The Price Is Right, But Are the Bids? An Investigation of Rational Decision Theory

By JONATHAN B. BERK, ERIC HUGHSON, AND KIRK VANDEZANDE*

The television game show The Price Is Right *is used as a laboratory to conduct a preference-free test of rational decision theory in an environment with substantial economic incentives. It is found that contestants' strategies are transparently suboptimal. In response to this evidence, simple rules of thumb are developed which are shown to explain observed bidding patterns better than rational decision theory. Further, learning during the show reduces the frequency of strategic errors. This is interpreted as evidence of bounded rationality. Finally, there is no evidence that a concern for fairness significantly alters bidding behavior.* (*JEL* C93, D44, C72, L10)

A common criticism of experimental economics is that the incentives in experiments are not large enough to induce subjects to incur the costs required to act optimally.[1] In response, experimentalists have designed ingenious techniques to address this issue. Ultimately, however, the best way to answer this criticism is to design an experiment in which the stakes are high.

In this paper we attempt to do just that. We conduct an experimental test of rational decision theory using the laboratory offered by the television game show *The Price Is Right.* We contend that this environment has all the advantages of a controlled experiment with an additional feature—the stakes on *The Price Is Right* are high enough to ensure that contestants have an economic incentive to play optimally.

In his Nobel Prize address, Herbert A. Simon (1979 p. 496) considers rational decision theory to be a paradigm in which fully rational agents optimize subject to the constraints placed on them by their environment. He contrasts this with theories of bounded rationality in which agents' capabilities are assumed to be much weaker: they may have incomplete knowledge of all available alternatives or computational abilities too limited to solve for the theoretical optimum.

The Price Is Right offers a unique experimental setting in which to test these theories. In this environment, by making the strong assumption of perfect rationality, we derive empirically verifiable predictions of the classical model of rational decision theory. Since models of bounded rationality do not rely on this assumption, deriving a similar empirically verifiable prediction of this paradigm is more dif-

* Berk: School of Business Administration, University of Washington, Box 353200, Seattle, WA 98195-3200; Hughson: David Eccles School of Business, University of Utah, Salt Lake City, UT 84112; Vandezande: Simon Fraser University, Faculty of Business Administration, Burnaby, BC V5A 1S6, Canada. The authors would like to thank Dan Bernhardt, Jim Brander, Rick Green, Steve Heston, David Hirshleifer, Burton Hollifield, Roger Ibbotson, Jon Ingersoll, Keith Murnighan, Daniel Loeb, Steve Ross, Tom Ross, Richard Sansing, and two anonymous referees for their insightful comments and suggestions, Venessa Brown, Chris Musto and Kate Ross for providing support in the data collection phase and especially Jim Storey both for his suggestions and his research assistance. Berk gratefully acknowledges financial support from the Richard D. Irwin Foundation and SSHRC (Grant No. 410-94-0725). The data on which this paper is based are freely available on Berk's home page: http://weber.u.washington.edu/~bschool/alpha.html.

[1] There is conflicting evidence on how the value of the prize affects the behavior of subjects in laboratory experiments. For example, Friedel Bolle (1990) contends that for small payoffs (that is, less than $100) changing the size of the prize only marginally affects behavior. However, Martin Sefton (1992) finds that even for small payoffs changing the size of the prize influences behavior and Steven J. Kachelmeier and Mohamed Shehata (1992) show that for large payoffs (on the order of three months' salary) changing the size of the prize has a significant effect on behavior.

ficult. However, we argue that whenever an agent who exhibits bounded rationality observes a strategy that is transparently welfare improving, he should adopt it. Presumably, the reason the agent did not initially use such a strategy was that he could not deduce it himself. Once the strategy is observed, this barrier is removed. Thus an empirically verifiable prediction of any theory of bounded rationality is that agents should adopt welfare-improving strategies upon observing them. The nature of the show makes this prediction testable.

The game show is described in the next section. To play for prizes on *The Price Is Right,* each contestant must win an auction that requires guessing the retail price of a consumer durable. This simple contest, which occurs six times on every show, is the focus of this study. In each auction, four contestants sequentially guess the retail price of a durable good worth about $1000. The contestant whose bid is closest to the retail price *without going over* wins the prize and plays in subsequent games for prizes worth as much as $60,000.

Three theoretical predictions of the fully rational model are derived in Section II. First, *any* contestant who bids fourth and who unconditionally prefers more to less must choose one of four bids. Next, in any equilibrium, there is a last mover advantage: the last bidder must win with a probability of at least one third. Finally, under symmetric information, the unique equilibrium features descending bids, that is, the first contestant bids highest, the next contestant bids second highest and so on.

The auction is also interesting because *The Price Is Right* is isomorphic to a variant of the classic location model in which firms compete by locating sequentially along a one-way street.[2] Thus, one may also interpret our results as an experimental test of this location model.

The empirical test of the fully rational model is in Section III. The most striking result is that in over half the sample, the last bid is transparently suboptimal. As long as bidders unconditionally prefer more to less, this result cannot be explained by appealing to any set of alternative preferences or informational assumptions.

In Section IV, we attempt to account for the observed behavior by postulating that contestants use simple rules or heuristics. To avoid data mining, we test these heuristics using new data. The behavior predicted by the heuristics better explains the data than does the fully rational model. We also find evidence that not all contestants use the same heuristic: contestants may use different strategies.

To examine why contestants might use strategies not supported by rational decision theory, in Section V we test two conjectures: 1) contestants exhibit bounded rationality, and 2) they derive utility from sources in addition to wealth maximization, in this case, a desire to be "fair." We do not find evidence that the strategies used on *The Price Is Right* result from contestants' desires to be "fair" in the sense of Matthew Rabin (1993). That is, there is no evidence that contestants' actions are motivated by a desire to increase others' material well being at their own expense. In contrast, we find convincing evidence that contestants improve their bidding after observing previous strategies: contestants appear to learn. This is consistent with the hypothesis that suboptimal bidding stems from the computational limitations of contestants, that is, bounded rationality.

Other researchers have also used television game shows as an experimental setting to study economic behavior. Robert Gertner (1993) uses *Card Sharks* to study risk aversion. Andrew Metrick (1995) uses *Jeopardy* to determine the frequency of best-response strategies. Finally, Randall W. Bennett and Kent A. Hickman (1993) investigate how economic agents process information by determining whether contestants on *The Price Is Right* optimally use their priors in formulating their bids.

[2] A review of the location literature is beyond the scope of this paper. The interested reader can consult Harold Hotelling (1929), B. Curtis Eaton and Richard G. Lipsey (1975) and Edward C. Prescott and Michael Visscher (1977). For an experimental test, see Jamie Brown-Kruse et al. (1993).

I. Description of the Game Show

For the hypotheses we examine, *The Price Is Right* has advantages over more conventional experimental settings. The rewards to participants are far greater than any prizes that have ever been awarded in a standard laboratory setting. In addition, contestants on the show are encouraged to weigh all available information. The host, Bob Barker, imposes no explicit time limit for placing a bid and he even reiterates the preceding bids if asked. Contestants are allowed to look to friends in the audience for guidance and the studio audience may shout out suggestions.

Contestants on *The Price Is Right* are chosen from the studio audience. Becoming a member of the studio audience requires waiting in line for about six hours. Screening of the audience is limited to a 15 second "interview" before taping begins.[3] Each audience watches only one show.

At the beginning of the show, four people are selected from the audience to sit on "contestants' row" and bid in an auction for a consumer durable worth about $1000. Contestants bid sequentially. A contestant must round his bid to the nearest dollar and may not submit the same bid as a previous contestant. The winner, who bids closest to the actual retail price *without going over,* receives the prize and the opportunity to continue on to the "pricing games." Prizes in the pricing games range in value from about $2,000 to $30,000. In the event that all four contestants overbid, they bid again in the same order for the same prize. This process is repeated as many times as necessary until a winner emerges. Unlike standard auctions, contestants are not required to purchase the prize—all prizes are free.

After each pricing game, a new contestant is selected from the audience. The new contestant replaces the last winner on contestants' row and bids first. The contestant on his left bids second, and so on. Thus, unless the first bidder wins, the order of bidding changes in the next auction. Six auctions are held in each one-hour show. All contestants who win an auction spin a wheel to decide which two will play in the final round of the show, the "showcase showdown," with prizes worth between $10,000 and $30,000. Ultimately, a contestant can win prizes worth as much as $60,000.

[3] Unfortunately the producers of *The Price Is Right,* Marc Goodsen Productions, refused to cooperate with us, so all information in this paper was either revealed on the air or observed during a live taping.

It is important to note that the nature of the show eliminates the problem that monetary values may not reflect a contestant's utility for the prize. Since contestants do not pay for their prizes, a bid should not be influenced by the bidder's private valuation. Furthermore, since the prize from the auction represents only a small fraction of each contestant's potential winnings, the nature of the product is unlikely to affect the contestant's motivation to win. Therefore, as long as contestants prefer more to less, they should bid to maximize the probability of winning the prize, regardless of their tastes or attitudes toward risk.

II. Theoretical Analysis of *The Price Is Right*

We now derive the behavioral predictions from the fully rational model. Define a round as one or more auctions where the four contestants bid for the same prize. In an auction, each contestant bids and the winner is the contestant whose bid is closest to the retail price of the prize without going over. If all contestants overbid, another auction commences and bidding continues until a winner is determined.

Define a "rational" bidder as any contestant who maximizes his probability of winning the round. Since the winning bid is closest to the price without going over, a contestant who bids $1 above a previous bid effectively reduces to zero the prior bidder's probability of winning. We refer to this strategy as "cutting off" another bid.

For simplicity, in the following analysis we assume that no contestant can rule out any positive price with certainty, and that bids may be arbitrarily close together rather than a minimum of one dollar apart. The following proposition characterizes a property that any equilibrium must have.

PROPOSITION 1: *A rational last bidder either cuts off a previous bid or bids zero.*

PROOF:

For any strategy in which the fourth bidder neither cuts off a previous bid nor bids zero, his probability of winning is increased by lowering his bid.

For convenience, we refer to any bidding strategy that is consistent with Proposition 1 as ''optimal.'' We next derive an implication, in this environment, of the assumption that contestants have rational expectations. We use as the definition of rational expectations what Steven M. Sheffrin (1983 p. 11) states is the most general statement of the rational expectations hypothesis: that ''... the subjective probability distributions of economic actors equal the objective probability distributions in the system.'' Here, this assumption implies that in equilibrium, any agent's subjective expectations about 1) the distribution of the price and 2) both his own and other agents' chances of winning, are the objective expectations. For concreteness, consider the following example.

Imagine an auction for a dishwasher in which contestants are asymmetrically informed. Suppose that the best-informed contestant has recently purchased a dishwasher, so that he knows the price of every model at every store in Los Angeles (where the prizes are priced). Since the exact model number of the prize is not usually revealed on the show,[4] this contestant's distribution is the distribution of prices of different dishwasher models at different stores in Los Angeles. Suppose that the next-best-informed contestant knows only the distribution of prices of kitchen appliances that appear on *The Price Is Right.* For her, the fact that the prize is a kitchen appliance is important, but she derives no additional information from knowing that it is a dishwasher. Finally, suppose that the other contestants are uninformed; they know only the distribution of the prices of products that appear on *The Price Is Right.* These contestants derive no additional information by observing that the prize is a dishwasher. If contestants have rational expectations, then 1) for each model dishwasher at each store in Los Angeles, the best-informed contestant must know the true probability that the producers use that dishwasher on the show, 2) for each kitchen appliance available in Los Angeles, the next-best-informed contestant must know the true probability that the producers use that kitchen appliance on the show, and 3) the uninformed contestants' distributions must be the true distribution of prizes used on the show. Observe that all three distributions are consistent in the sense that the next-best-informed contestant's conditional distribution (conditioned on the prize being a dishwasher) is the same as the best-informed contestant's distribution. Similarly, the uninformed contestants' conditional distributions (conditioned on the prize being a dishwasher (kitchen appliance)) is the same as the best-informed (next-best-informed) contestant's.

Thus, each contestant's belief about the distribution of the price is the true distribution of prices on *The Price Is Right,* conditioned on his own information. Similarly, each contestant's belief, conditioned on his own information, about both his own and other contestants' probabilities of winning must equal the true winning probabilities. Note that the rational expectations assumption places no further restrictions on contestants' beliefs. Contestants may be arbitrarily differentially informed, and may update from other bidders' actions, or even from the studio audience. The only requirement is that their beliefs be consistent. The following proposition, which is proved in the Appendix, relies exclusively on this assumption.

PROPOSITION 2: *Suppose that contestants have rational expectations. Then in equilibrium,*

1) *The fourth bidder must win at least as often as the third bidder; and the third bidder must win at least as often as either the first or second bidder.*
2) *The fourth bidder must win at least* $^1/_3$ *of the time.*
3) *The first and second bidders together cannot win more than* $^4/_9$ *of the time.*

When it is common knowledge that contestants are identically informed so that they all have the same distribution of the price of the prize, we show (in the Appendix) that there exists a unique equilibrium.

[4] For example, one contestant who was to bid on a VHS player asked: ''How many heads does it have?'' Barker did not provide this information.

TABLE 1—WINNING PERCENTAGE AND OVERBIDDING

	Contestant			
	1	2	3	4
Winning percentage	19.4	18.2	22.8	39.5
Percent of bids that exceed actual retail price	33.6	35.2	36.6	27.2

PROPOSITION 3: *If it is common knowledge that all contestants are identically informed, there is a unique equilibrium. In equilibrium, contestants bid in descending order. The first three bidders each win with probability $^2/_9$ and the fourth bidder wins with probability $^1/_3$.*[5]

Finally, it is worth noting that *The Price Is Right* is isomorphic to a particular problem in the location literature, that of firms that locate on a one-way street in a linear city. In such a city, consumers, who can only move up the street, purchase the product from the firm immediately up the street, even if the firm immediately down the street is closer. Consumers who are not serviced relocate randomly to a part of the street that is serviced. The decision problem for contestants on *The Price Is Right* is the same as that of firms locating to maximize market share. Although location theory is limited to games of complete information, the results are sometimes applied to real-world situations in which the complete information assumption may be violated. In light of this, if Proposition 3 is rejected, the appropriateness of using location theory in such situations could be questioned.

III. Empirical Results

An initial sample of 55 broadcasts was recorded on videotape and the results were then manually transcribed. A few shows were interrupted by news stories which left a total of 372 auctions, including 48 auctions in which all the contestants overbid.

TABLE 2—BIDDING-ORDER FREQUENCY

Bidding order (descending)	Occurrence frequency (percent)
1234	3.76
1243	2.42
1324	4.84
1342	2.69
1423	3.76
1432	3.23
2134	2.96
2143	4.30
2314	2.69
2341	2.69
2413	1.34
2431	4.03
3124	4.30
3142	3.50
3214	6.99
3241	3.76
3412	2.96
3421	4.57
4123	4.57
4132	3.76
4213	5.91
4231	3.23
4312	5.65
4321	12.10

We begin by testing the implications of Proposition 2, which relies on the rational expectations assumption. As shown in Table 1, the fourth bidder won more than $^1/_3$ of the rounds (39.5 percent) and both the first and second bidders won less than $^2/_9$ of the rounds (19.4 percent and 18.2 percent, respectively).[6] These frequencies are within the bounds specified in the proposition, so we cannot reject the hypothesis that contestants have rational expectations.

Proposition 3 predicts that if bidders have the same information and bid rationally, the first bid is highest, the rest follow in descending order, and the last is zero. This proposition is rejected by casual observation of Table 2: the observed ordering is strictly descending in only 3.8 percent of the auctions. Surprisingly, the strictly *ascending* ordering occurs in 12.1 percent of the auctions, significantly more frequently than the 4.17 percent that one would

[5] Daniel E. Loeb (1995) has independently derived the same existence result.

[6] We ignore auctions in which all contestants overbid.

TABLE 3—OCCURRENCE FREQUENCIES

Event	Frequency	Percentage
All bids less than (or equal to) the price[a]	156	41.94
First three bids less than (or equal to) the price[a]	163	43.82
Fourth bid is the highest[a]	131	35.22
Fourth bid is suboptimal[b]	183	56.48
A suboptimal fourth bid won[c,d]	56	30.60
Fourth bidder's likelihood of winning had he bid "optimally" instead of suboptimally[c,d]	79	43.17
Fourth bid is suboptimal and less than the other three bids:[b]	38	11.73
This fourth bid won[c,e]	4	10.53
This fourth bidder's likelihood of winning had he bid $1 instead of suboptimally[e]	9	23.68

[a] The percent is the frequency divided by the total number of auctions, 372.
[b] We restrict attention to the first auction in each round. The percentage is therefore calculated by dividing by the total number of rounds, 324.
[c] This includes cases in which the fourth bidder won in a later auction of the same round.
[d] The percentage is calculated by dividing by 183, the frequency that the fourth bid is suboptimal.
[e] The percentage is calculated by dividing by 38, the frequency that the fourth bid is suboptimal and is the lowest bid.

expect if all orderings were equally likely. Consequently, the hypothesis that the bidding orders occur equally often is convincingly rejected ($\chi^2(23) = 72.96; p < 10^{-6}$).

These results imply that the observed bids are unlikely to have been generated randomly. If bids were random, each bidder would win with the same frequency, and all bidding orderings would be equally likely. Yet the fourth bidder won significantly more auctions than the first three bidders,[7] and as shown above, the bidding orderings do not occur equally often. Thus the contestants appear to follow some strategy.

Since the implications of Proposition 2 cannot be rejected while those of Proposition 3 are convincingly rejected, it is tempting to conclude that the source of this rejection is the common information assumption. However, the most surprising result in this study comes from the test of Proposition 1 shown in Table 3. In more than half the sample auctions, the fourth bidder neither cuts off a prior bid nor bids zero.[8] This behavior is difficult to explain under the fully rational model because even under the most general assumptions, it is transparently suboptimal.

One might suspect that suboptimal bidding occurs so often because our definition of suboptimal is too strong: depending on the distribution of prices, a bid a few dollars above a previous bid might be no different from a bid of $1 above. However, the observed suboptimal bids had real consequences. Had fourth bidders who initially bid suboptimally reduced their bids to $1 above the next highest bid or zero, their winning frequency would have increased by almost half, from 30.60 percent to 43.17 percent (see Table 3). Note also that bidding a relatively low value is not equivalent to bidding zero. When fourth contestants' bids are lower than the others', but not zero, they win a mere 10.53 percent of the time. Had they bid zero instead, their winning percentage would have more than doubled to 23.68 percent (see Table 3). This shows that when the fourth contestant bids lowest but not zero, he is not bidding optimally.[9]

One might conjecture that if contestants derive utility from being on television, there

[7] The hypothesis that each contestant wins with equal probability is rejected ($\chi^2(3) = 34.7; p < 10^{-7}$).

[8] We operationalize cutting off as a bid not more than $1 above a previous bid and bidding zero as a bid of $100 or less. Operationalizing cutting off as a bid of not more than $5 above a previous bid did not qualitatively affect the results. We use $100 rather than $1 as the lower bound because it is not unreasonable to believe that prices are never below $100.

[9] Contestants also seem to make an interesting recurring mistake. Upon observing an optimal fourth bid sometimes contestants bid $1 *below* a preceding bid. Conditional on a previous optimal fourth bid, this mistake occurred 1 percent of the time. In contrast, it never occurred if there were no previous optimal bids.

might exist a disincentive to bid optimally. Losing allows a contestant to bid in subsequent rounds and so remain on television, and hence in the limelight, longer. However, winning an auction entitles a contestant to play a pricing game. Not only do contestants appear alone in pricing games, but they are on television for significantly longer than had they intentionally lost all six rounds. Therefore, a preference to be in the limelight should provide a stronger, not weaker, incentive to bid optimally.

In light of this evidence, the rejection of Proposition 3 cannot be merely a consequence of the common information assumption. Since Proposition 1, which does not require this assumption, is also rejected, we are compelled to infer that contestants use strategies that are not supported by rational decision theory.

IV. Alternative Bidding Strategies

Since rational decision theory cannot explain bidding even after allowing for the possibility that contestants are differentially informed, it is natural to ask what accounts for contestants' bidding strategies. In this section, we attempt to explain observed behavior by postulating that on *The Price Is Right,* differentially informed contestants bid using a few simple heuristics rather than fully rational strategies.

Although we have demonstrated that some contestants do not use fully rational strategies, they do appear to use previous bids as an input to their strategies. To see this, observe that the data in Tables 1 and 3 show that contestants underbid: they bid less than the actual retail price twice as frequently as they bid higher than the price. However, all four bids exceed the price in 12.9 percent of the auctions. If bids are independent, and contestants do not underbid on average, all four contestants should overbid in only 6.25 percent of the auctions. Systematic underbidding would reduce this probability still further. For overbidding by all four contestants to occur in 12.9 percent of the auctions, bids must be positively correlated. We take this as evidence that contestants use previous bids in making their bidding decisions.

The following important "stylized facts" therefore emerge from the data:

1) Contestants do not bid optimally, but neither do they bid randomly. They use some strategy.
2) Contestants use previous bids as an input to their bidding strategies.
3) Judging by the suboptimal behavior of the fourth bidder, at least half the time contestants do not use strategies that increase their likelihood of winning.

We use the first stylized fact to propose that contestants use one of three heuristic bidding rules: *simple* bidding, *sincere* bidding or *smart* bidding. Simple bidding is motivated by the last stylized fact. Simple bidders' assessments of the price of the prize require only their own information. Simple bidders ignore the information contained in other bids: their bids are simply what they expect the price to be.

We use the second stylized fact to construct the sincere bidding heuristic. Like simple bidders, sincere bidders bid the expected price given their information. Unlike simple bidders, as the auction proceeds, sincere bidders update their beliefs by using Bayes' rule to combine their own information with that contained in previous bids. Further, we assume that a fourth bidder sophisticated enough to bid sincerely is also savvy enough to cut off the next highest bid: a sincere fourth bidder reduces his bid from his expectation of the price to just above the next lowest bid or to zero.

Finally, we use the third stylized fact to justify the smart heuristic. Since the fourth bidder's problem is much easier than the others', we assume that smart fourth contestants bid to maximize the probability of winning the auction. Such a contestant computes the posterior distribution of the price as would a sincere bidder, but then chooses the bid that maximizes her probability of winning from the set detailed in Proposition 1. More detailed descriptions of the heuristics can be found in the Appendix.

On *The Price Is Right,* a contestant only wins if her bid is closest to the price without going over. Ignoring strategic considerations like cut-off bids, a contestant does better by

TABLE 4—WINNING PERCENTAGE AND OVERBIDDING FOR TWO HEURISTICS

	Winning percentage				Significance		Overbidding
Strategy followed	1	2	3	4	χ^2	p value	percentage
Actual (new sample)	16.0	17.5	25.2	41.3			11.6
No shading	19.5	19.9	19.4	41.2	8.9	0.031	14.9
25 percent shading	17.9	20.3	21.9	39.9	3.8	0.283	12.0

Notes: The first four columns show the frequency with which each bidder wins the auction. The statistics reported in the significance column provide a test whether the observed winning frequencies could have been generated by the heuristic. The last column shows the frequency with which all contestants overbid. The top line gives the frequencies from the new sample; the following lines are the simulation results.

bidding below the expected price. We therefore allow for the possibility that a sincere bidder *shades* her bid, that is, bids her expectation of the price less a fixed percentage of the standard deviation of her posterior distribution. It would make little sense to allow shading to drive a contestant's bid below the next highest bid, so we assume that in such a case, the shader cuts off the next highest bid.

We simulate the behavior of contestants who use these heuristic rules. Bidders start with the same prior distribution for the price. We envision this prior as representing contestants' beliefs before they observe the prize. The prize is then revealed to contestants and they each update the prior based on their knowledge of the product. Since the prior represents contestants' beliefs about the prize and since we assume that contestants have rational expectations, the price is drawn from the prior. Further details of the simulations can be found in the Appendix. For each simulation, we calculate the likelihood that each bidder wins, the likelihood of each bidding order, and the likelihood of overbidding. We then evaluate these strategies by testing their predictive power on *new* data.

To avoid data mining, we designed the simulations and picked the parameter values based on inferences made exclusively from the initial sample. We then created a second sample by recording 57 new broadcasts, a total of 381 new bidding auctions, including 44 auctions in which all contestants overbid. The predictions based on the simulations were then evaluated using the new data. Once testing started with the new data, no changes were made to the simulation programs or the parameter values.

The results of the simulations are given in Table 4 and Figure 1. None of the heuristics alone adequately explain the behavior observed on *The Price Is Right.* Better results are achieved by assuming that different contestants use different heuristics. The first simulation in Table 4 combines the sincere/smart strategy with no shading and the simple strategy. The first three contestants bid sincerely (without shading) with probability $1/2$ and simply with probability $1/2$. The fourth contestant bids smartly with probability $1/4$, sincerely with probability $1/4$ and simply with probability $1/2$. The winning percentages, as well as the frequency of overbidding, are quite close to the observed values. While the bidding order frequencies are also close, descending bids occur too often and ascending bids too infrequently (Figure 1). The fit improves with shading.[10] Indeed, a likelihood ratio test fails to reject the hypothesis that the actual winning frequencies were drawn from the distribution implied by this heuristic combination ($\chi^2(3) = 3.8, p = 0.283$). Unfortunately, we still reject the hypothesis that the actual bidding-order frequencies are drawn from this distribution. However, in light of the simplicity of the heuristics, the results of the simulations agree surprisingly well with the new data.

Although we provide evidence supporting the hypothesis that contestants use simple

[10] The first three contestants bid simply with probability $1/2$ and shade with probability $1/2$. They shade their bids by $1/4$ of a standard deviation. The fourth contestant bids smartly with probability $1/4$, sincerely with probability $1/4$ and simply with probability $1/2$.

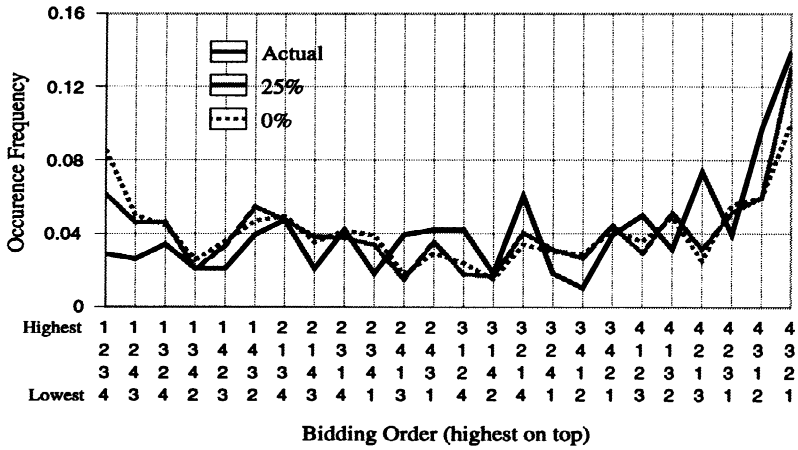

FIGURE 1. COMBINATION OF SINCERE AND SIMPLE BIDDING

Notes: The black line represents the observed bidding-order frequencies. The grey lines are the simulation results. In the simulation indicated by the broken grey line, contestants do not shade. In the simulation indicated by the solid grey line, they shade by $^1/_4$ of a standard deviation of their posterior distributions.

rules of thumb rather than the strategies predicted by rational decision theory, we have not yet investigated why contestants behave this way. In the next section we examine two alternative paradigms that might account for the observed behavior.

V. Bounded Rationality or Fairness?

Our empirical results show that contestants do not use fully rational bidding strategies. Our simulation results provide evidence that the bidding strategies that are actually employed derive from simple heuristics. Further, there appears to be a "round" effect: the fourth bidder is increasingly likely to bid optimally, that is, cut off an earlier bid or bid zero, later in the show.[11] To demonstrate this, we create a dummy variable, OPTBID, which is set to one in each auction where the fourth contestant makes an optimal bid. When OPTBID is regressed on six round dummy variables (Regression 1 in Table 5) the probability that the fourth contestant bids optimally differs across rounds ($\chi^2(5) = 19.34$, $p = 0.002$). Despite the presence of a selection bias,[12] this probability increases as the show proceeds.

There are (at least) two models of contestant behavior consistent with this evidence. In the first model, contestants exhibit bounded rationality. They initially use suboptimal decision rules, but after observing successful bids, learn to improve their strategies. Thus, contestants learn to be more rational. In the second model, contestants use optimal strategies from the start, but they derive utility from more than mere wealth maximization, in particular, they desire to be "fair."

In a recent paper, Rabin (1993) defines the notion of a fairness equilibrium. In this equilibrium, agents are willing to sacrifice their material well-being to help those who behave fairly and to punish those who behave unfairly. In the context of *The Price Is Right,* we take fairness to mean that each contestant has an equal chance of winning, regardless of his position in the bidding order. Hence, we assume that bids that ap-

[11] This effect was independently confirmed by Bennett and Hickman (1993).

[12] The selection bias results because contestants who play better are more likely to win and eliminate themselves from subsequent auctions. This effect should reduce the number of optimal bids later in the show.

TABLE 5—TESTS FOR LEARNING AND FAIRNESS

	1	2	3	4	5
Independent Variable	OPTBID	CUTOFF	Dependent Variable OPTBID	OPTBID	OPTBID
Constant			0.3643		0.3643
ROUND1	0.3243	0.1107		0.3243	
ROUND2	0.4603	0.1287		0.3893	
ROUND3	0.5366	0.1549		0.4110	
ROUND4	0.5349	0.1653		0.3662	
ROUND5	0.5581	0.1447		0.3722	
ROUND6	0.5641	0.1735		0.3729	
UNFAIR		0.0160[a]			
PREVOPT			0.2181	0.2033	
POPTW4O					0.2639
POPTW4N					0.1885
POPTL4O					0.2539
POPTL4N					0.2357
Observations	735	2205	735	735	735

Notes: Each column is a probit regression of the dependent variable that heads the column on the listed independent variables. The coefficients are reported as probabilities. All regressions use both the original data and the new data used in Section IV. All rounds following an interruption in a show were dropped. The variable definitions are as follows. CUTOFF is 1 if contestant is cut off, 0 otherwise; OPTBID is 1 if current bid is optimal, 0 otherwise; PREVOPT is 1 if at least one previous optimal bid by a fourth contestant on the same show, 0 otherwise; ROUND*n* is 1 if *n*th round, 0 otherwise. POPTW4O is 1 if at least one previous optimal winning bid by a fourth contestant on the same show and a contestant who previously bid optimally is bidding fourth, 0 otherwise; POPTW4N is 1 if at least one previous optimal winning bid by a fourth contestant on the same show and a contestant who has not previously bid optimally is bidding fourth, 0 otherwise; POPTL4O is 1 if at least one previous optimal bid but no previous optimal winning bid by a fourth contestant on the same show and a contestant who previously bid optimally is bidding fourth, 0 otherwise; POPTL4N is 1 if at least one previous optimal bid but no previous optimal winning bid by a fourth contestant on the same show and a contestant who has not previously bid optimally is bidding fourth, 0 otherwise; UNFAIR is 1 if the contestant had either (a) cut off another contestant who had never bid unfairly or (b) bid $1, 0 otherwise.

[a] $t = 0.703$.

pear to take advantage of a contestant's position in the bidding order, cut-off bids or bids of $1, for instance, are construed by other contestants to be unfair.[13] Under Rabin's concept of fairness, such bids invite retaliation against the unfair bidder in subsequent auctions. This retaliation may reduce the unfair bidder's subsequent chance of winning.

If the bidding behavior on *The Price Is Right* is influenced by fairness considerations, the first cut-off bid observed on a show should invite retaliation. Subsequent cut-off bids should then be observed that are retaliation against contestants who have previously bid unfairly. The implication is that a contestant who previously bid unfairly (that is, cut off a previous bidder who had not bid unfairly) is more likely to be cut off than a bidder who had not previously bid unfairly.

We test Rabin's fairness concept by comparing the frequency with which an unfair contestant is cut off to the frequency with which any other bidder is cut off. Two dummy variables for each contestant, UNFAIR and CUTOFF, are created. These variables are set to zero at the beginning of each show. Then, when any contestant either 1) cuts off another contestant who had never bid unfairly or 2) bids $1, we set UNFAIR to one for that contestant for the remainder of the show. When any contestant is cut off, we set CUTOFF to one for that contestant.

[13] Since the position of the fourth bidder is random, other contestants may not feel that the fourth bidder has a right to win when the other contestants overbid and therefore may view a $1 bid as unfair. A similar concept of fairness has been postulated to explain behavior in ultimatum games.

To increase the power of our tests, we combine our two samples for a total of 753 auctions. In each auction, we tally CUTOFF and UNFAIR for the first three contestants (the fourth contestant cannot be cut off). Then CUTOFF is regressed on UNFAIR and the six round dummy variables (Regression 2 in Table 5). The round variables are included to test whether fairness can account for the round effect. Because the complete history of the show is lost following an interruption, all rounds following an interruption in the show are dropped. This leaves 2205 observations.

The results of Regression 2 do not support the hypothesis that Rabin's fairness concept explains bidding observed on *The Price Is Right.* The coefficient on the UNFAIR dummy shows that bidding unfairly increases the probability of being cut off by a statistically insignificant 1.6 percent ($t = 0.70$).[14] This suggests that bidding on *The Price Is Right* does not reflect fair play.[15]

We next test whether a bounded rationality model is consistent with the empirical evidence. Simon (1979) notes that in classical models of rational choice, all alternatives, as well as the consequences of all alternatives, are assumed to be known. Under bounded rationality, contestants fall short of omniscience. It may be that they possess incomplete knowledge of the alternatives, or that they lack the computational ability to analyze the alternatives. However, relaxing the assumptions that characterize rational decision theory extracts a cost—the strong predictions of the theory. It is therefore difficult to test models of bounded rationality by deriving a specific set of predictions common to all. Instead, we argue that if contestants exhibit bounded rationality, when they observe strategies that improve their probability of winning, they should adopt these strategies. The theory of bounded rationality provides the explanation for why this strategy was not used initially: it is either beyond the pale of contestants' experience or computationally too difficult to derive. If agents observe a strategy that is transparently welfare enhancing it then becomes part of their experience. Under the bounded rationality paradigm, such a strategy should then be adopted.

On *The Price Is Right,* a welfare-improving strategy that not all agents initially adopt is optimal bidding by the fourth contestant. We therefore have the following testable prediction: the fourth contestant should bid optimally after having observed an optimal bid.

The round effect is consistent with the bounded rationality model because if contestants learn by observing better strategies, the frequency of suboptimal behavior should decrease during the course of the show. Hence, as the show proceeds, the fourth bidder should cut off other bidders more often. However, the model also makes a stronger prediction: this decrease in suboptimal behavior should result directly from contestants observing optimal fourth bidding. We therefore define a dummy variable, PREVOPT, which is set to one whenever there was a previous optimal bid by a fourth contestant on the same show. Regression 3 in Table 5 shows that the probability of observing an optimal bid by the fourth contestant is only 36.4 percent if no previous optimal bid has been observed, while it is 58.2 percent (constant + PREVOPT) when a previous optimal bid has been observed. This increase is statistically significant ($t = 5.73$, $p < 10^{-6}$).

There is additional evidence that the learning results mainly by observing optimal actions by other bidders. When the previous optimal bid dummy variable, PREVOPT, is included in a regression with the round dummies (Regression 4 in Table 5), the round effect disappears ($\chi^2(5) = 1.88$, $p = 0.86$). If the fourth contestant were to learn other than by observing previous optimal bids, then the round effect would persist.

Next, we examine whether contestants learn more from successful strategies than from unsuccessful strategies. That is, do they learn when they observe a strategy that wins or simply when they observe a strategy that they had previously not considered. In the former case, the probability of an optimal bid should increase only when a previous optimal bid wins. In the latter case, whether the previous optimal

[14] This result is not qualitatively sensitive to whether a bid of $1 is classified as unfair.

[15] Since the importance of fairness in Rabin's equilibrium concept declines as the payoffs increase, a different result may occur in a conventional laboratory version of the game.

bid won should not affect the subsequent probability of observing an optimal bid.

To test whether the success of previous optimal bids affects learning, we first partition rounds that follow an optimal bid into two sets. The set called POPTL contains the rounds in which the only previous optimal bid was a losing optimal bid; the set called POPTW contains all other rounds that followed an optimal bid. In either set it is possible that a contestant who bids optimally and loses will bid fourth again. This is much more likely to occur in the former case however, since by definition, a winning optimal bid implies that the fourth contestant will not participate in future rounds. Because a contestant making an optimal bid is likely to bid optimally again, it is important to control for the possibility that the fourth bidder might have bid optimally in the past. We therefore further partition each set into two subsets: POPTL4N and POPTW4N if the fourth contestant had never previously bid optimally; and POPTL4O and POPTW4O if the fourth contestant had previously bid optimally. Corresponding dummy variables for each of the four subsets are then constructed.

Regression 5 in Table 5 shows that contestants learn from observing previous optimal bids, but that they learn no more from winning bids than from losing bids. The probability of observing an optimal bid by a contestant who had not previously bid optimally is 36.4 percent if he had *never* seen an optimal bid (the constant in the regression), 60.0 percent after observing optimal losing bids only (POPTL4N + constant), and 55.3 percent after observing optimal winning bids (POPTW4N + constant). There is no evidence that contestants learn more from optimal winning bids than from optimal losing bids: the difference is −4.2 percent ($\chi^2(1) = 0.57$, $p = 0.45$). Merely seeing a better strategy is enough for contestants to adopt it, suggesting that contestants who did not use the better strategy originally did so because they were unable to deduce it themselves.

The importance of controlling for bidders who have previously bid optimally is clear. Presumably, a contestant who has previously bid optimally already understands the advantage of this strategy and so cannot learn by observing an optimal bid. Consequently, controlling for this effect provides a better test of the hypothesis that contestants learn from previous optimal bids. Contestants who have seen an optimal bid but have not previously bid optimally themselves bid optimally more often than contestants who have never seen optimal bids (that is, POPTW4N = POPTL4N = 0 is rejected, $\chi^2(2) = 24.77$, $p = 0.0001$).

Finally, one might conjecture that contestants learn better when they themselves are victims of cut-off bids. Therefore, after a cut-off bid has been observed, we examine the behavior of two groups of contestants who were bidding fourth for the first time. In the first group the contestant who was bidding fourth had been cut off in a previous round. In the second group the contestant who was bidding fourth had never been cut off. The probability of observing a cut-off bid in the first group is 63.1 percent, while the probability of observing a cut-off bid in the second group is 56.4 percent. Since this difference is statistically insignificant ($\chi^2(1) = 1.19$) we do not find evidence that contestants learn more when they themselves are cut off.

VI. Conclusion

We use the television game show *The Price Is Right* as a laboratory to conduct preference-free tests of rational decision theory. Although there is evidence that some contestants use optimal rules (the fourth contestant cuts off previous bids almost half the time), many fourth contestants use transparently suboptimal strategies. Moreover, there exist simple changes to the fourth contestant's behavior that would substantially increase his likelihood of winning. We also show that a few simple rules of thumb can better explain observed behavior than the fully rational game theoretic model.

We provide evidence that contestants tend to learn more sophisticated strategies as taping proceeds. This increase in strategic awareness occurs primarily because contestants observe other contestants using more sophisticated strategies. Interestingly, the success of the strategy is irrelevant. Merely observing the sophisticated strategy is enough for it to be adopted. This evidence is consistent with the hypothesis that behavior on *The Price Is Right* is explained by bounded rationality.

Our results indicate that rational decision theory cannot explain contestant behavior on *The Price Is Right.* Even when faced with relatively simple problems, we demonstrate that some (indeed most) contestants do not deduce the optimal strategy. This implies that we should observe situations in which agents with superior computational resources admit the possibility of suboptimal behavior by other agents when forming their strategies.[16]

APPENDIX

A. *Proofs*

PROOF OF PROPOSITION 2:

1) First note that if a bidder believes that his probability of winning exceeds that of another bidder, then since every bidder has rational expectations, this belief must be correct. Thus to show that a bidder wins more often than another bidder, it is sufficient to show that he believes that he wins more often.

First consider the last bidder. Since his optimal strategy is to pick the interval that he believes gives him the largest probability of winning, he must believe that he does at least as well as any previous bidder.

Next assume that the third contestant believes that a previous contestant, say i, has a strictly higher equilibrium probability of winning than he does. Let P_j be the third bidder's belief about the jth contestant's equilibrium probability of winning. Thus, $P_i > P_3$. This implies that contestant i's bid is immediately below that of the third contestant: there is no intervening bid. If this were not the case, then the third contestant could do strictly better by cutting off contestant i's bid since by cutting off contestant i, the third contestant is cut off no more often than contestant i would have been in the original equilibrium. Let Π_j be the third bidder's belief about what the fourth contestant thinks the jth contestant's probability of winning is, conditional on not being cut off by the fourth contestant. Thus Π_i is the amount of "probability mass" the third contestant believes the fourth contestant believes is between i's bid and the third contestant's bid. Note that by raising or lowering his bid, the third contestant can ensure that the amount of probability mass between his bid and the next highest bid is Π_i. If he were to do this, the amount of probability mass between i's bid and the third contestant's would then be Π_3. That is, from the third contestant's perspective, the fourth contestant's optimization problem is identical except that the ith bidder and the third bidder have switched positions (in terms of their respective winning probabilities). In light of this, since the third bidder has rational expectations, he must believe that by following this strategy, he will win with probability P_i. This contradicts the assumption that $P_i > P_3$.

2) The frequency with which the fourth bidder wins is the frequency which he believes that he wins the first auction of the round plus the frequency with which he believes that he wins subsequent auctions should all contestants overbid. Let x be the probability that he believes that all contestants overbid. Then, if the fourth bidder cuts off the previous contestant with the largest probability of winning, the fourth bidder wins the first round with probability not less than $(1 - x)/3$, and his probability of winning if he repeatedly cuts off a previous contestant is at least

$$\frac{1-x}{3} + x\left(\frac{1-\tilde{x}_2}{3} + \tilde{x}_2\left(\frac{1-\tilde{x}_3}{3} + \tilde{x}_3\left(\frac{1-\tilde{x}_4}{3} + \tilde{x}_4(\cdots)\right)\right)\right)$$

$$= \frac{1}{3} + \frac{x}{3} - \frac{x}{3} + \frac{x\tilde{x}_2}{3} - \frac{x\tilde{x}_2}{3} + \frac{x\tilde{x}_2\tilde{x}_3}{3} - \frac{x\tilde{x}_2\tilde{x}_3}{3} \cdots = \frac{1}{3}$$

where $\tilde{x}_i$ is what the fourth contestant believes is the probability that all contestants overbid on the ith repeated round of the auction. The fourth bidder chooses not to use this strategy

[16] Evidence of such behavior has been detected by Alvin E. Roth and Francoise Schoumaker (1983) and by Paul G. Straub and J. Keith Murnighan (1995). In a laboratory setting, they demonstrate that more sophisticated subjects can anticipate suboptimal play.

only if he can do weakly better, so his probability of winning must be at least $^1/_3$.

3) Suppose that the first two bidders together win more than $^4/_9$ of the time. Then, by (1), the third bidder must win at least $^2/_9$ of the time. But this implies that the first bidder must win less than $^1/_3$ of the time, a contradiction.

PROOF OF PROPOSITION 3:

Suppose the contestants bid as shown in Figure 2—that is, the first contestant bids c, the second b, the third a and the fourth L. We henceforth refer to each bid by its corresponding fractile. For example, we refer to the first contestant's bid as "$^7/_9$" and this bid gives her a $^2/_9$ probability of winning. We will show that these bidding strategies constitute a Nash equilibrium.

First we show that no contestant has an incentive to deviate by either cutting off a previous bid or bidding between two previous bids.[17] Consider possible deviations by the fourth contestant. Let p be the fourth contestant's probability of winning if she cuts off a previous bid. When she cuts off a previous bid, she wins the current round with probability $^2/_9$ and all contestants overbid with probability $^1/_3$. Therefore,

$$p = {}^2/_9 + {}^1/_3 p$$

so, solving for p, the fourth contestant wins with probability $^1/_3$ by cutting off a previous bid. The fourth contestant can therefore do no better by cutting off an earlier bid.

The third contestant, since she does not have to worry about being cut off, wins with probability $^2/_9$ if she bids $^1/_3$. She cannot do any better by cutting off a previous bid or bidding between two previous bids. Were she to cut off a previous bid or bid between two previous bids, the fourth contestant's optimal strategy is to continue to bid zero, and so the third contestant's probability of winning does not increase. Similar logic shows that the second contestant does not have an incentive to cut off the first contestant or bid between the first contestant's bid and U. Thus no contestant has an incentive to deviate by cutting off a previous bid or bidding between two previous bids. Since no contestant has an incentive to bid above the first contestant, he wins the auction with probability $^2/_9$ by bidding $^7/_9$.

Next note that the first contestant's probability of winning cannot exceed $^2/_9$. If it did, by Proposition 2 this would mean that the second contestant's probability of winning would be less than $^2/_9$. But were this the case, the second contestant could bid $^7/_9$ and guarantee herself a $^2/_9$ probability of winning. This follows because for the first contestant's probability of winning to exceed $^2/_9$ he must bid less than $^7/_9$, and the following contestants can do strictly better by not bidding above $^7/_9$. Thus no deviation by the first bidder can make him strictly better-off.

Since the first contestant has no incentive to deviate, the second contestant wins with probability $^2/_9$ by bidding $^5/_9$. Because the first contestant wins with probability $^2/_9$, Proposition 2 implies that the second contestant can win with a probability of no more than $^2/_9$. Thus no deviation by the second bidder can make her strictly better-off.

Similarly, the third contestant wins with probability $^2/_9$ by bidding $^1/_3$. By Proposition 2 he can do no better and thus no deviation can make him strictly better-off. Given these bids, the fourth contestant also has no incentive to deviate. Therefore, the bidding order in which the contestants bid, in descending order, $^7/_9$, $^5/_9$, $^1/_3$ and 0 is a Nash equilibrium.

Since the first bidder can always bid $^7/_9$ and win with probability $^2/_9$, in any other Nash equilibrium he must also win with probability $^2/_9$. If the first contestant does not bid $^7/_9$, we have shown that the second contestant can win with probability $^2/_9$ by bidding $^7/_9$. Therefore, in any Nash equilibrium, the first three contestants must win with probability at least $^2/_9$ and the last contestant must win with probability at least $^1/_3$. Since the first three bidders win with positive probability, they are not being cut off so (by Proposition 1) the last contestant must bid 0. This implies that in any Nash equilibrium, the first three bids must be some permutation of $^7/_9$, $^5/_9$ and $^1/_3$. However, the other five permutations allow deviations by at least one bidder that increase her probability of winning. Therefore, the equilibrium

[17] We consider a bid between the highest previous bid and U to be between two previous bids.

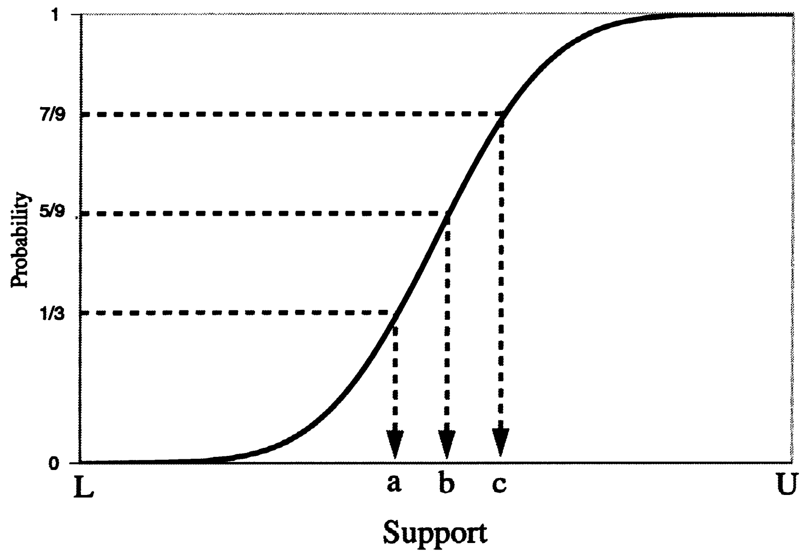

FIGURE 2. OPTIMAL BIDDING STRATEGIES

Notes: The graph shows the cumulative distribution function of the price of the prize. In the unique equilibrium, bids are strictly descending. The first contestant bids c, the second, b, the third, a, and the fourth, L. The fourth contestant wins with probability $^1/_3$ and others win with probability $^2/_9$.

described above is unique: there are no other pure strategy equilibria. Since no mixed strategy equilibria exist,[18] the equilibrium is unique.

B. *Heuristics*

Denote the price by random variable P with normal distribution $\mathcal{N}(\mu, 1/\tau)$, where τ is the precision. Let each contestant observe signal S_i, $i = 1, \dots, 4$ with distribution $\mathcal{N}(P, 1/\rho)$ where ρ is the precision. Assume that contestants use Bayes' rule to update their beliefs. Let V_i be the mean and ϕ_i be the precision of contestant i's posterior distribution for P. Let b_i be her bid. Finally, let B_i be the value of the maximum previous bid below V_i. If no such bid exists then let $B_i = 0$.

B.1. *Simple Bidding.*—Contestants use their prior beliefs and their signals optimally to determine V_i, but they ignore previous bids. Simple bidders bid their expectations of the price, $b_i = V_i$.

PROPOSITION 4: *Under simple bidding, contestant i bids*

$$b_i = V_i = \frac{\tau\mu + \rho S_i}{\phi_i},$$

where

$$\phi_i = \tau + \rho.$$

V_i is the expected price contestant i obtains by combining her prior information with her own signal using Bayes' rule.

PROOF:

Follows directly from Morris H. DeGroot (1970) Section 9.5, Theorem 1.

B.2. *Sincere Bidding.*—Having observed bids $\{b_1, \dots, b_{i-1}\}$, contestant i infers what $\{S_1, \dots, S_{i-1}\}$ must have been, assuming that

[18] A proof is available from the authors on request.

bids $\{b_1, \ldots, b_{i-1}\}$ were also sincere. The first three contestants each bid the mean of their posterior, which is what they expect the price to be. The last contestant cuts off the highest bid below his posterior mean if such a bid exists, or else bids zero.

PROPOSITION 5: *Under sincere bidding,*

$$V_i = \frac{\tau\mu + i\rho\bar{V}_i}{\phi_i}, \qquad i = 1, \ldots, 4$$

where

$$\bar{V}_i = \frac{1}{i}\left(S_i + \sum_{j=1}^{i-1} E\left[S_j \middle| \sum_{j=1}^{i-1} b_j\right]\right)$$

$$\phi_i = \tau + i\rho$$

$$\sum_{j=1}^{i-1} E\left[S_j \middle| \sum_{j=1}^{i-1} b_j\right] = \frac{b_{i-1}(\tau + (i-1)\rho) - \tau\mu}{\rho}.$$

For $i = 1, 2,$ or 3, the ith *contestant bids $b_i = V_i$. The fourth contestant bids $b_4 = B_4 + \varepsilon$, where ε is any small positive number.*

PROOF:

Follows directly from DeGroot (1970) Section 9.5, Theorem 1.

B.3. *Shading.*—

1) Contestant i, having observed $\{b_1, \ldots, b_{i-1}\}$ and S_i, calculates V_i as a sincere bidder would, assuming that previous contestants did not shade their bids.
2) If the contestant is not the fourth to bid, then:
   a) If $V_i - s_i\sqrt{1/\phi_i} > B_i$, she bids $b_i = V_i - s_i\sqrt{1/\phi_i}$.
   b) If $V_i - s_i\sqrt{1/\phi_i} < B_i$, she bids $b_i = B_i + \varepsilon$, where ε is a small positive number.
3) Since a sincere fourth bidder already bids $B_4 + \varepsilon$, shading leaves her strategy unaffected.

B.4. *Smart Bidding by Contestant Four.*— First, the fourth contestant assumes that all previous contestants bid sincerely and uses Bayes' rule to calculate the posterior distribution for P, that is,

$$\mathcal{N}\left(\frac{\tau\mu + 4\rho\bar{V}_4}{\tau + 4\rho}, \frac{1}{\tau + 4\rho}\right),$$

with a cumulative distribution function $\Phi_4(P)$. Next, let the previous three bids (sorted in descending order) be $\hat{b}_1$, $\hat{b}_2$, and $\hat{b}_3$. Then the fourth contestant wins with probability

$$\max\{1 - \Phi_4(\hat{b}_1), \Phi_4(\hat{b}_1) - \Phi_4(\hat{b}_2), \Phi_4(\hat{b}_2) - \Phi_4(\hat{b}_3), \Phi_4(\hat{b}_3)\}$$

by cutting off the bid at the lower end of the maximal interval.

C. *Simulations*

The price, P, is drawn independently from a normal distribution $\mathcal{N}(10, 1)$.[19] Signals S_i, $i = 1, \ldots, 4$ are drawn from a normal distribution $\mathcal{N}(P, \sigma^2)$, where σ^2 is itself a random variable drawn from a uniform [0, 1] distribution. That is, for some prizes, contestants have very precise signals, and for others, the signals are less precise and contestants rely more on their prior information. Each simulated auction was repeated 10,000,000 times to provide accuracy to three significant digits.

REFERENCES

Bennett, Randall W. and Hickman, Kent A. "Rationality and 'The Price Is Right'." *Journal of Economic Behavior and Organization,* May 1993, *21*(1), pp. 99–105.

Bolle, Friedel. "High Reward Experiments Without High Expenditure for the Experimenter." *Journal of Economic Psychology*, June 1990, *11*(2), pp. 157–67.

Brown-Kruse, Jamie; Cronshaw, Mark B. and Schenk, David J. "Theory and Experiments

[19] The mean of the prior is irrelevant to the analysis. We chose 10 to ensure that the probability of a price below zero was very small.

on Spatial Competition." *Economic Inquiry*, January 1993, *31*(1), pp. 139–65.

DeGroot, Morris H. *Optimal Statistical Decisions*. New York: McGraw-Hill, 1970.

Eaton, B. Curtis and Lipsey, Richard G. "Some New Developments in the Theory of Spatial Competition." *Review of Economic Studies*, January 1975, *42*(1), pp. 27–50.

Gertner, Robert. "Game Shows and Economic Behavior: Risk-Taking on 'Card Sharks'." *Quarterly Journal of Economics*, May 1993, *108*(2), pp. 507–22.

Hotelling, Harold. "Stability in Competition." *Economic Journal*, 1929, *39*(153), pp. 41–57.

Kachelmeier, Steven J. and Shehata, Mohamed. "Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the Peoples Republic of China." *American Economic Review*, December 1992, *82*(5), pp. 1120–41.

Loeb, Daniel E. "Bidding Optimally on The Price Is Right." *UMAP Journal*, 1995, *16*(2), pp. 123–35.

Metrick, Andrew. "A Natural Experiment in 'Jeopardy!'." *American Economic Review*, March 1995, *85*(1), pp. 240–53.

Prescott, Edward C. and Visscher, Michael. "Sequential Location among Firms with Foresight." *Bell Journal of Economics*, Summer 1977, *8*(2), pp. 378–93.

Rabin, Matthew. "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, December 1993, *83*(5), pp. 1281–302.

Roth, Alvin E. and Schoumaker, Francoise. "Expectations and Reputations in Bargaining." *American Economic Review*, June 1983, *73*(3), pp. 362–72.

Sefton, Martin. "Incentives in Simple Bargaining Games." *Journal of Economic Psychology*, June 1992, *13*(2), pp. 263–76.

Sheffrin, Steven M. *Rational Expectations*. New York: Cambridge University Press, 1983.

Simon, Herbert A. "Rational Decision Making in Business Organizations." *American Economic Review*, September 1979, *64*(4), pp. 493–513.

Straub, Paul G. and Murnighan, J. Keith. "An Experimental Investigation of Ultimatum Games: Common Knowledge, Fairness, Expectations, and Lowest Acceptable Offers." *Journal of Economic Behavior and Organization*, August 1995, *27*(3), pp. 345–64.