

American Economic Association

The Origin of the Winner's Curse: A Laboratory Study

Author(s): Gary Charness and Dan Levin

Source: *American Economic Journal: Microeconomics*, Vol. 1, No. 1 (February 2009), pp. 207-236

Published by: [American Economic Association](#)

Stable URL: <http://www.jstor.org/stable/25760353>

Accessed: 25/07/2014 19:25

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *American Economic Journal: Microeconomics*.

<http://www.jstor.org>

The Origin of the Winner's Curse: A Laboratory Study[†]

By GARY CHARNESS AND DAN LEVIN*

The Winner's Curse (WC) is a robust and persistent deviation from theoretical predictions established in experimental economics and claimed to exist in field environments. Recent attempts to reconcile such deviation include "cursed equilibrium" and level-k reasoning. We design and implement a simplified version of the Acquiring-a-Company game that transformed the game to an individual-choice problem that still retains the adverse-selection problem. We further simplified the problem so that simple ordinal reasoning could replace both Bayesian updating and contingent thinking. Our results suggest that the WC reflects bounded rationality in that people have difficulties performing contingent reasoning on future events. (JEL D81, D82)

Common-value auctions (CVA), where all bidders have the same (ex post) value for an item but receive different private (ex ante) signals about this value, have long been a major field of study in economics. One of the most robust findings regarding bidding behavior in CVA is known as the *winner's curse* (WC), where bidders systematically overbid, resulting in an expected loss. Explaining this persistent behavior remains an important challenge for theoretical models. In this paper, we present an experiment designed to isolate and focus on particular explanations of the WC in an environment where several recent theoretical resolutions have been proposed, as well as to test some approaches that might serve to ameliorate it.

In this type of environment, fully rational decision makers are expected to condition upon the critical future event (e.g., winning the auction) and correctly infer and incorporate the relevant posterior in their current bidding decisions. A systematic failure to do so in CVA results in overbidding and systematic negative (or below normal) payoffs.¹

* Charness: Department of Economics, University of California at Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106-9210 (e-mail: charness@econ.ucsb.edu); Levin: Department of Economics, Ohio State University, 1945 N. High Street, Columbus, OH 43210-1172 (e-mail: levin.36@osu.edu). We wish to thank Colin Camerer, Juan Carrillo, Yan Chen, David Cooper, Vincent Crawford, Erik Eyster, Yitzhak Gilboa, Jacob Goeree, Brit Grosskopf, Glenn Harrison, Nagore Iriberri, John Kagel, Eric Maskin, Tom Palfrey, Andy Postlewaite, Matthew Rabin, David Schmeidler, Bob Slonim, Lise Vesterlund, Mark Walker, Preston McAfee, Paul Pezanis-Christou, Roberto Weber, and Bill Zame for helpful discussions, as well as seminar participants at the SEA meetings in New Orleans and Washington D.C., the New and Alternative Directions Workshop at Carnegie-Mellon, ESA-Tucson and Montreal, Tel Aviv University, the Hebrew University, Yale University, University of Magdeburg, Humboldt University, University of Denmark, Vanderbilt University, University of Illinois, Max-Planck-Institute of Economics, the ASSA meetings in Boston, the University of Arizona, the University of Nottingham, the University of Amsterdam, and the University of California at Santa Cruz. All errors are our own.

[†] To comment on this article in the online discussion forum visit the articles page at:
<http://www.aeaweb.org/articles.php?doi=10.1257/mic.1.1.207>

¹ Unfortunately, economists often use the term WC to refer to the difference between the expected value of the item conditional on the event of winning and the "naïve," unconditional on winning, expectation of the CV. Such a difference may serve as a measure to the extent of the adverse selection in equilibrium.

What are the core elements that lead to the prevalence of the WC? Can the WC be rationalized as an equilibrium phenomenon or must it imply a more fundamental problem in individual decision making? Two recent innovative theoretical papers offer different but unifying stories to explain the persistence of the WC, with considerable success in explaining the winner's-curse data found in many laboratory experiments. Both of these papers permit a more relaxed belief structure than is required in standard sequential games, but maintain rationality in terms of agents (players) best responding to their (relaxed) belief structure.

Erik Eyster and Matthew Rabin (2005) introduce the notion of "cursed equilibrium," in which bidders correctly predict and best respond to the distribution of the other bids, but do not perceive correctly how these other bids depend on signals. This model, describing steady states of a cognitively-constrained learning process, permits flexible levels of value adjustment, depending on the degree of "cursedness." Vincent Crawford and Nagore Iriberri (2007) propose a model concerned with initial responses in incomplete-information games using "level- k " thinking, where each player's behavior is drawn from a common distribution of "types" that vary in their degree of sophistication. There are completely naïve and nonstrategic players, denoted as $L0$ types. Level- k type players, where $k = 1, 2, \dots$, best respond to their beliefs that all other players are of types $k-1$ and so are less sophisticated.

Since optimal behavior in a multiperson auction requires complicated calculations of one's best response, involving both beliefs about rivals' rationality and strategic uncertainty, it might be too much to expect that Nash equilibrium predictions will organize experimental data well.² Furthermore, given the complexity of the environment it is difficult to ascertain clearly the cause(s) of departures from the Nash norm. It is therefore useful to test behavior in a simpler environment where we can observe an individual's decision process in isolation, rather than as part of a more complex interaction involving other parties. Such a simpler environment may help isolate origins of the WC that cannot be explained by failure of beliefs, common knowledge of rationality, and/or best-response considerations.

One such environment is the "Acquire a Company" game, which is based on a game-theoretic analog of the famous "lemons" problem studied by George A. Akerlof (1970) and was first described in William F. Samuelson and Max H. Bazerman (1985). The game has two players: a bidder (firm) and a responder (the target firm). The bidder makes a bid and the responder either accepts the bid or rejects it. In either case, the game is over. The bidder is faced with the task of making an offer for a firm with a current value known to the firm but unknown to the bidder. The bidder does know the current value is an integer between 0 and 100, inclusive, with each value equally likely. The bidder also knows that the firm's value will increase by 50 percent if acquired. The responder transfers the firm to the bidder if and only if the bid is at least as high as the value to the current owner. In spite of the simplicity of this game, which abstracts from many of the complications embodied in an auction context, it is critical to emphasize that overcoming the WC (the adverse-selection problem) is

² Note that cursed equilibrium is meant to describe steady states of a cognitively-constrained learning process, while level- k is concerned with initial responses for learning.

still the key difficulty.³ In practice, the bid chosen is typically somewhere between the unconditional expected value to the seller of 50 and the unconditional expected value to the acquirer of 75.⁴ The WC has been found to be quite robust to a variety of alternative formulations of this task.⁵

Our experimental treatments abstract from beliefs about the behavior and cognition of other people by transforming the "Acquire a Company" game into an individual-choice task.⁶ We do so with a "robot" design in which there is no responding seller, not even any mention of such a role. A participant chooses a bid and then chooses one of 100 "cards" that are displayed face down on the computer screen (this is our "100VT"). The same rules apply, with the participant receiving 150 percent of the current value of the card (but paying the bid made), if the bid made is at least as large as the value indicated, with no payoff consequences otherwise. But here there is no other human agent's beliefs involved, either in the sense of Eyster and Rabin (2005) or Crawford and Iriberri (2007).

Thus, it is straightforward to reject these belief-based models when the decision is one of individual choice. We infer that the fundamental roots of the WC are not found in mistaken beliefs about other agents or in neglecting the cognitive process of others. But if the source of the winner's curse does not lie in one's belief structure, what could be the cause? We go deeper into the problem by developing a variation of the task in which there are only two values (this is our "2VT").⁷ In this case, the problem is greatly simplified so that ordinal dominance is sufficient to rule out all but two alternatives, and one of these is a rather unattractive gamble. In this treatment, even if we "bend backward" and allow the belief-based models to apply to individual-choice decisions, their predictions coincide with the standard Nash equilibrium. We do find that people are, by and large, capable of mastering this "ordinal reasoning" as dominated choices are rare. As we shall see, however, people often do not perform the final step and the winner's curse persists to a substantial degree.

We use the standard range of values in the 100VT for the sake of comparability to previous studies. Nevertheless, there are potential problems with a range in which the optimal bid is at one end of the range and the optimal choice is to essentially stay out of the auction. In order to confront this issue, we also run treatments where the range shifts to [20, 119], so that the optimal bid is no longer the minimum of the range (this is our "shifted-100VT"). We also conduct treatments where the only two values are 20 and 119 (our "shifted-2VT").

³ Upon suitable reflection, the bidder should realize that whenever the responder accepts the bid, B , it implies that the owner's value does not exceed the bid. This conclusion eliminates the possibility of high values (the adverse-selection issue) and necessitates forming a new and relevant expectation regarding the firm's value, incorporating the posterior distribution. With a uniform original distribution, the posterior is also a uniform distribution on $[0, B]$. Thus, its expected value is $B/2$. This leads to the conclusion that any nonnegative bid (smaller than 100) yields an expected value of $3/2 \times B/2 = 3B/4$ for an expected loss of $B/4$, and at this point it is clear that the optimal bid is 0.

⁴ Sheryl B. Ball, Bazerman, and John S. Carroll (1991) and Avishalom Tor and Bazerman (2003) use a protocol where they ask subjects to write down their reasoning. A typical explanation resembles: "Well, on the average the value is 50, which means it is 75 for me, so if I offer (say) 65, we both make handsome profits."

⁵ Other papers that examine this game include Yoella Bereby-Meyer and Brit Grosskopf (2002) and Grosskopf, Bereby-Meyer, and Bazerman (2003).

⁶ One explanation, put forth by Bazerman and a series of co-authors, is that the WC results from a failure to consider the cognition of other bidders. In a sense, this is a precursor of the belief-based theoretical models.

⁷ Camerer (1998) did pilot experiments in a 3VT design with analytically-skilled Caltech students, and found a reduction, but not an elimination of the Winner's Curse.

It is well known that people often have difficulty performing Bayesian updating or including all of the relevant information when making their decisions. For example, people may forget (as is illustrated in footnote 2) that whenever the bid is sufficient, it implies that the value does not exceed the bid, making it necessary to form a new and relevant expectation regarding the firm's value, incorporating the posterior distribution. In this sense, we might say that people have difficulty with contingent reasoning on future events. To confront this possibility, we use a two-pronged approach. First, we expand the experimental instructions to offer considerably more detail, providing multiple examples of combinations of bids and values, asking people to consider the outcomes. This has a considerable effect, particularly in the treatment where there are only two possible values, but we still see many high bids. Our final design involves transforming the bidding task into simple lotteries, where the need to do contingent reasoning is eliminated. Here, we do see that about 80 percent of the population makes the correct choice, with the residual percentage perhaps reflecting a taste for "action." These lotteries demonstrate that the framing is critical, since people do better when helped to overcome the contingent-reasoning problem. It appears that people didn't really frame the problem as lotteries in the first place.

The insightful, theoretical models put forward by Eyster and Rabin (2005) and Crawford and Iriberri (2007) definitely help to bridge the gap between laboratory behavior and equilibrium predictions in standard auctions and complex decision problems. However, we find it is not a simple matter of relaxing the very demanding belief rationality, mainly in the form of nonoptimal adjustments made for the information revealed by winning. Our study indicates that a more general cognitive error, which encompasses something like cursedness or its level- k analogs in games and certain kinds of failures to condition in individual decision-making tasks, is involved when people fall into the trap of the winner's curse.

I. Previous Empirical and Experimental Evidence on the Winner's Curse

The first formal claim of the WC was made by E. C. Capen, R. V. Clap, and W. M. Campbell (1971), three petroleum engineers, who argue that oil companies had fallen into such traps and thus suffered unexpected low profit rates in the 1960s and 1970s on Outer Continental Shelf (OCS) lease sales "year after year."^{8,9}

Not surprisingly, many economists greeted this claim with a degree of skepticism. After all, such a claim implies repeated errors and a departure from rationality, or at least from equilibrium predictions.¹⁰ As a result, that paper sparked a number of investigators to estimate rates of return from accounting data. Walter J. Mead, Asbjorn Moseidjord, and Phillip E. Sorensen (1983, 45) examine the before-tax rate of return on oil leases in the Gulf of Mexico and conclude that bidding "appears to reflect excessive enthusiasm for the amount of oil likely to be found." Kenneth Hendricks, Robert H. Porter, and Bryan Boudreau (1987) are more skeptical about

⁸ Of course Groucho Marx's statement: "I wouldn't want to belong to any club that would accept me as a member" can be viewed as an earlier recognition of the WC.

⁹ Capen et al. conclude, "He who bids on a parcel what he thinks it is worth will, in the long run, be taken for a cleaning."

¹⁰ For such a claim made by others see J. Lorenz and E. L. Dougherty (1983) and references therein.

the prevalence of the winner's curse in auctions for oil and gas leases. Nevertheless, while firms make positive profits on average, there is also evidence for the winner's curse. Most firms could have obtained larger profits by reducing their bids by a constant factor. For many of the firms the difference between the actual profits earned and those that could have been obtained with optimal bidding amounted to hundreds of millions of dollars. The authors conclude that "this result suggests that some firms may have ... failed to fully anticipate the impact of the 'winner's curse'" (Hendricks, Porter, and Boudreau 1987, 529).¹¹

The WC is a phenomenon that also seems to be present in a variety of environments other than auctions. James Cassing and Richard W. Douglas (1980) and Barry Blecherman and Colin F. Camerer (1998) find it in baseball's free agent market. John P. Dessauer (1981) suggests it exists in bidding for book publishing rights. Richard Roll (1986) suggests the WC exists in corporate takeovers. Kevin Rock (1986) presents a model in which the observed under pricing of initial public offerings is a direct result of the winner's curse problem facing uninformed investors, and Mario Levis (1990) provides empirical evidence (from British IPOs) that suggests the winner's curse and interest rate costs explain this under pricing. Finally, Cade Massey and Thaler (2005) find evidence from the National Football League (NFL) draft that is consistent with the presence of the winner's curse in that environment.

Thus, the evidence suggests that the WC is often present in field and market settings. Nevertheless, given the empirical nature of the WC and the complexities involved with field data, laboratory experiments are a natural and appealing approach for studying the WC in a carefully controlled environment. Kagel and Levin (1986), and Kagel and Levin in collaboration with others, study the existence of the WC and closely related issues resulting from such existence. The main finding is that the WC is a robust phenomenon in many CVA forms such as the first-price auction, the second-price auction (Kagel, Levin, and Ronald M. Harstad 1995), and the English auction (Levin, Kagel, and Jean-Francois Richard 1996).¹² Clearly, such persistent losses (or below normal profits) are not part of any equilibrium behavior with fully rational bidders. The typical experimental findings have been that behavior does not seem to converge to the relevant equilibrium and, to the extent that there is learning, it is very slow and situation specific.¹³

II. Theoretical Models and Behavioral Considerations

Eyster and Rabin (2005) formalize and generalize earlier ad hoc attempts by Kagel and Levin (1986) and Charles A. Holt and Roger Sherman (1994) to explain

¹¹ Hendricks and Porter (1988) also present evidence suggesting that winning bids may have been too high, but argue that it may be due to a failure to correctly anticipate the number of rivals. For more detailed discussions of the issues and evidence from these oil and gas studies, we refer the interested reader to John H. Kagel and Dan Levin (1986), Richard H. Thaler (1988), and Kagel and Levin (2002).

¹² See also Kagel and Levin (1991) and Kagel and Levin (1999). Kagel and Levin (2002) present a detailed survey of the large body of experimental work related to the WC.

¹³ Kagel and Levin (1986) find that "super-experienced bidders" do not fall prey to the winner's curse. In a similar vein, Glenn W. Harrison and John A. List (2005) find that experienced agents bidding in familiar roles are not greatly susceptible to the winner's curse, but that they do frequently fall prey to it when bidding in an unfamiliar role.

deviations from game-theoretic predictions based on full rationality. Their concept of cursed equilibrium provides an elegant formalization of the psychological principle that one tends to under appreciate the cognitive processes and/or the informational content of others' decisions. They define and apply this notion to a variety of economic environments ranging from CVA, Akerlof's lemons market, the take-over game, and voting in juries.

The cursed-equilibrium concept attempts to address the overwhelming evidence, coming in particular from experimental data, that game-theoretic predictions based on fully rational decision makers and common knowledge often miss the mark completely. They model players who do not fully take into account how other people's actions are contingent on these others' private information. As we explain below, such imperfect accounting may help explain systematic overbidding. This is the WC in both common-value auctions and the takeover game.

In the cursed equilibrium, "each player correctly predicts the distribution of other players' actions but underestimates the degree to which these actions are correlated with these other players' information." Applying this (for simplicity) to the continuous uniform $[0, 100]$ version of the takeover game, a (fully) cursed bidder correctly recognizes that for any bid $B \geq 0$, a fraction (probability) of $B/100$ (only those types with firm values that are smaller than or equal to B) will transfer the company. However, to the extent that she is cursed, the bidder then ignores the specific types that will (refuse to) transfer the firm and indiscriminately applies the probability $B/100$ to any type. Solving the bidder maximization problem under such fully cursed beliefs yields the following problem:

$$\max_{B \geq 0} \int_0^{100} \left[\frac{3v}{2} - B \right] (B/100) dv.$$

The solution to this problem is $B^* = 37.5$. And although 37.5 is still significantly lower than observed data in Ball, Bazerman, and Carroll (1991) experiments, where the dominant range of bids in most studies is between 50 and 75, the cursed-equilibrium model helps to close the gap between theory and evidence in the two-player takeover game (in fact, the 37.5 prediction comes quite close to the average bid in most of our 100VT cases).¹⁴

In our shifted-100VT, we increase the range of values by 20, so that the optimal NE bid becomes 40 instead of the "corner solution" of 0, and the fully cursed solution to the problem is $B^* = 62.5$. In this case, the bidder maximization problem under fully cursed beliefs yields

$$\max_{B \geq 0} \int_{20}^{120} \left[\frac{3v}{2} - B \right] [(B - 20)/100] dv.$$

¹⁴ The Eyster and Rabin model completes the spectrum from a fully rational player to a perfectly cursed player by introducing a parameter χ . For example, in our design, $\chi = 0$ means that the player is fully rational and thus bids 0; $\chi = 1$ means that the player is fully cursed in his or her beliefs and will bid 37.5, as in the text. Intermediate values such as $\chi = 1/2$ mean that the player is partially cursed, and thus, depending on the degree of cursedness, bids between 0 and 37.5. This rationalizes only bids lower than 37.5. Since, here, even $\chi = 1$ does not move predictions enough, for the sake of brevity, we omit discussing intermediate values in the text.

We shall see that this is not far from the average observed bid in our shifted-100VT case.

Since the model is based on beliefs about other players, the cursed-equilibrium story does not apply when we convert the takeover game to an individual-choice problem, with no other players. However, Eyster and Rabin (2005, 1632) argue that, in a less literal sense, an alternative interpretation of cursed equilibrium is not that each player is certain that every other player's action does not depend on her type but rather that they don't look for the correlations this mechanism implies. So when subjects hear the rules of the game, they understand that higher offers are more likely to be accepted but don't think through the mechanism behind this relationship. Thus, while our individual-choice results violate the model's predictions, they are arguably not inconsistent with the underlying motivation. While it is not clear (i.e., there are likely several ways) how to incorporate such an idea into an individual choice model, it seems that this might be a fruitful exercise.

Here is where our 2VT is particularly useful. Even if we effectively bend backward and consider the computer to be a player in a two-player game, the Eyster and Rabin model predicts that people will bid 0 when the value can only be either 0 or 99. To see this, recall that a cursed bidder gets the distribution of acceptances right, but doesn't take into account the correlation between the acceptance and the seller's value. So, a bidder in the 2VT knows that bids less than 99 are accepted with probability $\frac{1}{2}$, and bids at least as large as 99 are accepted with probability 1. However, he doesn't take into account what this implies for value, which he always assumes to be $\frac{3}{2} \times (0 + 99)/2 = 74.25$. So, his payoff from bidding $B < 99$ is $\frac{1}{2} \times (74.25 - B)$, and his payoff from bidding $B \geq 99$ is $1 \times (74.25 - B)$. Hence, his optimal bid is zero.¹⁵ Parallel reasoning in the shifted-2VT case yields an optimal bid of 20.

Crawford and Iriberri (2007) provide an alternative approach to explaining deviations from game-theoretic predictions based on full rationality developing a model based on level- k thinking. There is a nonstrategic (naïve) type, $L0$, in the population. The next level is a more sophisticated type, $L1$, who best responds to her beliefs that all other players are of type $L0$. The next most sophisticated type, $L2$, in turn best responds to his beliefs that all other players are of type $L1$, and so on. It is clear that matters depend critically on the specification of $L0$, at least in terms of the model's explanatory power, as this sets in motion and determines the best responses for the more "sophisticated" types. Crawford and Iriberri find that this approach, with only a small number of types, can be quite successful in a wide variety of auction environments.

Once again it is not clear how this model applies to our experimental design, however. What is the meaning of $L0$ when there are no other players? For example, suppose that type $L1$ does explain the experimental data. Who are the $L0$ types to whom they are making best responses? Certainly, if taken literally, the level- k approach does not have any bite in our environment.¹⁶ Once again, however, (as with Eyster and Rabin) the argument can be made that the underlying notion can be extended.

¹⁵ We are indebted to Jacob Goeree for pointing this out.

¹⁶ In fact, the conclusions for both the 100VT and the 2VT also extend to the Crawford-Iriberri "truthful" level- k model, in which $L0$ responders bid truthfully, making $L1$ and higher proposer types coincide with equilibrium.

Crawford and Iriberri argue that our analysis “leaves open the possibility that something like the level- k model accurately describes initial responses to environments, interactive or not, that pose cognitive difficulties that are isomorphic to those of predicting other players’ strategic decisions.”

In fact, even if we consider the computer to be a player in a two-player game, the level- k model predicts bids of 0 in both the 100VT and the 2VT. To see this, consider first that we are only interested in the choices of the bidder who anticipates responses from the “seller.” For the level- k model to have meaning in our context, we must presume the computer to be a random player (L_0), even though it has a dominant strategy. But then the computer accepts or rejects any bid with a 50 percent (or any other percentage) chance. Since each bid is accepted with the same likelihood, each bidder should “best respond” by bidding zero.¹⁷ Thus, the level- k model cannot explain the WC in the takeover game.¹⁸

Summing up, although these models may indeed provide insight, and may explain (at least partially) departures from equilibrium behavior in games, we demonstrate below that the WC is alive and well even in an individual-choice environment. It suggests that the origin of this phenomenon must stem from some form of bounded rationality, such as the decision maker’s failure to recognize that a “future” event, per se, is informative and relevant for their current decisions, compounded by poor updating when this idea is even considered.

Our experimental test totally strips away any element of another party’s decision process. In short, we “eliminate the middle man.” As previous studies (e.g., Gary Charness and Levin 2005) have shown, the complexity of the environment may be an important factor in departures from optimal behavior. People who may be poor Bayesians in a complex environment may do quite well in a simple one. Thus, we simplified the decision task by introducing the 2VT, which (in addition to stress-testing a more general version of the Eyster and Rabin model) provides the decision maker a way to reach the correct answer without having to invoke Bayesian updating.

We are not the first to use a computer to interact with a bidder in this task. To the best of our knowledge, James C. Cox, Vernon L. Smith, and James M. Walker (1987) was the first study to use robot bidders in a first-price auction experiment. Reinhard Selten, Klaus Abbink, and Ricarda Cox (2005) do frame a form of the 100VT task as an individual-choice problem.¹⁹ In their design, there is no graphical display on the screen, which we feel makes it completely transparent to the bidder that there is no other entity involved, just one’s bid and one’s choice of card on the screen—no middleman. More importantly, their focus is on learning direction theory rather than the winner’s curse, per se. They do not consider the theoretical models of the winner’s curse. We also conduct a variety of other treatments (such as the 2VT and the lottery treatment) that shed more light on the origin of the winner’s curse.

¹⁷ Parallel reasoning in our shifted-100VT and shifted-2VT treatments yields a best response of 20 in each case.

¹⁸ Why does this model fail in the takeover game while performing so well in auctions? The main difference between the takeover game and auctions is the information structure. In the takeover game, where one player has unambiguously better information (and this fact is common knowledge), there may be a different menu of possible heuristics available to the players than there is in auctions, where each player has some private information. Thus, it appears that when informational asymmetries are present, the level- k specification requires some adjustment.

¹⁹ Our study is independent from theirs, in the sense that we began our experiments in 2003, and were not aware of their study at that time.

As we shall see, there are additional behavioral considerations that apply to behavior in our experiments. For example, there is always some element of confusion, which we attempt to address by providing more detailed instructions and a larger number of examples. We also introduce a treatment (the 2VT) where the choice is, in essence, greatly simplified, as it involves only ordinal logic. In addition, there are issues of whether people are actually able to update properly, but don't always choose rationally, given their posteriors. Charness and Levin (2005) find that people are affected by emotional considerations such as psychological effect in more complex environments. Perhaps people are quite myopic, so they have difficulty envisioning their circumstances after a future event, so they are effectively unable to perform contingent reasoning on future events. The results from the 2VT, in combination with the lotteries, where no updating is required, shed light on these questions, as we do find a difference in behavior across both the shifted and unshifted cases.

III. Experimental Design

We conducted experimental sessions at the University of California at Santa Barbara. Participants were recruited from the general campus population by an e-mail message. Some 388 people participated in our experimental sessions. We distributed written instructions to the students and read these aloud, taking questions as they arose. The experiment itself was performed using a Web-based computer interface. We also conducted simple lotteries in which 86 people participated. No person participated in more than one experimental session or lottery.

Participants received \$5 for showing up on time and, since we expected losses, they also received a \$15 endowment. At the end of the session, we randomly chose one period from each 10-period block for payment, so that only six periods actually counted toward monetary payoffs. This device helped us to avoid (or at least minimize) possible intra-session income effects that could readily affect behavior, including potential bankruptcy issues.²⁰ After netting out the experimental gains and/or losses from the six chosen periods, we converted these to actual dollars at the rate of \$0.10 for each unit and added the result to the original endowment. While one could lose the entire \$15 endowment, we mandated that no participant could leave with less than the \$5 show-up fee, regardless of the outcomes.

After the instructions were read and discussed, 100 "cards" were displayed on each person's computer screen in a 10-by-10 array. Each card had an integer value between 0 and 99, inclusive, with one card having each of the possible 100 values. In

²⁰ In fact, the bankruptcy constraint should not affect behavior. First of all, at the time of bidding the subject does not know the profits or losses that have accrued in the past, as the periods used for payoff are not chosen until the end. Nevertheless, suppose we are faced with a subject who might wish to take advantage of the limited-liability issue. For simplicity, consider the 2VT case. First, consider behavior in a 20-period session. The subject thinks that if he bids 99, 20 times in a row, there is a 25 percent probability that he will gain 99, a 50 percent probability he will lose 49.5, and a 25 percent probability he will lose 198 but will only be liable for 150. This lottery has an expected value of -37.5 , even worse than the one-shot bid of 99, which has an expected value of -24.75 . More generally, applying this logic to 60 bids of 99 in the 2VT, calculations show that the expected value of this strategy is -79.40625 . Since this is not a profitable strategy, the limited liability should not be an issue here, barring the possibility of confusion.

the “standard” 100VT version, the cards were initially displayed in sequence (0 to 9 on the first row, 10 to 19 on the second row, etc.), so it was easy to verify that there was exactly one of each of these integers.

After viewing these cards, a participant clicked to flip the cards over and start the game, knowing that the cards would be reshuffled after every action (as he or she could verify during the session). Each period the participant chose a bid (a nonnegative integer) and then selected one of the cards. The card’s value was then revealed. If the bid was at least as large as the value displayed, the participant paid the bid and received 1.5 times the value displayed on the card. If the bid was (strictly) less than the value displayed, there was no bid payment or value received. This outcome was displayed on the screen.

In our 2VT, we reduce the complexity of the Bayesian updating problem while simultaneously increasing the variance of the outcome and feedback.²¹ We note that bidders in both treatments were free to choose any nonnegative integer as a bid. In this version a card can have only two possible values: 0 or 99. A participant sees 50 cards with a value of 0 and 50 cards with a value of 99.²² Thus, the only “reasonable” bids are either 0 or 99,²³ as any positive bid less than 99 is dominated strictly by a lower bid. A higher bid (but still less than 99) does not increase the probability of a transaction occurring, but does raise the cost to obtain it. Note that this is indeed a vast simplification as it involves only ordinal dominance logic. The only element of cardinality in the process is to compare bidding 0 to bidding 99, since bidding 0 yields a certain payoff of 0, while bidding 99 produces a rather unattractive lottery where the bidder wins 49.5 or loses 99 with equal probabilities of $\frac{1}{2}$. Note also that the variance of the value is much greater in the 2VT.²⁴

In contrast, the analysis is much more complex in the 100VT. Why is a bid of 60 worse than a bid of 59? If the firm’s value is 61 or more, it does not matter. If the firm’s value is precisely 60, the bidder is better off, by 30, by bidding 60 and winning than unsuccessfully bidding 59. In all of the other 60 cases, the bidder is worse off with a bid of 60 than a bid of 59, but by only 1. Thus, recognizing that bidding 59 dominates (in expectations) bidding 60 requires some nontrivial calculations, cardinal reasoning, that also involves the use of probabilities.

Each of our experimental sessions consisted of 60 periods. In some sessions, participants were first presented with cards having 100 distinct values. They were told that the distribution of the cards’ values would change after 30 periods, and the new values would be displayed at that time. In the subsequent 30 periods, cards had only two possible values {0, 99}, as discussed above. In some other sessions, participants

²¹ Colin Camerer (1999) did pilot experiments in a 3VT design with analytically-skilled CalTech students and formed a reduction but not an elimination of the winner’s curse.

²² We also conducted an intermediate four-value treatment, in which there were 25 cards each with values 0, 33, 66, and 99. Obviously, this is less complex than the 100VT, but is nevertheless qualitatively different than the 2VT, and could provide further insight into the nature of the WC in our setting. We chose the 4VT in part to examine the degree to which the increased proportion of zero bids observed in the 2VT was an artifact of there being only two “reasonable” bids. We report some details below, but in the interest of clear exposition, we refer the interested reader to our working paper for more information.

²³ Or 100 if one is uncertain if the bid needs to exceed the value or just equal it.

²⁴ In the original game, a bidder makes a profit conditional on winning whenever the firm’s actual value is between the bid and two-thirds of the bid (a probability of $\frac{1}{3}$) and make a loss whenever the firm’s actual value is between 0 and two-thirds of the bid (a probability of $\frac{2}{3}$).

were first presented with the 2VT, and the 100VT began after 30 periods.²⁵ In all cases, a history of the results of all previous periods in a segment (showing period, bid, value, profit) was displayed on the screen.²⁶

As mentioned earlier, one potential problem with the parameterization in both the 2VT and the 100VT is that the optimal bid in each case is zero. Since this is an action on the boundary of the action space, there could well be other reasons (e.g., boredom) that would lead people to avoid this action. To address this concern, we also conducted sessions in which the range of values was shifted to the integers between 20 and 119 (in the 2VT) and to {20, 119} in the 2VT. The optimal bid is now 40 in the 100VT and 20 in the 2VT. Each optimal bid yields a positive expected profit, and we can observe errors from underbidding.

Finally, it could be argued that our 2VT results stemmed from many (risk-seeking) people who wished to bid 99 even though they understood the consequences.²⁷ To test this, we employed two main strategies. First, we initially gave participants instructions with some detail and a minimal number of examples. We also ran additional two-segment sessions with more detailed instructions and more examples. Our less detailed (hereafter “LD”) instructions are shown in Appendix A, while the more detailed (hereafter “MD”) instructions are shown in Appendix B. Second, we abstracted from the multiple bidding choices available in the 2VT and framed four possible bids in terms of logically equivalent lotteries, producing two sets of four simple lotteries. The lotteries for the 2VT are shown in Appendix C.²⁸ In order to avoid the possibility of a limited liability issue, we endowed each person with an additional \$10, against which any gains or losses were netted out. We conducted these lotteries at the end of another experiment (involving the prisoner's dilemma), choosing five people from each session for actual additional payoffs.²⁹

Table 1 summarizes our experimental treatments.

IV. Experimental Results

This section presents the results of our experimental treatments, with a particular focus on theoretical issues and predictions in relation to these results. The first question we address is whether transforming the Acquire-a-Company game into an individual-choice problem eliminates the winner's curse in the “standard” 100VT.

A. 100VT Results

Table 2 and Figure 1 present the results in the 100VT.

²⁵ In a third set of sessions, there were 4 segments of 15 periods each, in the order 2VT-100VT-2VT-100VT.

²⁶ At the end of many of the sessions, we also passed out four questions that involved making estimates of probability and required some ability to do Bayesian updating. These questions are an attempt to substitute for the often-used demographic questions about a student's major, SAT score, and so forth, and are shown in Appendix D. For a thorough examination of the effect of such demographic factors on behavior see Marco Casari, John C. Ham, and Kagel (2007).

²⁷ See, for example, Ellen Garbarino and Robert Slonim (2005).

²⁸ Robert Dorsey and Laura Razzolini (2003) use a similar approach in converting bids in a first-price auction to lottery choices.

²⁹ All participants in these sessions knew that they had already earned positive sums of money prior to this exercise.

TABLE 1—EXPERIMENTAL TREATMENTS

Session sequence	Instructions	Number of participants
100VT–2VT	Less detail	54
2VT–100VT	Less detail	48
2VT–100VT–2VT–100VT	Less detail	47
4VT	Less detail	57
100VT–2VT	More detail	38
2VT–100VT	More detail	33
Shifted 100VT–2VT	Less detail	52
Shifted 100VT–2VT	More detail	59
Lottery Set 1 (original)	—	46
Lottery Set 2 (shifted)	—	40

TABLE 2—AGGREGATE BIDDING STATISTICS FOR 100VT
BY SEQUENCE AND TYPE OF INSTRUCTIONS

Session sequence, instructions	Observations	Avg. bid ^a	Zero-bid rate
100VT first, LD	1,620	38.86	0.0753
100VT last, LD	1,440	35.91	0.2090
Four segments, LD	1,410	33.47	0.1823
100VT last (4VT), LD	1,710	32.64	0.3099
100VT first, MD	1,140	35.17	0.2579
100VT last, MD	990	29.12	0.4010

Notes: By “100VT first,” we mean that the first segment was the 100VT, which was then followed by the 2VT. “LD” means less detailed instructions and “MD” means more detailed instructions.

^aThe average bid is computed excluding bids over 100, as a few very high bids otherwise distort it.

In contrast to the strict predictions of the Eyster-Rabin and Crawford-Irriberri models, the winner’s curse survives even when one cannot form beliefs about the behavior and motivations of other players. It is clear that, in order to address our results, these models would need to be able to incorporate probabilistic inferences that do not involve the decisions of others. If we relax this requirement, and simply pretend that there is a seller in this “game,” then the average bid made is not too far from the Eyster-Rabin prediction for fully cursed bidders, coming closer to our data than that in previous studies, where the “dominant range of responses ... is between \$50 and \$75” (Tor and Bazerman 2002, 8). With the specification we have chosen for L_0 , which seems to us to be the most natural one, the level- k model predicts that L_1 types will bid zero, so our data reject this model even if we pretend that there is a seller.

In fact, while the average bid does not differ greatly across the treatments in Table 2, we do see substantial variation in the proportion of zero bids made. There is a 13–15 percentage point increase in this proportion when the 100VT is last, so that subjects have first had the experience of playing the 2VT. This difference is significant on a Wilcoxon-Mann-Whitney rank sum test, using each subject’s average zero-bid rate as one observation, with $Z = 2.29$, $p = 0.011$, one-tailed test). In

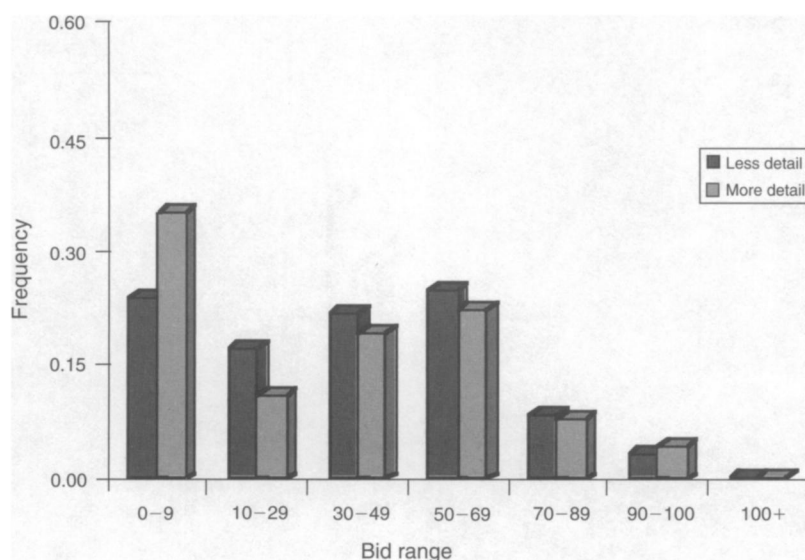


FIGURE 1. BID FREQUENCY IN 100VT, BY INSTRUCTION TYPE

addition, we find that providing more-detailed instructions increases the zero-bid rate by 18–19 percentage points (also a significant difference, using the same test: $Z = 2.16$, $p = 0.015$, one-tailed test) even though we have presented no theoretical explanation for this difference. We return to this point later.

Shifted 100VT.—We started with the design that has $\{0, 1, 2, \dots, 99\}$ valuations as it is the one closest to the standard Acquiring-a-Company game used in previous studies. However, two possible objections to this design are the fact that the optimal bid is at an end point of the action space, possibly biasing the results, and that the payoff for the optimal choice is zero (and although better than a negative payoff)

TABLE 3—AGGREGATE BIDDING FOR SHIFTED-100VT, BY SEQUENCE AND TYPE OF INSTRUCTIONS

Treatment, instructions	Observations	Avg. bid ^a	Overbid	Bid 40	Underbid
100VT Shift, LD	1,560	62.57	1,294 (82.9%)	29 (1.9%)	237 (15.2%)
100VT Shift, MD	1,770	68.44	1,561 (88.2%)	46 (2.6%)	163 (9.2%)

^a The average bid is computed excluding bids over 120, as a few very high bids otherwise distort it.

and may induce boredom. What happens when we shift the range to $[20, 119]$, where the optimal bid becomes 40 with expected earnings of 5 upon winning? Table 3 and Figure 2 present these results.

Once again, the WC seems present, as the proportion of overbids overwhelmingly exceeds the proportion of underbids, whether or not more detailed instructions have

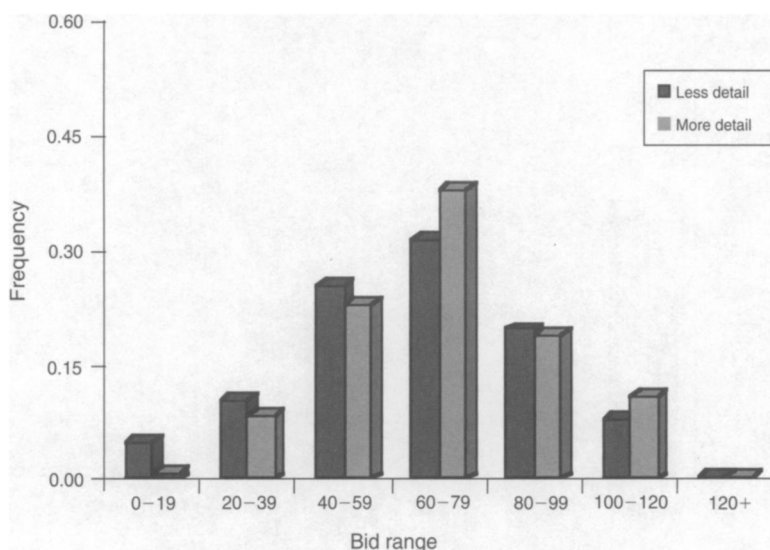


FIGURE 2. BID FREQUENCY IN SHIFTED 100VT, BY INSTRUCTIONS

been provided.³⁰ We note that, if we pretend that there is a seller, the Eyster-Rabin prediction for a fully cursed bidder is not that far from the average bid in both cases.

In Figure 2, we see the most common bids are in the 60–79 range, with the distribution resembling the normal distribution.³¹

B. 2VT Results

Next, we consider bidding behavior in our 2VT. Recall that in this case, both the Eyster-Rabin and the Crawford-Iriberri models predict bids of zero, whether or not there is an actual seller. We also anticipated a great preponderance of zero bids, as simplifying the task to require only ordinal reasoning limits the plausible choices to the two possible values, and the gamble implied by bidding the higher value seemed quite unattractive to us. Table 4 and Figure 3 summarize our overall results for the case where there are two possible values.

Around 30 percent to 60 percent of all bids are zero when we limit the number of possible values to two, so that the winner's curse is lessened in the less complex environment. Comparing individual tendencies, the zero-bid rate was higher in the 2VT than in the 100VT for 146 people, while this was reversed for 18 people (the rates were the same for the other 56 people). The binomial test shows that this is significantly different from random behavior ($Z = 10.00$, $p = 0.000$, two-tailed test). Note, the average bid is typically substantially higher in the 2VT than in the 100VT

³⁰ One caveat to this interpretation is that the optimal bid is closer to the lower end than the higher end, so that we would observe more overbids than underbids with random bidding. Nevertheless, with random bidding, we should expect 306 underbids out of the 1,531 non-40 bids in the LD case, and 345 underbids out of the 1,724 non-40 bids in the MD case. We observe substantially fewer underbids (and more overbidding) in the actual data.

³¹ We see a little confusion in the LD treatment, as nearly 5 percent of all bids are below 20 in that case.

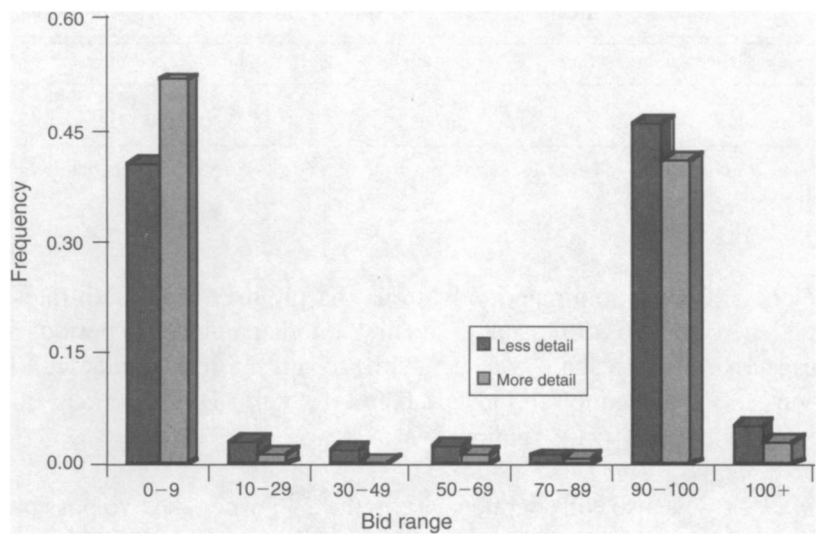


FIGURE 3. BID FREQUENCY IN 2VT, BY INSTRUCTION TYPE

TABLE 4—AGGREGATE BIDDING STATISTICS FOR 2VT, BY SEQUENCE AND TYPE OF INSTRUCTIONS

Session sequence, instructions	Observations	Avg. bid ^a	0-bid rate
2VT first, LD	1,440	57.08	0.3042
2VT last, LD	1,620	59.87	0.3346
Four segments, LD	1,410	37.07	0.4929
2VT first, MD	990	52.93	0.3848
2VT last, MD	1,140	36.21	0.5851

^a The average bid is computed excluding bids over 100, as a few very high bids otherwise distort it.

(although the difference is diminished when there are four segments and with the detailed instructions when the 2VT is the final treatment), so that participants tend to lose more money with the simpler design.³²

We see that the distribution of bids in the 2VT is bimodal, with either very low bids or bids near 100 predominating. It appears that most participants possessed some degree of sophistication, understanding that it was foolish to bid anything other than 0 or 99 (or 100, if one is uncertain about the tie-break rule). In fact, these bids comprised 82.5 percent of all bids in the 2VT with the LD instructions, and 90.8 percent of all bids in the 2VT with the MD instructions.³³ Thus, people seem able to master ordinal reasoning and avoid dominated (interior) bids. We note further that

³² The fact that bids are significantly higher in the 2VT than in the 100VT, so that people are actually losing more money in this case, illustrates that a little sophistication may be worse than none. A naïve bidder, who bids 68 in the 100VT, loses 11.56 on average, while a fully rational bidder breaks even with a 0 bid. A “partially sophisticated” bidder in the 2VT, who recognizes that bidding 68 makes no sense, but can’t see that 0 is better than 99, instead loses 24.75 on average—more than twice the naïve bidder’s average losses!

³³ In the 4-value treatment, 0 was bid 31 percent of the time; 33 was bid 23 percent of the time; 66 was bid 15 percent of the time; and 99 was bid 7 percent of the time. These bids comprised 76 percent of all bids (85 percent if we include bids of 34, 67, or 100).

TABLE 5—AGGREGATE BIDDING STATISTICS FOR SHIFTED 2VT, BY TYPE OF INSTRUCTIONS

Session sequence, instructions	Observations	Avg. bid ^a	20-bid rate
Less detail	1,560	52.01	0.4923
More detail	1,770	55.38	0.5757

^a The average bid is computed excluding bids over 120, as a few very high bids otherwise distort it.

the frequency of these dominated bids drops sharply over time, with rates of 16.7 percent, 5.8 percent, and 3.7 percent in the first, middle, and last 10 periods of a segment, respectively. Once again, we see a shift toward more low bids with the MD instructions, as 40 percent of all bids are below 10 with LD instructions, compared to 52 percent with the MD instructions.

Shifted 2VT.—We also consider behavior in the 2VT when the two possible values are 20 and 119. In this case, the lower value is still the optimal bid. However, bidding 20 does not mean that there is no “action” for the subject, since 50 percent of the time bidding 20 gives a profit of 10 units. Since nonnegative bids below 20 are allowed, we can also potentially observe underbidding. Our results in the shifted 2VT are summarized in Table 5.

We see that the proportion of optimal bids is around 50 percent or higher, as high as is seen in the standard 2VT. Nevertheless, it is clear that shifting the end points does not eliminate the WC. Figure 4 shows that, once again, people largely avoid dominated (interior) bids, suggesting that ordinal reasoning is not too difficult for our subject pool.

In the 2VT with less detailed instructions, 276 bids out of 1,560 (17.7 percent) were either less than 20 (55 bids), between 21 and 119 (160 bids), or greater than 120 (61 bids). Having more detailed instructions again reduced the confusion that people had with this choice, as in this case, 138 bids out of 1,770 (7.8 percent) were either less than 20 (20 bids), between 21 and 119 (98 bids), or greater than 120 (20 bids). The rate of dominated bids again drops steadily over time, with rates of 17.8 percent, 11.0 percent, and 8.5 percent in the first, middle, and last 10 periods of a segment, respectively.

C. Problems with Contingent Reasoning on Future Events

We now consider some possible explanations for the persistence of the WC that rely on cognitive issues and not on the (relaxed) belief structures proposed by the Eyster-Rabin and Crawford-Iriberri models. First of all, it seems clear that the 100VT environment is more complex than most subjects can handle, so simplifying to the 2VT environment makes the decision problem more approachable. Few people choose dominated bids frequently,³⁴ and yet a not inconsequential proportion

³⁴ In the original 2VT, 126 out of 220 (57 percent) subjects never choose a dominated bid, while 161 out of 220 (73 percent) subjects chose dominated bids no more than 10 percent of the time; 56 percent of these bids were chosen by 10 percent of the population. In the shifted 2VT, 67 of 111 (60 percent) subjects never choose a dominated

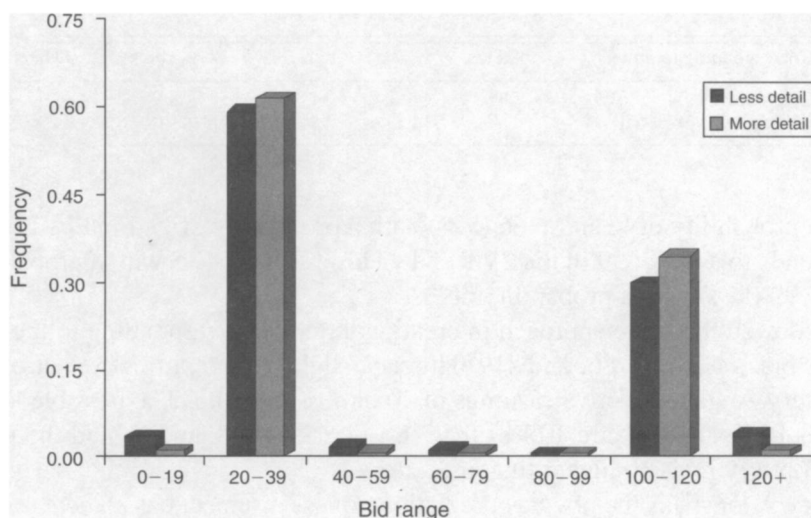


FIGURE 4. BID FREQUENCY IN SHIFTED 2VT, BY INSTRUCTIONS

of the participants chose the risky gamble with negative expected payoffs over either the assured result of zero (in the original 2VT) or a small, positive expected profit (in the shifted 2VT).

So we can see that complexity is a hurdle for many people in this decision task. One of our approaches to ameliorating this problem is to provide more detailed instructions and a greater number of examples. This device did have a moderate degree of success, decreasing the rate of dominated bids in all treatments, increasing the 0 bid rate in the original 100VT and 2VT, and increasing 20 bids in the shifted 2VT. Thus, helping people work through the contingent reasoning problem bears some fruit, perhaps particularly with the less complex task. However, we are still left with many people who choose to make large bids, even in the simpler 2VT.³⁵

Lottery treatments.—In order to proceed further in our investigation, we asked people who had not been in our 60-period sessions to choose a preferred lottery from a set of four possibilities labeled A, B, C, and D. Each of these lotteries corresponds to a bid in the 2VT. Lottery A is a safe “Lottery” that assures zero earnings and corresponds to a zero bid in the original 2VT task. Lottery B corresponds to a bid of 33 in this task and yields 0 with a probability of $\frac{1}{2}$ and -33 (loss) with a probability of $\frac{1}{2}$. Lottery C corresponds to a bid of 66 in the original task and yields

bid, while 82 of 111 (74 percent) subjects chose dominated bids no more than 10 percent of the time; 62 percent of these bids were chosen by less than 10 percent of the population.

³⁵ It has been suggested to us that one could consider that there is a cognitive cost to performing calculations, thus helping to explain why people overbid in the 100VT and the 2VT. In principle, this seems reasonable. However, we are a bit skeptical about its application here. In the 100VT, bidding randomly results in an expected loss of $49.5/4 = 12.375$ units in each period. Three periods (out of the 30 in this treatment in a session) count toward payoffs, at the rate of \$0.10 per unit, so that the expected cost of not performing the calculations is \$3.71. In the 2VT, the argument would be that there is very little cognitive cost when ordinal dominance is involved, so bids should be either 0 or 99. Bidding 0 or 99 randomly leads to an expected loss of 12.375, or \$3.71 for the three periods that count. This is a substantial cost in a lab experiment, where people are usually motivated by far smaller amounts of money.

TABLE 6—LOTTERY CHOICES

Corresponding treatment	A bids	B bids	C bids	D bids
Original 2VT (N = 46)	39	3	0	4
Shifted 2VT (N = 40)	30	1	0	9

0 with a probability of $\frac{1}{2}$ and -66 (loss) with a probability of $\frac{1}{2}$. Finally, Lottery D corresponds to a bid of 99 in the 2VT and yields a gain of 49.5 with a probability of $\frac{1}{2}$ and -99 (loss) with a probability of $\frac{1}{2}$.

We followed the same approach in creating lotteries corresponding to the shifted 2VT for bids of 20, 53, 86, and 119. This is a slightly more attractive set of lotteries. Lottery A' gives positive earnings of 10 half of the time. The possible losses in Lotteries B', C', and D' are 10 less than in Lotteries B, C, and D, and the possible gain in Lottery D' is 10 higher than in Lottery D.

Lottery A (or A') is the obvious prediction, unless a subject has some strong preference for making a risky choice. Lotteries B and C (or B' and C') correspond to dominated (interior) bids, while Lottery D (or D') is at least consistent with ordinal reasoning, but one might have expected a simple calculation would lessen the attraction of this high bid.

We had two important motivations for conducting these Lottery treatments. First, some reviewers felt that the relatively high proportion of bids of 99 in the 2VT might reflect strong preferences for risk seeking, so that people understood the gamble involved and found it attractive. Since each positive bid may be presented as a choice of a relevant lottery, one might argue that the WC is nothing more than an expression of risk loving. Second, there is a critical framing/contingent reasoning issue. Although the Acquire-a-Company task is logically and mathematically equivalent to a choice among 100 (relevant) lotteries, behaviorally these may not be equivalent, since the presentation as lottery circumvents the need for contingent reasoning.

The results of this treatment are shown in Table 6.

Overall, we see that 85 percent of the subjects chose the equivalent of bidding 0 in the original 2VT, compared to 75 percent in the shifted 2VT. Four people out of 86 (5 percent) made confused choices in this one-shot decision task. Note, choosing 119 in the shifted 2VT is a better gamble than choosing 99 in the original 2VT, yielding an expected loss of -14.75 compared to -24.75 . Perhaps this has something to do with why we see more D choices in the second lottery set.

We take this overall 80 percent result to be a rough upper limit for how far one can be able to eliminate the winner's curse in the laboratory, given certain characteristics of the environment such as the fact that people are playing with "house money." This treatment would seem to be about as transparent as is feasible, without overtly waving a flag in the participants' face. The contingent-reasoning problem is solved for the decision maker here.³⁶

Thus, our evidence indicates it is not the case that the half (or more) of the subjects in the 2VT (or shifted 2VT) who made high bids were risk seeking per se,

³⁶ In fact, evidence from Charles A. Holt and Susan K. Laury (2002) shows that people make very few risk-seeking choices in lotteries. Their Figure 6 indicates that fewer than 10 percent of all choices were risk seeking.

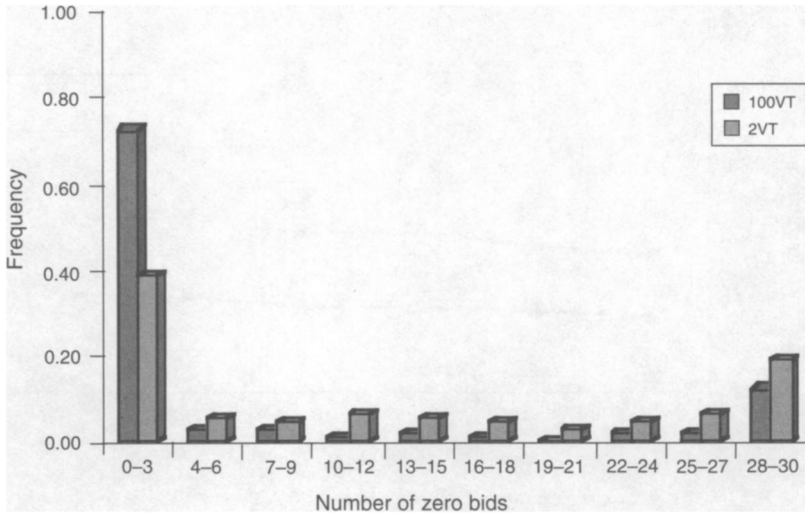


FIGURE 5. ZERO BIDS IN 100VT AND 2VT, BY INDIVIDUAL

being sufficiently attracted to the gamble and fully understanding the consequences when making the high bid. We see a reduction of 25–50 percentage points in the 0-bid (20-bid) rate in the 2VT (shifted 2VT) depending on the level of detail provided in the instructions. It appears that people do indeed have some problem with contingent reasoning regarding future outcomes.³⁷

D. Heterogeneity, Time Patterns, and Characteristics of Bid Changes

The data we have shown are aggregations. However, this presentation masks the fact that there is considerable heterogeneity in behavior.

Figure 5 shows the number of 0 bids (out of 30 possible) made by each individual in the 100VT and the 2VT. We see that few participants bid zero very often in the 100VT. More than three-quarters of the participants bid zero less than 10 times, while one-sixth of the participants bid zero more than 20 times. In the 2VT, half of the participants bid zero less than 10 times, while one-third bid zero more than 20 times. Some people mix between bids of 0 and 99 in the 2VT.³⁸

Economists often wonder if a given phenomenon is simply ephemeral or seems to have some staying power. We find, perhaps surprisingly, little change in our sessions over time. Figure 6 shows the zero-bid rate in the standard 2VT and 100 VT. Figure 7 shows the overbidding rate in the shifted 100VT and the shifted 2VT.

³⁷ Robert Dorsey and Laura Razzolini (2003) find that people with high values bid considerably and significantly more in the auction than in the corresponding lottery, but that this difference in behavior is substantially reduced by showing people the probability of winning the auction. Thus, there is evidence in their study that a misperception of probabilities is a factor in overbidding.

³⁸ It is possible that such people (and perhaps those who make small bids in the 100VT) understand that bidding zero is correct, but would nevertheless like to be active to some extent anyway. To the extent that this is true, we might expect that our observed zero-bid rates understate the level of understanding.

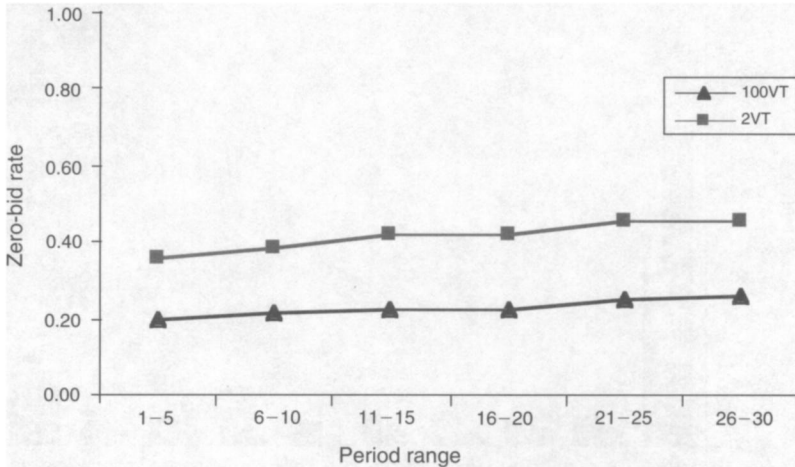


FIGURE 6. ZERO-BID RATE OVER TIME IN 100VT AND 2VT

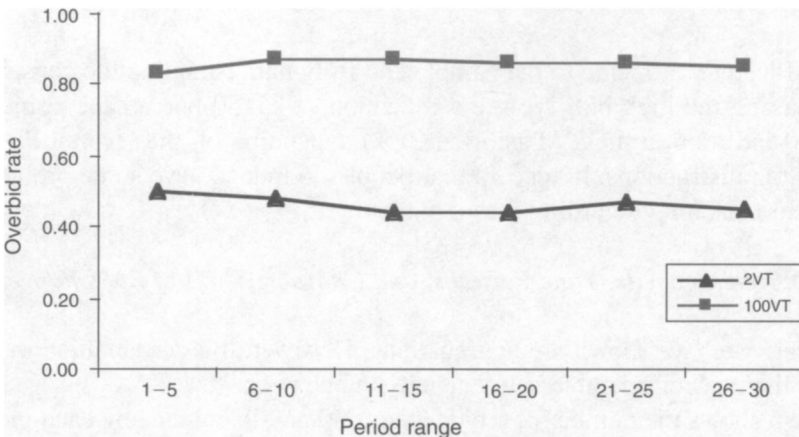


FIGURE 7. OVERBID RATE OVER TIME IN SHIFTED 2VT/100VT

As can be seen, there is no strong pattern over time, although the zero-bid rate is increasing slightly in Figure 6.

While learning is not a focus of this paper, and we see little change over time, we wish to point out a behavioral regularity present in both the original 100VT and the shifted 100VT. This is illustrated in Tables 7 and 8.

After a “no outcome,” there is a strong tendency to raise one’s bid, perhaps because the subject wanted some “action.” What is particularly interesting is that subjects had a tendency to lower their bids after an outcome, whether positive or negative. This is strongest in the shifted 100VT, where more people lower their bid than raise their bid by 52 percent to 17 percent after a negative outcome, and 43 percent to 16 percent after a positive outcome. In both the original and the shifted 100VT, the total number of times that the bid was raised and lowered was roughly equal, helping to explain why we do not see much drift over time.

TABLE 7—OUTCOMES AND BID CHANGES IN 100VT

	No. of times bid raised	No. of times bid unchanged	No. of times bid lowered
Positive outcomes	255 (27.2%)	364 (38.8%)	320 (34.1%)
No outcomes	1,583 (30.3%)	2,899 (55.5%)	742 (14.2%)
Negative outcomes	366 (19.6%)	483 (25.8%)	1,021 (54.6%)
Overall	2,204 (27.4%)	3,746 (46.6%)	2,083 (25.9%)

TABLE 8—OUTCOMES AND BID CHANGES IN SHIFTED 100VT

	No. of times bid raised	No. of times bid unchanged	No. of times bid lowered
Positive outcomes	102 (16.1%)	261 (41.2%)	271 (42.7%)
No outcomes	777 (45.7%)	603 (35.5%)	320 (18.8%)
Negative outcomes	135 (16.8%)	248 (30.8%)	422 (52.4%)
Overall	1,014 (32.3%)	1,112 (35.4%)	1,013 (32.3%)

On an individual level in the 100VT (shifted 100VT), patterns are somewhat more revealing. We observe 104 people in both treatments who decrease their bids more frequently (compared to the alternatives) after either positive or negative results, while nevertheless increasing their bids after no outcome. However, there are also 55 people who increase their bids after a positive outcome or no outcome, but reduce their bids after a negative outcome. There are also 42 people who keep their bids the same after a positive outcome, but lower them after a negative outcome and keep them the same after no outcome. Thus, different people respond differently to the various outcomes. Nevertheless, the result that people tend to raise their bids after not receiving an outcome is quite robust across the population. After no outcome, 195 people are more likely to raise their bids than to lower them. This is reversed for only 15 people.

Following up on these tables, we perform a formal analysis of the effect of the previous period's outcome on participant bidding. We present random-effects specifications in Table 9, for the change in the bid depends on the previous outcome in the original and shifted treatments.

The payoff outcome from the previous period affects one's bid. The worse the previous outcome, the more the current bid is reduced. In addition, bidders do tend to raise their bids significantly when these were "unsuccessful" in the previous period.

We also perform detailed regressions on the determinants of bidding behavior, which can be found in Appendix E. Summarizing these results, the optimal, 0-bid (20-bid) rate is higher in the 2VT, males make more optimal 0-bids (20-bids) and lower bids on average than females. More detailed instructions help with the optimal, 0-bid (20-bid) rate, but not with average bids in the 100VT. People who do better on our Bayesian updating questionnaire, which demonstrates a better feel for probability, have a higher 0-bid (20-bid) rate.

V. Conclusion

It seems that presently only a few economists doubt the existence of WC-type of behavior in the lab, and more and more are willing to accept that such behavior may

TABLE 9—BIDDING BEHAVIOR AND LAGGED OUTCOMES

Independent variables	Original Bid change	Shifted Bid change
Constant	−10.964*** (0.740)	−7.198*** (1.085)
2VT	4.191*** (0.542)	2.984*** (0.776)
Male	−1.823*** (0.559)	−1.527** (0.768)
Lagged outcome	0.033*** (0.008)	0.143*** (0.015)
Male × lagged outcome	−0.024* (0.014)	−0.151*** (0.019)
No profit or loss in previous period	20.573*** (0.565)	15.573*** (0.785)
Period	−0.043*** (0.016)	0.019 (0.046)
<i>N</i>	12,598	6,422
<i>R</i> ²	0.106	0.073

Note: Standard errors are in parentheses.

*** indicates significance at $p = 0.01$,

** indicates significance at $p = 0.05$, and

* indicates significance at $p = 0.10$ (all two-tailed tests).

exist outside the lab in real markets. It is important to study whether, and how, such departures from equilibrium and/or rationality are eliminated (mitigated) over time with experience, information structure (feedback), and institutional variations (rules of the auction or market).

We find that the winner's curse is alive and well, persisting even when we cast the Acquiring-a-Company game as an individual-choice task, thereby creating a one-person, adverse-selection problem. This device shows that the origin of the winner's curse does not lie in failing to take account of the cognition of other players or in current theoretical explanations involving beliefs about other players. Instead, it stems from a form of bounded rationality, as to a large extent decision makers fail to recognize that a future contingency is relevant for their current decisions. This is compounded by difficulties with Bayesian updating. We find that people who perform better on our questionnaire regarding such updating do significantly better in our data. In line with the notion of difficulties with contingent reasoning, providing more informative instructions has an appreciable effect on behavior. Finally, we observe a great deal of heterogeneity across the population.

As we noted earlier, both the Eyster and Rabin (2005) and Crawford and Iriberri (2007) models predict that people will bid zero in (a two-player version of) the 2VT (or 20 in the shifted 2VT), and so these predictions are partially correct. Avoiding the WC in the 2VT involves two steps. First, one should realize that bids other than 0 or 99 do not make sense; and the great majority of participants understand this. Second, one might expect decision makers to recognize that by bidding 99 rather than 0 (or 119 rather than 20 in the shifted 2VT), they are choosing a quite unattractive lottery,

with a 50 percent chance of losing 99 and a 50 percent chance of winning 49.5 (or, in the shifted 2VT, a 50 percent chance of losing 89 and a 50 percent chance of winning 59.5). Here is where many people seem to be unable to see the light and perform the contingent reasoning necessary to understand the meaning of these bids. When we test choices in our Lottery treatments, we find considerable improvement in terms of optimal bids. We find a limit of about 80 percent. It is possible that 20 percent of the population is actually risk seeking in this environment.³⁹

Our aim, in this paper, is primarily to demonstrate the persistence of the winner's curse in an individual-choice problem and the effect of simplifying the decision environment. We do not provide a comprehensive theoretical model, but we hope that our results will provide some insight into the origin of the winner's curse and stimulate better theoretical explanations of this pervasive economic problem.

APPENDIX A—BASIC INSTRUCTIONS

Welcome to our experiment. You will receive \$5 for showing up, regardless of the results. In addition, you will receive a \$15 endowment, with which you make bids on cards with some distribution of values.

You will see 100 cards on your screen, in a 10×10 array. They will have values that range between 0 and 99, inclusive. You will see all of these values before beginning, and will then click a button to flip the cards over and start the game. The cards are reshuffled after every action.

Each period you will choose a bid (a nonnegative integer) and will then select one of the cards; its value will then be shown.

If your bid is greater than or equal to the card's value, you receive 150 percent of this card's value. If your bid is less than the card's value, nothing happens; that is, you earn zero in such rounds.

If, for example, you bid 42:

1. Suppose the value of your selected card is 36, then its worth to you is $36 \times 1.5 = 54$; since your bid was 42, your profit is 12.
2. Suppose the value of your selected card is 20, then its worth to you is $20 \times 1.5 = 30$; since your bid was 42 your profit is -12 .
3. Suppose the value of your selected card is 67, then no transaction takes place and your profit is 0.

Your results will be displayed on the right side of the screen after you choose the card.

³⁹ However, anecdotal evidence (from talking with the subjects after the experiment) suggests that fewer of these risky gambles would have been chosen if the subjects were facing actual losses, rather than deductions from the "windfall" provided by the additional \$10 endowment. See Thaler and Eric J. Johnson (1990) for a discussion of this "house money" effect.

There will be a total of 60 periods in the session, consisting of two segments of 30 periods. The distribution of the card's values will change after 30 periods, and the new values will be shown to you. You will then click the button to flip them over and begin the second segment.

A history of the results of all previous periods (showing period, bid, value, and profit) will be displayed at the bottom of the screen.

At the end of the 60 periods, you will enter a name and then submit your results.

We will randomly choose one period from each 10-period block for payment, so that only six periods will actually count towards monetary payoffs. These periods will be chosen randomly at the end of the session, and will be shown on your screen. We will then net out your gains and/or losses from these six periods against your original endowment, converting them to actual dollars at the rate of \$0.10 for each unit.

At the end of the experiment, we will pay each participant individually and privately.

Thank you again for your participation in our research.

APPENDIX B—MORE DETAILED INSTRUCTIONS

(For the 2VT–100VT case, the instructions for the case of 100VT–2VT are similar, but reversed)

Instructions (1)

Welcome to our experiment. You will receive \$5 for showing up, regardless of the results. In addition, you receive a \$15 endowment, with which you make bids on cards with some distribution of values.

There will be 60 periods in this experiment. These instructions cover the first 30 periods. We will pass out further instructions at the end of 30 periods.

You will see 100 cards on your screen, in a 10×10 array. 50 of these cards have a value of 0, and 50 of these cards have a value of 99. You will see all of these values before beginning, and will then click a button to flip the cards over and start the game. The cards are reshuffled after every action.

Each period you will choose a bid (a nonnegative integer) and will then select one of the cards; its value will then be shown.

If your bid is greater than or equal to the card's value, you receive 150 percent of this card's value. If your bid is less than the card's value, nothing happens. That is, you earn zero in such rounds.

Your results will be displayed on the right side of the screen after you choose the card.

A history of the results of all previous periods (showing period, bid, value, and profit) will be displayed at the bottom of the screen.

SOME EXAMPLES

Suppose, for example, you bid 99:

1. Since the selected card's value cannot exceed 99, a transaction will take place in all cases. You earn positive profits whenever 150 percent of this card's value is

larger than your bid of 99, and you make negative profits (lose money) whenever 150 percent of this card's value is smaller than your bid of 99.

What would the outcomes be if the value drawn were 0 or 99?

Suppose, for example, you bid 42:

1. Suppose the value of your selected card is less than or equal to your bid of 42. In such an event a transaction will take place. You would earn positive profits whenever 150 percent of this card's value is larger than your bid of 42 and you would make negative profits (lose money) whenever 150 percent of this card's value is smaller than your bid of 42.
2. Suppose the value of your selected card is higher than 42. In this case, no transaction takes place and your profit is zero.

What would the outcomes be if the value drawn were 0 or 99?

Suppose, for example, you bid zero.

1. If the value of your selected card is greater than zero, no transaction takes place, and so your profit will be zero. If the value of your selected card is zero, then its worth to you is $0 \times 1.5 = 0$; since your bid was zero, your profit is zero.

What would the outcomes be if the value drawn were 0 or 99?

At the end of the 30 periods, please wait for further instructions.

We will randomly choose one period from each 10-period block for payment, so that only three periods of these 30 will count towards monetary payoffs. These periods will be chosen randomly at the end of the session and will be shown on your screen. The same will be done for three of the 30 periods in the second half of the session. At the end, we will net out your gains and/or losses from these six periods against your original endowment, converting them to actual dollars at the rate of \$0.10 for each unit.

At the end of the experiment, we will pay each participant individually and privately.

Thank you again for your participation in our research.

Instructions (2)

In this segment of 30 periods, you will once again see 100 cards on your screen, in a 10×10 array. However, now they will have values that range between 0 and 99, inclusive. You will see all of these values before beginning and will then click a button to flip the cards over and start the game. The cards are reshuffled after every action.

Once again, if your bid is greater than or equal to the card's value, you receive 150 percent of this card's value. If your bid is less than the card's value, nothing happens; that is, you earn zero in such rounds.

SOME EXAMPLES

Suppose, for example, you bid 99

1. Since your card's value cannot exceed 99, a transaction will take place. You earn positive profits whenever 150 percent of this card's value is larger than your bid of 99; and you make negative profits (lose money) whenever 150 percent of this card's value is smaller than your bid of 99.

What would the outcomes be if the value drawn were 0 or 22 or 50 or 99?

Suppose, for example, you bid 42.

1. Suppose the value of your selected card is less than or equal to your bid of 42. In such an event a transaction **WILL** take place. You earn positive profits whenever 150 percent of this card's value is larger than your bid of 42 and you make negative profits (lose money) whenever 150 percent of this card's value is smaller than your bid of 42.
2. Suppose the value of your selected card is higher than 42. In this case, no transaction takes place and your profit is 0.

What would the outcomes be if the value drawn were 0 or 22 or 50 or 99?

Suppose, for example, you bid 0:

1. If the value of your selected card is greater than 0, no transaction takes place, and your profit will be 0. If the value of your selected card is 0, then its worth to you is $0 \times 1.5 = 0$. Since your bid was 0, your profit is 0.

What would the outcomes be if the value drawn were 0 or 22 or 50 or 99?

APPENDIX C—LOTTERIES

I. Instructions

Below, you are asked to make a decision regarding which of four possible choices you would most prefer. Each choice involves one or more possible outcomes, and each of these outcomes occurs with some probability. These respective probabilities are clearly indicated. We ask that you indicate your choice by making a mark to the right of your preference.

Please note that there is a number in both the upper right- and left-hand corners of this sheet, which is unique for each copy. Please tear off one of these numbers and hold on to it. After we have collected all the questionnaires, we will randomly pick five of these numbers, and those whose decision sheets are selected will receive \$10 plus (or minus) their gains (or losses) from the selected choice. The chosen decision sheets will be given to the experimenter in the payment room. Please bring your number with you. The outcome for the indicated decision will be determined by a coin flip made privately in the payment room. The leftmost outcome will occur if the coin comes up heads, while the rightmost outcome will occur if the coin comes up tails.

Please do not look around at other people's decision sheets or engage in conversation with the other participants. Should you have any questions, please raise your hand and we will be happy to assist you.

Here are the four alternatives from which you must choose:

- A: Receive \$0.00 whether the coin comes up heads or tails. _____
- B: Receive \$0.00 with 50 percent probability and lose \$3.30 with 50 percent probability. _____
- C: Receive \$0.00 with 50 percent probability and lose \$6.60 with 50 percent probability. _____
- D: Lose \$9.90 with 50 percent probability and gain \$4.95 with 50 percent probability. _____

APPENDIX D—QUESTIONS ON BAYESIAN UPDATING

Consider two machines placed in two sides of large production hall, left side = *L* and right side = *R*.

The two machines produce rings, good ones and bad ones. Each ring that comes from the left machine, *L*, has a 50 percent chance of being a good ring and a 50 percent chance of being a bad ring. Each ring that comes from the right machine, *R*, has a 75 percent chance of being a good ring and a 25 percent chance of being a bad ring.

In each of the following four questions, you will observe some ring(s) that are the output of one of the two machines. After the information is given, you are asked to mark one (and only one) of the probabilities that you think is closest to the true probability that the *L* = left machine produced the ring(s), given your observations.

Each correct answer pays \$0.50.

Q1. You observe one *good ring*. What is your best assessment of the probability that the left machine produced this ring?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

Q2. You observe one *bad ring*. What is your best assessment of the probability that the left machine produced this ring?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

Q3. You observe six rings, and *all six are bad*. What is your best assessment of the probability that the left machine produced these rings?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

Q4. You observe six rings, of which *three are good* and *three are bad*. What is your best assessment of the probability that the left machine produced these rings?

1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0

APPENDIX E—ADDITIONAL REGRESSIONS

TABLE E1—REGRESSIONS FOR DETERMINANTS OF BIDDING BEHAVIOR

Independent variables	Dependent variable					
	Zero-bid (Probit) (1)	Zero-bid (Probit) (2)	Zero-bid (Probit) (3)	Bid made (Tobit) (4)	Bid made (Tobit) (5)	Bid made (Tobit) (6)
2VT	0.660*** (0.067)	0.636*** (0.065)	0.769*** (0.137)	12.37*** (0.78)	14.20*** (0.80)	12.37*** (1.51)
Male	0.387*** (0.138)	0.386*** (0.138)	0.471** (0.220)	−6.90** (2.93)	−13.50*** (1.33)	−15.29*** (2.04)
Segment not first	0.326*** (0.062)	0.022 (0.081)	0.047 (0.127)	−7.08*** (0.80)	5.56*** (1.43)	9.21*** (2.32)
More detailed instructions	0.468*** (0.149)	0.542*** (0.158)		1.43 (3.40)	−5.11 (3.78)	
Period		−0.004 (0.004)	0.003 (0.009)		0.450*** (0.063)	0.256** (0.124)
Period × last 30-period segment		0.018*** (0.004)	0.009 (0.009)		−0.924*** (0.087)	−0.634*** (0.167)
Period × four segments		0.027*** (0.009)	0.015 (0.013)		−1.175*** (0.094)	−1.039*** (0.158)
No. of questions correct			0.273** (0.118)			−15.95*** (0.90)
Constant	−1.362*** (0.118)	−1.224*** (0.132)	−1.649*** (0.248)	29.34*** (1.02)	38.32*** (1.26)	49.30*** (2.42)
ρ				0.385*** (0.010)	0.475*** (0.011)	0.350*** (0.015)
N	13,200	13,200	4,440	12,876	12,876	4,331
(Pseudo) LL	−7,456	−7,385	−2,331	−45,653	−45,584	−15,675

Notes: Standard errors are in parentheses. Period refers to the period in the regime. Last 30-period segment = 1 if it is the last half of a two-segment session, but is 0 otherwise. Bids > 100 were excluded in regressions (4)–(6). Regressions (1)–(3) are Probit regressions, using clustered standard errors at the level of the individual subject. Regressions (4)–(6) are two-sided Tobit regressions, with a lower limit of 0 and an upper limit of 100. We used clustered standard errors at the level of the individual subject.

- *** Significance at $p = 0.01$.
- ** Significance at $p = 0.05$.
- * Significance at $p = 0.10$ (all two-tailed tests).

TABLE E2—REGRESSIONS FOR DETERMINANTS OF BIDDING BEHAVIOR IN SHIFTED TREATMENT

Independent variables	Dependent variable			
	Twenty-bid (Probit) (1)	Twenty-bid (Probit) (2)	Bid made (Tobit) (3)	Bid made (Tobit) (4)
2VT			−17.58*** (1.55)	−17.56*** (1.55)
Male	0.549*** (0.181)	0.471** (0.220)	−9.06*** (1.70)	−7.81*** (1.76)
More detailed instructions	0.246 (0.181)	0.238 (0.179)	2.76* (1.67)	8.80*** (1.94)
Period	0.005 (0.004)	0.005 (0.004)	0.192*** (0.045)	0.192*** (0.045)
# Questions correct		0.308** (0.103)		5.49 (4.65)
Constant	−0.397** (0.168)	−0.538*** (0.175)	73.04*** (1.75)	68.24*** (2.27)
ρ			0.276*** (0.014)	0.265*** (0.036)
N	3,330	3,330	6,575	6,575
(Pseudo) LL	−2,209	−2,141	−31,571	−31,557

Notes: Standard errors are in parentheses. Period refers to the period in the regime. Last 30-period segment = 1 if it is the last half of a two-segment session, but is 0 otherwise. Bids > 120 were excluded in regressions (3) & (4). Regressions (1) & (2) are Probit regressions, using clustered standard errors at the level of the individual subject. Regressions (3) & (4) are two-sided Tobit regressions, with a lower limit of 20 and an upper limit of 120. We used clustered standard errors at the level of the individual subject.

- *** Significance at $p = 0.01$.
- ** Significance at $p = 0.05$.
- * Significance at $p = 0.10$ (all two-tailed tests).

REFERENCES

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84(3): 488–500.

Ball, Sheryl B., Max H. Bazerman, and John S. Carroll. 1991. "An Evaluation of Learning in the Bilateral Winner's Curse." *Organizational Behavior and Human Decision Processes*, 48(1): 1–22.

Bazerman, Max H. and William F. Samuelson. 1983. "I Won the Auction But Don't Want the Prize." *Journal of Conflict Resolution*, 27(4): 618–34.

Bereby-Meyer, Yoella, and Brit Grosskopf. 2008. "Overcoming the Winner's Curse: An Adaptive Learning Perspective." *Journal of Behavioral Decision Making*, 21(1): 15–27.

Blecherman, Barry, and Colin F. Camerer. 1998. "Is There a Winner's Curse in the Market for Baseball Players?" Unpublished.

Camerer, Colin F., 1999 "Rappresentazione mentale dei giochi: Framing e editing." *Sistemi Intelligenti*, 11, 113–43.

Capen, Edgar C., Robert V. Clapp, and William M. Campbell. 1971. "Competitive Bidding in High-Risk Situations." *Journal of Petroleum Technology*, 23(6): 641–53.

Casari, Marco, John C. Ham, and John H. Kagel. 2007. "Selection Bias, Demographic Effects, and Ability Effects in Common Value Auction Experiments." *American Economic Review*, 97(4): 1278–1304.

Cassing, James, and Richard W. Douglas. 1980. "Implications of the Auction Mechanism in Baseball's Free Agent Draft." *Southern Economic Journal*, 47(1): 110–21.

Charness, Gary, and Dan Levin. 2005. "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect." *American Economic Review*, 95(4): 1300–1309.

Cox, James C., Vernon L. Smith, and James M. Walker. 1987. "Bidding Behavior in First Price Sealed Bid Auctions – Use of Computerized Nash Competitors." *Economics Letters*, 23(3): 239–44.

- Crawford, Vincent P., and Nagore Iriberri. 2007. "Level-K Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 75(6): 1721–70.
- Dessauer, John P. 1981. *Book Publishing*. New York: Bowker.
- Dorsey, Robert, and Laura Razzolini. 2003. "Explaining Overbidding in First Price Auctions Using Controlled Lotteries." *Experimental Economics*, 6(2): 123–40.
- Eyster, Erik, and Matthew Rabin. 2005. "Cursed Equilibrium." *Econometrica*, 73(5): 1623–72.
- Garbarino, Ellen, and Robert Slonim. 2005. "Preferences and Decision Errors in the Winner's Curse." Unpublished.
- Grosskopf, Brit, Yoella Bereby-Meyer, and Max Bazerman. 2007. "On the Robustness of the Winner's Curse Phenomenon." *Theory and Decision*, 63(4):389–418.
- Hendricks, Kenneth, Robert H. Porter, and Bryan Boudreau. 1987. "Information, Returns, and Bidding Behavior in OCS Auctions: 1954–1969." *Journal of Industrial Economics*, 35(4): 517–42.
- Holt, Charles A., and Susan K. Laury. 2002. "Risk Aversion and Incentive Effects." *American Economic Review*, 92(5): 1644–55.
- Holt, Charles A., and Roger Sherman. 1994. "The Loser's Curse." *American Economic Review*, 84(3): 642–52.
- Kagel, John H., and Dan Levin. 1986. "The Winner's Curse and Public Information in Common Value Auctions." *American Economic Review*, 76(5): 894–920.
- Kagel, John H., and Dan Levin. 1991. "The Winner's Curse and Public Information in Common Value Auctions: Reply." *American Economic Review*, 81(1): 362–69.
- Kagel, John H., and Dan Levin. 1999. "Common Value Auctions with Insider Information." *Econometrica*, 67(5): 1219–38.
- Kagel, John H., and Dan Levin. 2002. *Common Value Auctions and the Winner's Curse*. Princeton, NJ: Princeton University Press.
- Kagel, John H., Dan Levin, and Ronald M. Harstad. 1995. "Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-Price Common Value Auctions." *International Journal of Game Theory*, 24(3): 293–319.
- Levin, Dan, John H. Kagel, and Jean-Francois Richard. 1996. "Revenue Effects and Information Processing in English Common Value Auctions." *American Economic Review*, 86(3): 442–60.
- Lorenz, J., and E. L. Dougherty. 1983. "Bonus Bidding and Bottom Lines: Federal Off-shore Oil and Gas." Society of Petroleum Engineers 12024, 58th Annual Fall Technical Conference. October 5–8, 1983, San Francisco, CA.
- Levis, Mario. 1990. "The Winner's Curse Problem, Interest Costs and the Underpricing of Initial Public Offerings." *Economic Journal*, 100(399): 76–89.
- Massey, Cade, and Richard H. Thaler (1995), "The Loser's Curse: Overconfidence vs. Market Efficiency in the National Football League Draft." Unpublished.
- Mead, Walter J., Asbjorn Moseidjord, and Phillip E. Sorensen. 1983. "The Rate of Return Earned by Leases Under Cash Bonus Bidding in OCS Oil and Gas Leases." *Energy Journal*, 4(1): 37–52.
- Rock, Kevin. 1986. "Why New Issues Are Underpriced." *Journal of Financial Economics*, 15(1–2): 187–212.
- Roll, Richard. 1986. "The Hubris Hypothesis of Corporate Takeovers." *Journal of Business*, 59(2): 197–216.
- Samuelson, William F., and Max H. Bazerman. 1985. "The Winner's Curse in Bilateral Negotiations." In *Research in Experimental Economics*. Vol. 3., ed. Vernon Smith, 105–37. Greenwich, CT: JAI Press.
- Selten, Reinhard, Klaus Abbink, and Ricarda Cox. 2005. "Learning Direction Theory and the Winner's Curse." *Experimental Economics*, 8(1): 5–20.
- Thaler, Richard H. 1988. "Anomalies: The Winner's Curse." *Journal of Economic Perspectives*, 2(1): 191–202.
- Thaler, Richard H., and Eric J. Johnson. 1990. "Gambling with the House Money and Trying to Break Even: The Effect of Prior Outcomes of Risky Choice." *Management Science*, 36(6): 643–60.
- Tor, Avishalom, and Max H. Bazerman. 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): 353–74.