האוניברסיטה העברית בירושלים

THE HEBREW UNIVERSITY OF JERUSALEM

PRIOR DIVERGENCE: DO RESEARCHERS AND PARTICIPANTS SHARE THE SAME PRIOR PROBABILITY DISTRIBUTIONS?

By

CHRISTINA FANG, SARI CARP,
and **ZUR SHAPIRA**

Discussion Paper # 587 **Aug 2011**

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 35 (2011) 744–762

ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/j.1551-6709.2011.01177.x

Prior Divergence: Do Researchers and Participants Share the Same Prior Probability Distributions?

Christina Fang,[a] Sari Carp,[b] Zur Shapira[a]

[a]*Department of Management and Organization, New York University*
[b]*Jerusalem School of Business, Hebrew University, Israel*

Received 12 October 2009; received in revised form 23 September 2010; accepted 24 November 2010

Abstract

Do participants bring their own priors to an experiment? If so, do they share the same priors as the researchers who design the experiment? In this article, we examine the extent to which self-generated priors conform to experimenters' expectations by explicitly asking participants to indicate their own priors in estimating the probability of a variety of events. We find in Study 1 that despite being instructed to follow a uniform distribution, participants appear to have used their own priors, which deviated from the given instructions. Using subjects' own priors allows us to account better for their responses rather than merely to test the accuracy of their estimates. Implications for the study of judgment and decision making are discussed.

Keywords: Prior probability distributions; Decision making; Experimental design

1. Introduction

In a classic article, Tversky and Kahneman (1983) reported that when naïve respondents were asked to reflect on their failure to use the conjunction rule properly, one of them said: ''I thought you only asked for my opinion'' (p. 300). This quote describes vividly how a participant responds to experimental instructions according to the way she perceives the problem. Furthermore, her perception seems very different from the intent of the experimenter. In other words, participants appear to rely on their own priors or models, which do not necessarily coincide with those taken for granted by researchers designing the experiments.

Correspondence should be sent to Christina Fang, Department of Management and Organization, New York University, New York, NY 10012. E-mail: cfang@stern.nyu.edu

This has important implications for how researchers interpret their experimental results in the behavioral decision-making tradition, where participants' responses are compared to an ''optimal'' response based on some normative model, and the gap between the responses is taken as evidence of a bias. Much less examined is the possibility that participants' priors may differ from those of the researcher (Suppes, 2007). In this article, we explore a possible divergence in priors between those assumed by the researchers and those brought to the experiment by participants. We ask: Do participants share the same priors as the researchers who design the experiment? If not, does this divergence matter?

To answer these questions, we explicitly ask participants to indicate their own priors in estimating the probability of a variety of events. We then compare their choices not only to optimal choices that are prescribed by some normative model (e.g., Griffiths & Tenenbaum, 2006) but also against the most likely response that is derived from their own prior probability distribution. Our experimental settings consist of (a) a version of the ''Acquiring a Company'' (AAC) problem (Samuelson & Bazerman, 1985); and (b) the ''Beauty Contest'' problem (Keynes, 1936). We chose these two problems because they are often used to illustrate judgmental biases. In both problems, there is a well-defined optimal behavior prescribed by normative theories, and deviations from those established norms are taken as evidence of judgmental bias or error. In addition, intelligent action in both problems requires that a subject think about how others will act, and specifically about the distribution of others' beliefs.

We find that in our studies, the participants did not share the same priors as the researchers who designed the experiment. Furthermore, these differences matter. Specifically, in both problems, the participants overbid in comparison with the optimal responses based on the normative models. Yet, when we compare their responses to the most likely response conditional on their prior probability distributions, we find that we can better account for their responses.

Our work has important implications for research on judgment and decision making. Existing research mostly concentrates on the ''computation'' aspect of participants and often assumes that suboptimal responses to these problems arise from participants' assigning improper weights to their priors. We show, however, that these deviations may be partially due to participants' adoption of a set of priors which differs from those assumed by the researcher and are based on normative theories. This is a different source of error—the fact that participants themselves may have a very different idea about the experiment. To the extent that participants respond to experimental instructions according to the way they perceive the problem, our work implies that their priors should be collected, studied, and analyzed.

It is important to note that we argue that we may be able to predict subjects' behavior better if we know they rely on priors derived from their own subjective distributions. We do not claim that such knowledge allows us to explain why they perform suboptimally on such tasks, or that it eliminates suboptimal performance. Indeed, our results replicate prior findings on suboptimal judgment. However, it is clear that people inject subjectivity into their interpretation of the problem, and that this predicts the choices they make. Thus, knowledge about subjects' priors can help us understand how subjects arrive at their judgments. In this article, we are interested in the *descriptive* aspect of behavior, not in how behavior departs from the *normative* perspective.

2. Study 1

Do subjects share the same priors as the researchers who design the experiment? Our setting was a version of Samuelson and Bazerman's (1985) ''AAC'' problem, which asks participants to determine an offer price for a company based on the uncertain outcome (uniformly distributed between \$0 and \$100/share) of its oil exploration project. Since the target has access to the true value, the distribution only represents the acquirer's uncertainty. Due to potential synergies, the target is worth 50% more to the bidder than to the current owner. Informational asymmetries between the target and the acquirer dictate that the rational offer price is zero. To see this, consider a participant who bids any positive amount x between 0 and 100. If the offer is accepted by the seller, the true value must be between 0 and x. Effectively, the fact that an offer is accepted narrows the range of possible values from [0,100] to $[0, x]$. As a result of this *truncation*, assuming that the participant actually does assume a uniform prior (i.e., he or she believes that all the outcomes of the oil exploration are equally likely), the expected value of the firm is $0.5x$ (and $0.75x$ to the bidder). Thus, the bidder is expected to lose an amount equal to $(x - 0.75x)$. The profit-maximizing (or loss-minimizing) and rational bid should be zero (i.e., not to bid at all). Therefore, a subject's ''overbid'' is defined by the entire amount of x (the difference between the actual bid and the optimal bid).

AAC remains one of the most persistent puzzles in the behavioral decision-making literature. In most studies (e.g., Carroll, Bazerman, & Maury, 1988; Ball, Bazerman, & Carroll, 1991; Selten, Abbink, & Cox, 2005), participants (over)bid a sum between \$50 and \$75/share. This behavior is typically attributed to a failure to consider the incentives of other players and/or rules of the game, or to risk taking. Though participants were shown to do better after seeing an unrelated set of puzzles such as the Monty Hall Problem (Idson et al., 2004), overbidding seems to be robust to a variety of decision support features (e.g., Bereby-Meyer & Grosskopf, 2008; Grosskopf, Bereby-Meyer, & Bazerman, 2007).

A key aspect of the problem, however, is that the optimal offer of zero is conditional on respondents' assuming a uniform distribution of potential values for the target. Subjects are explicitly instructed that ''all values between 0 and 100 are *equally* likely.'' Even without explicit instructions, respondents should assume uniformity in priors because in the AAC context, uniformity represents ''ignorance priors'' (cf. Fox & Rottenstreich, 2003). According to the *principle of insufficient reason* advanced by Leibniz and Laplace, ''if we see no reason why one case should happen more than the other, then these cases should be treated as equally likely'' (cf. Laplace, 1776, Leibniz, 1678, quoted in Hacking, 1975, pp. 132; chap. 14, respectively).

To our knowledge, no studies have explicitly asked participants for their own priors or tested whether these are truly uniformly distributed. Participants are simply assumed to follow this instruction and use it in their decisions. We set out to elicit participants' prior probability distributions directly and to examine how these distributions are used in determining an offer.

2.1. Participants and procedures

Seventy-nine undergraduate business students from New York University responded to one of the two versions of the AAC scenario (we kept substantive details the same as Samuelson and Bazerman's, while condensing and simplifying the wording). In the ''*informed*'' version, we retained and underlined the phrase (''All values within the range are equally likely'') employed in previous studies to instruct respondents to assume a uniform prior distribution with respect to the outcome of the oil exploration. In the ''*uninformed*'' version, we eliminated this sentence, providing no instruction as to the nature of the prior probability distribution.

We asked participants to indicate their priors about the target's values by assigning probabilities to values between $0 and $100, in $10 increments (see Appendix A). We then asked them to decide how much, if anything, to offer for the target. Finally, we added a second stage, in which participants bidding over $50 were told that their offers were accepted and the rest told that their bids were rejected. We then asked them to come up with probability distributions and offers for a second, similar target company.

2.2. Results

We first classify subjects' prior distribution into four categories: (a) Uniform: if the subject's probability assessment is 10% for each of the 10 intervals; (b) left skewed: if the sum of the subject's probability assessment for intervals *below* $50 is *less* than the sum of his or her probability assessment for intervals *above* $50; (c) right skewed: if the sum of the subject's probability assessment for intervals *below* $50 is *more* than the sum of his or her probability assessment for intervals *above* $50; and finally (d) symmetric: if the subject's probability assessment is symmetric around the intervals $40–$50 and/or $50–$60 (see Appendix B for details).

Table 1 exhibits the frequencies of different priors according to the two experimental conditions: whether subjects are told explicitly to assume uniformity or not.

As seen from Table 1, contrary to experimenters' expectation, only about 41% of the participants assume uniform priors. The remaining 59% do not. Even when respondents were

Table 1
Classification of participants' prior distributions

	Left Skewed	Symmetric	Right Skewed	Uniform
Not told	10	7	12	16
Told	5	7	6	16
Total	15	14	18	32
Percentage	19%	18%	23%	41%

explicitly instructed to assume uniformity, only 47% actually did so. Apparently, they either ignored or did not understand the simple instruction that all values were equally likely. Without explicit instruction (i.e., in the ''uninformed'' version), the proportion of respondents assuming uniformity fell still further, to 36%. Even in an ''ignorance prior'' situation, respondents did not naturally assume a uniform distribution, and their priors differed considerably from the instructions. This implies that as experimenters, we can no longer assume that uniformity is the implicit prior, which *all* subjects actually use in their decision-making process. Other, nonuniform, priors clearly play a role. Furthermore, we find that that there is no significant relationship between priors and instruction (i.e., whether the subjects were told or not told that they should have assumed uniformity), χ^2 (3, n = 79) = 2.1772, p = .54.

Consistent with prior studies, respondents overbid (average bid was $56.87). However, can we explain the extent of overbidding by subjects' actual prior probability distributions? Our conceptual argument is that part of the subjects' apparently irrational behavior may be subjectively rational, that is, consistent with their subjective priors. Thus, what appears to be irrational may in fact be ''rational'' in the sense that the behavior can be explained by priors differing from those assumed by normative theories and experimenters. Implicit in our argument is a shift in the underlying benchmark of ''rationality''—from an objective, normative definition to one which is subjective and individualistic. To answer this question, we compute the optimal bid implied by a participant's prior probability distribution. If this implied optimal is close to the actual offer, we may infer that participants' choices are at least consistent with their own priors, though they may still be overbidding according to a normative model. According to the normative theory, any positive bid above zero is not rational. Thus, the amount of (objective) overbidding, as compared to the rational bid of zero, is comprised of the entire positive bid. However, from our perspective, attributing any positive bid to irrationality ignores the possibility that subjects may have come to the problem with different priors from uniformity. For instance, a subject assumes a left-skewed prior, then his or her higher bid is consistent with his or her beliefs about B's value (which is believed to be high). We therefore conjecture that at least part of the positive bid may be attributed to subjects' priors diverging from those assumed by the experimenters. Because priors are by definition subjective, there is no right or wrong prior.

Specifically, for each subject, we empirically examine the degree of (subjective) overbidding, which is the difference between the actual bid and an implicit, subjectively optimal bid *conditional* on the subject's prior. Instead of using an objective benchmark in the form of the normative, rational bid of zero, we compare the actual bid with a subjective measure/benchmark for rationality. In other words, instead of computing overbidding with respect to a normative level of zero, we now compute an implied rational bid, *conditional* on each subject's priors.

To measure the degree to which participants' offers were consistent with their actual priors, we calculated the implied expected value of the target firm. We took each participant's actual prior probability distribution, and then truncated the distribution by the offer price the participant made. Since one should not expect the value of the target to be higher than the offer price, it is reasonable to infer that the true value of the target should lie below a participant's offer. We then calculated the implied optimal bid (conditional on each participant's priors) by taking the average (expected) value of the target given the truncated prior distribution for each participant, and multiplying this by 1.5. If this implied optimal bid was less than the actual offer, participants clearly overbid and their choices were not consistent with their own prior beliefs.

Here is an example. Suppose that a subject's prior is right skewed and has the following probabilities: .3, .3, .1, .1, .1, .1, 0, 0, 0, 0 over the intervals 0–10, 10–20, 20–30, 30–40, 40–50, 50–60, 60–70, 70–80, 80–90, 90–100, respectively. The subject's actual offer is $40. This offer indicates that the subject believes that the actual value of the target is from [$0, $40], that is, to the left of the cutoff. We then calculate the expected value of the target given this truncated prior distribution = [0.3*5 + .3*15 + .1*25 + .1*35)]/(.3 + .3 + .1 + .1) = $12/.8 = $15. Since the target's value is worth 50% more, the subjective optimal value is $15*1.5 = $22.50, resulting in subjective overbidding of $40–$22.50 = $17.50.

Fig. 1 plots subjects' actual bids, which we now decompose into two components: (a) a first part due to incorrect priors; and (b) a second portion attributable to incorrect computation or insufficient understanding of probability theory. The first component of the bid is the difference between the implicit rational bid and zero. This difference arises from priors, specifically from subjects not believing the kind of uniform priors experimenters assume. The second component, defined as the difference between a subject's own expected value of B and his or her bid, is the only component attributed to irrationality, per se. If a subject's implicit optimal bid is *X* (inferred from his or her priors), and he or she bids *Y*, then the difference between *X* and *Y* is the amount of ''subjective'' overbidding. This can result from many failures of rationality, including failure to account for others' motives as a result of asymmetric information (i.e., game-theoretic reasoning) and inability to compute expected

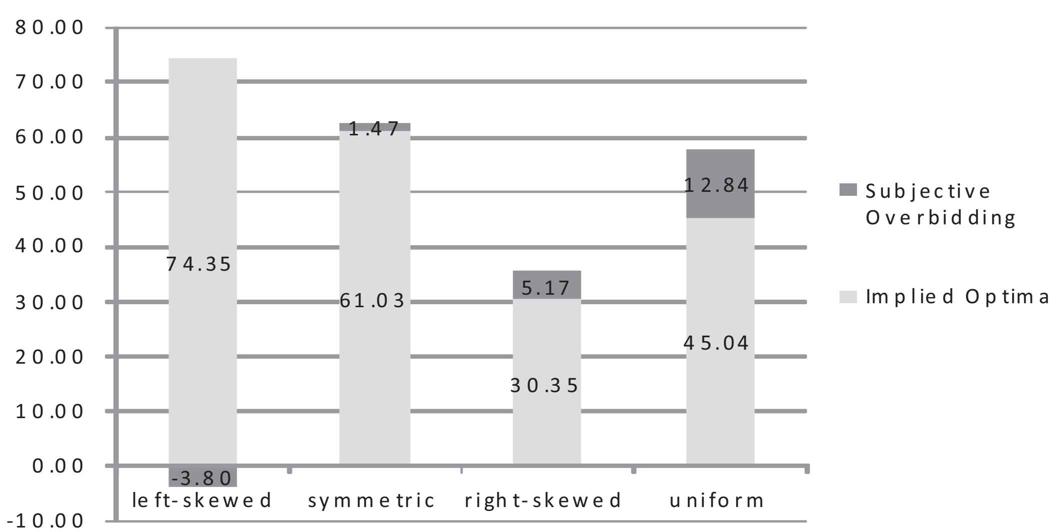

Fig. 1. (1) Average implied value of the target; (2) actual bids; and (3) the magnitude of overbidding (actual offer—implied value).

value or to apply basic principles of probability theory. These issues have been examined extensively in prior literature (e.g., Bereby-Meyer & Grosskopf, 2008; Grosskopf et al., 2007). The overbidding resulting from these failures represents the residue portion of the actual bid that cannot be explained by the subject's (potentially incorrect) priors. We consider this the true extent of ''irrationality.''

As shown in Fig. 1, there is a significant effect of priors in (a) the chosen numbers ($F(3, 75) = 10.72$, $p = .00$); (b) subjective optimal bids ($F(3, 75) = 22.60$, $p = .00$); and (c) overbidding ($F(3, 75) = 9.25$, $p = .00$). First, both implied optimal bids and actual bids vary sensibly as a function of the particular prior distribution the participant assumed. Specifically, those who have left-skewed priors (i.e., their priors are centered on high values) tend to extend significantly higher bids than the rest, while those with right-skewed priors have implicit rational bids that are predictably and significantly lower than the rest of the participants.

Second, participants assuming a prior distribution skewed toward higher values ended up with the smallest overbidding. They believed, whether reasonably or not, that the outcome of the oil exploration was likely to be highly successful. Thus, to them, some positive offer price was justified. Those who assumed a left-skewed distribution of oil exploration outcomes (thus believing that the outcomes are most likely concentrated on higher values) exhibited an average implied optimal bid of $74.35, which is very high compared to other participants. The actual bids of this group were slightly lower, averaging $71, which means that they actually ''underbid'' slightly (by $3.80). The magnitudes of overbidding were −$3.80, $1.47, $5.17, and $12.84, respectively, for participants assuming left-skewed (higher value), symmetric (but nonuniform), right-skewed (lower value) and uniform prior distributions.

Overall, subjects' implied optimal bids account for 70% of the variance in their actual offers ($B = 0.79$, $p < .00$) and are significantly and positively correlated: $r(77) = .81$, $p = .00$. These overall patterns demonstrate that divergent priors explain to some extent the pattern of ''irrational'' behavior we observe. Even though subjects may still be overbidding objectively (relative to the optimal bid of zero), their choices are consistent with their own priors.

The average amount of (subjective) overbidding, conditional on each subject's prior, is about $3.92. The average amount of (objective) overbidding, on the other hand, is the entire amount of the actual bid (i.e., $57). Clearly, the extent of irrationality (having accounted for subjects' beliefs) is much less than commonly believed where experimenters presume uniformity. Subjects may simply be coming to the problem with different assumptions and these differences matter as they potentially explain behavior previously labeled as ''irrational.''

How rationally will subjects respond to new information in the second part of the experiment? Of the total 79 subjects, we decided randomly whether their first-round offers were rejected or accepted—34 subjects were told that their offers were rejected, and the remaining 45 told otherwise. They were then asked to generate probability distributions and offers for a second, similar target company.

We find that 19 out of the 34 subjects with rejected offers (56%) changed their priors, while only 16 out of the 45 with accepted priors (26%) did. As seen in Table 2, of the former

Table 2
Frequencies of subjects who changed priors in the second round

	Rejected	Accepted
No. who changed priors	19	16
To the left skewed	12	3
To the right skewed	2	8
To the symmetric	4	4
To the uniform	1	1
Total number of subjects	34	45

group, 63% (12 out of 19) shifted their priors to be left skewed, reflecting a belief that the real value of the target is more concentrated in the higher values. In contrast, of the latter group (i.e., those who changed priors after their offers were told to be accepted), the dominant tendency (i.e., 8 out of 16, 50%) was to shift their priors to be right skewed. This change reflects, again, a sensible shift in beliefs which are concentrated in the lower ranges.

Even when subjects did not indicate a change of priors, their bidding behavior often changed to take into account the new information (i.e., accept or reject). We tabulate below the contrast in (a) actual offers; (b) implied optimal; and (c) overbidding, between those with rejected and those with accepted first offers. As can be seen in the average numbers and standard deviations in Table 3, there is a significant effect.

First, for those with rejected first offers, their second actual bid is higher ($t(33) = -5.73$, $p = .00$, two-tailed) and their second implied optimal is higher ($t(33) = -5.28$, $p = .00$, two-tailed). Subjects responded to rejected initial offers by increasing their actual offers, which were backed up by their actual beliefs, as shown in higher implied, subjective optimal numbers. We observe not only that subjects changed their actual bids sensibly but also revised their underlying priors (as reflected in their implicit, subjective rational bids conditional on their respective priors) in the expected direction. However, there was no difference between the degree of overbidding in the first and second rounds ($t(33) = 0.98$, $p = .33$, two-tailed). Recall that the degree of overbidding can be considered a proxy of the degree of subjective irrationality. Consistent with our overall argument, there is no reason to expect that subjects become more or less irrational during the course of the game, and the lack of significant change in the amount of overbidding confirms it.

Table 3
Tabulation of offers across the two conditions—those whose first offers were either accepted or rejected

	Rejected	Accepted
1st actual offer	55.63 (SD = 3.71)	56.89 (SD = 3.10)
2nd actual offer	70.82 (SD = 4.61)	45.88 (SD = 3.25)
1st implied optimal	47.70 (SD = 3.83)	47.69 (SD = 3.83)
2nd implied optimal	65.05 (SD = 4.92)	65.05 (SD = 4.92)
1st overbidding	7.93 (SD = 2.16)	4.99 (SD = 1.82)
2nd overbidding	5.77 (SD = 2.67)	3.15 (SD = 2.29)

Subjects with accepted first offers extended lower actual bids in the second round ($t(44) = 4.04$, $p = .00$, two-tailed), while their second round optimal bids were higher ($t(44) = 3.25$, $p = .00$, two-tailed). They responded to the accepted offers by decreasing their second offers, which was consistent with their actual beliefs as reflected in lower implied, subjective optimal numbers. As expected, there was again no difference between the degree of overbidding in the first and second rounds ($t(44) = 1.04$, $p = .31$, two-tailed).

These additional analyses provide rather conclusive evidence that subjects responded rationally to new information, updating their prior beliefs accordingly for the next bid.

2.3. Discussion

In tackling a generic probability-based problem such as the AAC, it is helpful to consider varying degrees of rational behavior exhibited by human subjects. First, subjects may have different priors from those believed by researchers to be normative or appropriate. Second, even assuming that a subject does adopt the normative prior, she may or may not incorporate the inherent information asymmetry. Of course, a rational subject should realize that, if her offer is accepted by the target who knows the true value of the firm (a situation of asymmetric information in game-theoretic terms), then the true value must lie below the offer. However, subjects may not fully comprehend the nature of asymmetric information and thus fail to account correctly for the incentives of the bargaining partner. Third, even if we assume that the first and second criteria of rationality are satisfied, subjects may or may not compute expected value correctly, given a prior probability distribution—they may simply not know how, may multiply incorrectly, etc.

We consequently define these three—roughly nested and not mutually exclusive—sources of ''errors'': (a) incorrect priors; (b) failure to truncate as a result of a lack of game-theoretical knowledge; and (c) lack of computational knowledge (e.g., probability theory). The conventional wisdom in research on behavioral decision making (e.g., Bereby-Meyer & Grosskopf, 2008; Grosskopf et al., 2007) is to attribute any bias primarily to a failure to consider the incentives of other players and/or rules of the game, and secondarily to gambling/risk taking. In other words, existing research focuses on the latter two sources of error. Any deviation from expected, normative behavior (e.g., a bid of zero) is attributed to the failure of subjects to truncate their priors (error 2) and/or the inability to calculate expected value for a given prior distribution (error 3). Our results also lend support to these views. Even subjects who share the correct probability distribution (i.e., uniformity) for the task still fail to give the normative response. Among the 32 subjects who assumed uniformity in our study, the average bid was \$57.88. The average implied optimal bid for these subjects was \$45.04, leaving the amount of subjective overbidding as \$12.84. The behavior of these subjects can be attributed to failures detailed in the existing literature (e.g., errors 2 and 3).

Our work ventures beyond the existing research by exploring the existence of the first error (incorrect priors) and its potential efficacy as an explanation for apparently suboptimal behavior. Our first question is a descriptive one: Do people actually have divergent priors? We find that a substantial portion of them do indeed hold priors which differ substantially from those assumed by researchers. Secondly, we find that this fact contributes considerably to explaining subjects' behavior in the AAC game.

3. Study 2

A natural question arises: Can participants' priors help explain their experimental responses even when the normative response is not dependent on any predefined priors? To examine this issue, we employed the ''Beauty Contest'' problem, which is frequently used to study the power and limits of iterative reasoning. In this game, first proposed by Keynes (1936), participants pick a number between 0 and 100, and the winning number is that closest to *2⁄3 of the average* of all the numbers chosen. Since the payoff depends on one's beliefs about others' beliefs, the game provides a convenient setting in which to study priors.

The optimal bid for this problem is zero, a response that is independent of the kind of priors participants may assume. To see this, consider a rational player. She should never choose a number above 67 ($100 \times 2/3$). This is because even if one assumes that everybody else bids the maximum of 100, one should not bid more than 2/3 of 100, according to the rules of the game. If she believes that others are also rational, she will not pick a number above $100 \times 2/3 \times 2/3$; and if she believes that the others are rational and that they too believe all players are rational, she will not pick a number above $100 \times 2/3 \times 2/3 \times 2/3$ and so on, until all positive numbers are eliminated. However, studies find that although players do learn as the game is repeated, the optimal response of zero is almost never reached (Nagel, 1995; Ho & Weigelt, 1996; Thaler, 1988). If iterative thinking does not entirely explain subject behavior, though, what else can explain it? So far, no studies have directly investigated players' beliefs about probability distributions. Leaving aside the issue of whether subjects behave optimally or not, we show that an understanding of participants' priors may better illuminate the choices they make.

3.1. Participants and procedures

Seventy-seven undergraduate business students from New York University participated in this study for course credit. They were given the rules of the game, then asked to (a) choose a number; and (b) illustrate graphically their beliefs (on a grid marked in deciles 0–10,…, 90–100) about the distribution of numbers chosen by a class of 100 participants including themselves (see Appendix B). Thirty-nine participants were asked to estimate the

Table 4
Frequencies of different priors, across the two experimental conditions

	Left Skewed	Symmetric	Right Skewed	Uniform
Before	3 (7%)	3 (7%)	31 (74%)	5 (12%)
After	4 (11%)	4 (11%)	27 (77%)	0 (0%)
Total	7 (9%)	7 (9%)	58 (75%)	5 (6%)

distribution *before* choosing a number, and the other 38 *after* choosing the number. We did not permit those in the ''after'' group to change the number they chose before their priors were elicited.

3.2. Results and discussion

First, as seen in Table 4, we find that regardless of whether we elicited probability distributions before or after the subjects chose their numbers, the percentages of those assuming different priors did not vary, χ^2 (3, n = 77) = 4.9663, p = .17. Across the two versions, a majority (75%) of subjects overall assumed a positively skewed prior distribution.

The average number selected was 33.48 (SD = 21.11). Consequently, the winning number was 22, the integer closest to 22.32. Consistent with previous studies, it is clear that the average participant selected a number much higher than the optimal choice, zero. Can we partially explain this behavior using self-generated priors? To determine whether our results were consistent with the participants' priors, we calculated each participant's *implicit* winning number given her priors about the overall distribution of the numbers picked. To generate this ''implicit'' figure conditioned on her priors, we computed the expected value of her prior distribution (without truncation), then took two-thirds of this value. If a participant was coherent in her calculation, this implicit number should be very close to the number she chose.

First, as can be seen from by the frequencies cross-tabulated in Table 5, we find no significant effect of version in (a) the chosen numbers (F(1, 75) = 1.89, p = .17); (b) subjective optimal (F(1, 75) = 0.77, p = .38); or (c) overbidding (F(1, 75) = 1.49, p = .23). Respondents who specified a probability distribution *before* choosing a number and those doing so afterward do not appear to behave differently. This implies that even if they are not directly asked to define their priors (i.e., in the ''after'' version), participants may already have some probability distribution in mind when deciding upon a strategy.

Second, we find a significant effect of prior distributions. As seen in Table 6, subjects professing different priors have significantly different (a) chosen numbers (F(3, 73) = 5.43, p = .00); (b) subjective optimal numbers (F(3, 73) = 8.49, p = .00); and (c) amounts of overbidding (F(3, 73) = 3.30, p = .03).

In addition, we compute the correlation between the actual number and the subjective optimal for each type of prior. This correlation is significantly positive for right-skewed priors: r(56) = .60, p = .00. For both left-skewed and symmetric priors, this correlation is not significant, r(5) = .67, p = .09 and r(5) = .04, p = .94. For those with uniform priors,

Table 5
Bidding behavior of subjects, across the two experimental conditions

	Chosen Number	Average Subjective Optimal	Average Overbid
Before	30.48	27.34	3.13
After	37.08	29.12	7.95

Table 6
Average chosen numbers, subjective optimal, overbidding figures as well as the correlation between the chosen number and the subjective optimal, for each type of priors

	Left Skewed	Symmetric	Right Skewed	Uniform
(1) Chosen No.	61.43	30.29	30.25	36.20
(2) Subjective optimal No.	38.89	34.09	25.69	33.33
(3) Overbid	22.54	−3.81	4.56	2.87
Correlation (1) and (2)	0.67 (.09)	0.04 (.94)	0.60 (.00)	n/a

Note. p values are inside parentheses.

the correlation cannot be computed since the subjective optimal numbers do not vary (i.e., fixed at 33.33 = 2/3*50).

The average of the ''subjective optimal'' numbers was 28.15 (SD = 8.85). Respondents seemed to utilize their priors coherently, as the mean difference between this implicit winning number and the actual winning number was small (5.32), though the two still are different statistically (Mann–Whitney *U* Test: $z(76) = 3.75$, $p < .001$). This offers some preliminary evidence that even in a problem where optimal response does not depend on a prior distribution, a participant's prior helps to explain observed results.

We also analyze in detail subjects' responses to our post-experimental Q&A and find even more telling evidence of their reliance on priors. Recall that we had gathered respondents' verbal descriptions of their thought processes at the end of the experiment. Most subjects appear to have followed a two-step heuristic, first deciding on an ''anchor,'' which was either an expected or average number (e.g., 50), and then adjusting based on what others might have believed (e.g., multiplying the anchor by 2/3 in multiple iterations). Here, are several excerpts from our experiment to illustrate this staged strategy:

> I assumed the average would be around 50 and that 2/3 of that is 33. Therefore, I figured most people would choose to pick 33. I pick 2/3 of 33, which is 24.

> Theoretically, the average of all numbers would be 50. 2/3 of 50 is 33, so I assumed most would put 33. If this is the case, by writing 22, I would be meaning 2/3 of this average. Obviously if people continued this line of thinking, eventually, all numbers given would theoretically be negligible from 0.

> The class seemed large enough that it would have a normal distribution and the average number between 1 and 100 is 50. However, students would know this and most likely pick 2/3 of 50 putting the average at 33. Being at 2/3 of 33 yields my guess of 22.

> I figured that most people were going to choose the numbers around the middle so the average * 2/3 would be around 30.

From these accounts of their thought processes, it is clear that subjects first fixated on an average number (based on their priors), and then chose various degrees of iterations. As seen

Table 7
The frequencies of different thought processes

	Number of Iterations					
	0	1	2	3 or More	Total	Percentage
Anchor = 50	12	15	4	5	36	47
Anchor = 100	1	1	2	9	13	17
Specific	12	4			16	20
Random	12				12	16

in Table 7, about 47% of subjects started with an anchor of 50, whereas another 17% used 100 as a starting point. They then carried out rounds of iterations to arrive at their chosen numbers. For example, 15 subjects (41% of those with 50 as the anchor) documented their own thought process as involving one iteration whereas 9 subjects (70% of those with 100 as the anchor) carried out more than three iterations.

It is very interesting that almost 64% of the subjects anchored on 50 or 100 as the starting point. For some reason, 50 being the average and 100 being the maximum carried a lot of weight in initially shaping their responses. Only when this ''anchor'' was set did they start to take into account ''what others might do'' by iterating through various rounds of rationality.

About 20% of the subjects, however, followed some specific, idiosyncratic heuristic:

> Most people will choose their favorite number, statistically most likely to be 3 or 5, or their birthday. Thus, the majority of numbers will be below 31.

> My number is 99. It is the last number with 2 letters. I thought that if I need to choose one number between 1 and 100, I want to choose the number which is located at the last with 2 letters.

> My chosen number is my favorite number b/c of my birthday 07-19-1987.

The remaining 16% simply guessed or followed a nonspecific heuristic:

> I chose randomly and liked that number.

> I honestly guessed.

The evidence from participants' responses further confirms our belief that even in a problem where optimal response is not dependent on a prior distribution, knowledge of a participant's prior can shed significant light on observed empirical results.

One remaining question is whether subjects would respond to feedback and learn to update their priors to ones that are more sensible. To explore this question, we ran an iterated version of the Beauty Contest experiment, asking subjects to play the game three consecutive times, and telling them the average and winning bids after each round. After deleting subjects with incomplete rounds of data, we have 44 new subjects who completed

Table 8
Round by round data for an iterative version of the Beauty Contest game

	Round 1	Round 2	Round 3
Left skewed	4	1	
Symmetric	4		
Right skewed	35	43	44
Uniform	1		
Total	44	44	44
Chosen number	26.51	16.20	6.75
Subjective optimal	23.14	13.86	6.93
Overbidding	3.37	2.34	−0.18

all three rounds of Beauty Contest. From the upper panel of Table 8, we first observe that consistent with our original experiment, about 15% of the subjects (9 out of 44) revealed priors other than right skewed. After receiving feedback and having the opportunity to learn from their experience, the frequencies observed for different priors converge toward right skewed. At Round 3, all subjects now have priors that are right skewed. We also computed the subjectively optimal bid as well as the amount of overbidding. A one-way within subject (or repeated measures) ANOVA was conducted to compare the effect of round on the actual bid, implied optimal bid and overbidding. We find that there is a significant effect of round for the actual bids, Wilks' Lambda = 0.417, $F(2, 42) = 29.30$, $p = .00$. Similarly, the effect of round is also significant for the implied optimal bids, Wilks' Lambda = 0.242, $F(2, 42) = 65.84$, $p = .00$. As expected, the amount of overbidding does not differ across rounds, Wilks' Lamdba = .897, $F(2, 42) = 2.417$, $p = .101$.

4. General discussion

As we demonstrated in these two problems, the observed biases (or deviations from normative behavior) can be due to the adoption of a different set of priors. Participants' priors differ substantially from those assumed by experimenters, and these differences are shown to matter. In both problems, when we take into account the priors participants bring to the experiments, we find that the extent of bias is much reduced and we can explain subjects' behavior much better. Note, however, that we do not argue that participants' performance with regard to the goal of a particular study necessarily improves when their priors are employed. Indeed, we replicated earlier results.

Our article complements a growing body of work showing that Bayesian models are becoming more prominent in various cognitive sciences (Griffiths, Kemp, & Tenenbaum, 2008). Principles of Bayesian probabilistic inference have been applied to diverse research programs ranging from visual scene perception (Yuille & Kersten, 2006), language processing and acquisition (Chater & Manning, 2006), to symbolic reasoning (Oaksford & Chater, 2001). This recent body of literature has mostly focused on the degree to which probability updating follows a Bayesian rule (for a notable exception, see Griffiths & Tenenbaum, 2009). We explore here an alternative, and much less examined, possible source of bias: incorrect priors.

The reliance on one's priors is found also in the work of McKenzie and his colleagues, who argue that participants bring ''extra-experimental knowledge'' to the task at hand in the form of multiple, competing hypotheses. They show, for instance, that participants' priors regarding the rarity of events can help explain some observed biases (McKenzie, 2004; McKenzie & Amin, 2002; McKenzie & Mikkelsen, 2000, 2007).

Our work has important implications on research on judgment and decision making. First, our results suggest that the existence and characteristics of priors can represent an important element in evaluating judgment and choice. Just as anchor accessibility results in insufficient adjustment (Epley & Gilovich, 2006), participants' priors may be a cause of judgmental errors. Therefore, an awareness of the actual distributions participants hold may enable us to understand better the roots of any judgmental biases or errors.

Second, given the fundamental role of priors, soliciting participants' priors may provide further insight into the way they process information. In some cases, assuming ignorance priors (Fox & Rottenstreich, 2003) is appropriate, but in others, it may not be. In such cases, experimenters may need to elicit participants' priors, rather than assume participants use the information provided. While manipulation checks are administered post hoc to insure correct inference about causality in most studies, these cannot guarantee that participants operate *solely* according to the information provided by the experimenter. If the purpose of an experiment is to test a particular hypothesis of whether participants behave according to a certain rule, manipulation checks provide a good supporting mechanism. However, searching for the source of potential bias in participants' behavior may require additional investigation such as the solicitation of prior distributions, as demonstrated in the studies reported in this article. Methods such as think-aloud verbal protocols (Ericsson & Simon, 1993; Payne, 1994; Tor & Bazerman, 2003) provide an example of how such information can be collected.

Finally, our results may shed some light on a potential way to improve judgment under uncertainty. The normative perspective argues that people will learn from their failure and will eventually put the appropriate weight on their own priors as well as on objective data they are provided with. According to this perspective, researchers assume that the improvement will result from teaching them Bayes' formula. Our view is that better learning may result from focusing subjects' attention on the differences between their own and some larger sample-based prior probability distributions.

Acknowledgments

We are thankful for the helpful comments and suggestions of Maya Bar-Hillel, David Budescu, Yaakov Kareev, Rosemary Nagel, Yuval Rottenstreich, two anonymous reviewers, as well as the executive editor Arthur Markman. All errors remain our own.

References

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

Bereby-Meyer, Y., & Grosskopf, B. (2008). Overcoming the winner's curse: An adaptive learning perspective. *Journal of Behavioral Decision Making*, *21*, 15–27.

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

Chater, N., & Manning, C. D. (2006). Probabilistic models of language processing and acquisition. *Trends in Cognitive Sciences*, *10*, 335–344.

Epley, N., & Gilovich, T. (2006). The anchoring and adjustment heuristic: Why the adjustments are insufficient. *Psychological Science*, *17*, 311–318.

Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data*, revised ed. Cambridge, MA: The MIT Press.

Fox, C., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, *14*, 195–200.

Griffiths, T. L., Kemp, C., & Tenenbaum, J. B. (2008). Bayesian models of cognition. In R. Sun (Ed.), *The Cambridge handbook of computational cognitive modeling* (pp. 59–199). New York: Cambridge University Press.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, *17*, 767–773.

Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*, 661–716.

Grosskopf, B., Bereby-Meyer, Y., & Bazerman, M. (2007). On the robustness of the winner's curse phenomenon. *Theory and Decision*, *63*, 389–418.

Hacking, I. (1975). *The emergence of probability*. Cambridge, England: Cambridge University Press.

Ho, T. H., & Weigelt, K. (1996). Task complexity, equilibrium selection, and learning: An experimental study. *Management Science*, *42*, 659–679.

Idson, L. C., Chugh, D., Bereby-Meyer, Y., Moran, S., Grosskopf, B., & Bazerman, M. (2004). Overcoming focusing failures in competitive environment. *Journal of Behavioral Decision Making*, *17*, 159–172.

Keynes, J. M. (1936). *The general theory of employment, interest and money*. New York: Harcourt Brace and Co.

McKenzie, C. R. M. (2004). Framing effects in inference tasks—And why they are normatively defensible. *Memory and Cognition*, *32*, 874–885.

McKenzie, C. R. M., & Amin, M. B. (2002). When wrong predictions provide more support than right ones. *Psychonomic Bulletin and Review*, *9*, 821–828.

McKenzie, C. R. M., & Mikkelsen, L. A. (2000). The psychological side of Hempel's paradox of confirmation. *Psychonomic Bulletin and Review*, *7*, 360–366.

McKenzie, C. R. M., & Mikkelsen, L. A. (2007). A Bayesian view of covariation assessment. *Cognitive Psychology*, *54*, 33–61.

Nagel, R. (1995). Unravelling in guessing games: An experimental study. *American Economic Review*, *85*, 1313–1326.

Oaksford, M., & Chater, N. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences*, *5*, 349–357.

Payne, J. W. (1994). Thinking aloud: Insights into information processing. *Psychological Science*, *5*, 241–248.

Samuelson, W. F., & Bazerman, M. H. (1985). Negotiation under the winner's curse. In V. Smith (Ed.), *Research in Experimental Economics III* (pp. 38–105). Greenwich, CT: JAI Press.

Selten, R., Abbink, K., & Cox, R. (2005). Learning direction theory and the winner's curse. *Experimental Economics*, *8*, 5–20.

Suppes, P. (2007). Where do Bayesian priors come from? *Synthese*, *156*, 441–471.

Thaler, R. (1988). The winner's curse. *Journal of Economic Perspectives*, *2*, 191–202.

Tor, A., & Bazerman, M. (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*, 353–374.

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90*, 293–315.

Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences*, *10*, 301–308.

Appendix A

I. Study 1 Questions

PART I. You are CEO of Company A, which is considering acquiring Company B. The value of B depends directly on the outcome of a major oil exploration project. The more oil is found, the more Company B is worth, up to a maximum of $100/share. If no oil is found, however, B will be worth nothing – $0/share. According to the information, all values within this range are equally likely. You (Company A) have the most efficient technology in the oil business, so B would be worth 50% more under your management than under its current management. That is, if B is worth *X* under current management, it would be worth *1.5X* to you. You must determine the price (ranging from 0 to infinity) to offer per share of B. Your dilemma is this: B knows the results of the exploration project, but you do not. However, if you are going to make an offer, you must make it now. B will accept your offer as long as it is economically profitable from its own point of view. What do you think is the probability (in percentage terms, from 0 to 100) that the oil exploration will result in B's value (under its current management) being:

between $0/and $10/share?	______ %
between $10/and $20/share?	______ %
between $20/and $30/share?	______ %
between $30/and $40/share?	______ %
between $40/and $50/share?	______ %
between $50/and $60/share?	______ %
between $60/and $70/share?	______ %
between $70/and $80/share?	______ %
between $80/and $90/share?	______ %
between $90/and $100/share?	______ %

Remember your percentages should sum to 100! Given your assessment above of the possible results of the oil exploration, what offer would you make (per share) for Company B? Remember B is worth 50% more to you than it is to current management. **My offer is: $______ per share.**

PART II. Your offer for B has *not* been accepted. What do you think this means? In pursuance of your corporate expansion strategy, you are now considering making an offer for Company C. C is engaged in exploring the oil concession next to B's, and like B, its value (also ranging from $0 to 100) depends directly on the outcome of the oil exploration. Like B, if C is worth *X* under current management, it is worth *1.5X* to you. How would you assess the probabilities for C, and what, if any, offer would you make? Probability (in percentage terms, from 0 to 100) that the oil exploration will result in C's value (under its current management) being:

between $0/and $10/share? _______ %
between $10/and $20/share? _______ %
between $20/and $30/share? _______ %
between $30/and $40/share? _______ %
between $40/and $50/share? _______ %
between $50/and $60/share? _______ %
between $60/and $70/share? _______ %
between $70/and $80/share? _______ %
between $80/and $90/share? _______ %
between $90/and $100/share? _______ %

My offer is: $______ per share.

PART III

To what extent was your offer decision affected by the fact that all share values from $0 to 100 were equally likely?

(not at all) 1 2 3 4 5 (very much)

To what extent was your offer decision affected by the fact that B was worth 50% more to you than to current management?

(not at all) 1 2 3 4 5 (very much)

II. Study 2 Questions

What is your assessment of the distribution of the numbers chosen by the students in the class (including yourself)?

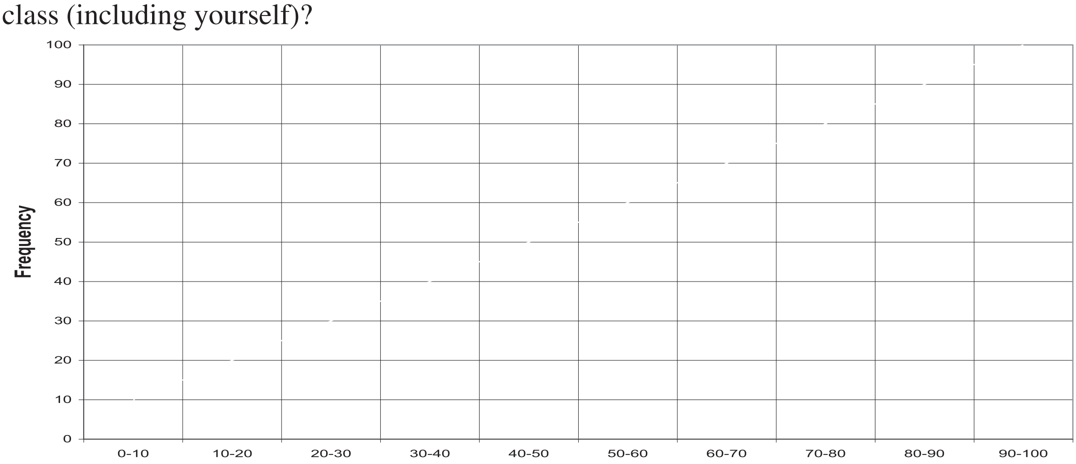

Appendix B: Classification of priors for data from Experiment 1

Our raw data for each subject is in the form of an array of numbers, each corresponding to the probability that B's value falls into one of the 10 equally spaced intervals. We classify the observed arrays into four prior categories following the simple heuristics below:

1. Uniform: if the subject's probability assessment is 10% for each of the 10 intervals.
2. Left skewed: if the sum of the subject's probability assessment for intervals *below* \$50 is *less* than the sum of his or her probability assessment for intervals *above* \$50. In this case, the mass of the subject's probability assessment is above the mid-point; that is, \$50.
3. Right skewed: if the sum of the subject's probability assessment for intervals *below* \$50 is *more* than the sum of his or her probability assessment for intervals *above* \$50. In this case, the mass of the subject's probability assessment is below the mid-point; that is, \$50.
4. Symmetric: if the subject's probability assessment is symmetric around the intervals \$40–\$50 and/or \$50–\$60.

Below we provide four quick examples for each category:

Intervals	Subject A (Left skewed)	Subject B (Right skewed)	Subject C (Uniform)	Subject D (Symmetric)
\$0 and \$10	0	0.2	0.1	0.02
\$10 and \$20	0	0.3	0.1	0.05
\$20 and \$30	0.05	0.4	0.1	0.03
\$30 and \$40	0.1	0.1	0.1	0.2
\$40 and \$50	0.15	0	0.1	0.21
\$50 and \$60	0.3	0	0.1	0.2
\$60 and \$70	0.3	0	0.1	0.2
\$70 and \$80	0.1	0	0.1	0.03
\$80 and \$90	0.05	0	0.1	0.05
\$90 and \$100	0	0	0.1	0.02