Available online at www.sciencedirect.com

ScienceDirect

Games and Economic Behavior 66 (2009) 830–840

GAMES and Economic Behavior

www.elsevier.com/locate/geb

When curiosity kills the profits: An experimental examination ☆

Julian Jamison [a,b,*], Dean S. Karlan [b]

[a] *University of Southern California, USA*
[b] *Economics Department, Yale University, USA*

Received 23 May 2005

Available online 27 September 2008

Abstract

Economic theory predicts that in a first-price auction with equal and observable valuations, bidders earn zero profits. Theory also predicts that if valuations are not common knowledge, then since it is weakly dominated to bid your valuation, bidders will bid less and earn positive profits. Hence, rational players in an auction game should prefer less public information. We are perhaps more used to seeing these results in the equivalent Bertrand setting. In our experimental auction, we find that individuals without information on each other's valuations earn more profits than those with common knowledge. However, given a choice between the two sets of rules, approximately half the individuals preferred to have the public information. We discuss possible explanations, including showing that there is a correlation between ambiguity aversion and a preference for having more information in the auction.

JEL classification: C91; D44; D81; D83

1. Introduction

There has been a recent surge of interest in economics concerning the study of different information structures. Consider, for example, gurus and advisors in the finance literature, cheap-talk and signaling in the game theory literature, and incompleteness in the contracting literature. A fascinating observation of the theory is that the value of information (to an informed party) can be negative in a strategic setting. While in a one-person decision problem it is necessarily the case that having more information increases one's expected payoff (at least weakly), this result can fail in strategic settings. It can be better to have strictly less information as long as the other players in the game know that this is the case. While not altogether surprising, this conclusion clearly runs counter to typical intuitions about the value of information. The purpose of this paper is to examine the public information version of this result in a specific experimental setting, a first-price auction (equivalent to a Bertrand duopoly). We test whether information makes

☆ We thank Jim Engle-Warnick, George Loewenstein, Richard Thaler, participants at an MIT Theory and Behavioral Lunch, an anonymous referee from the Russell Sage Behavioral Economics Roundtable, and especially two referees from this journal for useful suggestions concerning the design of the game. Jeff Butler was extremely helpful in programming the experiments. We thank the Russell Sage Behavioral Economics Program and Yale University for financial support. All errors are our own.

* Corresponding author at: University of Southern California, Los Angeles, CA, USA.
E-mail addresses: julison@gmail.com (J. Jamison), dean.karlan@yale.edu (D.S. Karlan).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2008.06.008

players worse off, and then we investigate individuals' preferences for the revelation of information. In particular, it is important for policy (including the design of institutions) to know if there is a preference for more information even when it is disadvantageous; and if so, to know whether this is a 'mistake' or if it is driven by the nature of the particular utility function.

When economics students first learn about Bertrand duopoly models, they often question the unique Nash equilibrium prediction,[1] which is for both firms to price at cost and earn zero profits.[2] Why not price somewhere above cost (which weakly dominates pricing at cost) and potentially make positive profits, with no risk of a loss? It is a legitimate question, and although the equilibrium stands, this illustrates the power of the assumption about common knowledge of other players' payoffs in such games. The same question appears in auction environments: if two bidders have the same value for the good (and this is commonly known), then they will end up bidding exactly that value. Under the more common assumption (mostly because it is theoretically more interesting) that values are not known, players bid below their value and both bidders earn positive payoffs in expectation. It may seem obvious in the auction setting that such common knowledge information is harmful to profits, but it is not always so transparent. Understanding similar environments is important to firms (and more generally to any players in these types of games), both when designing and influencing the institutions in which they will operate, and when making actual decisions about gathering and using information.

In this paper, we simulate a first-price auction game (formally equivalent to a unit-demand Bertrand oligopoly). By the logic above, subjects playing such a game should do better when they do not know each other's valuations versus when they do. We find that they earn higher profits with zero information, matching the theory, but that when asked their preferences, roughly half of the participants choose to play in the environment with more information and are willing to pay to do so. Hence they choose to decrease their earnings. We propose a hypothesis to reconcile this discrepancy: namely, that those particular subjects are ambiguity-averse. Ambiguity is distinct from risk, and applies not only when the state of the world is unknown, but also when the distribution over states of the world is unknown. Curious individuals presumably are averse to ambiguity since they seek information for the sake of information. The Ellsberg Paradox is the proto-typical example of ambiguity-aversion, though it focuses solely on a decision-theoretic setting. Support for our hypothesis comes from the fact that the more ambiguity-averse subjects (as measured by a sequence of Ellsberg-like decision problems) were likely to demand the full-information environment.

The paper proceeds as follows: Section 2 discusses some of the previous related literature, and Section 3 describes the experimental design. Section 4 presents the results, prefaced by some theoretical hypotheses to ground them. Section 5 concludes, and Appendix A provides details about the design parameters.

2. Literature

2.1. Auction theory

Auction theory is fairly well-developed for the familiar auction formats with basic assumptions (see, for instance, Milgrom and Weber, 1982). Recall that a first-price sealed-bid auction (FPA) is one in which bidders submit bids simultaneously and secretly; the highest bidder wins the object and pays his bid. Equilibrium bidding strategies involve bidding less than one's valuation in order to capture some surplus. Exact strategies depend on the expected distribution of the other bidders' values and on bidder preferences (e.g. risk-aversion). A second-price sealed-bid auction (SPA) is exactly the same, except that the winning bidder pays the second-highest bid rather than his own. Bidding exactly one's valuation is the weakly dominant strategy. The SPA is thus strategically equivalent to an English, or ascending-bid open outcry, auction, where bidders drop out at exactly their valuation. Furthermore, the SPA is also outcome-equivalent to a first-price auction in which bidders know each others' valuations (unlike above), since in that case the bidder with the highest value will simply bid at or marginally above the second-highest valuation.

The classic result in auction theory is the revenue equivalence theorem, which states that these standard auction formats produce equivalent (and optimal) expected revenue for the seller. Since they are all efficient as well, revenue equivalence from the seller's perspective implies that they are also cost equivalent for buyers. Revenue equivalence

[1] Indeed, this paper came to fruition because of an after class discussion between the authors when Jamison was the teaching assistant to Karlan in first year graduate game theory.

[2] This assumes equal and observable constant marginal costs.

holds under the following conditions: independent private values; symmetric prior distributions; and risk neutrality of the bidders. We maintain the assumptions of private values and symmetry, but we consider relaxing independence (and in some cases risk neutrality). In particular, if we drop independence and instead assume that values are "affiliated" (loosely speaking, this requires positive correlation to hold locally at every point in the support of the distribution), then the SPA produces more revenue than the FPA. Note that the SPA is still strategically equivalent (stronger than revenue equivalence) to the English auction even in this case.

For our purposes, since we are specifically interested in information *per se*, we run only first-price auctions, but in one case we inform players of each other's true values ('CK' for common knowledge) and in the other case we do not ('ZI' for zero information). As noted above, the CK model is outcome-equivalent to a SPA, while the ZI model is a FPA with no knowledge of the prior distributions. This latter assumption is unusual (again because theory has a limited amount that it can say concerning it), and one that we think warrants further study in general. In any case, this allows us to apply the theoretical results above to our setting. We point out formally here that a FPA is identical to a Bertrand oligopoly model with undifferentiated products and inelastic unit demand, with the same possible information structures as we have.

2.2. Auction experiments

For an initial survey of the vast and ever-expanding experimental work on auctions, see the book chapter by Kagel (1995). One of the main experimental results is that revenue equivalence does not seem to hold. More precisely, English auctions tend to converge quite quickly to the equilibrium outcome in repeated games, but there is systematic over-bidding in both first-price and second-price auctions (though it is considerably more pronounced in the SPA). Thus prices are higher in SPAs than they are in English auctions, so even strategic equivalence breaks down. Risk-aversion might help explain overbidding in the FPA, but nothing can explain overbidding in the SPA within the framework of the standard assumptions.

Experimental work has not focused yet either on the full ZI case (no information even about distributions of values) or on the full CK case (which is trivial theoretically). The case of affiliated private values has been studied by Kagel et al. (1987). Under risk neutrality, theory predicts that FPA prices should be lower than SPA prices, but risk aversion makes the effect ambiguous. Kagel et al. find that Nash equilibrium does a good job of organizing the data in the FPA, and find overall that seller revenue from the two formats is about the same. They find that public information about others' valuations does increase prices, but not by as much as would be predicted in a risk-neutral Nash equilibrium. Of course, our CK setting is not actually the same as a SPA experimentally, but that is certainly the closest environment that has been studied and we expect similar comparative statics relative to the FPA (our ZI). Kagel and Levin (1986) study public information in a common value environment, and find a mixed effect: it increases bids if there are few (3–4) bidders, but decreases bids with larger numbers (6–7) because it weakens a prevalent winner's curse bid error.

Few experiments have studied Bertrand competition directly. In the closest analogous environment ("posted-offers"; see Holt, 1995), the data support the Nash equilibrium outcome rather than the competitive outcome (Ketcham et al., 1984). Although theoretical auction predictions are not entirely borne out by experiments, there are empirical regularities. For instance, risk aversion appears to be present to some extent. Given risk aversion, affiliation moves revenue in the direction that theory predicts. Overall, Nash equilibrium appears to match the data more successfully than any simple ad hoc alternate models, however intuitively pleasing.

2.3. Ambiguity

As mentioned in the introduction, one possible reason that subjects chose the generally less profitable environment (i.e. CK) is that they place some inherent value on information *per se*, regardless of the implications for their payoffs. This can be formalized in the notion of ambiguity-aversion. Ambiguity was defined (ambiguously) by Frisch and Baron (1988) to be "uncertainty about probability, created by missing information that is relevant and could be known," while Camerer (1995) put it even more succinctly: "known-to-be-missing information." In essence, ambiguity aversion goes one step beyond risk aversion,[3] and in so doing poses a challenge for subjective expected utility theory (Savage,

[3] Sometimes ambiguity aversion is referred to as second-order risk aversion, as in, preferences over distributions of distributions.

1954). In a certain world, the state is known. In a risky world, the state is unknown but the probability of each state is known. In an ambiguous world, not only is the state unknown, but so is the distribution over states; possibly there are known probabilities for various distributions ('second-order' riskiness), but possibly not (e.g. no information at all).

The canonical thought-experiment dealing with ambiguity aversion is the Ellsberg Paradox (Ellsberg, 1961), one form of which is as follows: Urn 1 has 50 red marbles and 50 black marbles, for a total of 100. Urn 2 has 100 marbles that are either red or black, in some unknown distribution. One marble is chosen at random and the participant wins if red is picked. The subject chooses from which urn to draw. Ambiguity aversion predicts that the participant will prefer Urn 1, with a well-defined probability of winning of 50%. Furthermore, if the odds in Urn 1 are decreased, to 45% or even to 40%, many participants will still prefer the smaller but known probability for Urn 1 to the ambiguous probability of winning for Urn 2. Many decision-theoretic models have attempted to capture some aspects of ambiguity aversion, e.g. maxmin expected utility (Gilboa and Schmeidler, 1989) and non-additive models (Schmeidler, 1989 is one of several). Applications have been equally far-ranging, from finance to health to incomplete contracts. Of course, our auction game has more than one player, and less work has been done on understanding ambiguity aversion in strategic settings.

One counterexample to this trend is Chen et al. (2007) which is perhaps the closest to our study. They first present a theoretical model of auctions which allows for varying preferences toward risk and ambiguity, and then they approach the same question experimentally. In their theoretical model, ambiguity aversion leads to higher bids in a first-price auction.[4] Like us, they introduce ambiguity in the environment by considering a framework in which bidders do not know the distributions from which values are drawn.[5] Surprisingly, they find evidence of ambiguity-loving behavior (i.e. lower bids in the ambiguous auction). Unlike us, they do not directly consider common knowledge as one of the possible information structures, nor do they separately test subjects for ambiguity preferences.

Ellsberg's original paper presented his now-famous paradox as a thought-experiment only, but his intuition has been validated by many experiments since then.[6] These studies find that subjects tend to indeed be averse to ambiguity and are willing to pay an 'ambiguity premium' of roughly 10–20% in order to avoid it. This aversion is not a 'mistake' or lack of understanding of the question: Slovic and Tversky (1974) show that the result persists even after explaining the phenomenon to subjects. One interesting interpretation suggested by the work of Heath and Tversky (1991) is based on competence: expertise in the area of the ambiguous gamble tends to reduce ambiguity aversion (controlling for the level of ambiguity). This also has potential implications for ambiguity aversion in interactive settings with different perceived player skill levels.

In a world with ambiguity aversion, there can be a demand for information even if it is not going to affect the decisions that are made (i.e. simply for its own sake). For example, in medicine patients often want to know more about their conditions, but they do not want to make more decisions themselves: Strull et al. (1984) find that tests are often ordered that do not affect either the diagnosis or the treatment.

3. Experimental design

The experiment was conducted with 124 subjects in eight sessions, at the Experimental Social Science Laboratory (Xlab) at the University of California, Berkeley.[7] Participants were recruited by normal Xlab recruiting procedures (randomly chosen individuals from their database of interested participants). The experiment consisted of two parts: a series of incentivized questions that measured subjects' risk and ambiguity aversion, and then a simple two-player, sealed bid, first-price auction. All interactions were via computer, implemented using the z-Tree program.

In part I of the experiment, subjects were asked basic demographic information and then a series of decision problems to elicit their preferences (see Appendix A for the details). Subjects were asked several ambiguity questions and risk aversion questions. For the ambiguity questions, we used the standard Ellsberg urn questions: the participant had to choose between two urns, one with 50 red balls and 50 black balls and the other with an unknown proportion

[4] The intuition behind this conclusion is the same as is described in our Section 4.1 below.

[5] They also look at several other variations on standard auctions, including the use of reserve prices and subjects as auctioneers, and they directly compare first and second price auctions.

[6] See Camerer and Weber (1992) for an overview of the laboratory studies of ambiguity aversion.

[7] We previously ran a version of this experiment with 246 undergraduates at the University of Natal in Durban, South Africa. The design was not as careful, but the overall results were very similar.

of red and black balls. They were asked to bet on a red ball in Urn 1 or a red ball in Urn 2 (or they could chose "indifferent"). The question was repeated for 45 (and then 40) red balls and 55 (resp. 60) black balls in Urn 1.

At the end of part I, payoffs were calculated and shown on the subject's computer screen. Approximately 30 minutes had passed, and subjects were told that they could either take their current earnings (after waiting about 15 minutes for the checks to be processed) or stay and participate in a further experiment (adding to their earnings so far). Almost all subjects (87%) chose to remain.[8]

At this point, subjects were given detailed information (both verbally and on their monitors) on part II of the experiment—the auction task. In particular, they were told that they were bidding on behalf of their boss for a piece of art, and that in each round they would be given a maximum value for that particular round's artwork. They would be randomly matched with one other bidder but they would not know who the other person was. Each would place a bid, and whoever placed the higher bid won the object and kept the difference between their "valuation" (the amount their boss would recompense them) and their winning bid. Ties were broken randomly.

There were ten rounds, with new partners and valuations each round. No information about previous bids was made available. The valuations were drawn ahead of time from symmetric uniform distributions over a given range (e.g. from 45 to 65), although the specific range changed from round to round and was never known to the subjects (see Appendix A). Profits were computed as an average across all rounds.

At this point (still before playing the game), the subjects were told about the two possible information structures, which were referred to as A and B. In structure A, the valuations of both players were made known to both players, while in structure B each player was told only his/her own valuation. Examples were given, and subjects were asked if they had any questions.

Once everyone was ready to proceed, they were asked (via computer) to predict how many rounds they thought they would win (out of 10) and what percentile ranking they thought they would achieve in total profits (out of 100, with lower numbers corresponding to better performance). These two questions were not incentivized.

They were then asked if they preferred information structure A or B, and they were told that this would be incentivized so they should answer truthfully. The incentive system was explained (again both verbally and on the screens) as follows: Subjects were first asked to name a dollar value for the amount that their preferred structure was worth to them. They were told that they would be randomly placed into either environment A or B, but that if they got unlucky (in the sense of experiencing a different outcome from their previously expressed preference) they would be compensated by exactly the amount named; and if they got lucky they would be forced to pay the amount they stated out of their final earnings.

This created an incentive to truly express the strength of their preference, although it may have been difficult for them to accurately predict the size of their gross profits. However, when used as a relative measure the latter issue is less of a concern. They were allowed to enter amounts from $0 (if indifferent) to $10; we implemented a maximum for limited liability issues. All of this was carefully explained, and subjects were again given an opportunity to ask questions.

Finally, after entering their estimated strengths of preference (as measured in dollars), subjects proceeded to the auction task itself. As stated, they were split randomly into two groups (A and B) and were then randomly matched with partners in each of ten rounds, where of course the partner was constrained to be in the same group. After ten rounds, winnings were computed and participants were paid.

4. Results

4.1. Hypotheses

Our goal was not to test specific theories but to test and help to explain behavior in certain interesting environments. Hence this section is primarily meant to lay the groundwork for the interpretation of results, in addition to making various specific predictions.

[8] In the spirit of Levitt and List (2007), we gave individuals the opportunity to leave in order to analyze the selection decision to participate in an experiment (in this case, more experiments). In fact, across sessions we varied the amount of information available to the subjects regarding the exact nature of the second experiment. However, with only 16 individuals leaving, we did not have enough variation to analyze this second selection step.

We focus throughout on a two-bidder first-price auction with private values, where the latter phrase means only that bidders always know their own value. We will discuss below several possibilities regarding the underlying distributions that give rise to those values. Note that this setting is equivalent to one of Bertrand competition with inelastic demand, where firms know at least their own marginal cost. Unless specifically stated otherwise, we assume that agents are risk and ambiguity neutral.

We will refer to two distinct information environments in all sections of the paper. The first is ZI (or "zero information"; environment B above), which simply means that bidders are not told their opponent's valuation for the object. This is the standard assumption in the auction theory literature, although usually it is assumed that bidders at least know the mechanism according to which the values are derived (that is, the joint distribution used). It is likewise standard in the experimental literature to inform the subjects of the mechanism, although of course in the real world this is not often the case—and subjects have no trouble enacting actions (indeed, strategies) without a prior.

The second information environment is CK (or "common knowledge"; environment A above), in which both bidders are told both valuations for the object. The latter is obviously an unusual assumption, not because it is necessarily unrealistic but rather because it is theoretically uninteresting: if both agents have the same value, then the unique equilibrium outcome is for both to bid that value and earn no profits (again, as in Bertrand); if the values differ then the weaker player bids his value and the stronger player bids one minimal increment higher than that. We note here that this result does not depend on the distribution that was used to derive the values, nor on the risk/ambiguity preferences of the players. Further, it is the same as would be observed in a second-price auction.

Let us consider four mechanisms for generating the valuations:

(i) A single uniform draw from $[a, b]$ yields the common value for both agents;
(ii) Two independent uniform draws from $[a, b]$ yield the values for the two agents respectively.

In these two mechanisms we assume that there is no ambiguity, i.e. that in the ZI case the agents do at least know the true underlying distribution (as in the standard auction theory literature). If we added ambiguity (as discussed under (iii) below), it would only strengthen the case for ZI.

(iii) An ambiguous structure: one-half of the time the values are *iid* from $U[a_1, b_1]$, and one-half of the time *iid* from $U[a_2, b_2]$.[9] Here we assume $a_i < b_i$ and $b_1 > a_2$ so that there is true ambiguity in the intermediate range.

Hence mechanism (iii) is our analogue for the actual auction design used in our experiment. Given such a structure, we need to specify precisely how to interpret ZI, or rather how we model the players in the ZI setting. We will adopt the "principle of insufficient reason" enunciated by Luce and Raiffa (1957): since the agent has no information about the underlying distribution, they assume it to be uniform. The two natural possible domains for the perceived uniform distribution in this case are $[0, b_2]$ or $[a_1, b_2]$; we shall assume the latter since it is the worst-case scenario (i.e. highest bids and lowest profits) for ZI.

(iv) The same true underlying mechanism as in (i), but the agents know nothing about the distribution itself.

To be slightly more precise, we make no assumptions about how the agents model the distribution in (iv). This is an unusual case for which we can still draw clear conclusions without any such assumption, so it is worth mentioning. Again, recall that this is the most common environment used for analyzing Bertrand competition (simply adapted to an auction-theoretic setting).

We wish to compare profits in ZI versus profits in CK for each of these mechanisms, so that we can make some predictions about what might occur in the experiment and so that we can start to ask whether a potential player ought to prefer one or the other information environment.

Proposition. *Expected profits* (*for the bidders*) *are the same in ZI and in CK under mechanisms* (i) *and* (ii); *they are strictly higher in ZI than in CK under mechanisms* (iii) *and* (iv).

[9] Of course it doesn't matter that the probabilities are equal; any similarly ambiguous structure will yield the same results.

Proof. In (i), both players know that they have the same valuation all the time, and the distribution is irrelevant (as is the information). In all cases they each place bids equal to that value, as mentioned previously, and hence earn no profit.

In (ii), we use the fact that CK is outcome-equivalent to a second-price auction, while ZI is a standard first-price auction. (Formally, since we have continuous types, we invoke either infinitesimals or an endogenous tie-breaking rule to ensure that there is an equilibrium in the second-price auction, so that the stronger bidder reaps the full difference in valuations as profit.) Hence the revenue-equivalence theorem applies, since we are in the standard independent-private-values environment. But if the expected revenue is the same in the two cases then so must the expected profits of the bidders be.[10]

In (iii), we need to actually solve for some equilibria. First note that the symmetric Nash bidding function for a two-player first-price auction with *iid* values from $U[x, y]$ is to set $b = (v + x)/2$, where b is a player's bid upon observing a value of v. (This is straightforward to compute by assuming a linear function for one's opponent and solving for the best response.) So under CK, we can once more apply revenue equivalence to conclude that the total expected profits are the sum of the profits in the two underlying cases, i.e. what the bidders would have received from a first-price auction in which they were told the true domain and bid $(v + a_1)/2$ for draws from the lower distribution and $(v + a_2)/2$ for draws from the higher. Under ZI, however, we assume that they perceive a uniform distribution over $[a_1, b_2]$ and hence bid $(v + a_1)/2$ in all cases. Finally, note that all bid functions are monotonic and symmetric, so it is always the player with the higher valuation who wins the auction (i.e. all environments under discussion are efficient). Hence to compare profits we need only compare winning bids, and it is now clear that the winning bids under ZI are sometimes equal and sometimes strictly smaller than under the 'first-price equivalent' for CK. Thus ZI is indeed preferable from the bidders' point of view.

In (iv), CK again implies zero profits since the common valuation is known. Under ZI, we can't say with certainty what it is that players will decide about the world, but we do know that it is a dominant strategy to bid less than their value. So (in expectation), both bidders will make strictly positive profits. □

Of course, in (iv) it mattered that the underlying distribution produced common values. If instead the values were always quite far apart, then bidders might well prefer to know that (in order to reap the full benefits)—but this would naturally depend on what they believed (and hence how they bid) in the ZI case. Technically the comparison with no information at all is ambiguous for other mechanisms, but the qualitative conclusion is that if the values are positively correlated, then ZI is preferable. Since in most real-world situations values tend to be positively correlated (and agents certainly tend to believe this), we expect this result to hold generally. Similarly, although (iii) is only one instantiation of a specific ambiguous environment, the underlying rationale for preferring ZI will hold more generally.[11]

Thus risk and ambiguity neutral agents should generally prefer ZI to CK. As we are considering formally ambiguous environments, namely (iii) and (iv), we also consider ambiguity-averse agents. This has potentially two effects in each case: it changes bidding behavior (and hence profits), and it changes preferences over any given distribution of profits.

However, the results for mechanisms (i) and (ii) don't change, since there is no ambiguity in either of those environments. In fact, it could further be argued that mechanism (iii) is only a single complicated underlying joint distribution over valuations, but we consider it as presented: a distribution over distributions. One natural way of modeling ambiguity-averse preferences in this case is that agents will over-weight (relative to reality) the probability of the inferior (from their perspective) distribution.[12] This is precisely what seems to happen in the Ellsberg Paradox, and is roughly analogous to the fact that the certainty-equivalent of a risk-averse individual over-weights poor outcomes.

[10] Note that revenue-equivalence depends on risk neutrality, which we have indeed been assuming. With risk-aversion, standard results tell us that the seller prefers a first-price to a second-price auction, and hence in our setting that the bidders should prefer CK to ZI. This effect is ameliorated in environment (iii).

[11] Note that we can also solve for the equilibrium in (iii) under the assumption that players know the underlying distributions but nothing more (i.e. the standard auction environment). This introduces affiliation relative to the case where they know which uniform distribution is being used. The NE involves bidding intermediately for $v \in [a_2, b_1]$, but since players are more likely to win in that range when the true distribution is $[a_1, b_1]$—in which case they have 'overbid'—their expected profits are lower than in the baseline case (equivalent to CK). Which is to say: introducing more affiliation reduces profits, matching our overall trend.

[12] There is relatively little literature studying the implications of different forms of ambiguity-aversion in strategic situations, but this qualitatively fits, e.g., Gilboa and Schmeidler's (1989) 'set of priors' approach.

Table 1
Difference in average profits, by common knowledge or zero information.

	Played Common Knowledge game	Played Zero Information game	t-stat	Observations
Average profits	5.91	8.04	3.34***	105
Average profits if Common Knowledge preferred	5.59	8.01	3.24***	72
Average profits if Zero Information preferred	6.39	8.14	1.34	33
Profits in round one	8.3	14.93	2.98***	105
Profits in round ten	4.08	9.73	3.99***	105

Each observation in the above analysis is at the individual level. Profits are averaged over ten auction rounds. In Common Knowledge auctions, partners' valuations are revealed and Zero Information they are not. Common Knowledge (Zero Information) preferred indicates that player indicated a preference for Common Knowledge (Zero Information) before the auction began.
*** Significant at 1%.

In our case, it means that ambiguity-averse bidders will behave as if they believe their opponent's value to be higher than it (on average) actually is. We can imagine this as either a non-uniform distribution (e.g. for mechanisms (i) and (ii)) or as imputing a greater chance of a higher underlying uniform distribution (e.g. $[a_2, b_2]$ in mechanism (iii)). In either case, such bidders will bid higher than they would have had they not been ambiguity-averse, and will thereby make lower profits (since they are bidding closer to their valuations, all else equal). But they will 'expect' to do even worse than they actually do, since they imagine their opponents to have higher valuations than is in fact the case (so they think they won't win as often as they in fact do).[13]

All of this implies that the more ambiguity-averse an individual agent is, the less he will like the ZI setting (which is hardly surprising). Meanwhile, his preferences regarding CK (before the auction starts) incorporate only the second effect, namely that he calculates his utility as if overweighting the outcomes in which his opponent has a relatively high value. His actual bidding is unaffected, since he sees both valuations before making his bid (and indeed, we have already specified exactly how everyone bids under CK). So we predict that:

1. In ambiguous environments, profits will overall be higher in ZI than in CK;
2. But that ambiguity-averse individuals will, relatively speaking, prefer CK.

4.2. Experimental outcomes

Next we turn to the results from the experiment. Table 1 shows the basic summary results on average profits, and we find the expected result that ZI leads to higher average profits than CK. This is true for average profits as well as first round and last round profits. We also compare this difference for those who preferred CK and for those who preferred ZI and find similar differences (although the sample size of those who preferred ZI is sufficiently small that the differences are no longer statistically significant).

Table 2 presents the key results from the laboratory experiment. Columns 1 and 2 show that ambiguity averse individuals prefer CK information structures (unreported there, the correlation coefficient between the two measures from Table 1 is 0.234), and columns 6 and 7 show that the magnitude of this preference is also positively correlated, as expected, with ambiguity aversion. Columns 3 and 6 include controls for risk aversion, and the coefficients on ambiguity aversion remain unchanged, thus showing that our questions on ambiguity aversion are not simply picking up a correlation between risk aversion and preferences for more or less information.

Lastly, in Table 2 columns 4, 5 and 9, 10, we examine something quirky. We find that individuals who round off their bids to 5 or 10 also prefer common knowledge. An identical correlation was found in an initial study we conducted in South Africa with a similar framework. Although we do not want to over-interpret this result, it seems consistent with ambiguity aversion being a form of "thought-aversion." Common knowledge requires less thought, and just choosing simple round numbers is perhaps less-taxing to the mathematically challenged?

[13] The fact that ambiguity-aversion leads to higher bids is formalized in Chen et al. (2007).

Table 2
Ambiguity aversion as a predictor of preferences for Common Knowledge auction OLS.

	Prefers Common Knowledge					Strength of preference for Common Knowledge				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Most ambiguity averse	0.255**	0.380***	0.367**	0.399***		2.286*	4.110**	4.387**	4.422**	
	(0.016)	(0.008)	(0.010)	(0.005)		(0.096)	(0.026)	(0.018)	(0.015)	
Somewhat		0.221*	0.217*	0.233*			2.829*	2.917*	3.028*	
ambiguity averse		(0.084)	(0.090)	(0.066)			(0.091)	(0.080)	(0.066)	
Inconsistent ambiguity		0.026	−0.003	0.036			1.28	1.91	1.437	
preferences[a]		(0.859)	(0.983)	(0.807)			(0.511)	(0.336)	(0.452)	
Index of risk aversion[b]			−0.042					0.894		
			(0.360)					(0.134)		
Proportion subject's of				0.378*	0.327				6.131**	5.584**
bids ending in 5 or 10				−0.074	−0.135				(0.026)	(0.046)
Constant	0.625***	0.500***	0.600***	0.377***	0.589***	1.450**	−0.375	−2.51	−2.371	0.343
	(0.051)	0.000	0.000	(0.003)	0.000	(0.031)	(0.788)	(0.207)	(0.148)	(0.732)
Observations	105	105	105	105	105	105	105	105	105	105
R-squared	0.055	0.093	0.1	0.121	0.022	0.027	0.056	0.077	0.102	0.038

p values in parentheses. Prefers Common Knowledge indicates that before auction rounds that subjects preferred to play auctions where partners' bid were revealed. Valuation is their incentivized valuation of preferences for Common Knowledge (could be positive or negative).

[a] Inconsistent ambiguity preferences is a binary variable equal to one (19 out 108 respondents) if the responses to the three questions exhibited inconsistency, such as choosing the ambiguous urn over a known 50% probability of success, and then switching to the known urn when the probability of success is lowered to 45%, or expressing indifference between ambiguous and 45% probability of success, but then not switching to the ambiguous urn over a 40% probability of success.

[b] A lower number indicates a greater degree of risk aversion. The index is generated by counting the number of risky choices made in the risk aversion questions (and half a point for the first response in question two, see Appendix A).

* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

5. Conclusion

The fact that information can have a negative value in a strategic setting is well known, at least to economists. That is, it is sometimes the case that all players, if they behave optimally, would prefer less information on the table. In fact, it is possible that one player might individually prefer to have less information, as long as that fact is known to the other players. In this paper we explore a particular variant of this phenomenon experimentally. Specifically, in an auction game for which both players should theoretically prefer that private valuations not be common knowledge, we find experimentally that the players do earn higher profits without the information, but that many of them choose to have the information anyway. So the theory is confirmed, but either the players do not realize this or they have some reason to prefer the setting in which they enjoy lower profits. We suggest that ambiguity aversion explains this preference, and we provide evidence from standard behavioral games showing a link between ambiguity aversion and preference for full information in the competitive auction setting. Future experimental work may be able to better differentiate this rationale from competing hypotheses.[14] Here we put forward two potential measures, and find they are correlated as the theory predicted. This is both important substantively for understanding ambiguity aversion, and also useful methodologically for helping to assess the overlap of two potential measures of ambiguity aversion.

As far as the specific assumptions of our experimental model go, there are several limitations that we face. Our zero information framework gives the players no information about their rivals because we wanted the most extreme possible distinction from the public information case. Along the same vein, we were not interested in learning, which would confound knowledge of the distributions with pure preferences over the two environments. Finally, our ranking of the two possibilities only holds theoretically with positively correlated values. Certainly, a bidder who values the

[14] One possibility, suggested by a referee, would be to carefully conduct a repeated version of this game, with feedback on the profits from previous rounds. If subjects no longer choose CK after being told that on average they will earn less, this might imply that ambiguity aversion was not the original cause.

object considerably more than his rival may wish to know that in a first-price auction. We consider only the former environment, though we expect it to be empirically more relevant in the majority of cases.

Appendix A. Design parameters

After being assigned an ID number and consented, subjects entered their ID into their computers and were told to begin. They were first asked to give demographic data, and were told that they could skip any questions which made them uncomfortable. The requested data included: age; gender (male/female); college major (econ/other social science/humanities/math/physical sciences/business/other); father's education (less than HS/HS grad/some college/college grad/postgraduate); and mother's education.

Subjects then began part I of the experiment, the pure decision problems. The first set of problems was designed to measure risk aversion:

(1) [certain $9] vs [50/50 $3 or $20];
(2) [50/50 $18 or $5] vs [certain $10] vs [50/50 $25 or $2];
(3) [certain $2] vs [1/216 chance of $250, otherwise $0];
(4) [certain $8] vs [unknown chance of $23, otherwise $0];
(5) [certain $11] vs [50/50 $17 or $0].

Note that question (3) was designed to test for extreme gambles and question (4) involves ambiguity rather than risk.

The second set of questions was designed to measure social preferences. They all involved splits between the subject and a randomly chosen other participant in the experiment, with numerical values as follows:

(6) [$12/$0] vs [$10/$2] vs [$8/$4] vs [$6/$6];
(7) [$5/$5] vs [$4/$9] vs [$6/$2];
(8) [$4/$4] vs [$4/$5] vs [$4/$3].

Here question (6) is a standard dictator game; question (7) is a dictator game with multiplier greater than one; and question (8) is purely about relative payoffs.

The third and final set of decision problems in part I were Ellsberg-type ambiguity questions. Subjects were told that there were two urns with 100 balls, either red or black, and they were told the actual number of balls in Urn 1 (see below) but not Urn 2. They were then asked to bet in turn on red vs black in Urn 1; red in Urn 1 vs red in Urn 2; red vs black in Urn 2; and black in Urn 1 vs black in Urn 2. They could also choose "Indifferent" in each case. Winning a bet yielded $20, while losing yielded $4.

The number of red balls in Urn 1 was varied from 50 to 45 to 40 for each of the four question-types above, so there were 12 total questions in this section. Any subject who chose "Indifferent" was randomly identified with a response for that question when calculating payments. The actual number of red balls in Urn 2 (for the purposes of computing payoffs) was chosen uniformly between 0 and 100.

Subjects were paid the average of their earnings across the initial 20 questions, including (in the numerator of the average) any amounts received from their random partners in the social preferences questions. This came out to be between $9 and $10 for most participants, and was made known to them before they decided whether to stay for part II (as described in the text).

In the auction (part II), subjects were given new partners and valuations in each of the ten rounds. No information about previous bids was made available, including whether or not the subject had won in any given round. This was meant to reduce the possibility for learning (and for collusion) and to replicate a one-shot environment as closely as possible. The valuations were drawn once, ahead of time, from symmetric uniform distributions over a given range (e.g. from 45 to 65), although the specific range changed from round to round and was never known to the subjects, nor did they know that such a mechanism was even being used. Everything else about the environment was publicly known, including the fact that profits were again computed as an average across all rounds. The actual earnings were slightly lower in part II than in part I; see Table 1 for the exact numbers.

One advantage of using the same valuation-pairs for all subjects was that it increased the statistical power in the analysis, since more direct comparisons were possible. The ten instantiated pairs were as follows:

Round	$Value_1$	$Value_2$
1	65	63
2	32	38
3	83	77
4	50	49
5	36	23
6	42	40
7	24	20
8	63	78
9	25	29
10	71	75

References

Camerer, Colin, 1995. Individual decision making. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, NJ.

Camerer, Colin, Weber, Martin, 1992. Recent developments in modeling preferences: Uncertainty and ambiguity. J. Risk Uncertainty 5, 325–370.

Chen, Yan, Katuscak, Peter, Ozdenoren, Emre, 2007. Sealed bid auctions with ambiguity: An experimental study. J. Econ. Theory 136, 513–535.

Ellsberg, Daniel, 1961. Risk, ambiguity, and the Savage Axioms. Quart. J. Econ. 75, 643–669.

Frisch, Deborah, Baron, Jonathan, 1988. Ambiguity and rationality. J. Behavioral Dec. Making 1, 149–157.

Gilboa, Itzhak, Schmeidler, David, 1989. Maxmin expected utility with a non-unique prior. J. Math. Econ. 18, 141–153.

Heath, C., Tversky, Amos, 1991. Preference and belief: Ambiguity and competition in choice under uncertainty. J. Risk Uncertainty 4, 5–28.

Holt, Charles, 1995. Industrial organization: A survey of laboratory research. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, NJ.

Kagel, John, 1995. Auctions. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, NJ.

Kagel, John, Levin, Dan, 1986. The winner's curse and public information in common value auctions. Amer. Econ. Rev. 76 (5), 894–920.

Kagel, John, Harstad, R.M., Levin, Dan, 1987. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. Econometrica 55, 1275–1304.

Ketcham, Jon, Smith, Vernon, Williams, Arlington, 1984. A comparison of posted-offer and double-auction pricing institutions. Rev. Econ. Stud. 51, 595–614.

Levitt, Steven D., List, John A., 2007. What do laboratory experiments measuring social preferences reveal about the real world? J. Econ. Perspect. 21 (2), 153–174.

Luce, R. Duncan, Raiffa, Howard, 1957. Games and Decisions. Wiley, New York.

Milgrom, Paul, Weber, Robert, 1982. A theory of auctions and competitive bidding. Econometrica 50, 1089–1122.

Savage, Leonard J., 1954. The Foundations of Statistics. Wiley, New York.

Schmeidler, David, 1989. Subjective probability and expected utility without additivity. Econometrica 57, 571–587.

Slovic, Paul, Tversky, Amos, 1974. Who accepts Savage's axiom? Behavioral Sci. 19, 368–373.

Strull, William, Lo, Bernard, Charles, Gerald, 1984. Do patients want to participate in medical decision making? Journal of the American Medical Association 252, 2990–2994.