The Common-Probability Auction Puzzle*

M. Kathleen Ngangoué[†] Andrew Schotter [‡]

July 2022

Abstract

This paper presents a puzzle of the behavior of experimental subjects in what we call common-probability auctions. In common-value auctions, uncertainty is defined over values, while in common-probability auctions, uncertainty is defined over probabilities. We find that in contrast to the substantial overbidding found in common-value auctions, bidding in strategically equivalent common-probability auctions is consistent with Nash equilibrium. In our experiments, we isolate the different steps of reasoning involved in the bidding process and conclude that in competitive environments, the difference in bids across our two auctions stems from differences in the way that subjects estimate the objects' value they are bidding for rather than the way they bid conditional on these valuations.

JEL-Classification: D44, D81, C70, C90
Keywords: Common-value auction, winner's curse, uncertain values, uncertain probabilities, compound lotteries, bid, strategic uncertainty, motivated reasoning.

---

*This paper benefited from helpful discussions with David Budescu, Guillaume Frechette, Johanna Mollerstrom, Ryan Oprea, Debraj Ray, Ariel Rubinstein, Charles Sprenger, John Wooders, Will Zame and numerous seminar audiences. We thank Anwar Ruff for his programming expertise as well as Sofia Correa, Peter Li, Sisi Huang, Riley Conlon, Ian Leefmans, Xuemeng Yin for valuable research assistance.

[†]UCLA Anderson, 110 Westwood Plaza, Los Angeles 90095, kathleen.ngangoue@anderson.ucla.edu.

[‡]New York University, 19 West 4th St, 6th floor, NY NY 10012, andrew.schotter@nyu.edu

1 Introduction

In the typical common-value auction, bidders bid for a good whose value is not precisely known but is the same to all bidders. Bidding is then based on a privately informed appraisal of the good's expected value. Each bidder receives a signal about the good and, based on that signal, makes a bid. The key feature of this Bayesian game is that, when forming expectations, bidders face some uncertainty that is common to all bidders. Yet, there has been little discussion of auctions in which the main uncertainty refers to a probability rather than to a value. For example, consider firms bidding for bonds issued by a corporation under financial stress. Here, the value of the bond at maturity is known, but what is uncertain is the probability of default by the corporation. If investors do their due diligence, they will receive a signal about the common default risk drawn from a commonly known distribution and, based on this probability signal, make a bid for the bond. In such situations, what is uncertain is, first and foremost, a common probability and not a common value.[1]

In this paper, we ask whether in auctions, bidders process these two objects of uncertainty (values and probabilities) in the same way. In other words, do bidders, when facing two strategically equivalent common-value (CV) and common-probability (CP) auctions, submit identical bids? Or does the fact that one auction exhibits uncertainty in the value domain while the other exhibits it in the probability domain lead to differences in bidding behavior?

We find that subjects approach our two versions of the Bayesian game very differently. First, in contrast to the theoretical predictions, our subjects bid significantly more in CV than in equivalent CP auctions. More specifically, while our subjects in CV auctions tend to bid above the naïve bidding function (i.e., bidding the expected value given the signal), subjects in CP auctions tend to bid below the even lower Nash equilibrium bid function. As a result, winning bidders are less vulnerable to the winner's curse in the CP than in the CV auction. Why this difference exists is the puzzle we wish to unravel in this paper.

[1]Other auctions in which a probability constitutes a major uncertainty are auctions for artwork with dubious provenance, for which collectors may have a precise assessment if the object has a certificate of authenticity but where the risk of facing a counterfeit or stolen object may be the principal uncertainty.

In the different treatments of our experiment, subjects face either CV or CP auctions and bid for equivalent items. In CP auctions, we present subjects with a random asset whose positive value is known but where the probability of receiving that value is not, while in CV auctions, our subjects face a different asset whose failure risk is known but whose value in the case of success is uncertain. These assets define lotteries for which our subjects bid and are strategically equivalent, in that bidders should have identical expectations conditional on equivalent probability and value signals and, hence, bid identically.

To investigate our puzzle, we break down the bidding process into two stages, one dealing with fundamental uncertainty (Stage 1) and the other with strategic uncertainty (Stage 2), and run experiments attempting to identify which stage is responsible for our puzzle. In Stage 1, bidders have to calculate their subjective valuation for the good given a signal. This task requires them to solve the problem of fundamental uncertainty, as they have to estimate how much the good is actually worth. Once they have formed their subjective estimate, in Stage 2, they then have to convert their subjective valuation into a strategic bid, taking into account how others will bid. This second task deals only with strategic uncertainty. The difference in bidding may emanate from either or both of these two stages. In Stage 1, subjects may value the lotteries with uncertain values differently than those with uncertain probabilities, conditional on equivalent signals. Alternatively, given identical valuations in Stage 1, in Stage 2, bidders may view the strategic uncertainty across our two settings differently and bid accordingly.

Our experiments reveal that our puzzle is more complex than anticipated. Differences in bids originate by and large in the first stage of the bidding process, but fundamental uncertainty alone is not sufficient to generate the discrepancy. Concretely, subjects resolve uncertain values and probabilities differently–but only in a competitive setting. With uncertain probabilities, expectations about the lottery's payoff are mostly accurate, but of particular interest is the way that subjects contemplate uncertain values. Uncertain values give rise to misjudgments in a systematic manner: In non-competitive settings, subjects value CV lotteries close to their expected payoff but, in the auction, they expect them to pay off substantially more.

We came to this conclusion by running four experiments designed to study bidding behavior and its components in isolation. Experiment I establishes our main result on the asymmetry in bidding behavior across our CV and CP auctions. To sort out whether this difference emerges with fundamental uncertainty, we ran Experiment II, in which subjects were asked to evaluate (price) the same lotteries underlying our CV and CP auctions but not bid for them in an auction. Interestingly, stripping the auction game of its strategic elements revealed that the observed difference in bids cannot be attributed solely to fundamental uncertainty. When asked to price lotteries, subjects estimated both value and probability lotteries at their objective expected payoff.

We then ran two more experiments to study the auction context specifically. Experiment III investigated how bidders resolve the Stage-1 problem of fundamental uncertainty in an auction (competitive) as opposed to the (non-competitive) decision problem studied in Experiment II. This third experiment, in which we elicited a variety of fundamental estimates within and outside of the auction context, provides us with two main insights. First, differences in bids originate in Stage 1 of the bidding process, in that subjects come up with different estimates for equivalent CV and CP lotteries. Second, subjects must be placed in an auction context for these differences in expectations to emerge.

Experiment IV focused on Stage 2 and studied how bidders deal with the strategic uncertainty inherent in the auction game. Here, we found that strategic uncertainty does not offset, but further sustains, the differences that arise in Stage 1. Our subjects in CV auctions expected their competitors to overbid, which prevented them from bidding below their exaggerated estimates of the lottery's expected payoff. In contrast, in CP auctions, in which they had substantially lower estimates, they underestimated the competition. Finally, Experiment IV further confirmed that disparities between our CV and CP auctions emerge in Stage 1 because helping our subjects resolve fundamental uncertainty reduced the extent of overbidding in CV auctions and, hence, the differences in bids between auctions.

To appreciate the implications of our paper for auction design, it is important to understand that in the real world, it is hard to find examples of pure common-value, common-probability, or even pure private-value auctions. Bidders mostly face uncertainty over both values and probabilities, and, if an aftermarket exists for private goods, they too have a common-value component. Usually, however, one object of uncertainty prevails over the other. For example, consider a real estate developer who is thinking of bidding on a piece of rural land with the hope of building a shopping center. The developer knows the value of the land for that use, but the shopping center will be viable only if a highway being discussed by the state highway commission is built. If the highway is not built, the land remains farmland. The uncertainty here is primarily over the probability of the highway being built and not its value contingent on that happening (although that, too, may be uncertain).[2] To include such ex-post uncertainty into the sale of these contingent products, it is possible to auction off contingency contracts, and a recent and interesting literature on such auctions with contingent bids exists (see, e.g., DeMarzo et al., 2005). Still, in many cases, goods are auctioned with all-cash bids–i.e., without the use of such contracts.

Since most auctions have both probabilistic and value uncertainty, the way they are perceived by bidders can be influenced by the way the auction is presented to them. Descriptions that put relatively more weight on value uncertainty can be expected to increase revenues (with a higher incidence of the winner's curse), while emphasis on probabilistic uncertainty will tend to mitigate the winner's curse at the expense of revenues. Since bidders may shy away from auctions in which they repeatedly suffer from the winner's curse, stressing the probabilistic aspect of the auction, while reducing revenues in the short run, may increase entry and, hence, revenues in the long run.

The paper is divided in two main parts. In the first part, we briefly discuss the most closely related literature in Section 2 and then in Section 3, present the main Experiment I with its design, the theoretical hypotheses and the experimental results, which constitute our main puzzle. In the second half of the paper, Section 4, we tackle the puzzle by exploring potential explanations with the additional Experiments II to IV. We finally conclude with our interpretation of the findings in Section 5.

[2]Other contingent entities might be firms that create a new product whose value is contingent on their receipt of a patent. The worth of the firm can be estimated fairly accurately once the patent is granted, but the uncertainty over whether it will be is the vehicle driving the firm's value.

2 Related Literature

Our paper is connected to a number of different literatures. First, there is the obvious connection to the literature on common-value auctions and the extensive evidence on the winners' curse (Kagel and Levin 1986; Kagel et al. 1989; Charness and Levin 2009; Charness et al. 2014, i.a.; see also Kagel and Levin 2002 for an excellent review). This pervasive observation in the laboratory has spiked a wide experimental literature testing its robustness and, so far, the winner's curse effect has been found to decline with public information (Kagel and Levin, 1986; Grosskopf et al., 2018), learning in form of sufficient experience (Dyer et al., 1989; Kagel and Richard, 2001; Casari et al., 2007) or familiarity with the task in the field (Harrison and List, 2008).[3] Our contribution here is to investigate the extent to which the winner's curse is robust to having stochastic rather than deterministic objects auctioned off. Having lotteries as auction prizes allows for modeling the common uncertainty in two alternative ways, either as common probabilities or as common values. Our experimental findings are novel because they directly connect the incidence of the winner's curse to the type of object up for sale. We demonstrate that persistent overbidding is eliminated in an equivalent variant of the same game that requires identical skills of Bayesian updating, contingent reasoning and learning.

The main drivers of the winner's curse phenomenon are still subject to a debate. The experimental literature provides mixed evidence on the importance of emotions such as the thrill of winning (Cox et al., 1992; Holt and Sherman, 1994; Bos et al., 2008; Astor et al., 2013) or the fear of losing (Delgado et al., 2008). Other explanations offered relate more directly to strategic uncertainty. For instance, subjects may misidentify the connection between other bidders' actions and their private signals (Eyster and Rabin, 2005; Crawford and Iriberri, 2007; Eyster, 2019). Alternatively, subjects might have difficulties performing the type of contingent reasoning involved in equilibrium behavior. More precisely, in order to avoid over-

[3]Relatedly, overbidding in independent private value auctions has been attributed to a misperception of winning probabilities (to some extent) (Dorsey and Razzolini, 2003), learning dynamics (feedback information) (Neugebauer and Selten, 2006), ambiguity aversion (Salo and Weber, 1995), anticipated loser's regret (Engelbrecht-Wiggans and Katok, 2007; Filiz-Ozbay and Ozbay, 2007), spite (Morgan et al., 2003), reference-dependent utility with induced values (Lange and Ratan, 2010) and imperfect best response combined with risk aversion (Goeree et al., 2002).

bidding, subjects' bids should be conditional on their private signal being the highest among all signals, and they should shave their bid downward. Anticipating the informational content of winning is, however, a difficult task. It requires a sophisticated level of contingent reasoning that, in general, most bidders struggle with. Besides common-value auctions, difficulties related to contingent reasoning extend to other settings, such as "Acquiring-A-company" games (Bazerman and Samuelson, 1983; Charness and Levin, 2009; Martinez-Marquina et al., 2019, in an individual decision game variant); voting games (Esponda and Vespa, 2014); or asset markets (Carrillo and Palfrey, 2011; Ngangoue and Weizsacker, 2021). In all these settings, uncertainty appears to be a crucial factor in impeding contingent reasoning (Martinez-Marquina et al., 2019; Koch and Penczynski, 2018; Ngangoue and Weizsacker, 2021; Moser, 2019).

Note, however, that none of the behavioral explanations mentioned above hinge on a specific object of uncertainty. Our findings are surprising since it had been assumed that the winner's curse was primarily the result of faulty strategic (Stage-2) reasoning and not faulty Stage-1 valuations. This result could not be detected by previous auction studies because the two stages of the bidding process are typically not investigated separately, as we do here.

Our study also connects to recent decision-theoretic experiments that have exposed the fragility of attitudes toward uncertainty. Subjects' attitudes have been found to vary with different sources of ambiguity (Abdellaoui et al., 2011; Li et al., 2018), with different sources of risk (Chew et al., 2012; Armantier and Treich, 2016) and sometimes with different outcomes (e.g., with money vs. time in Abdellaoui and Kemel, 2014). These treatments highlight that it matters how subjects think about uncertainty, and we contribute to this literature by proposing an additional treatment variation that puts the domain over which uncertainty is resolved in the foreground. In our study, the final outcome is always some amount of money, and the mechanism generating uncertainty is similar (the source), but the object over which the uncertainty is defined varies. This matters substantially in our game with incomplete information because it implies that our subjects will have to form and update beliefs over different objects. This distinction is also picked up in a growing literature in economics and psychology on how people view lotteries with uncertain outcomes versus those with uncertain probabilities. The findings suggest mainly that when asked to choose between lotteries involving uncertain outcomes or uncertain probabilities, subjects appear to have no strong preference (Kuhn and Budescu, 1996; Gonzalez-Vallejo et al., 1996; Du and Budescu, 2005; Eliaz and Ortoleva, 2016). However, when asked to price them, subjects value lotteries with uncertain outcomes above those with uncertain probabilities (Schoemaker, 1991; Du and Budescu, 2005). The results of Experiments II and IIIb are consistent with this existing literature on decision problems, but our auction setting brings attention to the finding that these small differences observed in individual decision-making are substantially magnified in a strategic environment.

Finally, it is useful to point out that a number of things we do in this paper are novel to the experimental auction literature. For example, this paper is probably the first to separate the bidding process into two stages and investigate them separately. This separation leads us to elicit a number of objects typically not revealed in conventional experimental auctions, such as a subject's estimate of the value of the object on which they are bidding given their signal; their estimate of the value of the object in the event of having the highest signal among all competitors (something we elicit in Experiment III and that is needed to bid optimally); their beliefs about the highest bid they are likely to face in the auction given their signal; as well as their belief that they have the highest signal among all competitors, etc. All of these objects are vital for understanding bidding behavior, and we believe that they may be useful to others.

3 A Different Auction Game

Our study consists of four experiments with a total of 360 students from New York University.[4] Experiment I provides the main evidence on bidding behavior in auctions with 107 subjects. Experiment II, with 104 subjects, sheds some light on how, in a non-strategic setting, bidders value lotteries with uncertain values versus those with uncertain probabilities. Experiments III and IV, with 96 and 53 subjects, respectively, help us understand how subjects map their private signals into bids. All experiments were conducted in a between-subject design, in which subjects faced uncertainty either in values or in probabilities. In this section we present the two main auction

[4]Participants were recruited with the software hroot (Bock et al., 2014).

treatments CV and CP of Experiment I and relegate the description of Experiments II, III and IV to Section 4.

3.1 The Bayesian Game

At the beginning of an auction, subjects were randomly placed into groups of four bidders ($i = 1, .., 4$).[5] Subjects then bid for lotteries described as either common value (CV) or common probability (CP) lotteries. Both lotteries are defined by two parameters v and p, where v is a non-zero payoff of the lottery and p is the percentage probability of receiving that payoff (with $(100 - p)$ defining the percentage probability of receiving 0).[6] We define by k and $\tilde{u}$ the known and unknown components of the lottery, respectively (we use tildes in the following to denote random variables). In the CP lottery, the two outcomes $\{v, 0\}$ are known such that $k =: v$, but $\tilde{u} =: \tilde{p}$ is uncertain, while in the CV lottery, $k =: p$ is known, but $\tilde{u} =: \tilde{v}$ is uncertain.

The computer determined the exact lottery by randomly drawing $\tilde{u}$ from a uniform interval such that $\tilde{u} \sim U[\gamma_l, \gamma_h]$, $0 < \gamma_l < \gamma_h < 100$. The four bidders, who did not observe the realization of $\tilde{u}$, each received a private signal s_i independently from each other. The signal was informative about the unknown $\tilde{u}$ in that $s_i \sim U[\tilde{u} - \varepsilon, \tilde{u} + \varepsilon]$, $\varepsilon > 0$. Signals became more informative with a smaller support–i.e., with decreasing ε.

A Bayesian bidder would infer from observing a specific signal s_i that the unknown $\tilde{u}$ must lie within $[s_i - \varepsilon; s_i + \varepsilon]$. To help the subjects, we provided this Bayesian update to them before they bid. Given this information, subjects placed a bid for the lottery at the bottom of the decision screen (see Figure 1 for an example of a CV auction).

At the end of an auction, the unknown $\tilde{u}$ was revealed along with broad feedback: Every bidder observed the true lottery (with p and v), the lottery outcome (0 or v) and the highest bid.[7] The lottery was played, and the auction winner received the lottery's outcome, either 0 or v, and paid her bid.

---

[5] We chose $n = 4$ because in the experimental literature, the winner's curse has been extensively studied in auctions with four bidders (see Kagel and Levin, 2002).

[6] We deliberately focus on binary zero-outcome lotteries to keep the cognitive costs of computing expected values comparable. This lotteries describe an investment structure that is not unusual because in many auctions the possibility of a zero outcome in form of a failed investment is relevant. For instance, in the standard example of common-value auctions, bidding for an oil lease also entails the risk of hitting a dry well.

[7] The computer broke ties between maximum bids randomly.

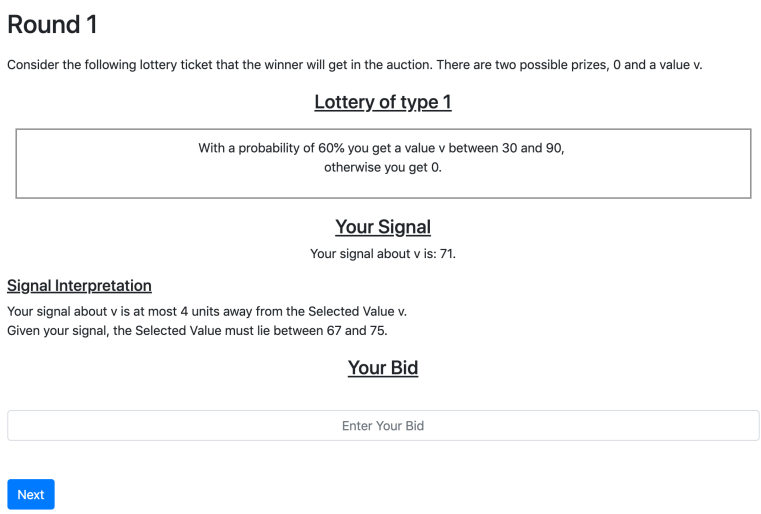

Figure 1: Example of Decision Screen in Treatment CV

3.2 Predictions Under Linear Expected Utility

There are three standard benchmarks under linear expected utility to which we can compare empirical bids. The derivation of the symmetric risk-neutral Nash equilibrium (RNNE) can be found in Wilson (1977) and Milgrom and Weber (1982). Following the standard experimental procedure of discarding observations with signals close to the lowest and highest possible values of $\tilde{u}$ (see, e.g., Kagel and Levin, 2002), we will constrain our data analyses to observations in the signal domain $(\gamma_l + \varepsilon < s_i < \gamma_h - \varepsilon)$ for which our three bidding benchmarks (multiplied by 100) take the following functional forms:

$$\texttt{Naive bid:} \quad 100 \cdot E[L|s_i] = k \cdot s_i \tag{1a}$$

$$\texttt{Break-Even bid:} \quad 100 \cdot E[L|s_i = \max_{\forall j}\{s_j\}] = k \cdot \left(s_i - \varepsilon \frac{n-1}{n+1}\right), \; j = 1, ..., 4 \tag{1b}$$

$$\texttt{RNNE bid:} \quad 100 \cdot b^*(s_i) = k \cdot \left[s_i - \varepsilon + \frac{2\varepsilon}{n+1} e^{-(\frac{n}{2\varepsilon})[s_i - (\gamma_l + \varepsilon)]}\right] \tag{1c}$$

A naive bidder will bid the expected payoff of the lottery given her private signal $E[L|s_i]$ (see Equation 1a). A more sophisticated bidder will take into account the winner's curse effect and will bid the expected payoff assuming that her signal is the highest. She will, therefore, shave her bid downwards to make, on average, zero profits with a break-even bid (see Equation 1b). A highly sophisticated bidder will shave her bid even more, assuming that, in a symmetric equilibrium, everyone else uses the same strategy (see Equation 1c). The break-even and the RNNE bid do not differ by much; the analyses will, therefore, focus mainly on the naive and the RNNE benchmarks as these represent the highest and the lowest bidding benchmark, respectively.[8]

3.3 Procedures and Parameters

In Experiment I, 55 and 52 subjects were assigned to the treatments CV and CP, respectively. Sessions lasted approximately 90 minutes and subjects earned, on average, \$23.78. The currency used in the experiment was credits (₵) with ₵6 corresponding to \$1.

Both treatments had identical procedures. The experiment was computerized with oTree (Chen et al., 2016) and consisted of two parts.[9] In the first part of both treatments, subjects engaged in eight different auction environments with ten separate auctions each. The eight environments determined the type of lotteries they bid for and were defined by an n-tuple $(k, \bar{\gamma}, \varepsilon)$, where $\bar{\gamma}(=\frac{\gamma_l+\gamma_h}{2})$ defined the interval $[\gamma_l, \gamma_h]$ (of fixed length) from which $\tilde{u}$ is drawn, and ε defined the signal precision $(\frac{1}{3}\varepsilon^2)^{-1}$. A 2x2x2 factorial design varied these three components across two sets of parameters: $k \in \{40, 60\}, \bar{\gamma} \in \{40, 60\}, \varepsilon \in \{4, 8\}$. Thus, as shown in Table 1, the eight environments differed with respect to the known parameter k (column 2), the support of the unknown parameter $[\gamma_l, \gamma_h]$ (columns 3 and 4), as well as

[8] In auctions with affiliated values, a prize in the form of a lottery ticket may generate some precautionary bidding if subjects have decreasing absolute risk aversion (Eso and White, 2004; Kocher et al., 2015). In the instructions, we did not specifically frame the lottery as an *ex post* risk, but the subjects probably perceived it that way. The observed bids in our CV treatment are, if at all, too high and do not suggest that the lottery ticket introduced a precautionary premium by lowering bids, although the general direction of a corresponding DARA effect in common-value auctions is not clear. More importantly, if bids exhibit a precautionary premium, there are no apparent reasons for premia to drastically differ across treatments.

[9] Subjects needed to pass a comprehension test before they could start the first part of the experiment. In the first part, subjects participated in a set of first-price auctions. In the second part, attitudes toward risk, compound risk and ambiguity were elicited (see Appendix C for a detailed description). At the end of the experiment, subjects learned their payoffs in the first and second parts and answered a small, unincentivized questionnaire. In the questionnaire, they provided some information on their socio-demographic background and their general approach to the auction game, and they took Frederick's cognitive reflection test (Frederick, 2005).

the signal precision given by ε (column 5).

Table 1: Lottery parameters

Lottery type	k	γ_l	γ_h	$\bar{\gamma}$	ε
1	60	30	90	60	4
2	40	10	70	40	4
3	40	30	90	60	4
4	60	10	70	40	4
5	60	30	90	60	8
6	40	10	70	40	8
7	40	30	90	60	8
8	60	10	70	40	8

Subjects then played ten different auctions in each of these eight auction environments. For each of these ten auctions, the computer randomly selected a true lottery on the basis of the environment's parameters such that the exact lottery (i.e., the true v of the CV or p of the CP lottery) and the corresponding signals possibly differed from auction to auction within an environment.

3.4 Results

Overall, bids significantly differed between the two auction formats.

Result 1 *Subjects generally overbid in common-value but bid according to Nash equilibrium in common-probability auctions.*

As predictions vary with parameters and signals, aggregate data will be described by bid factors defined as the difference between the subject's bid and the Nash equilibrium bid. Normalizing bids in this fashion allows us to focus on statistics that are independent of the private signals.[10,11] Bid factors are zero if subjects bid according to Nash equilibrium but are positive

[10]Empirical bid factors are usually defined as the difference between the naive bid and the subject's bid and reflect how much subjects shave their bid relative the naive expectations. We opted for a varying definition of bid factors that, in our opinion, offered better visualization of the data. Here, bid factors refer to deviations from the Nash equilibrium bid.

[11]We focus here on the comparison of bid factors across treatments, given that the empirical distributions of the randomly generated signals s_i were similar ($p = 0.113$ in Kolmogorov-Smirnov test).

(negative) if they bid above (below) the Nash equilibrium bidding function.[12]

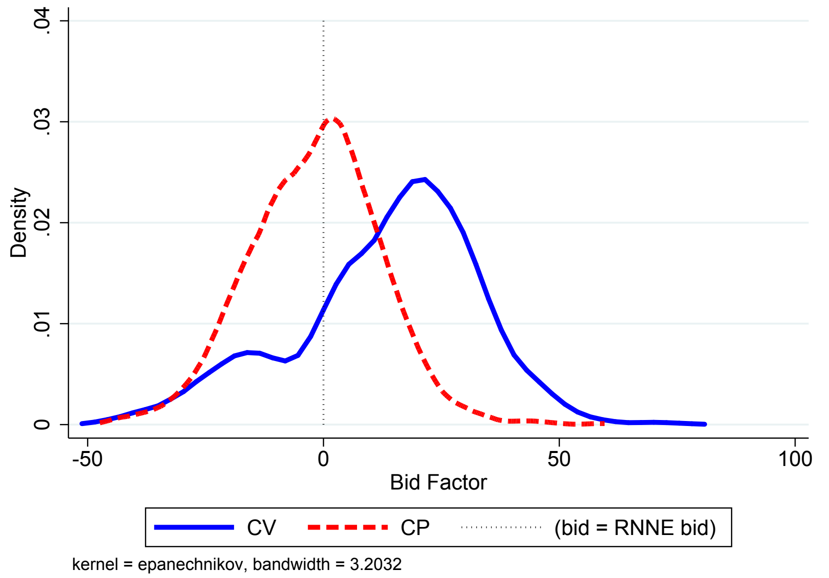

Figure 2: Bid Factors (=bid - RNNE bid) in Treatments CV (solid) and CP (dashed)

As Figure 2 shows, the distribution of bid factors significantly differed between the two treatments; Bid factors were predominantly positive in CV (indicating a fair amount of overbidding) but slightly negative, albeit consistent, with Nash equilibrium in CP. Appendix Table A1 presents the summary statistics of bid factors with respect to all three benchmarks. In CV, subjects bid more than the expected payoff of the lottery given their private signal, as mean and median bids are above the naive bid. In CP, by

[12]Despite passing the comprehension test, some subjects chose dominated bids that were above the highest possible outcome a lottery could pay off (given one's signal in CV or independent of the signal in CP). For subjects who made these dominated choices too frequently (more than 10% of all rounds), we interpret these choices as evidence of inattention and exclude 3 and 13 subjects in treatment CV and CP, respectively. In the Experiments II to IV, we exclude in a similar manner 5 and 8 subjects in CV and CP, respectively. Therefore, our analyses present the decisions of a reduced sample of 332 out of 360 subjects. It is important to note that, first, removing those subjects does not affect our main conclusion, as we continue to observe a substantial difference in bidding behavior with the entire sample. Second, this reduced sample is balanced in the sense that, across treatments, the remaining subjects do not substantially differ with respect to demographics and personal characteristics measured at the end of the experiment (see Appendix Table A7.)

contrast, subjects bid slightly below the RNNE bid. In sum, subjects significantly overbid for CV lotteries, but bid close to Nash equilibrium for CP lotteries. This difference between CV and CP occurred for all parameter combinations–i.e., in all eight auction environments. Appendix Figure A 3 shows the estimated median bid as a function of signals. In all eight environments, the median bids for CV lotteries lie substantially above, while those for CP lotteries are slightly below the RNNE curve.[13]

The question remains whether this difference in bid factors affects the incidence of the winner's curse. Figure 3a shows the distribution of bid factors in winning bids separately for the CV and CP treatments. Winners in both auctions fell prey to the winner's curse, as average winning bids were significantly above the break-even bid (see Appendix Table A2). Yet the difference in bid factors between the two auction formats remains substantial. Winners bid higher in CV and lost, on average, more than in CP (mean loss of ₵ −24.04 in CV vs. ₵ −10.04 in CP, p-value$<$ 0.001 in t-test of differences with cluster-robust standard errors). As shown in Figure 3b, the cumulative distribution function (CDF) of winning payoffs in CV auctions first-order stochastically dominates the CDF in CP auctions.

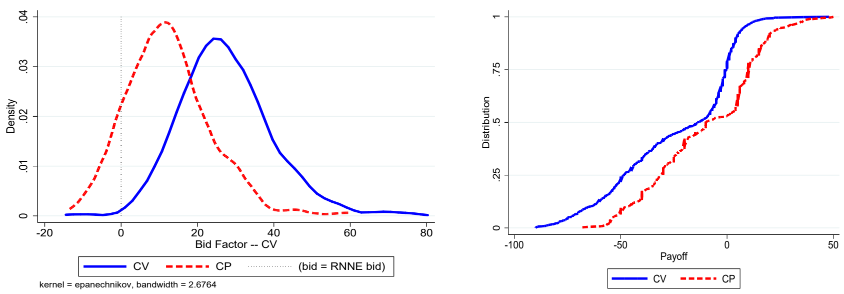

(a) Bid Factors in Winning Bids **(b)** CDF of Winners' Payoff

Figure 3: The Winner's Curse in CV and CP Auctions

Result 2 *The winner's curse effect is attenuated in the common-probability compared to the common-value auctions.*

[13]Previous experimental studies with a similar design exhibit some variety in the extent of overbidding with, for instance, a percentage of bids above the break-even bid ranging from 10% to 82% (see, e.g., Kagel and Levin, 2002). In our design that features a stochastic instead of a deterministic prize, we observe that the extent of overbidding is slightly in the upper end of this range (62% in Experiment I) but still in line with previous findings.

Several factors may have contributed to the winner's curse being less severe in CP than in CV auctions. The first possibility is that winners in CP lost less money because they shaved more as a result of more-sophisticated reasoning through the adverse selection problem; a second alternative is that CP auctions required less shaving in the first place because the empirical break-even bids differed across auctions. Specifically, subjects may have shaved differently because winning in CV and CP revealed unequal information. A winner's curse can arise only to the extent that winning reveals some relevant information that should have been considered in formulating a bid. Here, the information that a bidder needs to consider to properly correct for the adverse selection effect is given by the difference between her estimate of the object given her signal and her revised estimate upon winning the auction. In a symmetric equilibrium, these two estimates will differ because being the highest bidder implies having the highest signal. Yet, in our auctions, bidding strategies may not have been symmetric, and, thus, to assess how informative winning is, we estimate the predicted probability of having the highest signal conditional on winning. Winning in CV increases the likelihood of having the highest signal from 19% to 36%, which is three times more than in CP (24% to 28%, as shown in Appendix Figure A 5).[14] In other words, in the CP auction, a naive bidder could have easily won an auction without having the highest signal and, hence, without overestimating the object. This implies that, empirically, the CP auction requires bidders to shave less than the CV auction, which may explain why, in CP auctions, the winner's curse is less severe, but not why bids are lower. Appendix Section B.2 provides a more detailed analysis of the winner's curse effect in those auctions. We conclude that, in the aggregate, the winner's curse is mitigated in CP because, there, a winner is less likely to have the highest signal. This suggests that subjects in CP may not have been more sophisticated than those in CV, but it does not preclude the possibility that the ability to reason through the winner's curse differed.

[14]The probit estimation is done with the entire sample in Experiment I. The estimation with the reduced sample leads to more extreme results with a higher marginal effect of winning in CV, but a nonsignificant and weak effect in CP.

4 The Bidding Process and Its Constituents

4.1 Experiment II: Pricing of Lotteries

Differences between our CP and CV auctions may originate in the first stage of the bidding process in which subjects form subjective valuations for the good. In other words, subjects possibly bid differently simply because they resolved the fundamental uncertainty existing in Stage 1 differently, without differing in their assessment of the strategic uncertainty they faced in Stage 2. To investigate this in Experiment II, we stripped the auction game of its strategic elements (cf. Charness et al., 2014, i.a.) by having subjects submit their willingness to pay (henceforth WTP) for the same lotteries used in Experiment I. In treatment CVL, we elicited subjects' WTP for a series of CV lotteries, whereas in treatment CPL, we did the same for CP lotteries.

In Experiment II, subjects engaged in the same eight environments used in Experiment I but now made 12 decisions per environment. In the first round, they submitted an *ex-ante* WTP without observing a signal. This round was followed by ten rounds, in each of which a new lottery was generated. For each of these ten lotteries, subjects submitted an *interim* WTP *after* observing a signal. Lastly, later in the experiment, they stated another ex-ante WTP (without signals) for the first lottery they were presented with, but this time in its reduced rather than its compound version. This allows us to assess whether cognitive difficulties of compounding lotteries affected valuations because for this elicitation, we aggregated the compound probabilities for them. To this end, we presented lotteries defined over final payoffs by using a wheel that displayed all possible outcomes of the reduced lottery in a simple and condensed graph (see Appendix Figure A 1 for an example of a CV lottery). We reversed the order of eliciting WTP for compound and reduced lotteries across sessions: One third of the subjects first submitted a WTP for lotteries in their reduced version before pricing lotteries in their compound version.[15] The experimental interface remained

[15]To determine whether or not a lottery was bought, we endowed the subjects with ₵100 with which to bid, with any unspent credits paid to the subject. We then used the Becker-DeGroot-Marschak mechanism (1964) so that after a subject submitted her WTP, a random number between 0 and 100 set the lottery price. The subject bought the lottery if its price was weakly less than her WTP. In that case, any gains or losses were added to or subtracted from her endowment of ₵100. Otherwise, she did not engage in the lottery and ended the round with her initial endowment.

essentially the same as in Experiment I, with the only difference being that the subjects submitted WTPs for lotteries in a decision problem rather than bids in an auction game. To keep learning dynamics as similar as possible to those in Experiment I, in rounds without a signal there was no feedback after submitting an ex-ante WTP. In contrast, after every round with a signal, the subjects learned the actual lottery ticket, its price and its outcome, irrespective of whether or not they bought it.

4.1.1 Results: Experiment II

We collected data from 54 and 50 subjects in treatments CVL and CPL, respectively. We present our main results using the difference between w_i, the subject's willingness to pay for a lottery, and $E[L|s_i]$, the lottery's expected payoff given the subject's signal. We call this measure the *price factor* $PF = w_i - E[L|s_i]$, emphasizing its correspondence to the bid factor. In other words, the price factor is equivalent to the negative of a risk premium.

To compare the lottery valuations with the bids from Experiment I, we use the fact that $E[L|s_i]$ also corresponds to the naive bidding curve in the auction and compare the price factor with a *non-strategic* bid factor with respect to the naive benchmark ($BF^{ns} = \texttt{bid} - E[L|s_i]$). Figure 4a shows the distribution of non-strategic bid factors in the auctions of Experiment I. We juxtapose Figure 4b, which shows the distribution of price factors in Experiment II. The treatment effect that we found in the auction game is largely attenuated in the decision problem. The average difference of ₵ 4 (p-value < 0.001 in quantile regression, ₵ 6 in the means) is substantially smaller than the difference of ₵ 17.4 observed in the auctions.

Appendix Sections B.3 and B.4 show that subjects processed value and probability signals similarly in the valuation of lotteries. We obtain a similar picture with the ex-ante valuation of CV and CP lotteries, for which we find no differences in median WTP (see Appendix Figure A 4). Without any signal, subjects chose an average uncertainty premium of ₵ 4 regardless of whether uncertainty was represented by a range of values or a range of probabilities (see Appendix Table A3). We also find that, in the aggregate, subjects priced compound and reduced lotteries similarly (see Appendix Section B.5). In that sense, the general perception of lotteries with uncertain values versus uncertain probabilities does not explain differences in bids.[16]

[16]Note that incentives differ between the auction and the lottery treatments. In first-

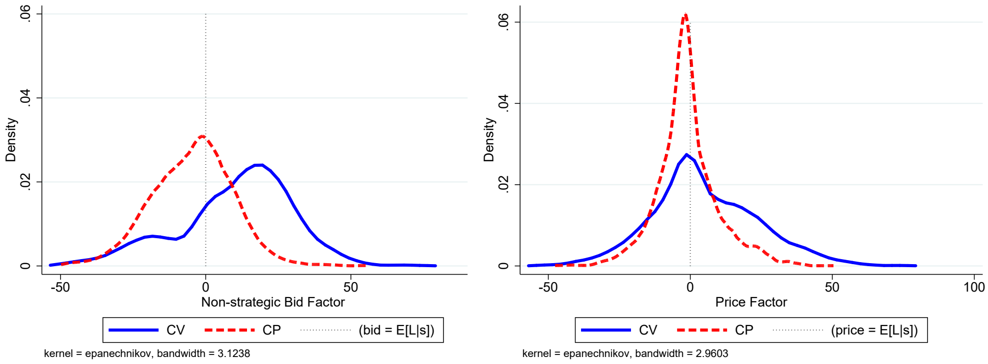

(a) Non-Strategic (Naive) Bid Factors **(b)** Price Factors

Figure 4: Bidding in Exp. I (Left) versus Pricing in Exp. II (Right)

We next investigate the discrepancy between the bidding in auctions and the pricing of lotteries in decision problems. Valuations for CP lotteries with and without strategic incentives are rather stable. On average, subjects priced CP lotteries almost ₵1.8 below their expected payoff, and in auctions, bid ₵2 less for the same CP lotteries ($p = 0.29$ for the comparison of bids and prices in quantile regression with clustered standard errors). In contrast, valuations for CV lotteries substantially differed across settings. In CV auctions, subjects bid, on average, ₵13.6 above the expected payoff of the lottery but priced the same lotteries close to the expected payoff in the decision problem (₵2.4 above expected payoff–i.e., on average, ₵11.2 less than in auctions, $p < 0.001$ in quantile regression).

Experiment II seems to suggest that the observed difference in bidding across our CV and CP auctions was a Stage-2 phenomenon involving strategic reasoning since subjects seemed to value CV and CP lotteries equivalently. This conclusion is premature, however, since, as we will see in Experiments III, IIIb, and IV, placing a subject in a competitive environment tends to change their Stage-1 valuations.

price auctions, a Nash equilibrium bidder pays his bid and in expectations makes small profits. In the lottery treatment, subjects pay the random price, which is, in expectation and conditional on buying, half the subject's WTP. Monetary incentives are, therefore, on average, higher in the lottery treatment and could partly explain the smaller differences in the lottery treatments. However, monetary incentives should have a similar effect across CV and CP treatments, but the asymmetry in findings between CV and CP treatments casts some doubt on monetary incentives being the main reason for the observed differences between auction and lottery treatments.

4.2 Experiments III-IV: The Bidding Process in the Auction

When bidding in an auction, it is not enough to estimate the common-value item by looking only at one's own signal. One must consider that others have also received signals and, if they bid in a monotonic fashion, then a bidder must condition her estimate of the common-value item on her signal being the highest. The persistent overbidding in our CV auctions suggests the opposite reflection, in that subjects may heavily weight the possibility of signals above their own as being relevant. Overestimating the item's value in the auction context (Stage 1) can then lead to overbidding in the absence of faulty strategic reasoning in Stage 2.

In Experiments III and IV, we look into the black box of bidding to identify what element in this mapping of signals into bids is different across our two auctions. Technically, the ingredients of the black box should not depend on signals being over values or probabilities; yet our results from Experiments I and II suggest that they do. Experiment III investigates the mapping from signals to fundamental estimates (Stage-1) *in an auction setting*, while Experiment IV explores strategic (Stage-2) behavior.

4.2.1 Experimental Designs

Procedures Experiments III and IV were conducted online on Zoom during the Spring and Fall of 2021. For more efficiency, we made use of the data in Experiment I by matching every participant in the new experiments with three participants from auctions that took place in Experiment I. In other words, every new subject replaced one randomly chosen subject from an auction in Experiment I. The new subject then bid for the exact same lottery ticket, observed the same signal and bid against the same three opponents as the replaced subject from Experiment I.[17]

[17]Hence, we held constant the environment our subjects faced by making it identical to that of Experiment I–except, of course, for the fact that they were bidding against subjects in a previous experiment, which the subjects knew. This allowed us to collect more independent data since we did not need to recruit four subjects to obtain data on one auction. In addition, this kept the feedback after each round as similar as possible because information about the winning bid changed only if a new participant won the auction. We also took into account that it was less risky to conduct an online experiment as an individual decision-making experiment rather than with groups because of the risk of losing group members due to internet connections or distractions. Each experiment had two parts with separate instructions to which subjects had access throughout the experiment. In addition, instructions were read aloud by the experimenter or made available through a video. In every experiment, a comprehension test before Part 1 ensured that subjects

Experiment III Intelligent bidding in an auction requires the ability to estimate the expected payoff of a lottery given one's own signal and, in Nash equilibrium, conditional on having the highest signal. We specifically wanted to know whether subjects form correct estimates of the lottery's expected payoff in the auction game (as opposed to the decision problem in Experiment II), in particular when they condition on having the highest signal and, more importantly, whether subjects perform these tasks differently across our value and probability auctions.

Experiment III consists of two parts. To assess the robustness of our results when matching participants to previous subjects, in Part 1, we let our subjects bid for a lottery after receiving a (value or probability) signal, just as in Experiment I. In Experiment III, however, because of a time constraint, subjects participated in only 40 different auctions. More specifically, we used the same parameters as in Experiment I, but we focused on the parameter sets with $\epsilon = 4$, for which, in Experiment I, overbidding in CV was more pronounced. As we show in Appendix B.6, this part of Experiment III replicates our main results of Experiment I.

In a second part, we elicited different estimates for 12 different lotteries. For each lottery we elicited three objects of interest:

1. The estimate conditional on a signal (which we call the **Naive Estimate**): Subjects stated their best estimate of the expected payoff of the lottery ticket given a signal. In the experiment, we referred to the expected payoff as the true *average worth* of the lottery ticket. We described it as the payoff that they would get, on average, if the lottery ticket with its Selected Value or Selected Probability was played very often. There was one important difference between the estimate elicited here and the valuation stated in Experiment II: While in Experiment II, we simply asked subjects to price the lottery given a signal in isolation, in Experiment III, we elicited these values after subjects had participated in 40 rounds of auctions. In other words, subjects' assessed the lottery in the context of an auction, forcing them to consider this auction context when estimating the lottery's expected payoff. For example, the expected payoff of a lottery conditional on one solitary signal may strike subjects as different from the one conditional on knowing that the signal is just one of four. If, for some reason, our

understood the information structure and the auction (in Experiment III).

subjects thought that their signal might have been the lowest, they might have increased their estimates to incorporate this fact. This is, of course, false strategic reasoning since, to avoid the winner's curse, one computes expectations under the assumption that one has the highest signal of all bidders.

2. The estimate conditional on having the highest signal (which we call the **Contingent Estimate**): We elicited subjects' estimates of the lottery's expected payoff in the hypothetical event of having the highest signal. Such contingent reasoning is required for bidding in Nash equilibrium. Differences between the Naive and the Contingent Estimate tell us whether subjects were able to construct Bayesian updates correctly–that is, whether they were able to compute correct fundamental estimates once they were nudged to consider the relevant conditional event of having the highest signal.

3. The belief about the likelihood of having the highest signal: We asked subjects to assess the probability that the signal that they received was the highest one among the four bidders' signals in an auction. This measure is necessary to interpret our data because someone who believes that she has the highest signal may not display any difference between her Naive and Contingent Estimates in the first place.

We also ran Experiment III with a slightly different variant, in which the two main parts were swapped: While in Experiment III, we elicited subjects' estimates about features of the lottery after they had already engaged in 40 auctions, in Experiment IIIb, we elicited this information *before* they bid in auctions. Reversing the order of the two parts in Experiments III and IIIb allowed us to elicit fundamental estimates both outside and inside the auction context. More precisely, in Experiment IIIb, we told our subjects that their signal was part of a vector of four signals, but we did not reveal that lotteries would be auctioned off in the second part. Hence, comparing Experiments III and IIIb controls for auction experience in the valuation of lotteries. The difference between Experiments II and IIIb, on the other hand, is that subjects in Experiment IIIb knew that their signal was one of four signals drawn, while in Experiment II they received a single signal in isolation.[18]

[18]All point estimates were incentivized with a binarized scoring rule (Hossain and Okui,

Experiment IV In Experiment IV, we isolated strategic from fundamental uncertainty by providing our subjects with the most accurate fundamental estimate for Stage 1, leaving them with mainly strategic uncertainty in Stage 2. This experiment also consisted of two main parts. In the first part, we explored subjects' strategic uncertainty by eliciting beliefs about their opponents' bids. Here, subjects saw the same information as their peers in Experiment I–that is, a lottery ticket with an uncertain value or probability–and then, based on their private signal, they stated what they thought the highest bid of the other three bidders was (we call this the **Belief about the Competitive Bid**).[19] Subjects stated their beliefs for 20 independent auctions without receiving any feedback between auctions so as to not bias their decisions in the subsequent Part II.

In the second part, subjects engaged in the same auctions as in Part 1, seeing the same lottery tickets with the same signal. However, in this part, we provided them with two pieces of information to allow them to weigh the fundamental risks against the strategic ones. The first one pertained to strategic uncertainty and reminded them of the belief about the competitive bid they had stated for that particular auction in Part I (cf. box in the upper right corner of Appendix Figure A 2). The second one pertained to fundamental uncertainty. Our goal was to reduce fundamental uncertainty as much as possible by telling our subjects that they had the highest signal among all four bidders. In addition, we explicitly told subjects what it

2013). A bonus payment of ₵ 12 was paid if the estimate was sufficiently close to the object of interest. More precisely, the bonus was paid with probability Pr[(Estimate - object of interest)$^2 < K$] with $K \sim [0, 36^2]$. To elicit the probability belief, a bonus of ₵ 24 was paid with probability $\Pr[(\mathbb{1}_{(s_i \geq s_{-i})})$ - guess$)^2 < K]$ with $K \sim [0, 1]$, where $\mathbb{1}_{(s_i \geq s_{-i})}$ takes the value 1 if the subject had indeed the highest signal (and zero otherwise). In addition, we incentivized subjects to report the *smallest* interval for which they were certain that it contained the lottery's Expected Payoff Conditional on the Signal and its Expected Payoff Conditional on Having the Highest Signal (Schlag et al., 2015; Enke and Graeber, 2021). To incentivize the choice of the interval, we penalized wider intervals with the following rule: subjects received a fixed bonus only if the object of interest was indeed in their chosen interval, but the magnitude of the bonus would depend on the interval width. The bonus was calculated as follows: $12 * [1 - (u - l)/(E[L(\gamma_h)] - E[L(\gamma_l)])]$, where l and u denote the chosen lower and upper threshold of the interval and $E[L(\gamma_l)]$ and $E[L(\gamma_h)]$ the minimum and maximum expected payoff of the lottery. Thus, the smallest interval–i.e., a point prediction– would lead to a maximum bonus of ₵ 12, while increasing the interval width would decrease the bonus. This measure was meant to reflect subjects' uncertainty about their estimates.

[19]Again, subjects submitted first a point belief and then the smallest interval they were sure would contain the highest bid of the other three bidders.

meant to have the highest signal by telling them what the expected value or probability of the lottery was given that their signal was highest, as well as the resulting expected payoff of the lottery if one would multiply values with probabilities (see Appendix Figure A 2). Hence, they had all the important information about the lottery, leaving them mainly with strategic uncertainty about the opponent.[20]

4.2.2 Results: Experiments III and IV

Stage 1 and Fundamental Uncertainty We start with Experiment III, which examined whether differences in bids can be captured by different fundamental estimates in Stage 1 of the bidding process. The distributions of estimates are presented in Figures 5a and 5b, which plot the densities of estimates conditional on the signal (solid lines for Naive Estimates) and conditional on the signal being highest (dashed lines for Contingent Estimates) relative to the objective expected payoff $E[L|s]$ in CV and CP, respectively.

Result 3 *Subjects in CV overestimated the lottery's expected payoff conditional on the signal they received and also conditional on the signal being the highest among all bidders, while subjects in CP did not.*

One striking feature in Figures 5a and 5b is that both Naive and Contingent Estimates of the lottery's expected payoff are significantly higher in CV than in CP. For equivalent signals, subjects believed that the CV lottery would, on average, pay more than the CP lottery. Furthermore, subjects did not condition correctly contingent on having the highest signal among the four bidders. The aggregate distribution of Naive and Contingent Estimates are practically identical. The same observation is made at the individual level, where subjects in both CV and CP chose, on average, the same estimates in both tasks. This is significant because it indicates

[20]Note that some of the information we collected in Experiments III and IV have, to our knowledge, not been collected in any previous auction studies, making our results relevant to auctions in general. While some previous studies have elicited estimates of the common-value item with or without incentives (Bazerman and Samuelson, 1983; Charness et al., 2019), and others have provided subjects with the opportunity to revise bids conditional on winning (Moser, 2019), to the best of our knowledge, ours are the first experiments to directly elicit estimates about the common-value item contingent on having the highest signal, the probability of having the highest signal and beliefs about the most competitive bid.

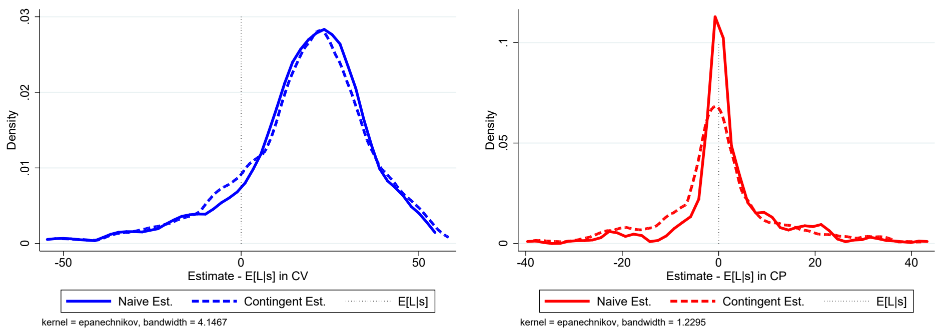

(a) CV **(b)** CP

Figure 5: Naive and Contingent Estimates in Experiment III

that subjects failed to do one of the most critical parts of Stage-1 processing, which is to condition correctly on their signal being the highest. If they had, the distribution of Contingent Estimates would have shifted to the left in both auctions, reflecting that having the highest of four signals lowers one's expectation about the lottery's payoff.[21] Thus, even when nudged to consider the right conditional event, subjects did not recognize that they needed to adjust their estimates downward.[22]

In Experiment IIIb, we asked subjects to assess lotteries before bidding in auctions knowing that the signal they received was one of four signals. The difference between Experiments III and IIIb should be informative of the impact of experiencing competition on Stage-1 valuations.

We find that the majority of subjects in Experiment IIIb estimated the lottery's expected payoff close to its actual expected payoff (median difference of 0.4 in CV, $p = 0.14$, 0 in CP, $p = 1$, in median regression with

[21]We come to the same conclusion when we take into account subjects' probabilistic beliefs of having the highest signal. We find no systematic correlation between these probabilistic beliefs and the Naive (and Contingent) Estimates, which, in turn, may not be surprising given subjects' failure to adjust Contingent Estimates in the right direction.

[22]Note that this finding contrasts with Moser (2019) and Esponda and Vespa (2021) which differ in two ways. First, Moser (2019) and Esponda and Vespa (2021) study actions (bids) rather than fundamental estimates contingent on relevant events; second, in their studies, the relevant event that subjects are nudged to consider is winning. We believe that the two events, winning and having the highest signal, are not empirically equivalent. Conditional on winning, increasing the bid is a strictly dominated action, making it easy for subjects to recognize that bids can only be adjusted downward. However, in Nash equilibrium, the relevant informational event is the one of having the highest signal, thereby also allowing for the possibility of non-monotonic or asymmetric bidding. Conditional on having the highest signal, it may not be obvious to profit-maximizing subjects that estimates need to be adjusted downward. This is what we in fact find.

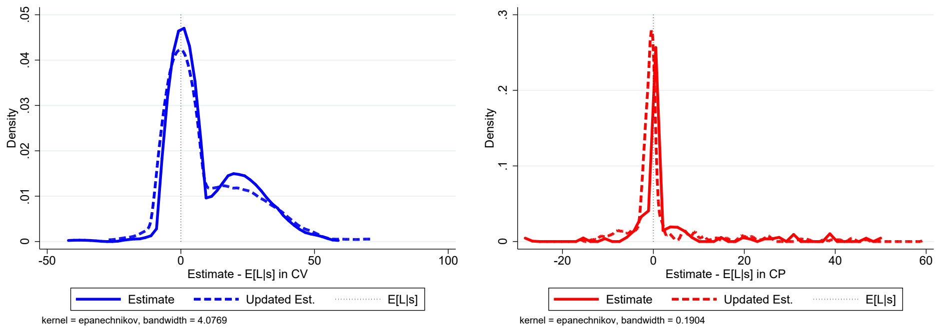

(a) CV

(b) CP

Figure 6: Naive and Contingent Estimates in Experiment IIIb

cluster robust standard errors (CRSE)). In the CP variant of Experiment IIIb, subjects continued to evaluate the lotteries at their expected payoff, albeit with a smaller noise. In CV, as can be seen in Figure 6a, the density of estimates is shifted to the left compared to that in Figure 5a. While a small hump remains at higher values, the mode is centered around the expected payoff $E[L|s]$.[23] This finding is consistent with the finding in Experiment II that subjects evaluated CV lotteries correctly in a non-competitive setting with one signal. Our findings here show that it does not matter whether signals are presented individually or as part of a vector as long as they are not placed in a competitive context. More importantly, these results suggest that if, conditional on their expectations, subjects shaved their bids equivalently in our CV and CP auctions, then the overbidding we observe in the CV auction stems from a problem of overestimating fundamental values in Stage 1 and not a strategic problem in Stage 2.

Stage 2 and Strategic Uncertainty When looking at bidding behavior in Experiment III, we see that contingent on their estimate, subjects shaved their bid in a manner that was not significantly different across auctions.[24]

[23]We note that in Experiment IIIb, subjects formed better Contingent Estimates. In both treatments, subjects lowered their estimates by, on average, ₵1 ($p = 0.02$ and $p = 0.001$ in CP and CV, in median regression with CRSE) contingent on having the highest signal (compared to a theoretical benchmark of lowering by, on average, ₵1.7).

[24]In absolute numbers, subjects actually shaved their bids more in CV than in CP, but between auctions, differences in shaving were nonsignificant. In CV, they bid, on average, ₵-2.20 below their estimates ($p = 0.235$, median regression with CRSE). In CP, bids were, on average, ₵-0.8 ($p = 0.6$) below estimates. We obtain the same conclusion if

This suggests that differences in bids do not stem from systematic differences in the way subjects convert their estimates into bids. Of course, we can get a complete picture of how subjects translated their estimates into bids only once we know their beliefs about their opponents' strategies. In Figure 7, we place the objective expected payoff of the lottery given a signal, $E[L|s]$, on the horizontal axis, and for each such expected payoff, we plot two observed values: (i) the median belief about the competitive bid (black line); and (ii) the actual, median competitive bid (from Experiment I, dashed line). The dotted line is the 45-degree line.

Result 4 *Subjects in CV expected their strongest competitor to bid highly, and their expectations were correct. By contrast, subjects in CP expected their strongest competitor to bid close to the lottery's expected payoff, thereby underestimating the competition.*

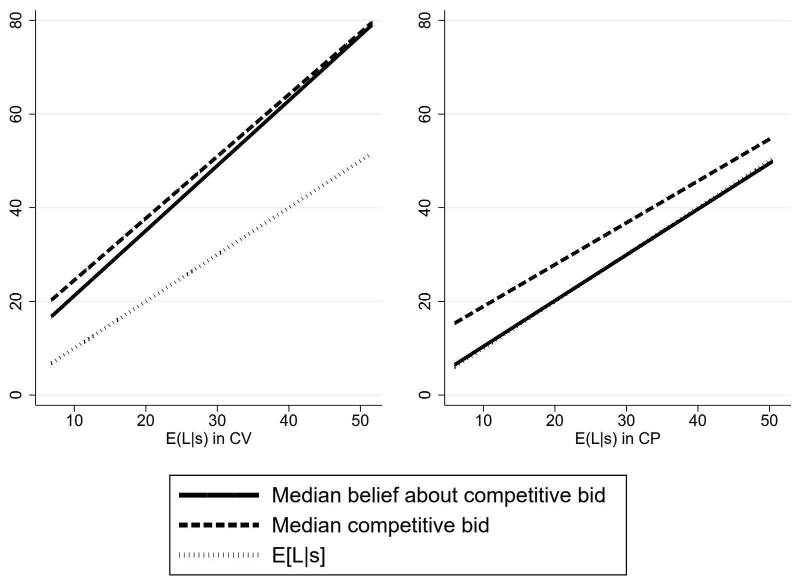

Figure 7: Beliefs and competitive bids

Figure 7 shows that, for a given expected payoff, beliefs about the competition subjects faced in the CV auction were significantly higher than

we estimate the bid function as the function of endogenous fundamental estimates in a two-stage least squares procedure with mixed effects. The estimates of the bid functions are not significantly different across auctions.

those in the CP auction. This observation is consistent with the fact that equivalent signals induced higher fundamental estimates in CV compared to CP.[25] In addition, in the CV auction, subjects were approximately correct in terms of their beliefs about the median competitive bid. Subjects expected high bids from their competitors (bids greatly higher than the expected payoff of the lottery), and these beliefs were confirmed. The pattern was different in CP. Here, subjects did not expect their competitors to bid above the expected payoff of the lottery and actually underestimated the competition. These beliefs further sustain the difference in fundamental estimates observed in Experiment III. In CV, the fact that subjects believed their competitors bid close to their inflated estimates left them no room to shave. For CP, the situation differed in that subjects had lower estimates and also underestimated the competition.

A natural question that arises is whether subjects would adjust their bids downward if they had correct estimates of the lottery's expected payoff when knowing they had the highest signal. Both Experiments IIIb and IV, in which subjects had access to correct fundamental estimates before bidding, can shed light on this question. While in Experiment IIIb, subjects computed correct fundamental estimates themselves, in Experiment IV, we provided them with correct fundamental estimates conditional on them having the highest signal. If subjects in CV failed to form accurate estimates, then providing them with such information should have changed their bidding behavior. In fact, we find this to be the case. We start by looking at the gentle nudge in Experiment IIIb. Figure 8 shows that having subjects enter the auction with correct fundamental estimates shifts the distribution of bids in Experiment IIIb to the left, relative to bids in Experiment III.

[25]This finding is akin to the type projection bias that induces one to believe that others have the same signal as one does (Breitmoser, 2019), although, here, subjects would be projecting their (biased) valuations rather than their type. Also note that there seems to be some endogeneity between valuations and beliefs about others, as excessive valuations were observed only in the presence of others.

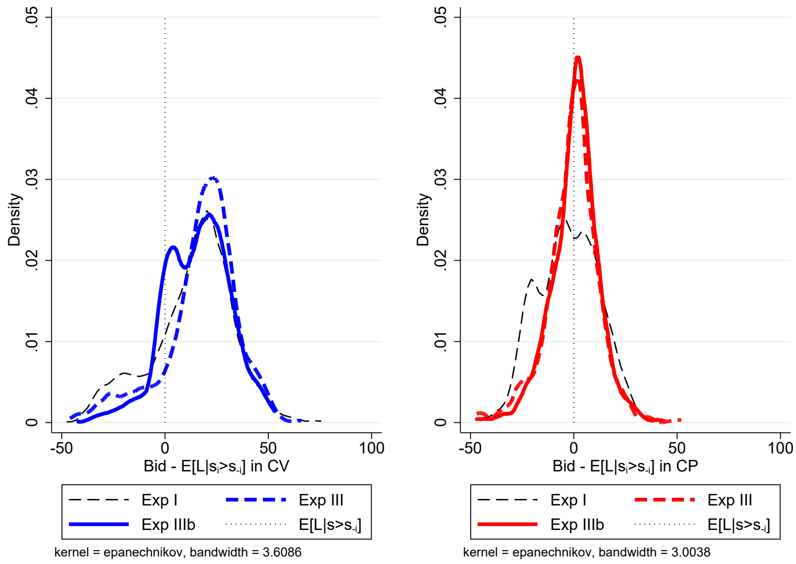

Figure 8: Bids in Experiments IIIb vs. I and III

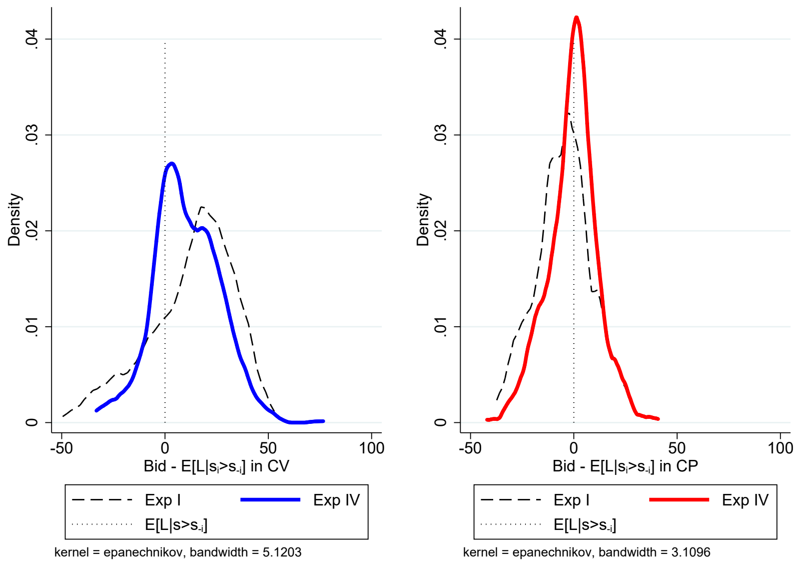

Figure 9: Bids in Experiment IV compared to Experiment I

The stronger nudge in Experiment IV led to an even stronger shift. Figure 9 shows subjects' bids in Experiment IV as deviations from the expected payoff given that they were told their signal was the highest. We juxtapose the bids (relative to the same benchmark) of their peers in Experiment I who had the exact same signal but were not informed of its rank. As Figure 9 suggests, providing subjects with the conditional expected payoff had a correcting effect on bids in both auctions. In CP, subjects adjusted their bids upward, around the conditional expectation. In CV, we observe that, relative to Experiment I, a substantial portion of bids were adjusted downward towards the conditional expectation, even if the correction was not enough to eliminate the aggregate evidence of overbidding. The shift to the left in the CV distribution is consistent with our conjecture that in both Experiments I and III, in which subjects in CV bid highly, a good part of this overbidding was the result of not computing proper valuations. This confirms our suspicion that receiving a signal in the value domain led to the failure to process values correctly. When the expected payoff was provided, behavior changed but could do so only partially, as subjects' competitors from Experiment I who were not given this information did, in fact, bid excessively.

5 Discussion and Conclusion

This paper has presented a puzzle concerning the object of uncertainty in auctions. We found that when we auctioned off items with an uncertain common value, subjects largely overbid, but when the object of uncertainty was a probability, the average bidding was close to the Nash equilibrium.

Our experiments revealed that our main puzzle consists of a combination of two different phenomena: the overvaluation of value lotteries and the fact that this overvaluation occurs only in a competitive context–i.e., in auctions, but not in decision problems.

We gained these novel insights by breaking down the process of bidding into two stages. In Stage 1, given the signal received, subjects form an estimate of the item up for sale (either conditional on their signal or on knowing their signal is the highest). In Stage 2, they transform their estimate into a bid, with the understanding that they will face other bidders. This two-stage process separates the fundamental uncertainty in the auction–i.e., how much the good they are bidding for is worth–from the strategic uncertainty–i.e., what others will bid given their signals.

This procedure allowed us to pin down the origin of our puzzle to the way in which subjects resolve fundamental uncertainty when competing with others. When subjects received a value or probability signal about the lottery in decision problems, they did not only value the underlying lottery equally across our two domains, but also valued it close to the lottery's expected payoff. These results were observed in Experiments II and IIIb. However, when subjects were presented with signals in an auction, or, more precisely, after having experience in an auction, (as was true in Experiment III), they overvalued the lottery in the CV but not in the CP auction. Thus, there was an asymmetry in how bidders resolved fundamental uncertainty about values and probabilities but only when subjects were placed in the auction. What this means is that the wall we have erected between assessing fundamental and strategic uncertainty is not as solid as we thought. There is a spillover. Our main goal in this discussion is to see whether we can make sense in integrating these two different pieces of our puzzle.

In the following, we first discard possible explanations that have been discussed in the literature, and then consider alternative explanations for why competition triggers misjudgments with value uncertainty. While we offer some suggestions to help us solve our puzzle, there is plenty still left to investigate.

Enhanced utility via affect. An existing literature has discussed how competitive settings may affect subjects' valuation for an item. As discussed in Section 2, one of the first mentioned explanation for the winner's curse was the joy of winning that inflates the utility payoff of winning and, as a result, entices a subject to bid higher. Somewhat related, Imas and Madarasz's (2020) "economics of exclusion" suggests that other peoples' desire for a scarce object makes it more valuable to the decision maker. In both of these mechanisms, competition affects outcomes via an affective channel, which is modeled as an increase in utility. Yet a surprising insight from our experiments is that, in our Bayesian game, competition affects outcomes via a cognitive channel: Our subjects directly misestimated the item's value (rather than having a higher utility for it). This implies that our puzzle's source is not in the utility domain but, rather, in additional cognitive challenges to belief formation. Thus, in the following, we abstract from mechanisms that distort valuations through utility and focus on how competition may complicate belief formation.

Misinterpretation of information in the auction. Bidding for lotteries is more complex than pricing them because it requires bidders to form rank-order statistics. Bidders know that their signal is only one out of many, and their final expectation of the good will depend on how, they believe, their signal compares to others' signals. In our Experiment IIIb, however, we do not find that computing rank-order statistics is the main issue: Subjects were able to interpret their signals as part of a vector as long as they were not associated with competition with others.

Desirability bias. Another way in which competition may affect belief formation is via motivated reasoning (Benabou and Tirole, 2002; Benabou, 2013). One essential motive in competitive settings is the desire to win the competition. In several experimental studies, the desire to win has been shown to distort beliefs (see Kunda 1990 for a review of experimental research). Interestingly, this desirability bias features some irregularities that fit our observed pattern. For example, Bar-Hillel and Budescu (1995); Windschitl et al. (2010); Krizan and Windschitl (2009) have noticed that predictions of outcomes are more prone to the desirability bias than are predictions of likelihoods. Why this bias arises with the contemplation of values but not of probabilities is unresolved, but our findings corroborate this observed discrepancy in the literature and expose the importance of modeling the right object of uncertainty in settings that trigger additional desires. One possible hypothesis to be tested would be that subjects weight outcomes that trigger a desire more heavily in their decision making. For instance, the desire to win may trigger a focus on the "best case scenario" that is worth "fighting" for, which, in the CV auction, means receiving a high prize in the lottery. In contrast, in the CP auction, exaggerating the probability of receiving the high prize offers less of an opportunity for wishful thinking (perhaps even less with probability weighting). Whatever the probability turns out to be, it must be multiplied by the prize to determine an expected value, which is guaranteed to be less than the highest value hoped for in the CV auction.

Complexity matters more when pursuing a second objective. An alternative explanation might be that stage 1 valuations could be affected by the additional complexity that Stage 2 reflections introduce into decision making. There is a basic difference in complexity between valuing a lottery in isolation and bidding in an auction for that same lottery. When valuing lotteries, subjects must choose a price that maximizes their profit, or an estimate that maximizes its accuracy. To excel at this task, it suffices to estimate the lottery's expected payoff. By contrast, bidding in the auction trades off two objectives: The optimal bid must be sufficiently high to maximize the chance of winning but, in the case of a win, sufficiently low to generate profits. Hence, in the competitive setting, the desire to win introduces a second objective in addition to the one of maximizing profits from buying. The difficulty of attending to multiple objectives in a competitive environment might interfere with estimating the item's value.

The consequences of this type of complexity are not obvious, however. Increased complexity can make the use of heuristics more appealing and, hence, possibly lead to less-sophisticated bidding, or it can entice subjects to exert more cognitive effort and, hence, enhance sophistication. What is not clear is which auction should be considered more complex and which reaction to this complexity is elicited.

In the following, we investigate reasons why the complexity of our CP and CV auctions might differ.

Complexity is determined by the (mis-)match between information and the response mode. One possibility for why one might consider the CP auction more complex than the CV auction is that the domains of the signal and the response differ in the CP auction. Subjects receive signals in percentage probabilities, but they have to bid in credit values. This additional computational task of converting probability signals into bid values might induce subjects to direct more attention and effort to computing correct estimates; it forces them to multiply their probability estimate with its associated value, which leads them to think in expected value terms. We call this expectation-based reasoning. In contrast, subjects in the CV auction do not engage in such expectation-based reasoning but, rather, use some other heuristic. In such auctions, a bidder might fantasize about a high prize that she might be able to obtain, inducing her to bid higher without taking the probability of that event occurring into account. In CP auctions, any fantasizing must occur over the probability of winning the high prize, requiring a bidder to compute expectations when converting her beliefs into a price. Note that our design highlights the importance of expectation-based reasoning because, in contrast to previous common-value auction experiments, in our design, simply bidding one's signal does not conform with naive expectations.

Previous experimental studies have commented on this hypothesis and shown that responses become more prone to biases when the response mode (a ₵ bid value) is aligned with the object of beliefs (a ₵ signal value) (Tversky et al., 1988; Chapman and Johnson, 1994). For example, Kagel and Levin 1995 have alluded to the importance of the response mode to explain the difference between the high bids in sealed-bid auctions, in which prices must be submitted, and the lower bids observed in Dutch auctions with a binary accept/reject response mode.[26] Whereas in Kagel and Levin's comparison, the object of uncertainty is the same (values) while the response mode differs (prices vs. accept/reject), in our experiments, the response mode is the same (price values), but the object of uncertainty is different (values vs. probabilities). An interesting test of this conjecture would be to investigate the extent of overbidding in CP auctions in which subjects bid for a dollar–a response scale that is identical to a probability scale. Interestingly, we note that the prevailing response mode in auctions on non-performing loans are bids on the dollar.[27]

Complexity is determined by the support of outcomes. Valuing compound lotteries with uncertain probabilities over binary values appears different

[26]In the words of Kagel and Levin, "The behavioral breakdown of the strategic equivalence of first-price and Dutch auctions and of second-price and English auctions is analogous to the preference reversal phenomenon, where theoretically equivalent ways of eliciting individual preferences do not produce the same preference ordering (see the Introduction, section III.F.l and Camerer, chapter 8, section 111.1).12. Psychologists attribute the preference reversal phenomenon to a breakdown in procedural invariance principles so that the weight attached to different dimensions of a problem varies systematically with the response mode employed. In the auctions, prices are higher when bidders must specify a price, as in the first and second-price auctions, compared to the open auctions where the decision is essentially to accept or reject the price that the auctioneer announces. Like the P(rice)-bets in the preference reversal phenomenon, the sealed bid auctions focus attention on the price dimension of the problem, and, like the P-bets, generates somewhat higher prices. On the other hand, the accept/reject decisions involved in the Dutch and English auctions focus attention on profitability, generating somewhat lower prices." Kagel (1995, p. 512)

[27]To the best of our knowledge, whether these auctions on non-performing loans are informationally efficient is still an open empirical question.

from valuing compound lotteries with multiple uncertain values. For example, under standard expected utility, the expected valuation of CP lotteries increases linearly in the average probabilistic belief. This is because with the CP lottery bidders know that, despite the uncertain probability distribution, the winner will get one of two possible lottery outcomes: either the positive value v or nothing. Thus, to assess the CP lottery subjects must simply weight the utility of the high value by a probabilistic belief. With CV lotteries, on the other hand, the order of aggregating uncertainty matters more: averaging utilities that are nonlinear over different values is surely more difficult than considering the utility of a single (average) value belief. Bidders may therefore find it difficult to convert values in utils and keep track of all possible utils in the aggregation process. Perhaps it is this additional cognitive load that in the auction makes reasoning with uncertain values more complicated and leads subjects to focus on a few salient values that may not be good summary statistics. This raises the question of whether overbidding generally decreases with a smaller support of outcomes–particularly in settings like ours, in which the signal is not a good summary statistic.[28]

Since complexity is in the eye of the beholder, we might want to consider which auction the subjects actually perceived to be more complex. In Experiment I's post-experimental survey, while we found a good deal of heterogeneity in how subjects perceived our two auctions, the majority of our subjects preferred the CV to the CP lottery (54% for CV vs. 25% for CP, with the rest being indifferent), and also preferred to participate in the CV over the CP auctions (47% vs. 25%). While preference cannot be considered synonymous with a lack of complexity, it is likely that our subjects preferred the CV lotteries over their CP counterparts because they were better able to comprehend them. What remains unclear is whether a better comprehension of CV lotteries implies that it is easier to deal with them in the bidding process. Further research targeting subjects' reflection and choice processes should be undertaken, and a recent literature has taken up this ambitious challenge (see, e.g., Agranov et al., 2015; Oprea, 2020; Kendall and Oprea,

[28]We have some suggestive evidence against this conjecture in Experiment I where we increased the signal precision in our CV auction, which reduced the support of possible outcomes, and found that subjects overbid more. Of course, a cleaner test is required than this since subjects may also intuitively recognize that shaving becomes less important with more-precise signals. As a result, the size of the support of values remains on the table as a possible explanation.

2021).

To conclude, we think that this paper has taken us some distance in presenting a phenomenon that may be more general than the auction puzzle discussed here. For example, many problems in the real world contain elements of both probabilistic and value uncertainty just as most auctions are hybrids. What our results suggest is that one may be able to influence behavior in such hybrid environments by simply stressing one aspect of the uncertainty at the expense of the other. For example, if focusing attention on the value uncertainty in an auction causes subjects to bid higher, then an auctioneer whose objective is revenue maximization may want to do that. Alternatively, if a social planner aims at maximizing social welfare, then in that same environment she might want to stress the problem's probabilistic uncertainty sine that would mitigate the winner's curse. A possible avenue for future research might be to investigate how, in general hybrid environments, behavior can be altered by stressing the different domains of uncertainty facing decision makers. Since many economic mechanisms are such hybrid entities, our results may have implications for how these mechanisms should be implemented or framed in the real world.

References

M. Abdellaoui and E. Kemel. Eliciting prospect theory when consequences are measured in time units: "time is not money". *Management Science*, 60:1844–1859, 2014.

M. Abdellaoui, A. Baillon, L. Placido, and P. P. Wakker. The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review*, 101:695–723, 2011.

M. Agranov, A. Caplin, and C. Tergiman. Naive play and the process of choice in guessing games. *Journal of the Economic Science Association*, 1:146–157, 12 2015. ISSN 2199-6776. doi: 10.1007/s40881-015-0003-5.

O. Armantier and N. Treich. The rich domain of risk. *Management Science*, 62:1954–1969, 7 2016. ISSN 15265501. doi: 10.1287/mnsc.2015.2215.

P. J. Astor, M. T. P. Adam, C. Jahnig, and S. Seifert. The joy of winning and the frustration of losing: A psychophysiological analysis of emotions in first-price sealed-bid auctions. *Journal of Neuroscience, Psychology, and Economics*, 6:14–30, 2013.

A. Baillon, Y. Halevy, and C. Li. Experimental elicitation of ambiguity attitude using the random incentive system. *Experimental Economics*, 2022.

M. Bar-Hillel and D. Budescu. The elusive wishful thinking effect. *Thinking Reasoning*, 1:71–103, 1 1995. ISSN 14640708. doi: 10.1080/13546789508256906.

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

G. M. Becker, M. H. DeGroot, and J. Marschak. Measuring utility by a single-response sequential method. *Behavioral Science*, 9:226–232, 1964.

R. Benabou. Groupthink: Collective delusions in organizations and markets. *Review of Economic Studies*, 80:429–462, 4 2013. ISSN 00346527. doi: 10.1093/restud/rds030.

R. Benabou and J. Tirole. Self-confidence and personal motivation. *Quarterly Journal of Economics*, pages 871–915, 2002.

O. Bock, A. Nicklisch, and I. Baetge. hroot: Hamburg registration and organisation online tool. *European Economic Review*, 71:117–120, 2014.

W. V. D. Bos, J. Li, T. Lau, E. Maskin, J. D. Cohen, P. R. Montague, and S. M. McClure. The value of victory: Social origins of the winner's curse in common value auctions. *Judgment and Decision Making*, 3:483–492, 2008.

Y. Breitmoser. Knowing me, imagining you: Projection and overbidding in auctions. *Games and Economic Behavior*, 113:423–447, 1 2019. ISSN 10902473. doi: 10.1016/j.geb.2018.10.004.

J. D. Carrillo and T. R. Palfrey. No trade. *Games and Economic Behavior*, 71:66–87, 1 2011.

M. Casari, J. C. Ham, and J. H. Kagel. Selection bias, demographic effects, and ability effects in common value auction experiments. *American Economic Review*, 97:1278–1304, 2007.

G. B. Chapman and E. J. Johnson. The limits of anchoring. *Journal of Behavioral Decision Making*, 7:223–242, 1994.

B. G. Charness and D. Levin. The origin of the winner's curse: A laboratory study. *American Economic Journal: Microeconomics*, 1:207–236, 2009.

G. Charness, D. Levin, and D. Schmeidler. A generalized winner ' s curse: An experimental investigation of complexity and adverse selection, 2014.

G. Charness, D. Levin, and D. Schmeidler. An experimental study of estimation and bidding in common-value auctions with public information. *Journal of Economic Theory*, 179:73–98, 1 2019. ISSN 10957235. doi: 10.1016/j.jet.2018.10.009.

D. L. Chen, M. Schonger, and C. Wickens. otree-an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.

S. H. Chew, R. P. Ebstein, and S. Zhong. Ambiguity aversion and familiarity bias: Evidence from behavioral and gene association studies. *Journal of Risk and Uncertainty*, 44:1–18, 2 2012. ISSN 08955646. doi: 10.1007/s11166-011-9134-0.

J. C. Cox, V. L. Smith, and J. M. Walker. Theory and misbehavior of first-price auctions: Comment. *American Economic Review*, 82:1392–1412, 1992.

V. P. Crawford and N. Iriberri. Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner 's curse and overbidding in private-value auctions? *Econometrica*, 75:1721–1770, 2007.

M. R. Delgado, A. Schotter, E. Y. Ozbay, and E. A. Phelps. Understanding overbidding: Using the neural circuitry of reward to design economic auctions. *Science*, 321:1849–1852, 2008.

P. M. DeMarzo, I. Kremer, A. Skrzypacz, K. Back, S. Board, S. Brusco, N. Garleanu, B. Hollifield, and J. Morgan. Bidding with securities: Auctions and security design. *The American Economic Review*, 95:936–959, 2005.

R. Dorsey and L. Razzolini. Explaining overbidding in first price auctions using controlled lotteries. *Experimental Economics*, 6:123–140, 2003. ISSN 13864157.

N. Du and D. V. Budescu. The effects of imprecise probabilities and outcomes in evaluating investment options. *Management Science*, 51:1791–1803, 2005.

D. Dyer, J. H. Kagel, and D. Levin. A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis. *The Economic Journal*, 99:108–115, 1989.

K. Eliaz and P. Ortoleva. Multidimensional ellsberg. *Management Science*, 62:2179–2197, 2016.

R. Engelbrecht-Wiggans and E. Katok. Regret in auctions: Theory and evidence. *Economic Theory*, 33:81–101, 10 2007. ISSN 09382259. doi: 10.1007/s00199-006-0180-9.

B. Enke and T. Graeber. Cognitive uncertainty, 2021.

P. Eso and L. White. Precautionary bidding in auctions. *Econometrica*, 72: 77–92, 2004. ISSN 00129682.

I. Esponda and E. Vespa. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6:180–202, 2014.

I. Esponda and E. Vespa. Contingent preferences and the sure-thing principle: Revisiting classic anomalies in the laboratory. *Working Paper*, 2021.

E. Eyster. Errors in strategic reasoning, 2019. ISSN 2352-2399.

E. Eyster and M. Rabin. Cursed equilibrium. *Econometrica*, 73:1623–1672, 2005.

E. Filiz-Ozbay and E. Y. Ozbay. Auctions with anticipated regret: Theory and experiment. *The American Economic Review*, 97:1407–1418, 2007.

S. Frederick. Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19:25–42, 2005.

B. Gillen, E. Snowberg, and L. Yariv. Experimenting with measurement error: Techniques with applications to the caltech cohort study, 2019.

J. K. Goeree, C. A. Holt, and T. R. Palfrey. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104:247–272, 2002.

C. Gonzalez-Vallejo, A. Bonazzi, and A. J. Shapiro. Effects of vague probabilities and of vague payoffs on preference: A model comparison analysis. *Journal of Mathematical Psychology*, 40:130–140, 1996.

B. Grosskopf, L. Rentschler, and R. Sarin. An experiment on asymmetric information in first-price common-value auctions: The blessed winner. *Games and Economic Behavior*, 109:40–64, 2018.

Y. Halevy. Ellsberg revisited: an experimental study. *Econometrica*, 75: 503–536, 2007.

G. W. Harrison and J. A. List. Naturally occurring markets and exogenous laboratory experiments: A case study of the winner's curse. *The Economic Journal*, 118:822–843, 2008.

B. C. A. Holt and R. Sherman. The loser's curse. *American Economic Review*, 84:642–652, 1994.

T. Hossain and R. Okui. The binarized scoring rule. *Review of Economic Studies*, 80:984–1001, 7 2013. ISSN 00346527. doi: 10.1093/restud/rdt006.

A. Imas and K. Madarasz. Mimetic dominance and the economics of exclusion: Private goods in public context, 2020.

J. H. Kagel. Auctions: A survey of experimental research, 1995.

J. H. Kagel and D. Levin. The winner's curse and public information in common value auctions. *American Economic Review*, 76:894–920, 1986.

J. H. Kagel and D. Levin. *Common value auctions and the winner's curse.* Princeton University Press, princeton edition, 2002. ISBN 0691016674.

J. H. Kagel and J.-F. Richard. Super-experienced bidders in first-price common-value auctions: Rules of thumb, nash equilibrium bidding, and the winnerâs curse. *Review of Economics and Statistics*, 83:408–419, 2001.

J. H. Kagel, D. Levin, R. C. Battalio, and D. J. Meyer. First-price common value auctions: Bidder behavior and the "winner's curse". *Economic Inquiry*, 27:241–258, 1989.

C. Kendall and R. Oprea. On the complexity of forming mental models, 2021.

C. Koch and S. P. Penczynski. The winner's curse: Conditional reasoning and belief formation. *Journal of Economic Theory*, pages 57–102, 2018.

M. G. Kocher, J. Pahlke, and S. T. Trautmann. An experimental study of precautionary bidding. *European Economic Review*, 78:27–38, 2015.

Z. Krizan and P. D. Windschitl. Wishful thinking about the future: Does desire impact optimism? *Social and Personality Psychology Compass*, 3: 227–243, 5 2009. doi: 10.1111/j.1751-9004.2009.00169.x.

K. M. Kuhn and D. V. Budescu. The relative importance of probabilities, outcomes, and vagueness in hazard risk decisions. *Organizational Behavior and Human Decision Processes*, 68:301–317, 1996.

Z. Kunda. The case for motivated reasoning. *Psychological Bulletin*, 108: 480–498, 1990.

A. Lange and A. Ratan. Multi-dimensional reference-dependent preferences in sealed-bid auctions - how (most) laboratory experiments differ from the field. *Games and Economic Behavior*, 68:634–645, 3 2010. ISSN 08998256. doi: 10.1016/j.geb.2009.08.008.

Z. Li, J. Muller, P. P. Wakker, and T. V. Wang. The rich domain of ambiguity explored. *Management Science*, 64:3227–3240, 7 2018. ISSN 15265501. doi: 10.1287/mnsc.2017.2777.

A. Martinez-Marquina, M. Niederle, and E. Vespa. Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109: 3437–3474, 2019. ISSN 19447981. doi: 10.1257/aer.20171764.

P. R. Milgrom and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis Policy*, 2, 2003.

J. Moser. Hypothetical thinking and the winner's curse: an experimental investigation. *Theory and Decision*, 87:17–56, 2019.

T. Neugebauer and R. Selten. Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets. *Games and Economic Behavior*, 54:183–204, 2006.

M. K. Ngangoue and G. Weizsacker. Learning from unrealized versus realized prices. *American Economic Journal: Microeconomics*, 13:174–201, 5 2021. ISSN 1945-7669. doi: 10.1257/mic.20180268.

R. Oprea. What makes a rule complex? *American Economic Review*, 110: 3913–3951, 12 2020. ISSN 19447981. doi: 10.1257/AER.20191717.

A. A. Salo and M. Weber. Ambiguity aversion in first-price sealed-bid auctions. *Journal of Risk and Uncertainty*, 11:123–137, 1995.

K. H. Schlag, J. Tremewan, and J. J. van der Weele. A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics*, 18:457–490, 9 2015. ISSN 13864157. doi: 10.1007/s10683-014-9416-x.

P. J. Schoemaker. Choices involving uncertain probabilities: Tests of generalized utility models. *Journal of Economic Behavior and Organization*, 16:295–317, 1991.

A. Tversky, S. Sattath, and P. Slovic. Contingent weighting in judgment and choice. *Psychological Review*, 95:371–384, 1988.

R. Wilson. A bidding model of perfect competition. *The Review of Economic Studies*, 44:511–518, 1977.

P. D. Windschitl, A. R. Smith, J. P. Rose, and Z. Krizan. The desirability bias in predictions: Going optimistic without leaving realism. *Organizational Behavior and Human Decision Processes*, 111:33–47, 1 2010. ISSN 07495978. doi: 10.1016/j.obhdp.2009.08.003.

A Experimental Interfaces

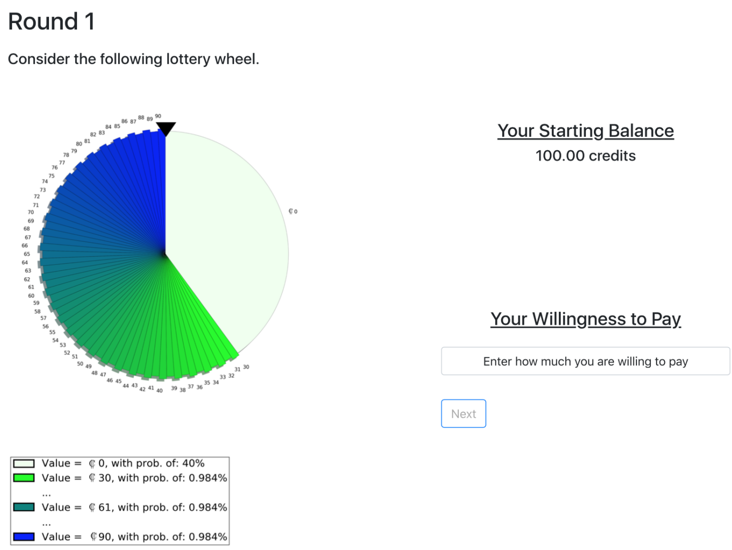

Figure A 1: Example for Reduced Lottery in Experiment II

Figure A 1 shows the display of a reduced CV lottery in Experiment II. The display of reduced CP lotteries was similar but with only two possible outcomes.

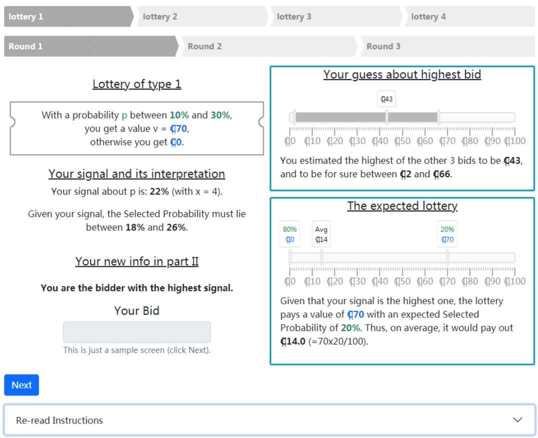

Figure A 2: Example of CP Interface in Part II of Experiment IV

Figure A 2 shows the interface in the second part of Experiment IV. The two framed boxes correspond to the two pieces of information given to the subjects. The upper right box reminds the subject of their belief about the competitive bid for the exact same lottery. The lower box renders information about the lottery to be expected conditional on having the highest signal in the market.

B Descriptive Statistics

B.1 Bid and Price Factors

Table A1: Bid Factors (BF) in Reduced Sample

		Naive bid (bid-bid^{Naive})		Break-Even bid (bid-bid^{BE})		Nash-Eq. bid (bid-bid^{RNNE})	
		mean	median	mean	median	mean	median
Exp. I	CV	10.53*** (1.872)	13.6***	12.26** (1.871)	15.32***	13.33*** (1.870)	16.45***
	CP	-4.60*** (1.424)	-3.80**	-2.84** (1.425)	-2.04	-1.77 (1.425)	-0.82
	Diff.	15.13*** (2.346)	17.4***	15.11*** (2.345)	17.36***	15.10*** (2.345)	17.26***
Exp. III	CV	16.14*** (2.570)	19***	17.80*** (2.559)	21***	17.53*** (2.791)	20.60***
	CP	-1.33 (1.225)	-0.40	-0.64 (1.178)	1	-0.31 (1.621)	0.80
	Diff.	17.47*** (2.819)	19.4***	17.86*** (2.789)	20***	17.84*** (3.199)	19.80***
Exp. IIIB	CV	12.95*** (1.550)	13.2***	14.55*** (1.558)	15***	15.9*** (1.689)	16.8***
	CP	-0.60 (1.099)	0	1.23 (1.141)	1**	0.84 (1.380)	1.80*
	Diff.	13.54*** (1.885)	13.2***	13.32*** (1.941)	14***	15.07*** (2.166)	15***
Exp. IV	CV	9.46*** (1.828)	8.4***	10.66*** (1.828)	9.36***	11.45*** (1.828)	10.00***
	CP	-1.61 (1.750)	-0.4	-0.41 (1.750)	0.76	0.39 (1.750)	1.60
	Diff	11.07*** (2.522)	8.80***	11.07*** (2.522)	8.60**	11.07*** (2.521)	8.40**

Note: Cluster robust standard errors (CRSE) clustered at subject level in parentheses. P-values: *: p-value<.1,**: p-value<.05, ***: p-value<.01. Clustering standard errors by sessions do not alter tests results and accounts for approximately 1% of the residual variance. In the remaining analyses standard errors are clustered at the subject level.

Tables A1 and A2 present the computation of bid factors relative to all three benchmarks: the naive, the break-even and the RNNE bid function. Positive (negative) bid factors imply that average bids are above (below) the benchmark.

Table A2: Bid Factors (BF) – Winning Bids in Exp. I

		CV	CP	Diff
Naive bid	mean	24.68***	9.50***	15.18***
(bid-bid^{Naive})		(0.731)	(1.076)	(1.299)
	median	23.40***	8.40***	5.00***
Break-Even bid	mean	26.41***	11.25**	15.17***
(bid-bid^{BE})		(0.727)	(1.099)	(1.316)
	median	25.16***	10.28***	14.88***
Nash-Eq. bid	mean	27.47***	12.33***	15.15***
(bid-bid^{RNNE})		(0.725)	(1.108)	(1.322)
	median	26.40***	11.52***	14.88***

Note: Cluster robust standard errors (CRSE) at subject level in parentheses. P-values: *: p-value<.1,**: p-value<.05, ***: p-value<.01.

Table A3: Price Factors in Exp. II

		CV	CP	Diff
Part CL with signal				
Price Factor	mean	5.12***	-1.03	6.16***
(bid-$price - E[L\|s]$)		(1.676)	(1.001)	(1.945)
	median	2.4***	-1.8***	4.2***
Part CL without signal				
Price Factor	mean	0.75	-3.00**	3.75*
(bid-$price - E[L\|s]$)		(1.679)	(1.410)	(2.188)
	median	-4***	-4***	0
Part RL				
Price Factor	mean	-2.93*	-1.37	-1.56
(bid-$price - E[L\|s]$)		(1.684)	(1.128)	(2.022)
	median	-4**	-1	-3

Note: CRSE in parentheses. P-values: *: p-value<.1,**: p-value<.05, ***: p-value<.01.

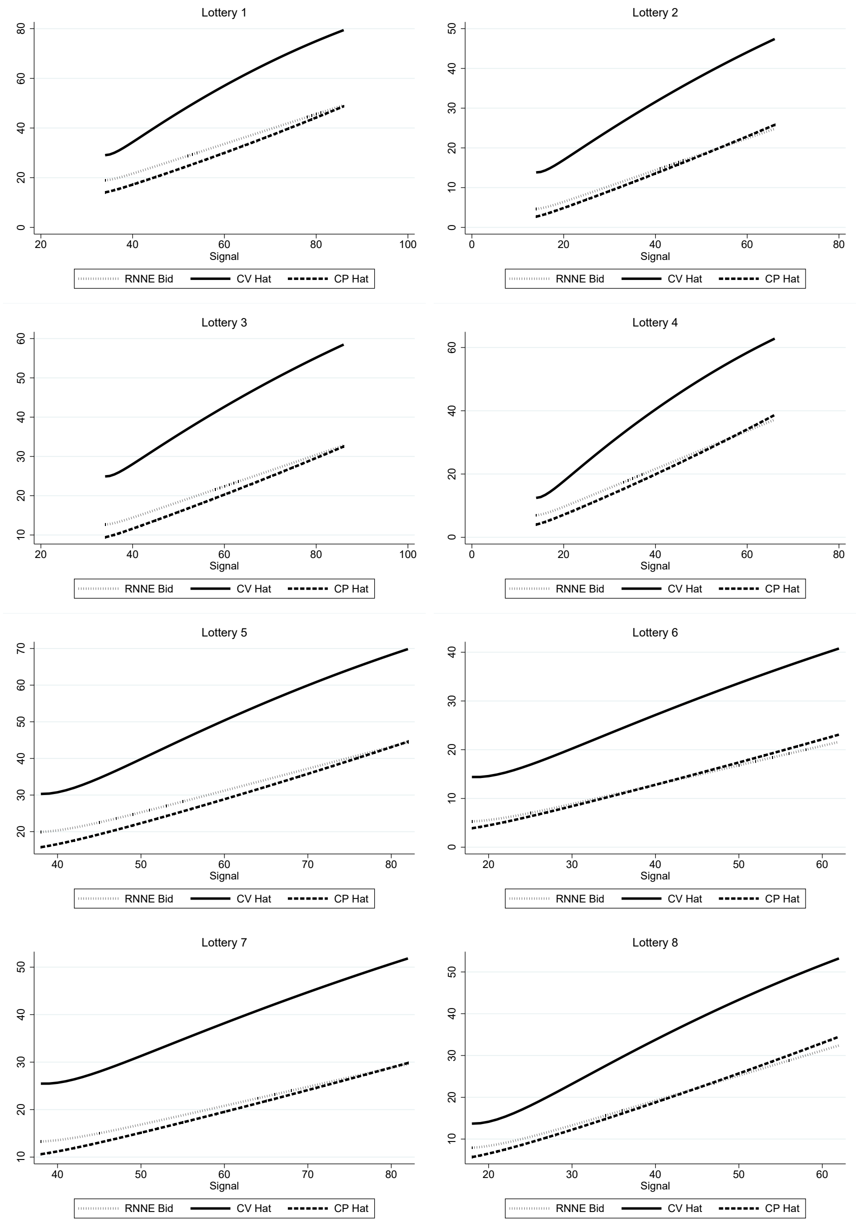

Figure A 3: Estimated Median Bids in CV and CP Auctions by Lottery Types (Experiment I)

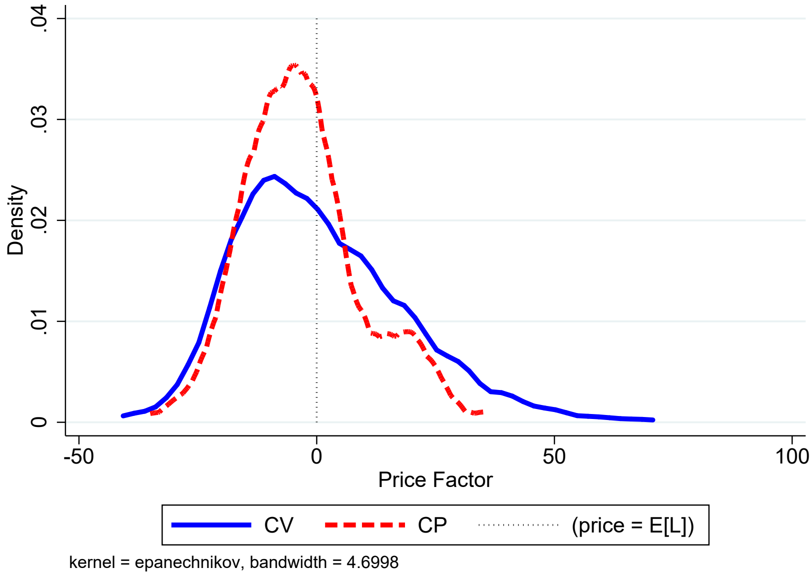

Figure A 4: Distribution of *Ex ante* Price Factors for Compound Lotteries ($=w_i - E[L]$) in Treatments CVL (solid) and CPL (dashed)

B.2 Adverse Selection

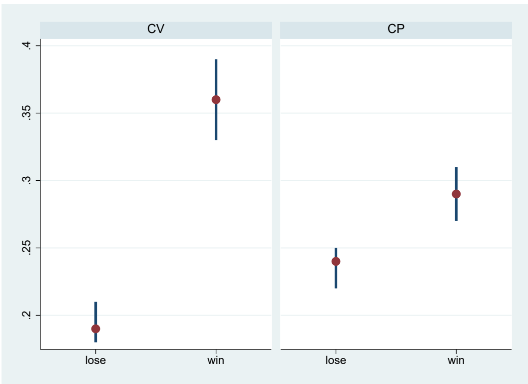

Figure A 5: Predicted probabilities of having the highest signal conditional on winning and losing

While the winner's curse is less severe in CP than in CV, it is not clear whether this is because subjects reason better through the adverse selection problem with probabilistic uncertainty or because winning reveals less information in the first place. To shed more light on the extent of the adverse selection problem, we juxtapose how informative two events are: the event of having the highest signal and the event of winning. The lottery's actual average payoff conditional on having the highest signal tells us how much, in each auction, subjects must actually shave their bids to break even. This is the empirical benchmark for updating in a Bayesian manner. To this end, we regress the lottery outcome on the signal and a dummy that takes the value one if the signal is the highest in the auction. The lottery's average payoff conditional on winning, on the other hand, tells us how much subjects can actually learn from winning the auction. In a similar manner, we regress the average lottery outcome on the signal and a dummy that takes the value one if the bidder with the same signal won the auction. Note that the two events of having the highest signal and of winning will convey the same information if the winner is always the bidder with the highest signal, thereby providing the need to account for an adverse selection effect in bidding. We find that in the data having the highest signal in CV (CP) requires adjusting expectations downward by an average of ₵-3.08 (₵-2.65, $p < 0.001$ in either case). In contrast, winning is not as informative, in particular in CP auctions. Winning an auction allows one to adjust expectations only by a fraction of what can actually be learned from having the highest signal: up to 35.71% in CV but only 10.07% in CP. Differences in the reduced sample are even more striking: 60.38% in CV versus 3.62% in CP. This is a direct result of a weaker correlation between signals and winning the auction. Appendix Figure A 5 plots the predicted probabilities of having the highest signal conditional on winning. The risk of falling prey to the winner's curse is present in CV but marginal in CP. Winning in CV increases the likelihood of having the highest signal from 19% to 36%, which is three times more than in CP (24% to 28%).

B.3 Decision weights

To assess the importance of signals relative to the known component of the lottery, we estimate the elasticity of the bid with respect to the known and the unknown (i.e., signal) component of the lottery. To this end, we use a simple Cobb-Douglas bidding function in the form of $b(s_i) = k^\alpha \cdot s^\beta$. A naive agent, for instance, would bid $E[L|s] = k^\alpha \cdot s^\beta$ with $\alpha = \beta = 1$.

We use the marginal rate of substitution (MRS) to compare the esti-

Table A4: MEDIAN REGRESSION COEFFICIENTS IN BIDDING

$ln(bid)$	(CV)	(CP)	(Diff)
$ln(k)$	$0.244^{*\dagger\dagger\dagger}$	$0.749^{***\dagger}$	-0.505^{**}
	(0.130)	(0.137)	(0.197)
$ln(s_i)$	$1.096^{***\dagger\dagger\dagger}$	1.254^{***}	-0.158
	(0.030)	(0.161)	(0.171)
$Cons$	-1.488^{***}	-4.818^{***}	3.331^{***}
	(0.570)	(0.687)	(0.901)
N	3253	2564	5817
$Subjects$	52	39	91
R^2	0.015	0.072	
$F-Test$	0.000	0.040	
MRS	$\approx 0.22\frac{s}{k}$	$\approx 0.60\frac{s}{k}$	

Note: Median regression with CRSE in parentheses. Significant difference from 0: *: p-value<.1,**: p-value<.05, ***: p-value<.01. Significant difference from 1: †: p-value<.1,††: p-value<.05, †††: p-value<.01. F-test refers to a test on equal weighting of known parameter and signal ($\alpha = \beta$).

Table A5: MEDIAN REGRESSION COEFFICIENTS IN PRICING

$ln(bid)$	CVL	CPL	Diff.
$ln(k)$	$0.546^{***\dagger\dagger\dagger}$	$0.810^{***\dagger\dagger\dagger}$	-0.264
	(0.157)	(0.056)	(0.067)
$ln(s_i)$	0.947^{***}	0.974^{***}	-0.027
	(0.069)	(0.044)	(0.066)
$Cons$	-2.500^{***}	-3.847^{***}	1.347
	(0.721)	(0.346)	
N	4256	4000	8256
$Subjects$	54	50	104
R^2	0.141	0.3919	
$F-Test$	0.0209	0.0017	
MRS	$\approx 0.57\frac{s}{k}$	$\approx 0.83\frac{s}{k}$	

Note: Median regression with cluster robust standard errors (CRSE) at subject-level in parentheses. Significant difference from 0: *: p-value<.1,**: p-value<.05, ***: p-value<.01. Significant difference from 1: †: p-value<.1,††: p-value<.05, †††: p-value<.01.

mated bidding functions. The MRS represents here how much units of the signal subjects are willing to trade against a unit of the known parameter to maintain the same bid. For a naive bidder, the MRS equals $\frac{\alpha s}{\beta k} = \frac{s}{k}$. For our parameter variation, MRS under Nash equilibrium should be close to $\frac{s}{k}$. In both auction formats, the estimated MRS is smaller than $\frac{s}{k}$ ($\approx 0.23\frac{s}{k}$ in CV vs. $\approx 0.60\frac{s}{k}$ in CP in Appendix Table A4), indicating that subjects overweighted their private signal but underweighted the known component. Subjects in CV auctions put relatively more weight on the signal compared to those in CP auctions. Similar results are obtained with the pricing data, where MRS are closer to the naive benchmark $\frac{s}{k}$ (see Appendix Table A5). While subjects put more attention on signals in both CV and CP formats it is important to keep in mind that these signals are about different components of the lotteries. In CV treatments, subjects paid more attention to values in the lottery whereas in CP treatments they rather focused on the probabilities. In a nutshell, it appears that the uncertain component determines how subjects allocate their attention to different features of the auctioned item.

B.4 Information processing in Experiment II

We study the importance of information processing in the decision problem. The empirical value of a signal is obtained by comparing subjects' willingness to pay before and after receiving signal s_i. To this end, we regress subjects' willingness to pay w_i on objective measures like the prior expected value $E[L]$ and the information content of the signal given by the change in expectations $(E[L|s_i] - E[L])$. We also include a dummy D_{signal} that equals one when the willingness to pay was submitted after observing a signal.

As shown in Table A6, we do not find substantial differences in the way subjects processed these value and probability signals (consistent with our results in Table A5). Under risk-neutral expected utility, pricing occurs at the expected value. That is, an increase of ₵1 in prior and interim beliefs is reflected in an equivalent increase of ₵1 in prices, while uncertainty premia (captured by the constant and the dummy variable) should be zero (cf. first column of Table A6). In treatment CVL, subjects reacted reasonably to variations in both prior parameters and signals as the corresponding coefficients do not substantially differ from the RNEU benchmark. In treatment CPL, subjects slightly underreacted to variations in the parameters, but

Table A6: Median Regression Coefficients

WTP	Rational	CV	CP	Diff
$E[L]$	1	0.898***†	0.830***†††	0.068
		(0.053)	(0.039)	(0.069)
$E[L\|S] - E[L]$	1	1.051***	0.905***†††	0.146
		(0.084)	(0.030)	(0.103)
D_{signal}	0	5.292***	2.074**	3.218
		(2.027)	(0.947)	(2.419)
$Cons$	0	-0.335	0.074	-0.409
		(2.349)	(1.497)	(2.783)
N		4688	4400	9088
$Subjects$		54	50	104
R^2		0.234	0.390	
$F - Test$		0.0080	0.0001	

Note: Median regression with CRSE in parentheses. Significant difference from 0: *: p-value<.1,**: p-value<.05, ***: p-value<.01. Significant difference from NE coefficient: †: p-value<.1,††: p-value<.05, †††: p-value<.01.

more importantly coefficients do not differ from the ones in CVL. Hence, subjects processed value and probability signals similarly.

A striking observation is that in both CVL and CPL, the mere fact of observing a signal significantly increased WTP by ₵ 5 and ₵ 2, respectively. In other words, even when objective prior and interim expectations coincided, subjects were willing to pay more after observing a signal. This could be rationalized to some extent with a reduced uncertainty premium in interim beliefs, as seen in the treatment CPL where after getting a signal subjects bid closer to expected value. Rather surprising is that in CVL subjects bid, on average, even above expected values after seeing a signal, implying that the mere fact of getting a signal led subjects to move from an average positive to a negative uncertainty premium.

B.5 Reduction of Compound Lotteries in Experiment II

In CVL, subjects made no such distinction when valuing reduced and compound CV lotteries. The median premium for compound risk in values is zero, suggesting that compound risk in values may not necessarily have been perceived as such. In CPL, they chose a small average compound risk premium of ₵2 for CP lotteries, pricing the reduced CP lotteries slightly higher than their compound analog. There are some order effects in the comparison of reduced and compound lotteries. Whether subjects first saw reduced or compound lotteries matters but only in the CV treatments. On average, subjects chose similar WTP with and without compound risk when they valued the compound lottery before its reduced form version (median compound risk premium of 0 in CV lotteries). Seeing the reduced lottery first, on the other hand, *increased* (rather than decreased) their WTP for the compound version of CV lotteries by ₵3.5. In other words, the median premium for compound risk defined over values is even negative, implying that the average subject was more averse to the reduced than to the compound version of the CV lottery.

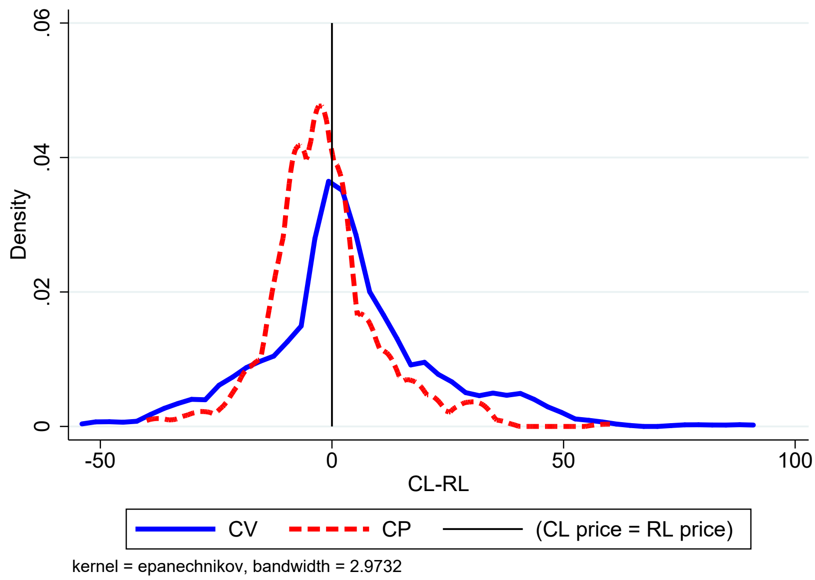

Figure A 6: Differences in Pricing of Compound and Reduced Lotteries in Treatments CVL (solid) and CPL (dashed)

B.6 Experiment III

Figures A 7a and A 7b show the robustness of our Experiment I results to our Experiment III design where subjects bid against previous Experiment I subjects. Like in Experiment I, in Experiment III subjects in CV bid significantly more than their peers in CP. While in both CV and CP auctions, subjects in Experiment III bid slightly higher than their peers in Experiment I, the difference in bids between CV and CP remains of similar magnitude: Subjects in CV bid, on average, ₵ 19.2 more than their peers in CP (compared to ₵ 17.60 for the same auctions in Experiment I). Hence, our results are robust to our design modification in Experiment III.

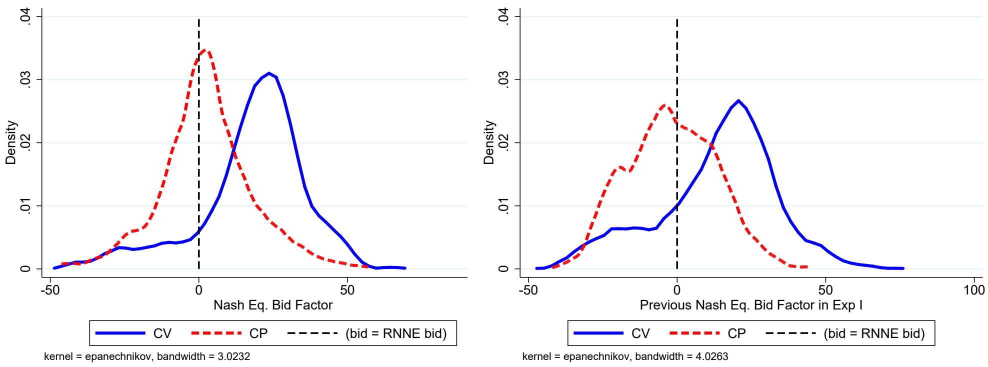

(a) Exp. III

(b) Exp. I with identical signals

Figure A 7: Comparing bid factors in Exp. III & I

C Individual Covariates

C.1 Attitudes toward risk, compound risk and ambiguity

In the last part of Experiments I & II, we elicited subjects attitudes toward risk, compound risk and ambiguity. Subjects started this part by first selecting the payoff relevant task. To this end, they threw a dice, knowing that the number on top of the dice would define the selected task. The correspondence between the dice numbers and the tasks were, however, revealed only at the end of the experiment (Baillon et al., 2022). The exchange rate remained the same ($1 for 6 credits), but payoffs from the main part of the experiment were weighted more heavily than those in this last part (3:1).

This part consisted of only six decision problems. The six decision screens corresponded to three types of decision problems with two replicate measurements each.

C.2 Elicitation

We elicited risk attitudes with a multiple price list akin to Abdellaoui et al. (2011) and Gillen et al. (2019). Subjects faced virtual bags with red and blue chips. First, subjects chose the color to bet on and then gave their certainty equivalent (henceforth CE) for their chosen bet. Risky bets were implemented with the following lottery (100:0.5;0) and (150:0.5;0) (i.e., a 50% chance of winning ₵100 / ₵150 or otherwise nothing).

To implement bets with compound risk, subjects were told that the computer would first randomly select one virtual bag from a set of virtual bags containing each a different mixture of red and blue balls (Figure A 8 shows an example of the screen for a bag with 20 chips), and would then randomly draw a chip from the selected bag. Subjects received ₵100 (₵150 in the replicate measurement) if the color of the drawn chip matched the color they bet on.

The implementation of ambiguous bets was similar, except that the mixture of red and blue chips was determined ex ante by a research affiliate and was not known to subjects.

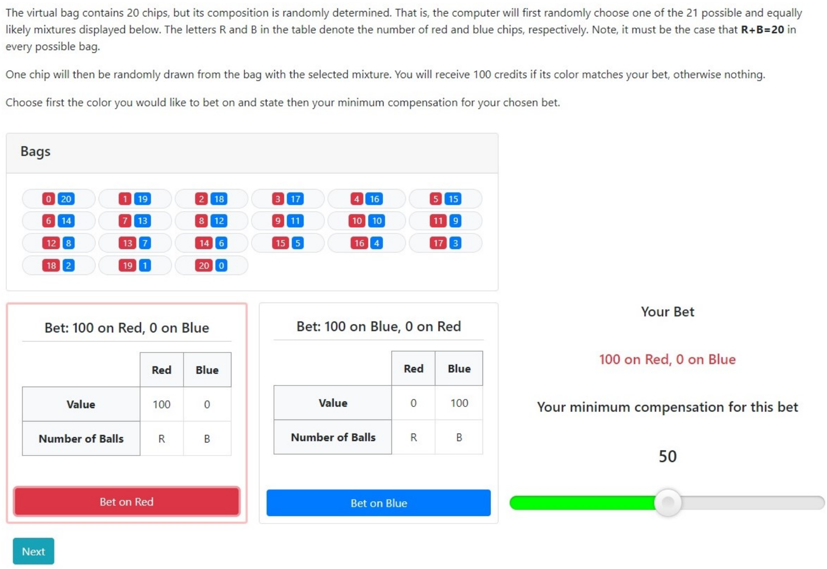

Figure A 8: Example for a decision screen to elicit attitudes toward compound risk (after selecting to bet on red and a certainty equivalent of 50 credits.)

C.3 Descriptive statistics

Methods. We classify attitudes as averse toward a type of uncertainty if subjects' prices display a premium for the lottery. The premium is given by the difference between the lottery's expected value and the subject's CE. A positive (negative) premium reflects aversion (proclivity).

We mitigate possible measurement errors by taking the mean of the two replicate measurements: To this end, we first normalize the CE by the lottery's expected value and average the normalized CE across the two replicate measurements.[29,30] Note that all decisions under uncertainty should be affected by a risk premium, if a subject is not risk-neutral. In a crude attempt to control for risk attitudes in decisions with compound risk and ambiguity, we subtract the subject's average risk premium from the chosen premium

[29] For the ambiguous bets, we assume uniform beliefs over possible probabilities to compute the lotteries' expected value.

[30] Most subjects were also consistent in their attitudes, especially in their attitudes toward ambiguity. The redundant measures yield the same classification for 71.15%, 75.96% and 79.81% of the subjects regarding attitudes toward risk, compound risk and ambiguity, respectively (in the full sample).

for lotteries with compound risk and ambiguity (cf. Gillen et al., 2019). This yields a conservative measure of the premia for compound risk and ambiguity since risk premia for binary lotteries should be highest when the success probability equals 50% (as in the risky lotteries). Thus, subjects who were less averse toward compound risk/ ambiguity than toward risk exhibit a negative premia for compound risk and ambiguity (applies to 59 (60) out of 195 subjects for the compound risk (ambiguity) premium).

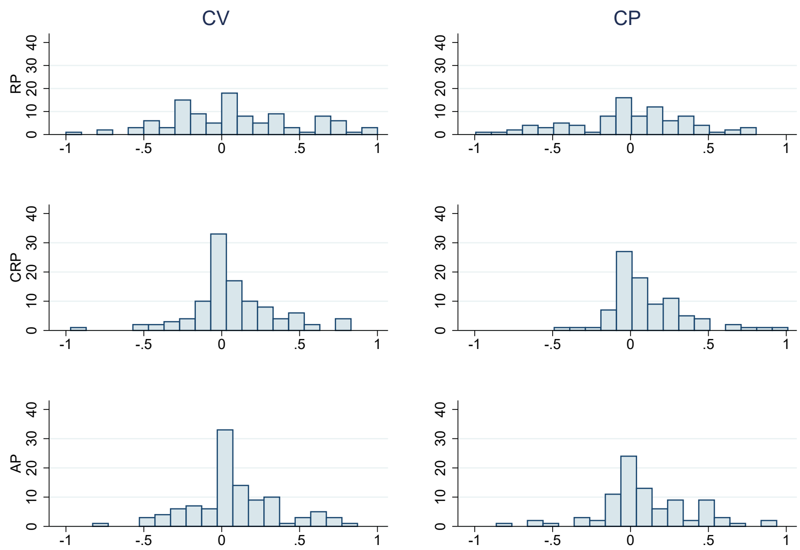

Figure A 9: Distribution of Premia in Exp. I II – by Treatments CV (left) and CP (right).

Results. Figure A 9 shows the distribution of risk, compound risk and ambiguity premia, averaged across the two duplicate measures. In general, most subjects were averse toward uncertainty.

Distributions of premia are not significantly different from each other across treatments (the Kolmogorov-Smirnov statistics yields p-values of $p = 0.21, p = 0.45, p = 0.89$ for risk, compound risk and ambiguity premia, respectively). Most subjects chose a premium close to zero, and attitudes toward compound risk and ambiguity are positively correlated (consistent with Halevy (2007)'s finding). The pairwise correlation coefficients are $\rho_{RC} = -0.24, \rho_{RA} = -0.10, \rho_{CA} = 0.54$.

C.4 Individual Characteristics

In general, individual characteristics do not significantly differ between the CV and CP treatments. The measures of the cognitive reflection tests (CRT) are higher in the treatments III-IV and have to be interpreted with caution because the experiment was conducted online.

Table A7: MEANS OF INDIVIDUAL CHARACTERISTICS BY TREATMENT IN REDUCED SAMPLE

		Male	Age	CRT	RP	CRP	AP
	CV	0.538***	21.962***	1.519***	0.006	0.095**	0.120***
		(0.070)	(0.355)	(0.149)	(0.053)	(0.041)	(0.036)
I	CP	0.615***	22.333***	1.308***	-0.076	0.153***	0.115**
		(0.079)	(0.460)	(0.160)	(0.065)	(0.033)	(0.050)
	Diff	-0.077	-0.372	0.212	0.082	-0.058	0.005
		(0.105)	(0.581)	(0.219)	(0.084)	(0.052)	(0.062)
	CV	0.434***	21.415***	1.389***	0.151**	0.057	0.043
		(0.069)	(0.386)	(0.164)	(0.059)	(0.035)	(0.043)
II	CP	0.480***	21.440***	1.5***	0.050	0.094**	0.090**
		(0.071)	(0.365)	(0.154)	(0.048)	(0.037)	(0.042)
	Diff	-0.046	-0.025	-0.111	0.101	-0.037	0.047
		(0.099)	(0.531)	(0.225)	(0.076)	(0.051)	(0.060)
	CV	0.450***	23.050***	2.300***	-0.028		
		(0.114)	(0.671)	(0.242)	(0.159)		
III	CP	0.389***	23.278***	2.500***	0.069		
		(0.118)	(0.645)	(0.217)	(0.124)		
	Diff	0.061	-0.228	-0.200	-0.097		
		(0.164)	(0.931)	(0.325)	(0.202)		
	CV	0.522***	21.913***	2.391***	0.068		
		(0.106)	(0.569)	(0.137)	(0.103)		
IIIb	CP	0.538***	21.346***	2.269***	0.223***		
		(0.100)	(0.474)	(0.172)	(0.724)		
	Diff	-0.017	0.567	-0.122	-0.155		
		(0.146)	(0.740)	(0.219)	(0.125)		
	CV	0.448***	23.279***	2.655***	-0.023		
		(0.094)	(0.660)	(0.114)	(0.177)		
IV	CP	0.333***	22.286***	2.048***	0.141		
		(0.105)	(0.492)	(0.243)	(0.169)		
	Diff	0.115	0.990	0.608**	-0.164		
		(0.141)	(0.825)	(0.268)	(0.219)		
	CV	0.480***	22.130***	1.865***	0.049	0.076***	0.081***
		(0.038)	(0.221)	(0.0084)	(0.043)	(0.027)	(0.028)
I-IV	CP	0.494***	21.981***	1.773***	0.062*	0.120***	0.101***
		(0.040)	(0.214)	(0.089)	(0.035)	(0.025)	(0.032)
	Diff	-0.013	0.150	0.092	-0.013	-0.040	-0.020
		(0.055)	(0.307)	(0.122)	(0.055)	(0.037)	(0.043)

Note: *: p-value<.1,**: p-value<.05, ***: p-value<.01. Robust standard errors clustered by subject in parentheses.