HBS Working Paper #05-037
Revised: 8/25/2005

Bounded Awareness: What You Fail to See Can Hurt You

Dolly Chugh
and
Max H. Bazerman

Economists and psychologists rely on widely divergent assumptions about human behavior in constructing their theories. Economists tend to assume that people are fully rational, while psychologists, particularly behavioral decision researchers, identify the systematic ways in which people depart from rationality. Yet, while the two disciplines have offered different predictions of how individuals react to external stimuli, behavioral decision researchers and economists have shared the implicit assumption that individuals will accurately perceive the stimuli available to them.

In this paper, we argue that people actually often fail to perceive and process stimuli easily available to them. In other words, we challenge the tacit assumption that awareness is unbounded and provide evidence that humans regularly fail to see and use stimuli and information easily available to them. We call this phenomenon "bounded awareness" (Bazerman and Chugh, 2005).

Herbert Simon introduced bounded rationality as a "behavioral model (in which) human rationality is very limited, very much bounded by the situation and by human computational powers" (1983, page 34). The concept of boundedness leads us to the important distinction between descriptive and normative views of human behavior. A normative view of behavior is grounded in the assumption that the operation of the human mind is not constrained. A descriptive view of behavior is grounded in the assumption that mental processes are constrained in ways that make normative operations impossible.

Traditional economic theory is based upon a normative view of human behavior,

while much of the work of today's behavioral and experimental economists reflects a descriptive approach. Within psychology, the goal of generating descriptive theories has been to accurately predict how individuals, groups, and organizations will behave.

Here, we argue that people have "bounded awareness" that prevents them from focusing on easily observable and relevant data. We define bounded awareness as the phenomenon in which individuals do not "see" accessible and perceivable information during the decision-making process, while "seeing" other equally accessible and perceivable information; as a result, useful information remains out of focus for the decision-maker.

A "focusing failure" results from a misalignment between the information needed for a good decision and the information included in awareness. Our less formal definition of a focusing failure is captured in the familiar "How could I have missed that?" reaction to the realization that important information that could have been easily seen was ignored. We are deliberate and normative in our labeling of these instances as failures, as they represent costly errors. Researchers are just beginning to understand the systematic effects caused by our bounded awareness but we can and will demonstrate the suboptimality that results when individuals systematically ignore important information. We will argue that the focusing failures that arise from bounded awareness are indeed quite costly.

As illustrated in Figure 1, we propose that awareness is "bounded" when individuals fail to see, seek, use, or share highly relevant, easily accessible, and readily perceivable information during the decision-making process. In early stages, decision-

makers may fail to see or seek out key information. Later, decision-makers may fail to use the information that they have within their awareness because the relevance of this information remains outside of awareness. Finally, individuals may be aware of key information that they fail to contribute to a group's shared awareness despite its relevance to a group decision.

We begin by first describing perceptual mental processes in which obvious information is missed – that is, simply not seen -- by the visual perceiver. Inattentional blindness and change blindness are examples of processes in which individuals do not see or use clearly visible information. We then extend this phenomenon to decision making and forecasting, using evidence about focalism to illustrate how people over focus on some information and fail to use other easily available information. We next examine how these processes of bounded awareness may extend to other important domains and across levels of analysis. We discuss the tendency for groups to show bounded awareness by focusing on shared information and ignoring unique or unshared information held by group members. We examine bounded awareness in negotiators, particularly the tendency for the rules of the game and the decisions of the other party to be ignored. In competitive bidding situations such as auctions, we look at how bidders tend to ignore the impact of uncertainty and the number of competing bidders on their decision, often succumbing to the winner's curse.

The diversity in the evidence we present will make it evident that bounded awareness is not a one-dimensional concept with a singular underlying mechanism. Rather, we are proposing that a wide range of psychological processes, including those of

perception, attention, and cognitive reasoning, are subject to the errors of bounded awareness. We recognize that robust literatures exist on many well-established and related concepts and we make no attempt to thoroughly review them here. Rather, we strive to tie those seemingly disparate examples together with the concept of bounded awareness, Bounded awareness is a phenomenon that encompasses a variety of psychological processes, all of which lead to the same error: a failure to see, seek, use, or share important and relevant information that is easily seen, sought, used, or shared.

INATTENTIONAL BLINDNESS

Neisser (1979) presented a videotape of two visually superimposed teams passing basketballs, one wearing light-colored shirts and the other wearing dark-colored shirts, and asked participants to count the number of passes made between the two teams. Because the players were superimposed on top of each other, the task was moderately difficult. To score accurately, participants had to pay close attention to the task. Yet only 21 percent of participants reported seeing a woman abruptly and clearly walked through the group of players carrying an open umbrella. A repeated viewing of the videotape, without the counting task, reveals that the woman is unambiguously visible in the middle of the screen for a significant part of the video.

We have used this video in the classroom, and find that an even smaller percentage notices the woman. After the first showing of the video, during which students count passes, we ask whether anyone saw anything noteworthy. In a large class,

it is common for just a few people to mention seeing a woman with an umbrella. Predictably, others scoff at the suggestion. Yet, when we show the video for the second time, everyone sees the woman, leading to significant laughter and disbelief. This information sits visible and available in the visual field, yet escapes awareness when competing with a task requiring other attentional resources.

Simons and Chabris (1999) have replicated and extended Neisser's work with a video in which a person in a gorilla costume walks through a basketball game, thumping his chest, and is clearly and comically visible for more than five seconds. Simons provides a series of such demonstrations on a video that is available at www.viscog.com. Recently, Simons and Chabris were awarded the "Ig Nobel Prize," an honor given for scientific achievement "which first make you laugh, and then make you think" (Abrahams, 2004). Since Neisser (1979) first observed this phenomenon, now called "inattentional blindness" (Simons and Levin, 2003), it has become an important area of study for cognitive and perceptual psychologists. Mack and Rock (1998) demonstrate the phenomenon using perceptual experiments in which participants miss some information while focusing on other information. The consequences of inattentional blindness extend to real, life-and-death activities. For example, an airplane pilot who is attending to his controls could overlook the presence of another airplane in his runway (Mack, 2003). Similarly, cell phones can diverts drivers' attention, making inattentional blindness a likely contributor to car accidents. Moore and Egeth (1997) connect inattentional blindness to neural regions in the brain, and Mack (2003) identifies many key independent variables that affect the probability of not seeing the obvious.

CHANGE BLINDNESS

Change blindness is a perceptual phenomenon related to inattentional blindness. Change detection researchers have demonstrated that people fail to notice changes in the information that is visually available to them (Simons, 2000). Interestingly, people often cannot describe the change that has taken place, but do demonstrate traces of memory of what they saw before the change. For example, an experimenter holding a basketball stopped pedestrians to ask for directions (Simons, Chabris, Schnur, and Levin, 2002). While the pedestrian was giving directions, a group of people (confederates in the experiment) walked between the experimenter and the pedestrian. During this interruption, the experimenter handed the basketball to one person in the group. After giving directions, the pedestrian was asked if he or she noticed any sort of change during the brief exchange with the experimenter. Most did not. However, when led to think about a basketball, the pedestrian did recall seeing it at the beginning of the exchange, and some even recalled specific features of the ball. So, while the participants failed to explicitly notice that a change took place, they did hold accurate implicit memory representations of both the pre- and post-change image.

Another example of change blindness is demonstrated in a study by Angelone, Levin, and Simons (2003). Participants watched a videotaped interaction in which easily visible clothing or objects were altered as the camera cut to different angles. Many participants failed to notice the change. Yet it was clear that they had perceived the initial visual stimulus, as they accurately selected the pre-change information from a

photographic lineup. Additional studies by Mitroff, Simons, and Franconeri (2002) found similar patterns.

The possible influence of change blindness in decision making is seen in a study in which participants are asked to choose the more attractive of two faces displayed on a computer screen (Johansson, Hall, and Olsson, 2004). As participants moved the cursor to indicate their choice, a flash on the screen distracted them, and the two pictures were reversed. Nonetheless, most subjects continued to move the cursor in the same direction, and selected the picture they originally viewed as the more attractive. Importantly, they both failed to notice the switch and provided reasons, post-hoc, to support their unintended decision.

The information that is missed in these studies is visual, and the mental processes at work appear to be perceptual. We believe that further research should explore how these processes operate when the information is not necessarily visual and the processes are not necessary perceptual. We anticipate those studies will find that other types of changes also go unnoticed, leading to important errors in decision-making. Bounded awareness, we believe, will extend from perceptual processes to decision-making processes.

For example, Cain, Loewenstein, and Moore (2005) describe the slippery slope of auditors becoming unethical. Essentially, they argue that auditors may be blind to changes made in corporate accounting practices as long as the changes are made slowly – on a slippery slope. Imagine that an accountant is in charge of the audit of a large corporation with a strong reputation. The auditor and client have an excellent

relationship, and the auditor receives tens of millions of dollars in fees from the client each year. For three years, the accountant has viewed and approved the client's high-quality, extremely ethical financial statements. Suddenly, the corporation begins stretching, and even breaking, the law in certain areas. If the accountant were asked if she noticed these transgressions, would she sign a statement certifying that the financial statements were acceptable according to government regulations? Cain et al. believe that most would notice and would refuse to sign.

Now suppose that the auditor saw and approved of high-quality, highly ethical financial statements for one year, after which the corporation begins stretching the law in a few areas, but does not appear to break the law. The third year, the firm stretches the ethicality of its returns a bit more. Some of the company's accounting decisions may now violate federal accounting standards. By the fourth year, the corporation has begun to stretch the law in many areas and occasionally to break it. Auditors are much more likely to notice and refuse to sign the statements in the first version than in the second version, even if the unethical behavior is the same in year four of both stories.

This type of change blindness illustrates the "slippery slope" theory of unethical behavior (Cain et al., 2005), which predicts that individuals are more likely to engage in unethical behavior that occurs in small increments than in unethical behavior that occurs suddenly. Decision makers are less likely to notice small changes and to code them as unethical than they are to code a dramatic change as wrong (Tenbrunsel and Messick, 2004).

Similarly, Gino and Bazerman (2005) show that people are less likely to perceive change if it occurs slowly over time rather then abruptly. They argue that recent business scandals such as the fall of Enron and Worldcom illustrate the *boiling frog syndrome*. According to this folk tale, if you place a frog in a pot of water, and raise the temperature ever so slowly, the gradual warming will comfort the frog into a state of relaxation. Eventually, however, the frog will die due to his inability to sense the gradual increase in water temperature.

FOCALISM AND THE FOCUSING ILLUSION

"Focalism" is the common tendency to focus too much on a particular event (the "focal event") and too little on other events that are likely to occur concurrently (Gilbert & Wilson, 2000; Wilson, Wheatly, Meyers, Gilbert, & Axsom, 2000). Thus, individuals overestimate the degree to which their future thoughts will be occupied by the focal event as well as the duration of their emotional response to the event. For example, individuals overestimate the impact of positive events, such as the win of a preferred sports team or political candidate, on their overall happiness. And even more dramatically, individuals overestimate the impact of negative events, such as a major medical condition, on overall happiness. This error of bounded awareness is a failure to use information that is available and known to the decision-maker, but not incorporated into the conscious decision-making process.

Using similar logic, Schkade and Kahneman (1998) define the "focusing illusion" as the human tendency to make judgments based on attention to only a subset of available

information, to overweight that information, and to underweight unattended information. College students in the Midwest and in Southern California evaluated their own life satisfaction and the perceived life satisfaction of others. While Californians and Midwesterners reported similar levels of life satisfaction, both groups predicted that Californians had greater life satisfaction than Midwesterners. The factor that led to this pattern was that salient difference between California and the Midwest, such as climate, heavily influenced non-residents' judgments of residents' life satisfaction. Schkade and Kahneman argued that when Californians imagined moving to the Midwest, and vice versa, weather became a salient factor, and all other life events affecting satisfaction were outside of awareness. But these factors played a much smaller role in determining the experienced life satisfaction of residents of either region.

Again, these examples of bounded awareness show how decision makers can overlook and fail to see easily available information when considering the implications of future possible events. This tendency also occurs when individuals are considering the probabilities of possible future events. When the 1995 National Basketball Association tournament reached the "elite eight," Fox and Tversky (1998) recruited basketball fans as participants. The subjects were asked to evaluate one of the following: (1) the probability that each team (Chicago, Indiana, Orlando, New York, Los Angeles, Phoenix, San Antonio, and Houston) would win the championship, (2) the probability that the winning team would come from each of the four divisions (Central [Chicago and Indiana], Atlantic [Orlando and New York], Pacific [Los Angeles and Phoenix], and Midwestern [San Antonio and Houston]), or (3) the probability that the winning team would come

from either the Eastern conference (comprising the Central and Atlantic divisions) or the Western conference (comprising the Pacific and Midwestern divisions). If the participants were well calibrated, the sum of the probabilities in all three conditions (for the eight teams, for the four divisions, and for the two conferences) should each have added up to 100 percent.

The combined probabilities for the two conferences summed to 102 percent, close to the calibrated norm. But the sum of the probabilities of the four divisions was 144 percent; even more striking, the sum of the probabilities of the eight teams was 218 percent. Fox and Tversky predicted and explained the finding by observing that when participants focus on a particular team, they see data to support that team winning the tournament; meanwhile, they do not see or use data to support the possibility of another team winning the tournament. In a similar study, Tversky and Koehler (1994) asked medical doctors to assess the probabilities of four mutually exclusive prognoses for a patient and found that the summed probability of the four diagnoses far exceeded 100 percent.

The implications of focalism are not limited to laboratory studies. The Challenger space shuttle tragedy, for example, can be better understood through the lens of focalism (see Vaughn [1996] for a deeper analysis of this disaster). On January 28, 1986, the Challenger was launched at the lowest temperature in its history, leading to a failure of the "O-rings" and an explosion that killed all seven astronauts aboard. Before the launch, the decision makers at NASA examined seven prior launches in which some sort of O-ring failure occurred. No clear pattern between O-rings and temperature emerged

from this data, and the launch continued as scheduled.

Critically, the decision makers failed to consider 17 previous launches in which no O-ring failure occurred. A logistic regression of all 24 launches would have led to an unambiguous conclusion: the Challenger had more than a 99 percent chance of malfunction. The scientists at NASA, however well meaning, caused a tragedy by missing information easily visible and accessible to them. More broadly, the bounded awareness often results when decision-makers and groups limit their analysis to the data at hand, rather than seeking out information most relevant to the question being considered.

BOUNDED AWARENESS IN GROUPS

In terms of decision making, one advantage of groups is that they collectively can possess, and utilize, more information than any individual. This advantage occurs in situations where individuals have and exchange unique information that is not held by other members of the group. However, this comparative advantage of groups is limited by the amount of unique information sharing that occurs in the group. In fact, a number of studies (Stasser and Titus, 1985; Stasser, 1988: Stasser and Stewart, 1992; Gruenfeld, Mannix, Williams, and Neale, 1996) show that groups tend to focus much more on shared information than on unique or unshared information, thus reducing the advantage of the group over the individual as a decision-making body. Group awareness occurs when unique information is distributed throughout the group, thus minimizing the amount of uniquely held information in the group.

For example, Stasser and Titus (1985) asked college students to choose between three candidates running for student council president. The data provided favored Candidate A, who was preferred by 67 percent of individuals and 83 percent of groups comprised of these fully informed individuals. An alternative version of the exercise simulated the nature of information in most real-world groups. Some of the information about the candidates was held by all group members and some of it was unique, including much of the positive information about Candidate A. Now only 23 percent of the individuals in the unshared condition chose Candidate A. However, when the individuals were placed in groups, where every piece of needed information was held by at least one member, a mere 18 percent of the groups chose Candidate A.

Surprisingly, groups miss easily visible and accessible information, demonstrating a bounded awareness similar to that of individuals. In fact, groups discuss shared information more often than unshared information (Stasser and Titus, 1985), despite the fact that groups are often created for the very purpose of increasing the amount of information available for a decision. A paradox emerges; groups are created to share unique information, yet devote their attention to commonly held information.

BOUNDED AWARENESS IN NEGOTIATIONS

Two types of information are critical for any effective negotiator: the decisions of others and the rules of the game. However, due to bounded awareness, these types of information are often not seen or used by negotiators, leading to critical errors. We will use three types of negotiation problems to illustrate the phenomenon: multi-party

ultimatum games, the Monty Hall problem, and the Acquiring a Company problem.

Multi-Party Ultimatum Games

In an ultimatum game, Player 1 divides a known, fixed sum of money any way he chooses by filling out a form stating, "I demand X." Player 2 either accepts the offer and receives her share of the money as allocated by Player 1 or rejects the offer and leaves both parties with nothing. Concerns for fairness often lead Player 1s to be more generous and Player 2s to demand more than economic models suggest.

In this section, we examine multiple-party ultimatum games (Messick, Moore, and Bazerman, 1997). In one multi-party version, six participants are assigned to the roles of A, B, C, D, E, and F. Player A is given $60 dollars to allocate to the six parties and told that the offers to B, C, D, E, and F must be equal and must be an integer. B, C, D, E, and F each record the minimum amount that they would accept from A.

Two versions of the problem exist, and differ in their decision rules. In the "largest amount" version, if the amount that A offers to B-F is equal to or greater than the largest amount requested by B, C, D, E, or F, then A's allocation is distributed; if it is not, then all parties receive 0. In the "smallest amount" version, if the amount that A offers to B-F is equal to or greater than the smallest amount requested by B, C, D, E, or F, then A's allocation offer is distributed; if it is not, then all parties received 0. As with the two-party ultimatum game, players B-F often respond in a bimodal fashion. Many B-F players accept $1, since $1 is better than the $0 they would receive from turning the offer down. But another large group of players B-F demand $10—they want their "fair" share. Since individuals underestimate disjunctive events (those that can occur independently)

and overestimate conjunctive events (those that must occur in conjunction with one another) (Tversky and Kahneman, 1974), player As tend to underestimate the ease of persuading at least one out of five people to accept $1 and overestimate the likelihood of all five individuals accepting anything less than $10.

Messick et al. (1997) demonstrated that players who offered anything less than 10-10-10-10-10-10 in the "largest amount" version tended to get $0 themselves (because the probability of getting even 15-9-9-9-9-9 is trivial). In addition, players who offered anything more than $1 or $2 to the other players in the "smallest amount" version did so because that wanted to be "fair" or because they made a bad decision; the expected payoff by player As fell dramatically as they increased their offers to B-F.

Bounded awareness leads to a failure to distinguish between these two problems. Yet, despite the surface similarities in the problems, the differences in the rules of the game and the likely decisions of the other party are sharply distinct. The effective negotiator focuses on these differences and uses this information to determine the rational response.

The Monty Hall Game

In the 1960s, Monty Hall was the host of the American television game show "Let's Make a Deal." On the show, Monty would ask contestants to pick one of three doors. He and the contestants knew that one of the doors led to the grand prize and that the other two doors led to "zonks," which were small prizes or gag gifts. After a contestant picked a door, Monty would often open one of the other two doors, reveal a zonk, and then offer the contestant the chance to trade their chosen door for the remaining unchosen and

unopened door. Most contestants preferred to stick with their original door.

Years after "Let's Make a Deal" went off the air, statisticians, economists, and journalists became intrigued with the analytical choice faced by the contestants (Selvin, 1975; Nalebuff, 1987; vos Savant, 1990a; vos Savant, 1990b; vos Savant, 1991). Their analysis argued that contestants erred by not switching to the remaining unchosen door. In addition, they argued against the common belief that, with only two doors remaining following the opening of one door by the host, the odds of winning the grand prize are 50-50.

The researcher proposed the following logic. When a contestant first chose a door, she had a 1/3 chance of winning the prize. Let's assume that Monty always opened an unchosen door (the "Monty always opens" condition) and then offered the contestant the opportunity to switch. When Monty opened one door to reveal a zonk, as he always could do, the 1/3 probability did not change. In other words, the contestant still had a 1/3 chance of having picked the grand prize from the start, and a 2/3 chance that the grand prize was behind one of the other two, unchosen doors. After Monty opens one zonk door, the 2/3 chance was now carried by the unopened, unchosen door. Switching doors is the winning strategy, as it increases the participant's odds of winning the grand prize from 1/3 to 2/3.

The critical assumption in this analysis was, "Monty always opens an unchosen door." A different assumption would be the "Mean Monty" condition, in which Monty knew where the grand prize was located and made decisions in order to minimize the contestant's chances of winning. So, after the contestant picked a door, "Mean Monty"

could either choose to end the game or open one door and offer a switch. Because Monty was actively trying to minimize the probability that the contestant would win the grand prize, the contestant should never have accepted an offer from Monty to switch. In fact, Monty's choice to offer a switch should have served as a signal that the contestant had already picked the winning door.[1]

Thus, the winning strategy in the "Monty always opens" condition is to always switch doors, while the winning strategy in the "Mean Monty" condition is to never switch doors. Seeing the importance of Monty's decision rule, and his likely decisions, is the key to the winning strategy. But bounded awareness can prevent individuals from seeing this readily available information.

Acquiring a Company

In the Acquiring a Company problem, an acquiring firm has the opportunity to buy out a target firm. The acquiring firm does not know the exact current value of the target, only that it falls between $0 and $100 per share, with all values equally likely. The target firm is expected to be worth 50 percent more under the acquirer's management than under the current ownership. Thus, it appears to make sense for a transaction to take place. The target firm can accept or reject the acquiring firm's offer.

Consider the logic that would lead to a rational response when considering whether to make an offer of $60 per share:

If I offer $60 per share, the target will accept the offer 60 percent of the time—

[1] In a dynamic game-theoretic equilibrium, the contestant would not know that she won, but should still keep her original choice.

whenever the firm is worth between $0 and $60 to the target. Since all values between $0 and $60 are equally likely, the firm will, on average, be worth $30 per share to the acquirer, for a loss of $15 per share ($45 to $60). Thus, a $60 per share offer is unwise.

Similar reasoning applies to any positive offer, such that, on average, the acquirer obtains a company worth 25 percent less than the price it pays when its offer is accepted. If the target accepts the acquirer's offer of $X, the company is currently worth anywhere from $0 and $X. As the problem is formulated, any value in that range is equally likely. Thus, the expected value of the offer equals $X/2. Because the company is worth 50 percent more to the acquirer, the expected value to the acquirer is 1.5($X/2) = 0.75($X), just 75 percent of its offer price. Thus, for any value of $X, the acquirer's best option is to not make an offer ($0 per share). The paradox of the problem is that although the firm is always worth more to the acquirer than to the target, any offer higher than $0 generates a negative expected return to the acquirer. The paradox results from the high likelihood that the target will accept the acquirer's offer when the firm is least valuable to the acquirer—in other words, when it is a "lemon" (Akerlof, 1970).

The Acquiring a Company problem demonstrates the "winner's curse," where one side, typically the seller, often has much better information than the other side, putting the buyer at a disadvantage. A structural asymmetry is built into the problem: A rational buyer will bid $0 despite the fact that the buyer's valuation of the company is higher than the seller's valuation. The strategic seller will not provide the buyer with information about the company's true value, especially when that value is low. As a result, to avoid

an expected value loss, game theory recommends that buyers not make any offer.

In all three of these seemingly very different negotiation problems (multi-party ultimatum games, the Monty Hall Game, and the Acquiring a Company problem), people tend to ignore the rules of the game and the decisions of the opposing party (Tor and Bazerman, 2003). The problems are not analytically difficult and all of the information the negotiator needs is easily seen. Yet, in the multi-party ultimatum game, Player As allocated an average of $8.15 to the other players in the "smallest amount" condition), while allocating $8.47 to the other players in the "largest amount"). These results reflect bounded awareness, as the optimal strategies for Player A should be dramatically different between the two conditions (offers of $1 vs. $10). Players are passing up significant profit due to this error.

In the "always open" condition of the Monty Hall problem, Friedman (1998) has found that most people do not switch doors, essentially opting for a 1/3 chance of winning rather than trading for a 2/3 chance. Tor and Bazerman (2003) found that 59 percent of participants behave in this suboptimal manner. But in the Mean Monty version, 79 percent made the right decision to keep the existing door. Finally, the majority of participants made the same decision in both versions of the game; only 24 percent answered both versions correctly. The key "Monty Hall" result in Tor and Bazerman (2003) is that the consistency across the two very different versions of the game demonstrate that the rules of the game and the decisions of others are outside the bounds of awareness for most decision-makers.

The game that most extensively demonstrates bounded awareness is "Acquiring a

Company." Extensive research on this problem suggests that bounded awareness leads decision makers to ignore or simplify the cognitions of opposing parties as well as the rules of the game (Tor and Bazerman, 2003; Carroll, Bazerman, & Maury, 1988). Across studies, the modal response range falls between $50 and $75. The common reasoning is: "On average, the firm will be worth $50 to the target and $75 to the acquirer; consequently, a transaction in this range will, on average, be profitable to both parties." Typically, less that 10 percent of participants offer $0 per share. Replications with accounting firm partners, CEOs, investment bankers, and many other skilled groups have produced similar results. Finally, even participants who were paid according to their performance and given many opportunities to learn through experience exhibited the same pattern of responses (Ball, Bazerman and Carroll, 1991; Grosskopf and Bereby-Meyer, 2005).

Most people can follow the logic that the optimal offer in this problem is $0 per share, yet when unaided, most people make a positive offer. Why? Because they systematically exclude information from their decision-making processes that they have the ability to include. They overlook the fact that their expected return depends on an acceptance by the other party, which in turn is affected by the rules of the game, which state that the other side knows its true value before accepting or rejecting the offer. This logic implies that acceptance by the target is most likely to occur when the acquirer is in the least desirable position.

Respondents overwhelmingly have provided solutions that yield a negative expected return. But in a recent adaptation of the Acquiring a Company exercise, Valley,

Moag, and Bazerman (1998) found that when parties talk face-to-face, they can make a trade at a mutually beneficial value. Thus, social interaction can overcome the inefficient outcomes predicted by game theory and behavioral decision theory. Valley et al. suggest that communication motivates parties to create benefits for the other side, facilitates trust, and allows for the type of information exchange not predicted by game theoretic models. Essentially, the social context of face-to-face interaction helps negotiators overcome bounded awareness.

These three problems illustrate how the rules of the game and the decisions of others, two critical and often accessible pieces of information in a negotiation, can be out of focus to negotiators. We believe that these documented focusing failures help to explain negotiation failures beyond the three problems presented here. Camerer and Lovallo (1999) use the term "reference group neglect" to describe the tendency of people to be insensitive to the quality of their competition. Along these lines, Moore and colleagues (Moore and Kim, 2003; Moore and Small, 2004) have found that people expect to perform better than average on easy problems, but worse than average on hard problems, ignoring the fact that easy or hard problems are easy or difficult for most people as well. Moore (2000) also has observed bounded awareness in the context of negotiation deadlines. In a negotiation in which both buyer and seller receive zero payoff if they fail to reach agreement, Moore imposes a publicly known deadline on party, which intuitively appears to put that party at a disadvantage. Of course, both sides should recognize that if one party has a deadline, the other does too. Yet negotiators tend to believe that a deadline put them at an asymmetric disadvantage. In another experiment,

Moore imposes time-related costs on one party, costs that do give an advantage to the party without time-related costs. Moore then gives the party with time-related costs the opportunity to impose a firm deadline on the negotiation, thereby eliminating the individual's asymmetric time-related costs and creating symmetric costs for the failure to reach agreement. Most people reject this option, despite its strategic benefit. By failing to consider how the rules of the game affect the other party, individuals suboptimize their outcomes.

A related phenomenon is "system neglect" (Massey and Wu, 2001), or the human tendency to undervalue the importance of the general decision-making context. One notable example of this type of bounded awareness is the longtime failure of U.S. citizens to treat campaign-finance reform as an important means of curbing the undue influence of special-interest groups on the political process (Bazerman, Baron, and Shonk, 2001). When asked whether they support the issue of campaign-finance reform, most people say "yes." Yet they tend to rank it very low in their list of priorities, relative to other issues. Bazerman et al. (2001) argue that voters undervalue campaign-finance reform because of their bounded awareness of the critical indirect impact of campaign finance reform on other issues important to them. Failing to think through this process, people tend instead to value issues that are more clearly seen as end states or outcomes, such as tax cuts or education. A broader awareness would direct their attention toward a set of outcomes that would have a large, positive effect on many issues (Bazerman et al., 2001).

Finally, bounded awareness can prevent negotiators from considering the impact of their decisions on outsiders. Decision and negotiation scholars have promoted

cooperation in prisoner dilemma games and social dilemmas. In a prisoner dilemma game, two or more parties would be jointly better off if they both cooperated rather than if they both defected (betrayed each other), yet each party would be better off defecting on the other, regardless of the behavior of the other party. The prisoner dilemma problem has been used to understand defection in the nuclear arms race, the failure of strategic alliances, and the overharvesting and overfishing crises. In these realms and others, negotiators make gains at the expense of others outside of the bounds of the problem or crisis.

BOUNDED AWARENESS IN AUCTIONS

Bounded awareness extends beyond individual perception and decision-making and into competitive situations. We return now to the "winner's curse," which Bazerman and Samuelson (1983) describe as a situation such as a two-party negotiation between buyer and seller in which the buyer errs by ignoring the perspective of the seller. In auctions, the winner's curse occurs when the winning bidder fails to consider that he or she has bid higher than everyone else, despite the fact that all bidders have an information disadvantage relative to the seller. Bazerman and Samuelson (1983) argue that the highest bidder may have significantly overestimated the worth of the item they have just purchased.

Figure 2 illustrates this argument. Curve E depicts the distribution of bidder estimates for the true value of the commodity. Curve B shows the distribution of bids. The depiction assumes that (1) the mean of the distribution equals the true value of the commodity—that no aggregate under- or overestimation is expected; and (2) bidders

discount their estimates a fixed amount when making bids, thus explaining the leftward shift of the estimate distribution. The figure suggests that a winning bid—one from the right tail of the distribution—will probably exceed the actual value of the commodity. The highest bidder is likely also to have been one of the highest estimators; unless they had reason to believe that they had better information than other bidders, overpayment is likely. In fact, we have found in our research (Samuelson and Bazerman, 1985) that the winning bidder in auctions of highly uncertain commodities with a large number of bidders typically pays more than the commodity is worth.

If an individual bidder or group of bidders assumes that their bid will win the auction, this assumption should tell them that they may very well have overestimated the value of the commodity relative to other bidders. Based on this reasoning, bidders on highly uncertain commodities who are competing against a large number of bidders should lower their estimates of the true value of the commodity as well as their bids. Thus, if they do end up winning, they are less likely to have overbid, or at least not by a smaller margin. But most people ignore the effects of uncertainty, even falsely viewing the presence of many bidders as a sign that they should be confident of the commodity's value and quality.

Insert Figure 2

We propose that bounded awareness is the psychological process that explains the information left out of the decision-making process, leading to the winner's curse. In the 1980s, Wall Street saw a blitz of acquisition activity, in which companies often competed destructively against each other to win a takeover battle and paid too much for what they

"won." Not surprisingly, two-thirds of acquisitions turn out to be failures or fall short of expectations.

CONCLUSION

Bounded awareness could viably be included under a broad definition of bounded rationality, one of the most powerful concepts to emerge in the social sciences. But we believe that bounded awareness research emphasizes a direction quite different from recent operationalizations of bounded rationality. In our view, most work that is positioned as part of modern-day bounded rationality focuses on imperfections in the integration of data in the decision-making process. As a result of this perspective, other aspects of decision perfection are missed.

In work parallel to our own, Thaler (2000) suggested that there are two additional ways in which decision making is limited, which he describes as "bounded willpower" and "bounded self-interest." The former describes the pervasive human tendency to give greater weight to present concerns than to future concerns, leading to inconsistencies between temporary motivations and long-term interests. The latter notes that our self-interest is bounded; unlike the stereotypic economic actor, we care about the outcomes of others. Similarly, Chugh, Bazerman, and Banaji (2005) have introduced the concept of "bounded ethicality" to refer to the limits on our ethics of which we are not even aware. We now add the term bounded awareness. Together, these concepts provide a specific roadmap of how human judgment departs from economic models and common intuition.

We are not advocating an unrealistic state of unboundedness, nor are we denying the potential adaptive benefits of cognitive efficiency. But, we propose that people tend

to overestimate their own awareness and underestimate the bounds on their awareness, demonstrating a stubborn resistance to the very existence of bounded awareness. Subsequently, they disbelieve that they have overlooked information crucial to making a successful decision. The failure to recognize these normal psychological limitations poses an even greater danger than the limitation itself. That is, the risks of the meta-error may exceed those of the error itself.

In fact, we suspect that people tend to over-search for information in lower priority contexts and under-search in higher priority contexts. The costs of bounded awareness are greatest, then, in the contexts where the decisions are of the highest priority. In lower priority situations, the adaptive benefits of bounded awareness may, in fact, outweigh the costs.

Our goal in this article has been to specify the nature of the ways in which our awareness is bounded. We are sure that there are many other bounds to human awareness, and we look forward to developments of this understanding. But if we have convinced you that readily available and important information exists that you and others miss on a systematic and predictable basis, this article has served its purpose.

FIGURE 1

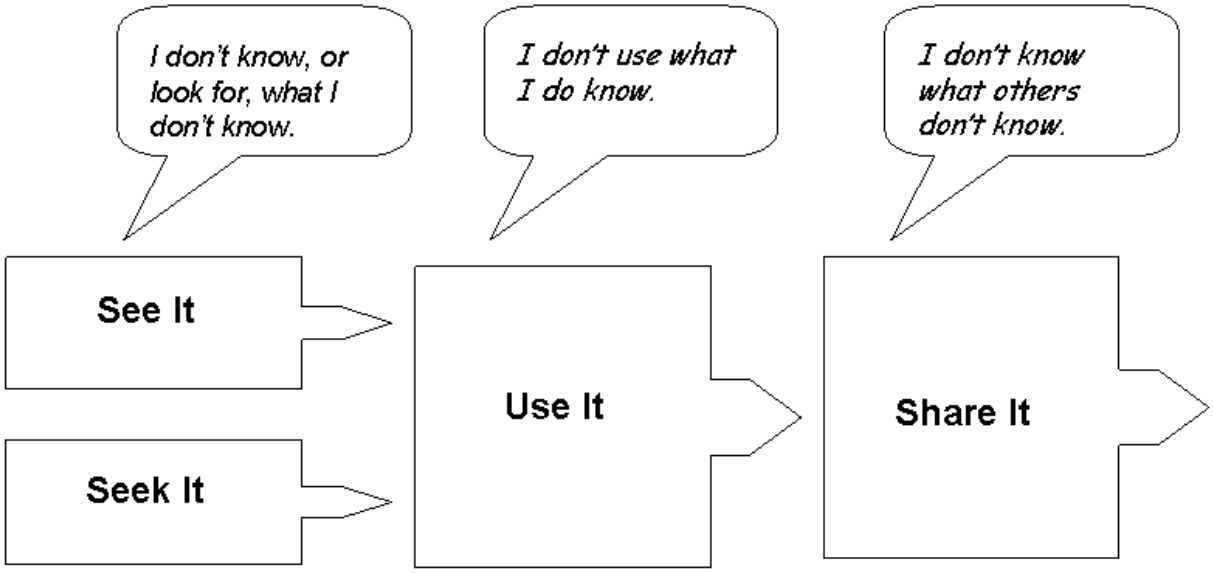
I don't know, or look for, what I don't know.
I don't use what I do know.
I don't know what others don't know.
See It
Seek It
Use It
Share It

FIGURE 2

Caption: Winner's Curse in Auctions

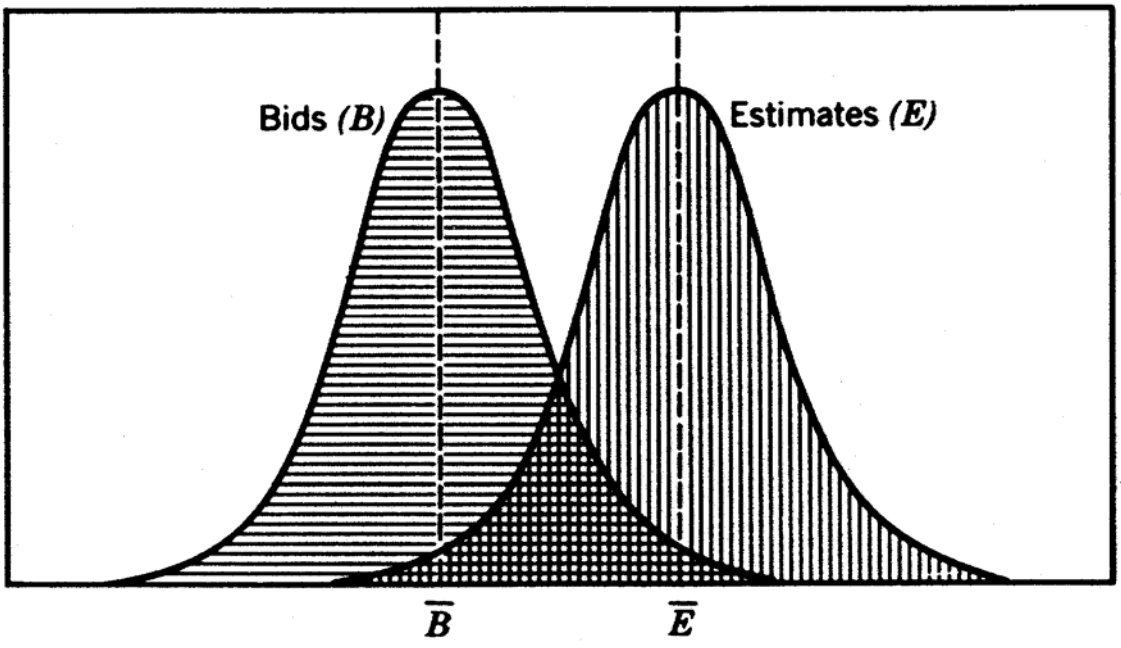

Variables	Assumptions
E = Estimates B = Bid D = (Amount of discounting) $= E - B$	1. True value $\approx \overline{E}$ 2. True value will be equal for all bidders

REFERENCES

Abrahams, M. (1999). What is this Ig? http://www.improbable.com/ig/what-is-this-ig.html (October 1, 2004).

Akerlof, G. (1970). The market for lemons. *Quarterly Journal of Economics, 89,* 488–500.

Angelone, B.L., Levin, D.T., and Simons, D.J. (2003). The relationship between change detection and recognition of centrally attended objects in motion pictures. *Perception, 32*(8), 947-962.

Ball, S. B., Bazerman, M. H., and Carroll, J. S. (1991). An evaluation of learning in the bilateral winner's curse. *Organizational Behavior and Human Decision Processes* 48, 1–22.

Bazerman, M.H., Baron, J., and Shonk, K. (2001). *You Can't Enlarge the Pie: Six Barriers to Effective Government*, Basic Books.

Bazerman, M.H. and Chugh, D. (2005). Bounded Awareness: Focusing Failures in Negotiation. Thompson, L. (ed.), Frontiers of Social Psychology: Negotiations. College Park, MD: Psychology Press.

Bazerman, M. H., and Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution* 27, 618–634.

Cain, D.M., Loewenstein, G. and Moore, D.A.(2005). The Dirt on Coming Clean: Perverse Effects of Disclosing Conflicts of Interest. *Journal of Legal Studies.*

Camerer, C., & Lovallo, D. (1999). Overconfidence and Excess Entry: An Experimental Approach. *The American Economic Review*, 306-318.

Carroll, J. S., Bazerman, M. H. and Maury, R. (1988). Negotiator cognitions: A descriptive approach to negotiators' understanding of their opponents. *Organizational Behavior & Human Decision Processes 41*(3), 352-370.

Chugh, D, Bazerman, M.H., & Banaji, M.R. (2005) Bounded Ethicality as a Psychological Barrier to Recognizing Conflicts of Interest. In D.A. Moore, D. M., Cain, G. Loewenstein, & Bazerman, M.H. (Eds.). *Conflicts of Interest: Problems and Solutions from Law, Medicine and Organizational Settings*. London: Cambridge University Press.

Fox, C. R. and A. Tversky (1998). A Belief-Based Account of Decision under Uncertainty. *Management Science 44*(7).

Friedman, D. (1998). Monty Hall's three doors: Construction and deconstruction of a choice anomaly. *American Economic Review, 88* (4), 933-946.

Gilbert, D. T. and Wilson, T. D. (2000). Miswanting: Some problems in the forecasting of future affective states. In J.P. Forgas (Ed) *Feeling and thinking: The role of affect in social cognition* Studies in emotion and social interaction, second series (pp 178-197) xvi, 421pp.

Gino, F., & Bazerman, M.H. Slippery slopes and misconduct: The effect of gradual degradation on the failure to notice unethical behavior, Harvard Business School working paper.

Gruenfeld, D., Mannix, E. A., Williams, K. Y., and Neale, M. A. (1996) Group composition and decision making: How member familiarity and information distribution affect process and performance, *Organizational Behavior and Human Decision Processes*, *67*, 1-15.

Johansson, P., Hall, L., and Olsson, A. (2004). *From Change Blindness to Choice Blindness*. Toward a Science of Consciousness Conference, Tucson, AZ.

Mack, A. (2003). Inattentional Blindness: Looking Without Seeing. *Current Directions in Psychological Science, 12*(5), 180-184.

Mack, A. and Rock, I. (1998). *Inattentional blindness,* Bradford Books.

Massey, C., & Wu, G. (2001). Detecting regime shifts: a study of over-and under-reaction. Unpublished manuscript, Chicago, IL.

Messick, D. M., Moore, D. A. and Bazerman, M. H. (1997). Ultimatum bargaining with a group: Underestimating the importance of the decision rule. *Organizational Behavior & Human Decision Processes, 69*(2), 87-101.

Mitroff, S. R., Simons, D. J., & Franconeri, S. L. (2002). The Siren Song of implicit change detection. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 798-815.

Moore, D. (2000). The unexpected benefits of negotiating under time pressure. Unpublished doctoral dissertation, Northwestern University, Evanston, IL, USA.

Moore, C. M. and Egeth, H. (1997). Perception without attention: Evidence of grouping under conditions of inattention. *Journal of Experimental Psychology: Human Perception & Performance*, *23*(2), 339-352.

Moore, D., Cain, D., Loewenstein, G., & Bazerman, M.H. *Conflicts of Interest: Problems and Solutions from Law, Medicine, and Organizational Settings*. London: Cambridge University Press, 2005.

Moore, D. A. and Kim, T. G. (2003). Myopic Social Prediction and the Solo Comparison Effect. *Journal of Personality & Social Psychology, 85*(6), 1121-1135.

Moore, D. A., Kurtzberg, T. R., Fox, C. R., and Bazerman, M. H. (1999). Positive illusions and biases of prediction in mutual fund investment decisions. *Organizational Behavior and Human Decision Processes, 79*, 95-114.

Moore, D. A., Loewenstein, G., Tanlu, L., & Bazerman, M. H. (2004).Psychological Dimensions of Holding Conflicting Roles: Is it possible to play advocate and judge at the same time? *Tepper Working Paper 2004-E40*, Pittsburgh, PA.

Moore, D. A., & Small, D. A. (2004). Error and bias in comparative social judgment: On being both better and worse than we think we are. *Tepper Working Paper 2004-E1*, Pittsburgh, PA.

Neisser, U. (1979). The concept of intelligence. *Intelligence, 3*(3), 217-227.

Samuelson, W.F., & Bazerman, M.H. (1985). Negotiating Under the Winner's Curse. In V. Smith (Ed.), Research in Experimental Economics, JAI Press, Volume III.

Schkade, D. A. and D. Kahneman (1998). Does Living in California Make People Happy? A Focusing Illusion in Judgments of Life Satisfaction. *Psychological Science, 9*(5), 340-346.

Selvin, S. (1975). [Letter to the editor]. *American Statistician 29,* 67.

Simon, H.A. (1983). *Reason in Human Affairs.* Stanford: Stanford University Press.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition, 7*(1-3), 1-15.

Simons, D. J. and Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, *28*(9), 1059-1074.

Simons, D.J., Chabris, C.F., Schnur, T., and Levin, D.T. (2002). Evidence for preserved representations in change blindness. *Consciousness & Cognition: An International Journal*, *11*(1), 78-97.

Simons, D. J. and Levin, D. (2003). What makes change blindness interesting? In D. E. Irwin and B. H. Ross (Ed.) *The psychology of learning and motivation*. San Diego, CA, Academic Press.

Stasser, G. and Stewart, D. (1992). Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of Personality & Social Psychology*, *63*(3), 426-434.

Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality & Social Psychology*, *48*, 1467-1478.

Tenbrunsel, A.E. and Messick, D.M., Ethical fading: (2004). The role of self deception in unethical behavior. *Social Justice Research, 17* (2), 223-236.

Thaler, R. (2000). From homo economicus to homo sapiens. *Journal of Economic Perspectives, 14*, 133-141.

Tor, A. and Bazerman, M. H. (2003). Focusing failures in competitive environments: Explaining decision errors in the Monty Hall game, the acquiring a company

problem, and multiparty ultimatums. *Journal of Behavioral Decision Making, 16*(5), 353-374.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*, 1124-1130.

Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review, 101*, 547-567.

Valley, K. L., Moag, J. S., and Bazerman, M. H. (1998). A matter of trust: Effects of communication on the efficiency and distribution of outcomes. *Journal of Economic Behavior and Organization, 35.*

Vaughn, D. (1996). *The Challenger Launch Decision : Risky Technology, Culture, and Deviance at NASA.* Chicago, IL: University of Chicago.

vos Savant, M. (1990a). Ask Marilyn. *Parade Magazine.* September 9, 1990. New York, NY.

vos Savant, M. (1990b). Ask Marilyn. *Parade Magazine.* December 2, 1990. New York, NY.

vos Savant, M. (1991). Ask Marilyn. *Parade Magazine.* February 17, 1991. New York, NY.

Wilson, T. D., Wheatley, T., Meyers, J. M., Gilbert, D. T. and Axsom, D. (2000). Focalism: A source of durability bias in affective forecasting. *Journal of Personality & Social Psychology, 78*(5), 821-836.