. The Economic Journal 100(402): 718-755.
- Fischhoff, B. (1975). Hindsight \neq foresight: The effect of outcome knowledge on judgment under uncertainty. Journal of Experimental Psychology: Human Perception and Performance, 104, 288-299.
- Hinds, Pamela 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. Journal of Experimental Psychology: Learning, Memory and Cognition, 15, 605-619.
- Koszegi, B. (2001). Who has anticipatory feelings. Working Paper. University of California, Berkeley.
- Lewis, T. R. and D. E. M. Sappington. 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. Economic Journal, 97, 666-684

Loewenstein, G. (1999). Because it is There: The Challenge of Mountaineering... For Utility Theory. Kyklos, 52, 315-44.

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

Osband, Kent. 1989. Optimal forecasting incentives. The Journal of Political Economy, 97(5): 1091-1112.

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

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

Rabin, M. & Schrag, J. (1999). First Impressions Matter: A Model of Confirmatory Bias. Quarterly Journal of Economics, 114(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. Psychological Science 8: 368-373.

Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self perception and social perception: Biased attributional processes in the debriefing paradigm. Journal of Personality and Social Psychology, 32, 880-892.

Simons, D.J., & Levin, D.T. 1997. Change Blindness. Trends in Cognitive Science 1: 261-267.

Appendix A. – Images from experiment 1



“City” – Images differ in the reflection of the building in the water



“Statue” – Shrub to the right of obelisk disappears