Damaging conflict: All-pay auctions with negative spillovers and bimodal bidding

James W. Boudreau*
Kristy Buzard†
Timothy Mathews‡
Lucas Rentschler§
Shane D. Sanders¶

February 21, 2025

Abstract

We investigate how the presence of a negative externality in an all-pay auction influences bidding behavior in the laboratory. In the standard risk-neutral model, Nash equilibrium predicts no difference in strategies between treatments with and without the externality. Our experimental results provide some support for this prediction, as average bids do not differ significantly across treatments and generally align with equilibrium benchmarks. However, bidding distributions in both treatments exhibit a pronounced bimodal pattern that is consistent with previous all-pay auction experiments but inconsistent with risk-neutral Nash predictions. To account for these features, we evaluate two models of bounded rationality: quantal response equilibrium (QRE) and noisy introspection, both incorporating prospect theory-inspired preferences. While both models can rationalize bimodal bidding, QRE provides the superior fit while also predicting the location of bid peaks across treatments. These results reinforce the growing evidence that while subjects may not mix strategies exactly as prescribed by Nash equilibrium, their bidding behavior remains broadly equilibrium-consistent. Our findings also support the role of reference dependence and loss aversion in explaining bimodal bidding and strategic behavior more generally.

JEL Classification Codes: C72, C92, D44, D62, D72.

Keywords: All-pay auctions; bimodal bidding; contests; rent-seeking; spillovers; prospect theory.

*Kennesaw State University, E-mail: jboudre5@kennesaw.edu.
†Syracuse University, Email: kbuzard@syr.edu.
‡Kennesaw State University, Email: tmathew7@kennesaw.edu.
§Utah State University, E-mail: lucas.rentschler@usu.edu.
¶Syracuse University, Email: sdsander@syr.edu

1 Introduction

We present the results of laboratory experiments on a type of all-pay auction game that incorporates negative externalities. We term these games *contests with damages*. Contests with damages are two-player games in which each player simultaneously chooses a bid to win a prize of common value. The player who bids the most wins the prize with certainty, but both players sacrifice their bids regardless of whether they win or lose as in the standard all-pay auction. The difference from the standard all-pay auction is that each player's bid also imposes some amount of "damage" in terms of increased cost upon the other player regardless of winning or losing. The players' own bid is therefore not only costly to themselves, but also injurious to their rival.

One interpretation of this model could be a political lobbying contest in which both parties spend to win office. The spending is sunk on both sides in pursuit of victory, but in the process each also damages the opposing side's reputation, win or lose. Similarly, the model could represent a battlefield or armed-conflict situation, with spending in the pursuit of victory by one side also felt (in payoff terms) in the form of damaging material consequences for the other. More generally, the model captures the idea that competition can be destructive in ways that are not internalized: as players spend more to win the prize they are interested in, they also impose costs on their opponent in payoff terms that have no bearing on their own payoffs.

A model of contests with damages is especially interesting for experimental analysis because when players have risk-neutral preferences it is a specific variant of the "rank-order spillover" contest model developed by Baye et al. (2012), which provides testable results. While the general spillover model's parameters allow for a wide variety of contests with both positive and negative externalities, we focus on the version that captures an all-pay auction with only negative externalities exerted equally by both players. Baye et al. (2012) established that this specific version has a symmetric Nash equilibrium (NE) identical to the standard all-pay auction, providing an opportunity for comparison in terms of subject behavior: players should exhibit the same levels of spending (on average) whether or not damages are present in the model. Our results support that specific prediction, but also lead to further questions.

Comparing the results of our experimental treatments of two-player all-pay auctions with and without damages, we find that subjects behave similarly in either case. Average bids were not significantly different across treatments, equal to approximately half subjects' valuation for the prize as predicted by the model with risk-neutral preferences. In addition, however, both treatments displayed similar patterns of *bimodal* bidding, with large proportions of subjects choosing minimal or maximal bids. This presents a discrepancy, as bimodal bidding is not consistent with equilibrium behavior in the all-pay auction model with risk-neutral preferences, with or without damages. According to the risk-neutral model, the perfectly discriminating nature of an all-pay auction means players should mix uniformly on the interval $[0, V]$, where V is the prize value.

Though we are the first (to our knowledge) to include damages in experimental treatments of all-pay auctions, previous lab experiments involving standard all-pay auctions have revealed two consistent results that are especially relevant to our own. First, average bidding does seem to be in line with equilibrium predictions when experiments involve only two bidders (Potters et al., 1998; Ernst and Thöni, 2013), though when experiments involve more than two subjects competing the tendency is to overbid relative to Nash equilibrium (Gneezy and Smorodinsky, 2006; Lugovskyy et al., 2010). Second, regardless of the number of bidders, experimental studies of all-pay auctions have consistently displayed bimodal bidding patterns, with subjects placing very low (close to zero) and very high (close to V) bids more frequently than others (Müller and Schotter, 2010; Ernst and Thöni, 2013; Dechenaux et al., 2015), contrary to the uniform distribution predicted by NE in the linear model. The question we then turn to is: why bimodal bidding? Again, the risk-neutral model can not reconcile this phenomenon.

The leading explanation for bimodal bidding comes from prospect theory (Kahneman and Tversky, 1979), as it can be shown that bimodal strategies are consistent with NE in an all-pay auction model where players have reference-dependent preferences and loss aversion (Müller and Schotter, 2010; Ernst and Thöni, 2013) and a survey by Dechenaux et al. (2015) cites this as the leading explanation for such behavior in the lab. The issue for our setting, however, is that incorporating such preferences into a contest with damages makes the model prohibitively difficult to solve, preventing us from obtaining closed-form NE predictions for the case with damages. We therefore turn to a more general form of equilibrium, quantal response equilibrium (QRE), which nests NE and can be solved for computationally.

Introduced by McKelvey and Palfrey (1995), QRE is similar in spirit to NE, but instead of requiring each player to best-respond to the other with perfect rationality, it posits that each player mixes pure strategies based on their relative payoffs according to a logit response function under the assumption that other players do the same. A noise parameter in the response function provides flexibility, allowing for a spectrum of behavior from perfect rationality (NE) to uniform randomness. Incorporating reference-dependent preferences and loss aversion, we numerically solve for QRE in our contests, fitting the model's parameters to our data using maximum likelihood. For comparison, we also estimate an alternative model of bounded rationality—noisy introspection (NI)—introduced by Goeree and Holt (2004).

Our results indicate that subjects' bidding behavior aligns closely with equilibrium predictions, with the QRE model providing the stronger fit. The estimated QRE strategies successfully capture the key features of the observed bidding distributions, including bimodality and occasional overbidding. Crucially, QRE also predicts a shift in the second peak of the bid distribution across treatments, a prediction supported in our data. In treatments without damages, the second peak occurs near the prize value of 5, whereas in treatments with damages, it shifts downward to around 4 as the QRE model anticipates. This result strengthens the case for QRE as a descriptive model of behavior, as it not only accounts for the presence of bimodal bidding but also correctly predicts the effect of damages on bidding distributions.

The NI model provides an interesting point of comparison. While its predicted distributions bear some visual resemblance to the empirical data in terms of bimodality, its weaker statistical fit and imprecise parameter estimates suggest that it does not provide a fundamentally better explanation. The standard error of the NI noise parameter relative to its magnitude in particular implies that the noisy introspection process adds little explanatory power beyond standard equilibrium play.

Overall, our findings support the conclusion that subjects are behaving in an equilibrium-consistent manner, and that QRE with reference-dependent and loss-averse preferences provides a compelling explanation for the observed bidding patterns. The model successfully accounts for the bimodal nature of bidding seen in previous all-pay auction experiments while also capturing the effect of damages on bid distributions. These results therefore reinforce the explanation for bimodal bidding suggested by the literature (Ernst and Thöni, 2013; Dechenaux et al., 2015). More generally, our findings are also consistent with prior experimental research on standard two-player all-pay auctions showing that while players may not mix exactly as predicted, average bidding levels tend to align with equilibrium benchmarks (Potters et al., 1998; Gneezy and Smorodinsky, 2006).

After reviewing related literature in the next section, Section 3 presents the risk-neutral model of contests with damages. This model yields a symmetric NE in terms of expected bids that is equivalent to the model without damages; however, it also predicts a uniform distribution of bids. We present our experimental setup in Section 4 and our experimental results showing bimodal behavior in Section 5. In Section 6 we first describe our models of bounded rationality, QRE and NI, and the procedures to estimate their parameters. We then present our results for those models suggesting QRE is the better explanation for the observed behavior in our experiments. We then finish with a short discussion of what our results mean for the explanation of behavior—including bimodal bidding—in all-pay auction experiments.

2 Related Literature

In addition to the works referenced above specifically relating to all-pay auction experiments, this paper also relates more broadly to the literature on externalities in all-pay contests. For a survey of the theoretical literature on contest models, including various treatments of externalities, see Konrad (2023). For surveys of the literature on experiments with contests, see Dechenaux et al. (2015) or Chowdhury et al. (2024).

The contests in our experiments are a special case of contests with rank-order spillovers by Baye et al. (2012) where the spillover doesn't depend on rank (who wins or loses). Rather, each player is symmetrically damaged as a linear function of the other's effort. We chose this specification because it means the presence of the (negative) externality does not change players' NE strategies with risk-neutral utility, providing us with a crisp prediction: bids will be the same in both settings.

Dechenaux and Mancini (2008) consider a different special case of Baye et al. (2012) in experiments where they vary the size of the externality parameters to model differences in legal systems in terms of the degree of winner- or loser-reimbursement. Park and Shogren (2021) also consider a contest model with linear rank-order externalities, but with a Tullock rather than all-pay contest success function, and focus on the difference in spending based on who is impacted by an aggregate externality. Chung (1996) and Lee (2000) also study spending levels in contests with aggregate externalities, meaning the impact of the externality is based on total (as opposed to individual) spending.

Other all-pay contest models that include externalities focus on those that are identity-dependent, meaning agents' valuation for winning changes depending on who wins (Linster, 1993; Esteban and Ray, 1999; Klose and Kovenock, 2015; Sela and Yeshayahu, 2022; Betto and Thomas, 2024). These models are more general than ours, as the externalities may or may not depend on effort as well as who wins, and focus on the existence of equilibria in such settings. The externalities we consider are very specific, with the (symmetric) relationship between each player's spending and it's impact on their opponent clearly defined and yielding more straightforward Nash predictions. Boosey and Brown (022a) and Boosey and Brown (022b) also study all-pay contests with identity-dependent externalities, but with effects that are network-based, so the extent of the externality on any one player depends on who they're connected to. Sacco and Schmutlzer (2008) and Jönnson and Schmutzler (2017) characterize equilibria in all-pay contests with a general form of negative externality such that the prize to the winner is reduced by the second-highest bid, whereas in our more specific model each player is harmed whether they win or lose.

Our framework appears to be more consistent with the irrevocability of an all-pay auction: Damages, like spending, are imposed whether or not one wins the prize. This appears apt, e.g., in the case of damages to physical and human capital regardless of outcome in armed conflict. Another relevant application includes political campaigning with negative campaigning, in which reputational costs can be imposed regardless of outcome.

3 Model

In the all-pay auctions we study, two players, $i = 1, 2$, simultaneously choose bids, $b_i \in [0, \infty)$, to win a prize commonly valued at V. The player choosing the highest bid wins with certainty (with ties broken by the toss of a fair coin), but both parties' bids are sunk as in the standard all-pay auction with complete information. In addition to the direct cost of the bids, however, their bids can also inflict damage on one another. When players have risk-neutral preferences, payoffs are specified as

$$u_i(b_i, b_j) = \begin{cases} V - b_i - \gamma b_j & \text{if } b_i > b_j \\ -b_i - \gamma b_j & \text{if } b_i < b_j \\ \frac{1}{2}V - b_i - \gamma b_j & \text{if } b_i = b_j \end{cases}, \tag{1}$$

where the parameter $\gamma \geq 0$ captures the size of the negative externality. If $\gamma = 0$, the game is a standard all-pay auction with complete information. If $\gamma > 0$, the game is a contest with damages; each player's bid is not only a cost they must bear themselves, but also damages the other party to a degree.

It is well known (Hillman and Samet, 1987; Baye et al., 1996) that the standard (risk-neutral) two-player all-pay auction with complete information (here the case of $\gamma = 0$) entails a unique symmetric NE of

$$F^*(b) = \frac{b}{V} \text{ on } b \in [0, V];$$

that is, players will randomize uniformly over $b_i \in [0, V]$. The expected payoffs in this equilibrium are zero.

It has also been established that the case of the all-pay auction with negative rank-order spillovers (here the case of $\gamma > 0$, what we refer to as a contest with damages) entails the same unique symmetric NE (Baye et al., 2012). This makes sense, as the "damages" aspect of the payoff function does not affect the players' individual choice process. Accordingly, γ is referred to as a "nuisance" parameter by Baye et al. (2012). Though payoffs differ for the case with damages, equal to $-(\gamma V/2)$ rather than zero, what is most important in terms of our experimental predictions is that expected spending is the same for both cases, equal to $V/2$. This leads us to the following initial hypothesis for our experiments.

Hypothesis 1. Average bids will be the same in experimental trials of standard all-pay auctions and all-pay auctions with damages.

Note that our experimental treatments presented subjects with payouts (the monetary value of payoffs) structured according to the risk-neutral model. That is, the monetary payout for any individual round of play was determined by equation (1), leading to our hypothesis above. This does not necessarily mean that subjects valued those payoffs in a risk-neutral fashion, however. Section 6 and Appendix B explore a richer set of preferences.

4 Experimental Design

In each of our sessions, subjects were randomly sorted into groups of two. Groups were fixed for all twenty rounds of the experiment because subjects are predicted to employ a mixed strategy in all-pay auctions. If we had used a random re-matching protocol, subjects could rely on the re-matching protocol to ensure that their bids were unpredictable. Thus, the fixed matching protocol encourages subjects to actively randomize, and is a key feature of our experimental design. Using fixed groups also allows us to use group-level data for our non-parametric tests.[1]

Subjects within each group played a symmetric and complete information all-pay auction with a prize of $5. Group members simultaneously chose unrecoverable bids between $0 and $7, the highest of which won the prize, with ties broken by fair randomization. After each round, subjects observed their own bid and payoff as well as the bid and payoff of the other member of their group member. Subjects also observed a history table, which contained this information for all preceding rounds.

Our experiments consisted of two treatments. The first was the standard all-pay auction described above. In the second treatment, each subject was penalized 20% of the bid of their counterpart, in addition to paying their own bid. We refer to the former as the "No damages" (ND) treatment, and the latter as the "With damages" (WD) treatment.

In the ND treatment we ran three sessions, with a total of 21 groups.[2] In the WD treatment, we ran four sessions, with a total of 31 groups.[3] All sessions were run at Utah State University. At the beginning of an

[1]One concern is that tacit collusion may be more likely to emerge with fixed groups. While we cannot definitively rule out that such tacit collusion was present, we note that our data in the treatment without damages is quite similar to what is typically observed. Further, any such effect would be present in both treatments, and would not affect any differences between treatments.

[2]In two sessions, there were sixteen subjects. In the third, there were ten subjects.

[3]In two of the sessions there were 20 subjects. One of the remaining sessions had ten subjects, and the other had twelve.

experimental session subjects were seated at computers separated by dividers to ensure privacy. Subjects were provided with a hard copy of the instructions, which were read aloud by an experimenter.[4] Subjects completed a short quiz prior to beginning the experiment. The computer interface was programmed in z-Tree (Fischbacher, 2007).

Subjects were paid a $7 show-up fee, and began the experiment with an additional $15 to cover losses. Four rounds were randomly selected for payment.[5] Subject payments were rounded up to the nearest $0.25, and paid in private. The average payment was $22.16, with a minimum of $7, and a maximum of $42.00.

5 Experimental Results

We take a pair of subjects as the independent unit of observation in our experiment. To conduct non-parametric tests, we average the data for each pair across all twenty rounds of the experiment. This approach allows us to account for the repeated nature of the interactions while maintaining the independence of observations. All statistical tests reported in this section are two-tailed.

[4] The instructions for the WD treatment can be found in Appendix A.

[5] Three subjects lost more than the starting balance in the four selected rounds, and were only paid their $7 show-up fee. We do not drop this data in our analysis, but our results are robust to doing so.

Table 1: Summary statistics

	Without damage	With damage
Bids	2.173	2.346
	(1.859)	(1.943)
Total expenditures	4.347	4.691
	(2.986)	(3.297)
Payoffs	0.327	-0.315
	(2.407)	(2.854)

Notes: Table contains means with standard deviations in parentheses.

Table 1 shows the key summary statistics for our experiment, including the average bid, the total expenditures (calculated as the sum of both players' bids), and the resulting payoffs in each treatment. While average bids are slightly higher when damages are present, we cannot reject the null hypothesis that average bids are the same (Robust Rank Order test, $p = 0.397$).[6] In other words, the hypothesized equivalence in average bids cannot be rejected.[7]

When comparing the observed bids to the risk-neutral theoretical predictions, note that the average bid in both treatments is slightly lower than the NE expected bid of 2.5. However, this difference is not statistically significant in either the treatment with damage (Wilcoxon Signed Rank test, $p = 0.735$) or the treatment without damage (Wilcoxon Signed Rank test, $p = 0.304$). This finding is consistent with the existing literature on all-pay auctions without damages and with two players, where average underbidding as a point estimate with no statistical difference from risk-neutral theoretical predictions is a common observation (Ernst and Thöni, 2013; Potters et al., 1998; Boudreau et al., 2022).

Our data is consistent with Hypothesis 1, indicating that the addition of a negative externality, in which

[6] The robust rank order test is closely related to the Mann-Whitney test, and is designed as a more resilient alternative to the Mann-Whitney test, particularly when higher moments of the underlying distributions are not likely to be equal. See Feltovich (2003) for a detailed discussion.

[7] It is important to note that it is possible that this inability to reject the null hypothesis of equivalence is a result of a relatively small number of observations.

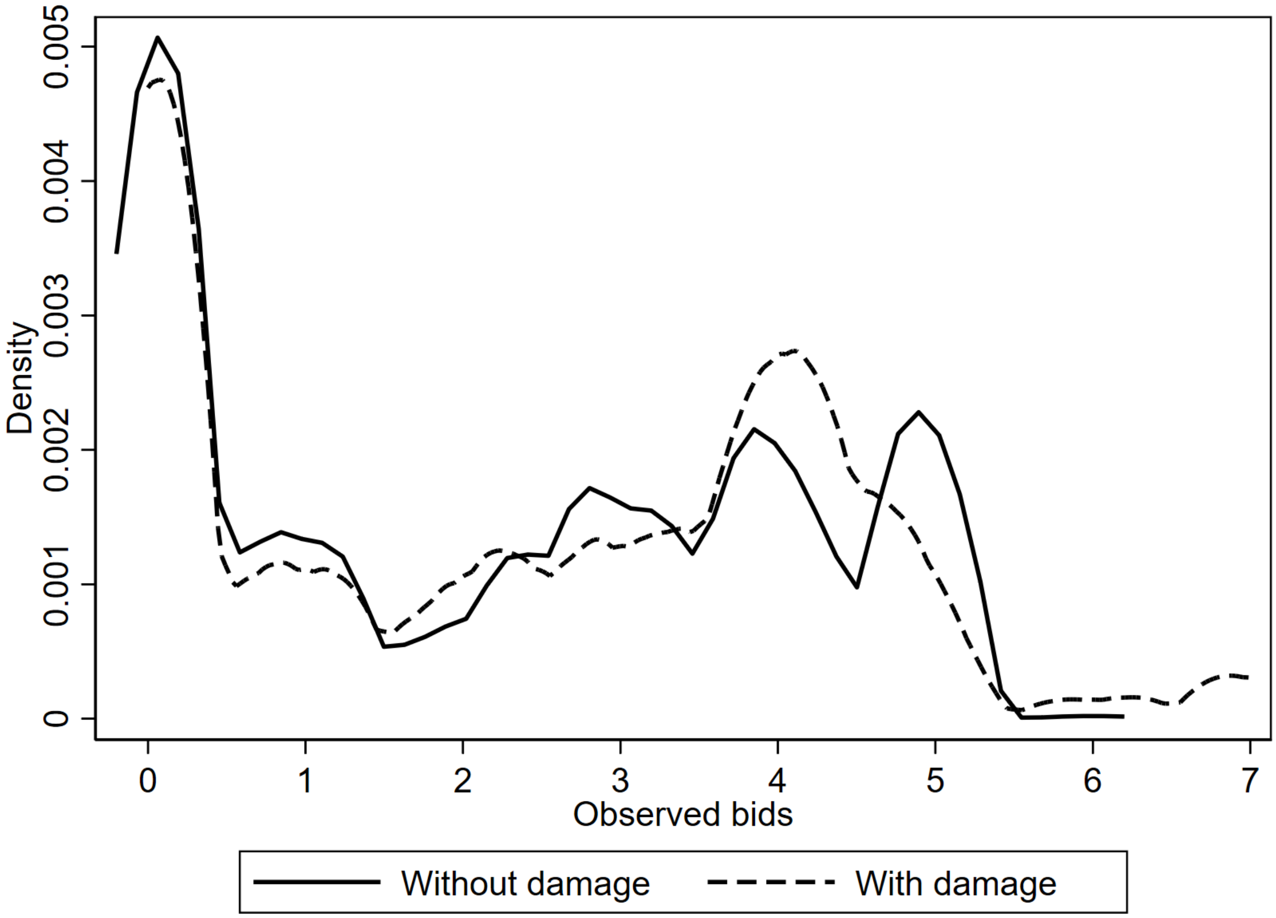

Figure 1: Bids with and without damage

a player's bid proportionally reduces their opponent's payoff, does not significantly affect bidding behavior on average. As expected, the results for individual bids extend to total expenditures and the level of prize dissipation.[8] We find no significant difference in either the sum of bids or the level of prize dissipation between the treatment with damage and the treatment without damage. Moreover, there is no significant difference between the observed total expenditures and prize dissipation levels and their corresponding risk-neutral theoretical predictions.[9]

Although risk-neutral theory predicts that average behavior in all-pay auctions with and without dam-

[8]The sum of bids, expressed as a percentage of the prize value, is commonly referred to as prize dissipation. See, for example, Aycinena et al. (2019).

[9]The test statistics and p-values for the sum of bids and prize dissipation are identical to those obtained from the corresponding tests on individual bids in all cases. This is because, in our experimental design, both the sum of bids and prize dissipation are simply linear transformations of the average bid in an auction, obtained by scaling the average bid by a constant factor.

ages will be the same, it also predicts that the introduction of damages will reduce player payoffs. In the absence of damages, players are expected to earn zero payoffs. While we observe slightly positive payoffs on average in our experiment, this difference is not statistically significantly different from zero (Wilcoxon Signed Rank test, $p = 0.304$). This result can be attributed to the minor underbidding observed in this treatment.

Since risk-neutral theory predicts that behavior will remain unchanged when damages are added, payoffs are expected to become negative. Specifically, in our context, expected payoffs are $-(\gamma V/2) = -0.5$. Indeed, we find that payoffs in the treatment with damages are not statistically different from the predicted negative payoff (Wilcoxon Signed Rank test, $p = 0.735$).[10] Consequently, we find that payoffs are lower in the treatment with damage compared to the treatment without damage (Robust Rank Order test, $p = 0.033$). This result is perhaps not surprising, given the similarity in average bidding behavior across the two treatments.

While the observed average bidding behavior in our experiment is close to the risk-neutral theoretical predictions, a closer examination of the bid distributions reveals two striking departures from the model's Nash equilibrium, in which players bid according to a uniform distribution on the interval $[0, 5]$.

First, a non-trivial number of bids are above the prize value of 5. In the no-damages treatment, there are twelve bids above the prize value, with a maximum of 6. In the damages treatment, there are 55 bids above the prize value, with a maximum of 7. Since these bids guarantee negative payoffs, there is no rational basis for them.

Second, as the kernel density plots of bids displayed in Figure 1 reveal, the bid distributions are clearly not uniform.[11] The observed distributions exhibit a bimodal pattern, which is a common finding in the literature on all-pay auctions with complete information and a common prize value (Sheremeta, 2013).[12] The bimodal distribution is characterized by a concentration of bids around zero, indicating cautious bidding, and another concentration slightly below the value of the prize, representing aggressive bidding.

[10] Although average payoffs are negative in our data, this difference is also not statistically significant from zero (Wilcoxon Signed Rank test, $p = 0.272$).

[11] A kernel density plot is a smooth, continuous curve that estimates the probability density function. It is important to note that all observed bids are between \$0 and \$7. Despite this, the smoothing associated with kernel density plots leads to a positive density below zero.

[12] Similar behavior has also been observed in all-pay auctions where the prize value is unknown (Grosskopf et al., 2010).

Interestingly, despite bids not corresponding to the predicted uniform distribution, the distributions are quite similar across treatments. While, using the Kolmogorov-Smirnov test, we are unable to reject the null that the two distributions are the same ($p = 0.375$), it is worth noting that the prevalence of cautious bids is slightly less when damages are present. Further, aggressive bids in the no damages care are most frequently close to $\$5$, while aggressive bids in damages treatment are concentrated around $\$4$.

It is important to note that the bimodal nature of bidding is present even after subjects have gained experience. Figure 2 illustrates this with kernel density plots of bids in the first and second half of the experiment. Interestingly, in both treatments, the concentration of cautious bids increases, while the concentration of aggressive bids increases.

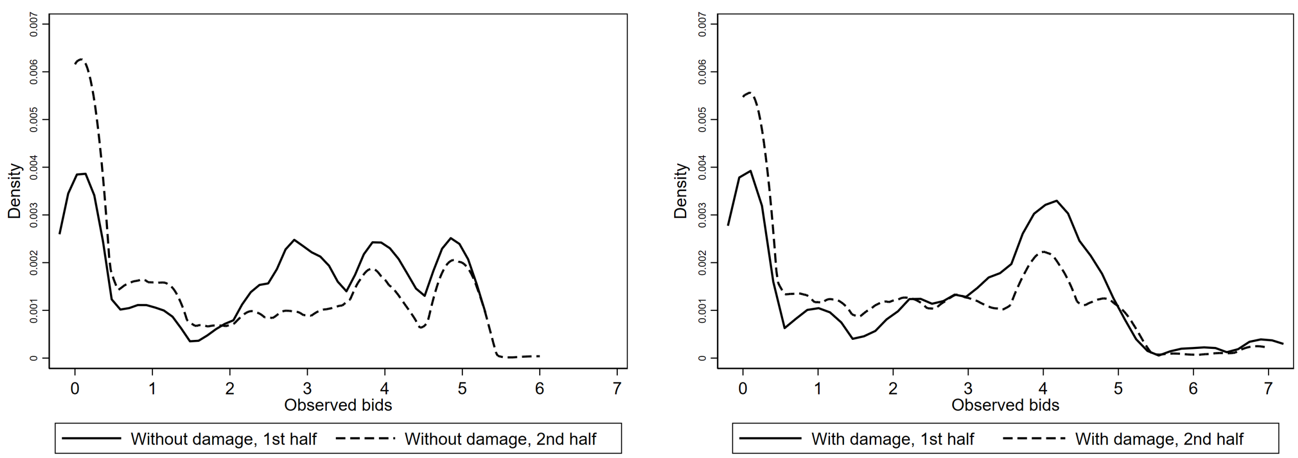

Figure 2: Kernel density plots of expenditures in the first and second half

These experimental results present an intriguing puzzle. On one hand, the introduction of damages to the model is not expected to alter bidding strategies, and indeed, our data are broadly consistent with this prediction. On the other hand, the distribution of bids in both treatments deviates substantially from the predictions of the risk-neutral Nash equilibrium. This discrepancy between theory and observation is not unique to our study, as similar bidding patterns have been documented in the literature on all-pay auctions with complete information and a common prize value. However, bidding behavior across treatments with and without damages adds a new layer of complexity to the problem. A comprehensive explanation for our findings would need to account for both the observed bimodal distribution of bids in both treatments, and that aggressive bids in the treatment with damages are frequently less than those observed in the treatment without damages.

6 Bounded Rationality with Reference-Dependent Preferences

6.1 Reference Dependence and Bimodal Bidding

The leading explanation for the bimodal distribution of bids in all-pay auctions without damages, introduced in Ernst and Thöni (2013), is allowing for reference-dependent preferences based on prospect theory (Kahneman and Tversky, 1979). Specifically, they employ the value function specified by Tversky and Kahneman (1992) :

$$v(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda(-x)^{\alpha} & \text{if } x < 0, \end{cases} \tag{2}$$

This functional form means that gains and losses are evaluated differently based on a reference point of zero. The parameter $\alpha \in (0,1)$ implies players are risk-averse in gains but risk-seeking in losses, while the parameter $\lambda > 0$ induces loss-aversion.

Ernst and Thöni (2013) convincingly argue that a Nash equilibrium with players who exhibit such reference-dependent preferences can explain the bimodal distribution of bids in all-pay auctions without damages. However, once damages are introduced, the resulting equilibrium differs from the no-damages case.[13] Since deriving a closed-form solution for Nash equilibrium with damages and reference-dependent preferences is infeasible, we rely on numerical methods.

In addition to allowing for reference-dependent preferences, we also consider a generalization of Nash equilibrium called quantal response equilibrium (QRE).[14] This approach makes sense for several reasons. First, Gneezy and Smorodinsky (2006) show that QRE can better explain observed bidding behavior in complete-information all-pay auctions than Nash equilibrium. Second, QRE nests Nash equilibrium, so we can let the data determine which notion of equilibrium fits best. Third, QRE is computationally feasible to estimate numerically.

[13] Appendix B discusses this issue in depth.

[14] QRE relaxes the assumption of perfectly rational players by allowing them to choose strategies probabilistically, with the probability of each strategy increasing in its expected payoff. See McKelvey and Palfrey (1995) for a detailed description.

For comparison, we also consider a similar model of bounded rationality known as Noisy Introspection (NI) introduced by Goeree and Holt (2004), which can be estimated in a similar fashion to QRE as explained below.

6.2 QRE estimation

To estimate the parameters of the QRE model with reference-dependent preferences, we employ a two-stage procedure. First, we determine the maximum likelihood estimates (MLE) of the parameters using the data from our experiments. Second, we implement a bootstrap procedure to compute standard errors by resampling the data and re-estimating the parameters.

The QRE computation follows a fixed-point iteration approach, where players' bidding strategies are intialized randomly, then iteratively updated according to a logit response function. Given a bid space discretized into increments of 0.01 to represent cents on the dollar, players update their bid distributions based on the expected utility of each bid, which depends on the opponent's bid distribution from the previous iteration.

The logit response function for player i specifies that the probability of choosing bid is given by:

$$f_i(b_i) = \frac{\exp(z\mathbb{E}u(b_i, f_{-i}))}{\sum_{b_i' \in [0,7]} \exp(z\mathbb{E}u(b_i', f_{-i}))},$$

where $\mathbb{E}u(b_i, f_{-i})$ represents the expected utility to player i of bidding b_i while their opponent uses the strategy f_{-i}, and z is the noise parameter governing the degree of rationality. As $z \to \infty$, players behave as if they fully best-respond to their opponents, approaching Nash equilibrium. Conversely, as $z \to 0$, choices become uniform, implying completely uniform-random behavior. In the intermediate range, QRE smooths best responses based on payoff differences. We set the bidding range to $[0, 7]$ to be in-line with our experimental setup, as some subjects did make bids above the prize value of 5.

The iterative process continues until convergence, defined as the maximum absolute difference in bid distributions between iterations falling below a given threshold. Here, we use the error threshold 10^{-6}. To ensure stability, we apply a damping factor to each iteration's update. If the procedure fails to converge within a set number of iterations (5000), it is terminated, and the parameter set is flagged as non-convergent.

The three parameters, a, λ, and z are then estimated by maximizing a regularized likelihood function. Given our data from the experiments with and without damages, we denote the two datasets $\mathcal{D}_1$ and $\mathcal{D}_2$ for with and without damages, respectively. The likelihood function is then constructed as the sum of log probabilities of observed bids under the equilibrium (converged) distribution:

$$\mathcal{L}(a, \lambda, z|\theta) = \sum_{i \in \mathcal{D}_1} \log f(b_i) + \sum_{j \in \mathcal{D}_2} \log f_d(b_j) - \theta \sum_{q \in \{a, \lambda, z\}} q^2,$$

where $f(b)$ and $f_d(b)$ represent the equilibrium bidding distributions for the cases without and with damages, respectively, and θ is a regularization parameter that penalizes large parameter values to prevent overfitting.[15] The logit response function defines these distributions based on expected utility differences.

To obtain the MLE, the first stage of our process performs an adaptive grid search. Beginning with a broad search over a coarse grid of parameter values, the likelihood function is evaluated at each grid point, and the top 10% of parameter sets (those with the highest likelihood values) are retained for further refinement. The retained parameter sets define a narrower search space, where a finer grid is applied around the highest-likelihood estimates. This process repeats iteratively, each time further zooming in on the most promising parameter regions.

After the first step of our procedure identifies a set of candidate parameters, to assess their accuracy we implement a nonparametric bootstrap approach with $B = 100$ resamples. For each bootstrap sample, we re-estimate the parameters using by solving the following constrained optimization problem:

$$\{\hat{a}, \hat{\lambda}, \hat{z}\} = \arg\max \mathcal{L}(a, \lambda, z|\theta) \quad \text{subject to} \quad \ell_a \leq a \leq u_a, \quad \ell_\lambda \leq \lambda \leq u_\lambda, \quad \ell_z \leq z \leq u_z.$$

[15]Due to the nature of the all-pay auction format, which prevents pure-strategy equilibria, as the noise parameter z increases it eventually leads to oscillatory behavior in the iterative QRE computation—i.e., cycles where bid probabilities shift back and forth rather than converging smoothly. Ridge regularization (also known as L2 regularization) helps stabilize the estimation by discouraging extreme parameter values that might amplify these cycles, improving numerical stability and convergence while also preventing model overfitting (Hoerl and Kennard, 1970; Hastie et al., 2009). We used a value of $\theta = 2$ to ensure convergence over a generous range of potential parameter values, but smaller values lead to the same estimates so long as convergence is possible over the relevant ranges.

where ℓ_q and u_q define the lower and upper bounds for each parameter $q \in \{a, \lambda, z\}$ as specified in the first stage. The optimization is carried out using MATLAB's fmincon function with numerical gradient estimation and strict convergence criteria. The standard error of each parameter is then computed as the sample standard deviation across bootstrap estimates:

$$\mathrm{SE}(\hat{q}) = \sqrt{\frac{1}{B}\sum_{b=1}^{B}(\hat{q}_b - \bar{q})^2},$$

where $\hat{q}_b$ represents the estimated value of parameter q in the b-th bootstrap iteration and $\bar{q}$ is the mean of the bootstrap estimates for parameter q. The final standard error for each parameter is simply the sample standard deviation of these bootstrapped estimates across all B resamples.

6.3 NI estimation

In the NI model of Goeree and Holt (2004), players update their beliefs about their opponent's strategies based on an internal reflection process. Each player starts with an initial belief[16] about the opponent's bidding strategy, which is represented as a probability distribution. At each iteration, players update these beliefs based on a noisy belief-formation process, with the degree of noise controlled by the parameter η. A random noise term, $\mathcal{N}(0, 1)$, is added to the opponent's strategy probabilities, with the magnitude of the noise being proportional to $\eta \in (0, 1)$. Higher values of η imply greater noise, making the players' beliefs more erratic. To ensure convergence, we follow Goeree et al. (2018) and keep the noise, η, constant for the first k updates, then allow it to decay exponentially in subsequent iterations.[17] This decay captures the idea that players rely more on past beliefs over time and become more certain in their expectations.

The belief updates are incorporated into the strategy choice process as follows: each player chooses their strategy based on their belief about the opponent's strategy, with the utility from their actions (using the same reference-dependent and loss-averse utility function) influencing the choice. The strategy is

[16] We follow Goeree and Holt (2004) and begin with uniform distributions, but randomized initial distributions lead to the same outcomes.

[17] For the results presented here we use $k = 5$, but alternative specifications, e.g. $k = 3$ or $k = 10$ have little impact, as the results below will explain.

updated using the logit response function, which incorporates the noisy beliefs and converts them into probabilities of bidding at each bid level based on relative expected payoffs. The estimation of the NI parameters then follows the same two-stage procedure as the QRE model, with a refinement process over the grid of candidate parameters followed by bootstrapping for standard error calculations.

6.4 Results

QRE	a	λ	z	$\mathcal{L}(a, \lambda, z)$
	0.397	4.590	10.020	-12913
	(0.029)	(0.256)	(0.286)	
NI	a	λ	η	$\mathcal{L}(a, \lambda, \eta)$
	0.200	1.234	0.535	-13452
	(0.089)	(0.838)	(0.217)	

Table 2: Estimated parameters for QRE and NI models, standard errors in parentheses

Our parameter estimates are presented in Table 2. As evident from the likelihoods, the QRE model provides a better fit to our data, with the resulting Akaike information criterion favoring QRE. Also, the noise parameter estimate for the NI model exhibits a relatively large standard error, suggesting the noisy introspection process adds little in terms of explaining the data.

The equilibrium strategies for the QRE model with the corresponding estimated parameters from Table 2 are depicted in Figure 3. Both exhibit bidmodal patterns with their first peak at zero. The model without damages ($\gamma = 0$) exhibits a second peak at the prize value, which is somewhat reflected in the no-damage treatments as seen in Figures 1 and 2, particularly for the second half data in Figure 2. The model with damages ($\gamma = 0.2$) has its second peak around 4, which is even more consistent with our data depicted in Figures 1 and 2 for the treatments with damage. Also of note, the boundedly rational nature of QRE allows us to account for bids above the prize value, which are not possible to reconcile with a true Nash equilibrium.

In Figure 4 we present the NI equilibrium strategies based on the estimates in Table 2. Visually, these do bear some resemblance to the features of our data as depicted in Figure 1, particularly in terms of the

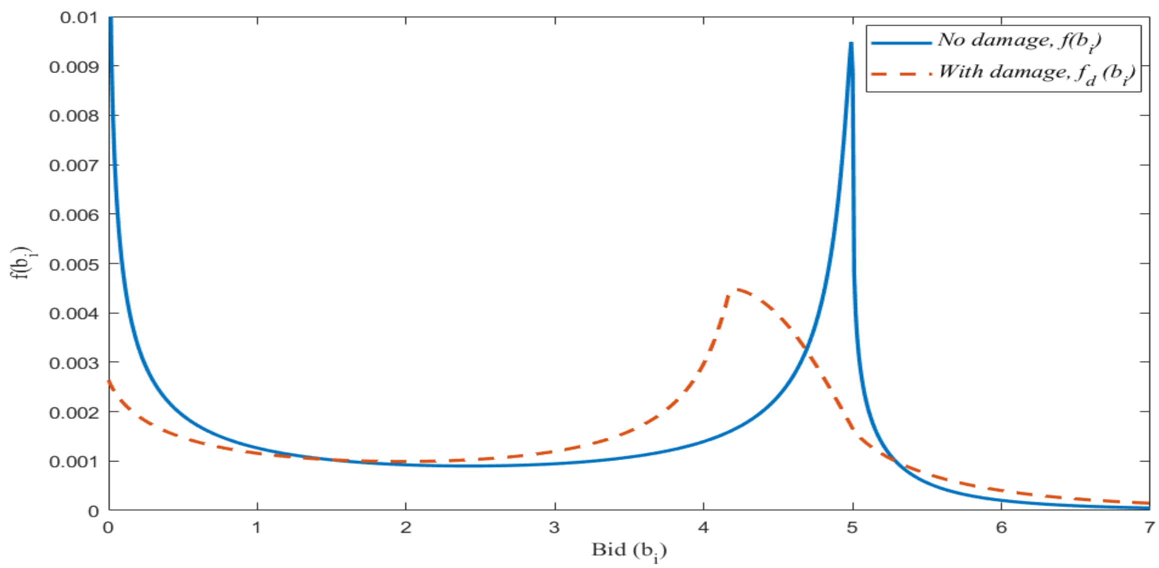

Figure 3: Estimated QRE Strategies

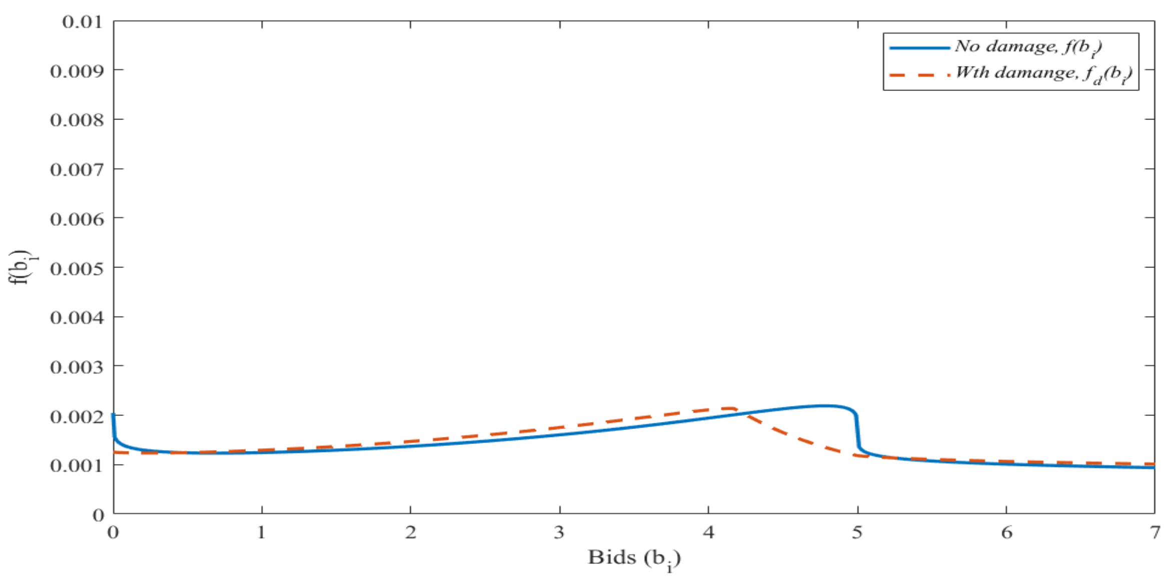

Figure 4: Estimated NI Strategies

overlapping of the strategies' second peaks just above bids of 4. In addition to yielding a less favorable likelihood and larger standard errors for the model's parameters, however, the strategies are much flatter than the QRE's, with $f_d(b_i)$ having a very slight first peak at zero. For comparison we also include Figure 5, which includes all four strategy distributions, to make the differences in shape between the two model's predicted strategies more clear.

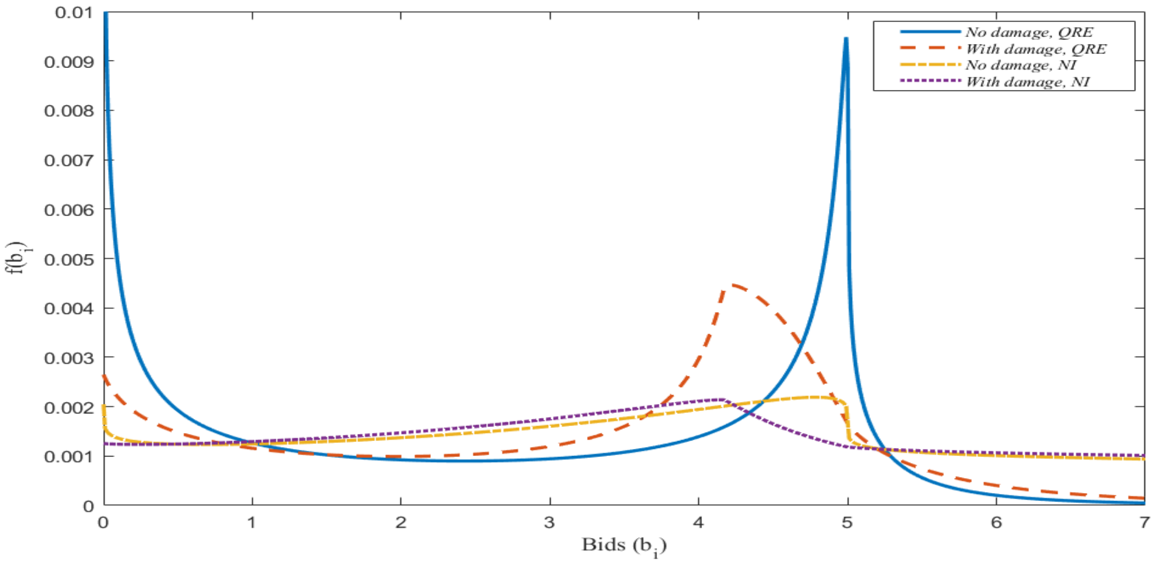

Figure 5: Estimated QRE and NI Strategies

7 Discussion

Our study provides evidence that bidding behavior in two-player all-pay auctions is largely consistent with equilibrium predictions, particularly those of the QRE model. By incorporating reference-dependent preferences and loss aversion, QRE effectively captures the bimodal nature of observed bidding behavior while also correctly predicting the effect of damages on bid distributions. The empirical validation of QRE's key predictions reinforces its value as a descriptive model of boundedly-rational behavior in strategic settings.

A key contribution of our analysis is demonstrating that QRE predicts not only the presence of bimodal bidding but also the locations of peaks in the bidding distributions across treatments. In line with the model's predicted strategies, the second peak of the observed bid distribution appears near the prize value of 5 in treatments without damages while shifting downward to around 4 in treatments with damages. This treatment-dependent shift is precisely what the QRE model anticipates, strengthening its validity as a framework for understanding strategic behavior in all-pay auctions. These findings add to the growing body of literature showing that while subjects do not always mix strategies exactly as prescribed by Nash equilibrium, their *average* bidding behavior generally aligns with theoretical benchmarks (Potters et al., 1998; Gneezy and Smorodinsky, 2006; Dechenaux et al., 2015).

Our comparison with the noisy introspection (NI) model provides additional insight into the nature of bidders' strategic reasoning. While NI predicts bimodal bidding and some level of behavioral noise, its weaker statistical fit and relatively large standard errors indicate it does not offer a fundamentally better explanation than QRE. In particular, the estimated NI noise parameter's imprecision implies that noisy introspection adds little explanatory power beyond standard equilibrium play. This result aligns with previous findings that equilibrium-based models, when augmented with bounded rationality assumptions, often provide the best descriptions of bidding behavior in experimental all-pay auctions (Ernst and Thöni, 2013; Dechenaux et al., 2015).

While our study provides support for QRE as a model of strategic behavior in two-player all-pay auctions, several avenues for future research remain open. One promising direction is to examine how bidding behavior changes as the number of players increases. Prior research suggests that multi-player all-pay auctions often exhibit higher rates of overbidding due to intensified competition, and it would be valuable to investigate whether QRE retains its explanatory power in such settings. Additionally, varying the size of the damage parameter could provide further insight into the robustness of our findings. A more systematic exploration of how different levels of damages influence bidding strategies could help refine our understanding of reference-dependent and loss-averse behavior in contests and auctions.

Overall, our findings contribute to the literature by reinforcing the established explanation for bimodal bidding in all-pay auctions and extending the empirical validation of QRE to settings with damages. Future research expanding on these results could further illuminate the interplay between bounded rationality, equilibrium behavior, and contest structure, providing deeper insights into strategic decision-making in competitive environments.

References

Aycinena, D., R. Baltaduonis, and L. Rentschler (2019). Valuation structure in incomplete information contests: Experimental evidence. *Public Choice 179*, 195–208.

Baye, M. R., D. Kovenock, and C. G. de Vries (1996). The all-pay auction with complete information. *Economic Theory 8*(2), 291–305.

Baye, M. R., D. Kovenock, and C. G. de Vries (2012). Contests with rank-order spillovers. *Economic Theory 51*(2), 315–350.

Betto, M. and M. W. Thomas (2024). Asymmetric all-pay auctions with spillovers. *Theoretical Economics 19*(1), 169–206.

Boosey, L. A. and C. Brown (2022a). Network-dependent externalities in contests. Working Paper. Available online at: https://myweb.fsu.edu/lboosey/network$_contests_09-29-22.pdf$.

Boosey, L. A. and C. Brown (2022b). Contests with network-based externalities: Experimental evidence. Working Paper. Available online at: https://myweb.fsu.edu/lboosey/network$_contests_experimental.pdf$.

Boudreau, J. W., H. N. Nagaraja, L. Rentschler, and S. D. Sanders (2022). The role of luck in political and economic competition: Noisy all-pay auctions. *Public Choice*, 1–21.

Chowdhury, S. M., D. Kovenock, and A. Mukherjee (2024). Experimental research on contests. Sheffield Economic Research Paper Series No. 2024004, Department of Economics, University of Sheffield.

Chung, T.-Y. (1996). Rent-seeking contest when the prize increases with aggregate efforts. *Public Choice 87*(1-2), 55–66.

Dechenaux, E., D. Kovenock, and R. Sheremeta (2015). A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics 18*(4), 609–669.

Dechenaux, E. and M. Mancini (2008). Auction-theoretic approach to modeling legal systems: An experimental analysis. *Applied Economics Research Bulletin 2*, 142–177F.

Ernst, C. and C. Thöni (2013). Bimodal bidding in experimental all-pay auctions. *Games 4*, 608–623.

Esteban, J. and D. Ray (1999). Conflict and distribution. *Journal of Economic Theory 87*(2), 379–415.

Feltovich, N. (2003). Nonparametric tests of differences in medians: Comparison of the wilcoxon–mann–whitney and robust rank-order tests. *Experimental Economics 6*, 273–297.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Gneezy, U. and R. Smorodinsky (2006). All-pay auctions–an experimental investigation. *Journal of Economic Behavior & Organization 61*, 255–275.

Goeree, J. K. and C. A. Holt (2004). A model of noisy introspection. *Games and Economic Behavior 104*(1), 247–272.

Goeree, J. K., P. Louis, and J. Zhang (2018). Noisy introspection in the 11-20 game. *Economic Journal 128*(611), 1509–1530.

Grosskopf, B., L. Rentschler, and R. Sarin (2010). Asymmetric information in contests: Theory and experiments. Technical report, Working Paper, Texas A&M University.

Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer series in statistics. Springer.

Hillman, A. L. and D. Samet (1987, January). Dissipation of contestable rents by small numbers of contenders. *Public Choice 54*(1), 63–82.

Hoerl, A. E. and R. W. Kennard (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics 12*(1), 55–67.

Jönnson, S. and A. Schmutzler (2017, July). Designing all-pay auctions. Working paper, University of Zurich. Available online at https://www.econ.uzh.ch/dam/jcr:e7df5df8-92ca-4bf7-8256-8eaebc2c62c9/JoenssonSchmutzler$_{2}017_{0}7_{0}3.pdf$.

Kahneman, D. and A. Tversky (1979, March). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–292.

Klose, B. and D. Kovenock (2015). The all-pay auction with complete information and indentity-dependent externalities. *Economic Theory 59*(1), 1–19.

Konrad, K. (2023). *Strategy and Dynamics in Contests*. LSE perspectives in economic analysis. Oxford University Press.

Lee, S. (2000). Contests with externalities. *Seoul Journal of Economics 13*(2), 123–134.

Linster, B. G. (1993). A generalized model of rent-seeking behavior. *Public Choice 77*(2), 421–435.

Lugovskyy, V., D. Puzzello, and S. Tucker (2010). An experimental investigation of overdissipation in the all pay auction. *European Economic Review 54*(8), 974–997.

McKelvey, R. D. and T. R. Palfrey (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior 10*(1), 6–38.

Müller, W. and A. Schotter (2010). Workaholics and dropouts in organizations. *Journal of the European Economic Association 8*(4), 717–743.

Park, S.-H. and J. F. Shogren (2021). A virtuous circle of governance contests with externalities. *Sustainability 13*, 7766.

Potters, J., C. G. de Vries, and F. van Winden (1998). An experimental examination of rational rent-seeking. *European Journal of Political Economy 14*, 783–800.

Sacco, D. and A. Schmutlzer (2008, May). All-pay auctions with negative prize externalities: Theory and experimental evidence. Working Paper No. 0806, University of Zurich. Available online at https://www.zora.uzh.ch/id/eprint/52393/1/wp0806.pdf.

Sela, A. and A. Yeshayahu (2022). Contests with identity-dependent externalities. Discussion Paper 22-03, Ben-Gurion University of the Negev.

Sheremeta, R. (2013). Overbidding and heterogeneous behavior in contest experiments. *Journal of Economic Surveys 27*(3), 491–514.

Tversky, A. and D. Kahneman (1992, November). Advances in prospect theory: Cumulative representations of uncertainty. *Journal of Risk and Uncertainty 5*(4), 297–323.

A Instructions

This appendix contains the instructions for the treatment with damages.

Introduction

Welcome to this experiment. The decisions you make during this experiment will determine how much money you earn. You will be paid in cash, privately, at the end of our experiment.

It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

The following instructions will explain how you can earn money. We will go over these instructions with you. After we have read the instructions, there will be time to ask clarifying questions. When we are done going through the instructions, each of you will have to answer a few brief questions to ensure everyone understands.

For today's experiment, you will receive an initial payment of $15, in addition to the $7 show-up fee.

Groups

In today's experiment you will be randomly matched with another participant. You will only interact with this other participant through the computer interface, and you will not know who this other participant is.

Available good

In each round, there is a good worth $5 available for each group of two participants. Only one of the two participants in each group will obtain this good; it cannot be divided.

BIDS

In each round, each participant will choose a BID, which is a number between $\$0$ and $\$7$, inclusive. The participant with the highest BID in a group will obtain the good.

If both participants in a group submit the same BID, then the tie is broken randomly, with each participant having equal probability of obtaining the good.

Each participant's BID must be paid, regardless of whether or not they obtained the good.

In addition, the participant who obtains the good pays a percentage of the BID of the other participant in their group. This percentage, which we call a LOSS PERCENTAGE, is 20%.

Similarly, the participant who does not obtain the good pays a percentage of the BID of the other participant in their group. This percentage (which we also call a LOSS PERCENTAGE) is 20%.

Participants will not know the BID submitted by the other participant in their group when they choose their own BID.

PAYOFF of a round

In each round, the participant in the group who submits the highest BID (or who is randomly chosen in the event of a tie), will obtain the good.

The PAYOFF for the round for this participant is

$$(\$5) - (\text{Their Own BID}) - (0.20) \cdot (\text{The BID of the Other Participant})\,.$$

The PAYOFF for the round for the participant who does not submit the highest BID (or who is not randomly chosen in the event of a tie), is

$$\$0 - (\text{Their Own BID}) - (0.20) \cdot (\text{The BID of the Other Participant})\,.$$

Example 1

Suppose that Participant 1 chooses a BID of $\$4.30$, and Participant 2 chooses a BID of $\$0.75$.

Since $4.30 > 0.75$, Participant 1 obtains the good.

His or her payoff for the round is $\$5 - \$4.30 - 0.20 \cdot \$0.75 = \$0.55$.

The payoff of Participant 2, who does not obtain the good, is $\$0 - \$0.75 - 0.20 \cdot \$4.30 = -\$1.01$.

Example 2

Suppose that Participant 1 chooses a BID of $\$2.95$, and Participant 2 chooses a BID of $\$3.60$.

Since $3.60 > 2.95$, Participant 2 obtains the good.

His or her payoff for the round is $\$5 - \$3.60 - 0.20 \cdot \$2.95 = \$0.81$.

The payoff of Participant 1, who does not obtain the good, is $\$0 - \$2.95 - 0.20 \cdot \$3.60 = -\$3.67$.

Participating in a round

At the beginning of each round, you will be asked to enter your BID for the round. Remember that this BID can be any number between $0 and $7, inclusive.

You will choose your BID without knowing the BID chosen by the other participant in your group.

Results of a round

At the end of each round the results of the round will be displayed on your screen. The results you will see are:

1. Whether or not you obtained the good.
2. Your BID.
3. Your LOSS PERCENTAGE.
4. Your Payoff for the round.
5. The BID of the other participant in your group.
6. The LOSS PERCENTAGE of the other participant in your group.
7. The Payoff of the other participant in your group.

In addition, the results of all previous rounds will always be displayed on your screen.

Selecting rounds for payment

Once all 20 rounds of the experiment have been completed, 4 rounds will be randomly chosen for payment. Each of the 20 rounds are equally likely to be chosen for payment.

Your payment for today's session will be the sum of your payoff in each of the 4 rounds randomly chosen for payment and the initial payment of $15. This is in addition to the $7 show-up fee.

Summary

1. In each round there is an available good which is worth $\$5$ to both participants in your group.
2. In each round, both participants in your group will choose a BID, which is a number between $\$0$ and $\$7$, inclusive. The participant in your group who chooses the highest BID will obtain the good. Ties are broken randomly.
3. The PAYOFF for the round of the participant who obtains the good is

$$(\$5) - (\text{Their Own BID}) - (0.20) \cdot (\text{The BID of the Other Participant})\,.$$

4. The PAYOFF for the round of the participant who does not obtain the good is

$$(\$0) - (\text{Their Own BID}) - (0.20) \cdot (\text{The BID of the Other Participant})\,.$$

5. At the end of the experiment 4 of the 20 rounds will be chosen randomly for payment.

B Nash Equilibrium with Reference Dependent Preferences

In Section 3, we see that when preferences are risk neutral, Nash Equilibrium in our all-pay auction predicts that players will choose a mixed strategy that randomizes uniformly over $b_i \in [0, V]$. However, our experimental results display player behavior that departs dramatically from this prediction, most notably with bidding that is bimodal. Below, we explore whether the reference-dependent preferences proposed by Ernst and Thöni (2013) can explain our experimental results.

B.1 Contests with Reference-Dependent Preferences

Ernst and Thöni (2013) suggest a model with reference-dependent preferences and loss aversion based on prospect theory (Kahneman and Tversky, 1979). Specifically, Ernst and Thöni (2013) employ the value function specified by Tversky and Kahneman (1992):

$$v(x) = \begin{cases} x^\alpha & \text{if } x \geq 0 \\ -\lambda(-x)^\alpha & \text{if } x < 0. \end{cases} \tag{3}$$

In this section, we establish the symmetric mixed-strategy Nash equilibrium of an all-pay auction when players have reference dependent preferences.

Any symmetric NE in fully mixed strategies for a two-player all-pay auction must satisfy the following:

$$\mathbb{E}u(b_i, F) = \mathbb{E}u_W(b_i, F) + \mathbb{E}u_L(b_i, F) = \mathbb{E}u(0, F) \text{ for all } b \in [0, V], \tag{4}$$

where $\mathbb{E}u_W(b_i, F)$ and $\mathbb{E}u_L(b_i, F)$ are the expected payoffs from winning and losing, respectively, for a player bidding b_i against the mixed strategy F.

For the case without damages, the following is a more concrete version for our problem, where we denote the NE mixed strategy cumulative distribution function as $\tilde{F}$ and the associated probability density function as $\tilde{f}$:

$$\mathbb{E}u(b_i, \tilde{F}(b_i)) = \int_0^{b_i} (V - b_i)^\alpha \tilde{f}(b_j) db_j + \int_{b_i}^{V} (-\lambda(b_i)^\alpha) \tilde{f}(b_j) db_j = 0 \text{ for all } b \in [0, V].$$

Since the utility from winning $(V - b_i)^\alpha$ and the utility from losing $-\lambda(b_i)^\alpha$ do not depend on b_j, this expression can be simplified considerably to

$$\mathbb{E}u(b_i, \tilde{F}(b_i)) = \tilde{F}(b_i)(V - b_i)^\alpha + (1 - \tilde{F}(b_i))(-\lambda(b_i)^\alpha) = 0 \text{ for all } b \in [0, V]. \tag{5}$$

Solving this equation for $\tilde{F}(\cdot)$ gives the following bidding strategy:

$$\tilde{F}(b_i) = \frac{\lambda(b_i)^\alpha}{\lambda(b_i)^\alpha + (V - b_i)^\alpha}, \tag{6}$$

which in turn yields a bimodal density as desired (Ernst and Thöni, 2013).

B.2 Contests with Reference-Dependent Preferences and Damages

Inserting damages into the model with reference-dependent preferences is less straightforward than the case of the standard model for two main reasons. First, the expected payoff for not bidding at all, $\mathbb{E}u(0, F_d)$, is no longer equal to zero. This can be seen upon inspection of the expected payoff of not bidding when there are damages

$$\mathbb{E}u(0, F_d) = -\lambda \int_0^V (\gamma b_j)^\alpha f_d(b_j) db_j,$$

where F_d is the cumulative distribution function used by both players, and f_d is its associated density.

Second, the gains and losses from a strictly positive bid no longer depend solely on the agent winning or losing. The expected payoff for losing with a bid of b_i is

$$\mathbb{E}u_L(b_i, F_d) = -\lambda \int_{b_i}^V (b_i + \gamma b_j)^\alpha f_d(b_j) db_j,$$

and the expected payoff for winning with a bid of b_i is

$$\mathbb{E}u_W(b_i, F_d) = \begin{cases} \int_0^{b_i} (V - b_i - \gamma b_j)^\alpha f_d(b_j) db_j & \text{if } b_i \leq V(1-\gamma) \\ \int_0^{\bar{b}} (V - b_i - \gamma b_j)^\alpha f_d(b_j) db_j - \lambda \int_{\bar{b}}^{b_i} (\gamma b_j + b_i - V)^\alpha f_d(b_j) db_j & \text{if } b_i > V(1-\gamma) \end{cases}, \tag{7}$$

where $\bar{b} \equiv \bar{b}_j(b_i) = (V - b_i)/\gamma$. This expression is piecewise because the presence of damages makes it possible for the payoff from a winning bid to be negative. This happens when both b_i and b_j are large enough, that is, both are larger than the threshold value $V(1-\gamma)$.

The analogue to Equation 4 for the model with reference-dependent preferences *and* damages is

$$\mathbb{E}u(b_i, F_d^*) = \mathbb{E}u_W(b_i, F_d^*) + \mathbb{E}u_L(b_i, F_d^*) = \mathbb{E}u(0, F_d^*) < 0 \text{ for all } b \in [0, V]. \tag{8}$$

This defines the symmetric NE strategy, F_d^*.

In general, there is no reason why $\tilde{F}$ and F_d^* should be the same. With damages present ($\gamma > 0$), players face a more significant threat of losses, meaning their behavior should differ if they possess reference-dependent preferences. We are unable to obtain closed-form solutions for the model with reference-dependent preferences and damages. However, in the next section, we provide intuition for how the symmetric NE strategies change in the presence of damages, and demonstrate numerically that the two equilibrium distributions cannot be the same. This is in contract to the risk-neutral model, which predicts that NE bid distributions should be identically uniform.

B.3 Intuition for Nash Equilibrium Strategies in the Presence of Damages

Following Ernst and Thöni (2013), we use maximum likelihood to fit the experimental data from the no damages case to the mixed strategy cumulative distribution function in Equation 6 and find that the best fit results from $\lambda = 2.3043083$ and $\alpha = 0.5623753$. That is, experimental participants display substantial aversion to both loss and risk. We next use these fitted parameters to demonstrate that the experimental evidence does not align with the predictions of the reference-dependent model of contests without damages. Specifically, we show that the Nash strategy for the model with reference-dependent preferences but without damages can not be the same for the model with both reference-dependent preferences and damages using the same parameters from our experimental setting. The same numerical exercise also provides intuition as to how the two strategies differ.

First, to fix ideas, Figure 6 shows the equilibrium mixed strategies (from Equation 6) when Player j has reference-dependent preferences and does not face damages from Player i's bid. The parameters are set at $\lambda = 2.304308$, $\alpha = 0.5623753$, $V = 5$, and $\gamma = 0.2$ (the latter two from the experimental implementation).

We see that neither risk-neutral or loss averse preferences lead to a bimodal bidding distribution. Risk aversion alone leads to bimodal bidding, and adding loss averse shifts the probability mass of bids toward lower values.

Figure 7 shows the expected payoffs for Player i when Player j uses the mixed strategy in Equation 6 in the model with reference-dependent preferences. The left panel displays Player i's expected value of winning (orange) and losing (blue) for each $b_i \in [0, 5]$ when there are no damages and Player j plays the no-damages mixed strategy (6). That is, the orange and blue lines show $\mathbb{E}u_W(b_i, \tilde{F})$ and $\mathbb{E}u_L(b_i, \tilde{F})$, respectively, from (4). We see that these two expected values net out to zero in the red line, which is the total expected value of bidding (i.e., the expected value of winning plus the expected value of losing). Since the expected value of not bidding (i.e., setting $b_i = 0$) is zero when there are no damages, we see that the condition for Player i to play a fully-mixed strategy—that the expected value of playing each bid is equal to the expected value of any other bid and $b_i = 0$ in particular—is met.

Note that in the left panel, three important expected values all overlap at zero: the expected value of bidding a positive number, the expected value of bidding zero, and the difference between these two. The difference between the expected value of a positive bid and the expected value of a zero bid is exactly what must be equal to zero in order to support a fully-mixed strategy. Importantly, this condition does *not* hold

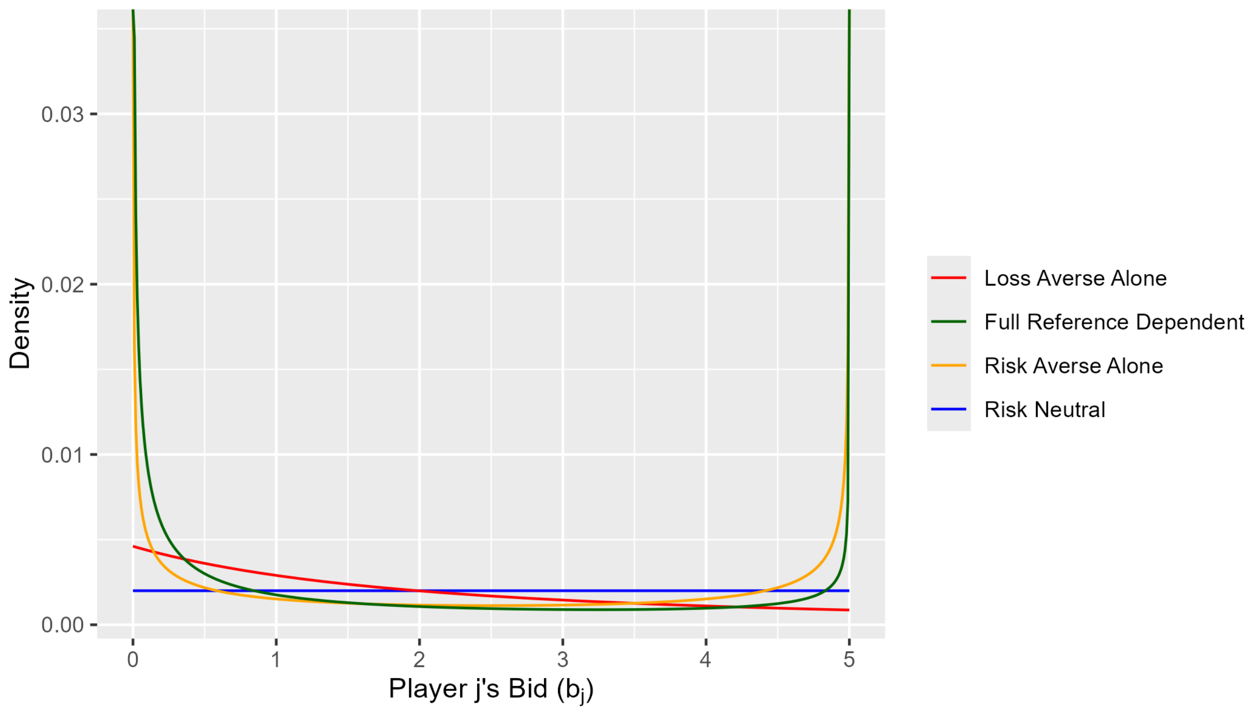

Figure 6: PDFs for Player j's equilibrium strategy under reference-dependent preferences and no damages

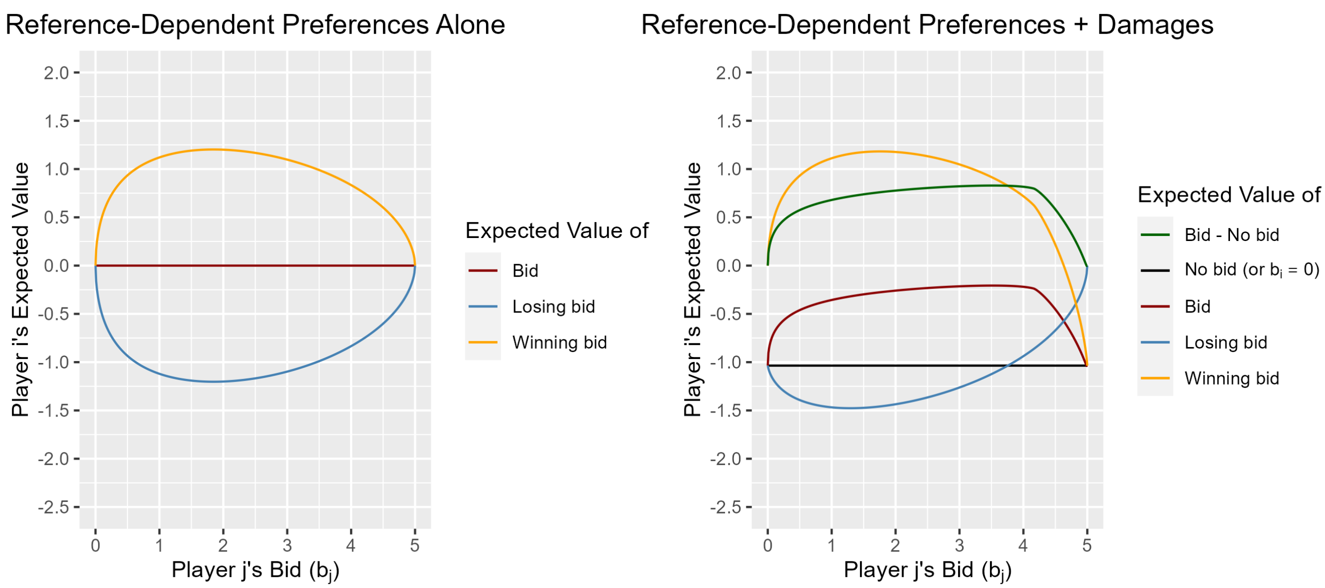

Figure 7: Expected Values against No-Damages NE Mixed Strategy, Reference-Dependent Preferences

in the right panel, where we change the utility to include damages but continue to use the mixed-strategy for the no-damages case ($\tilde{F}$) for Player j's mixed strategy. This shows that the fully-mixed strategy for the case with damages must be different from the strategy for the case with no damages. That is, Player i's best response to Player j playing $\tilde{F}$ is *not* to play *any* fully-mixed strategy.

In fact, we can say something even stronger. Because Player i's expected value from a positive bid (the red line) is greater than her expected value of not bidding (the black line) for all $b_i \in (0, 5)$, in order for Player i to play a fully-mixed strategy, the expected value of bidding must be reduced from the red line to the black line for each value of b_i. This means one of three things needs to happen: the expected value from winning (the orange line) must go down, the expected value from losing (the blue line) must go down, or both.

The value of Player j's bid enters negatively into Player i's expected values from both winning and losing through the damages term γb_j. To reduce either expected value for Player i, the mixed strategy for Player j must shift probability density from lower values of b_j to higher values of b_j compared to the mixed strategy for the no-damages case, $\tilde{F}(b_j)$. It is conceivably possible to also shift some weight to very small values of b_j at the same time as making a more significant shift from intermediate values to higher values so that $F_d^*(b_j)$ does not first order stochastically dominate $\tilde{F}(b_j)$, but $F_d^*(b_j)$ must have more probability mass on higher values of Player j's bids.

To further develop intuition, Figure 8 shows the expected values as in Figure 7, but for the case of a risk-neutral expected utility maximizer (i.e., $\lambda = \alpha = 1$). Here, consistent with the theory outlined in Section 4, we see that the presence of damages shifts the expected values, but does not disturb the mixed strategy. However, risk-neutral preferences cannot explain the bimodal bidding distribution.[18]

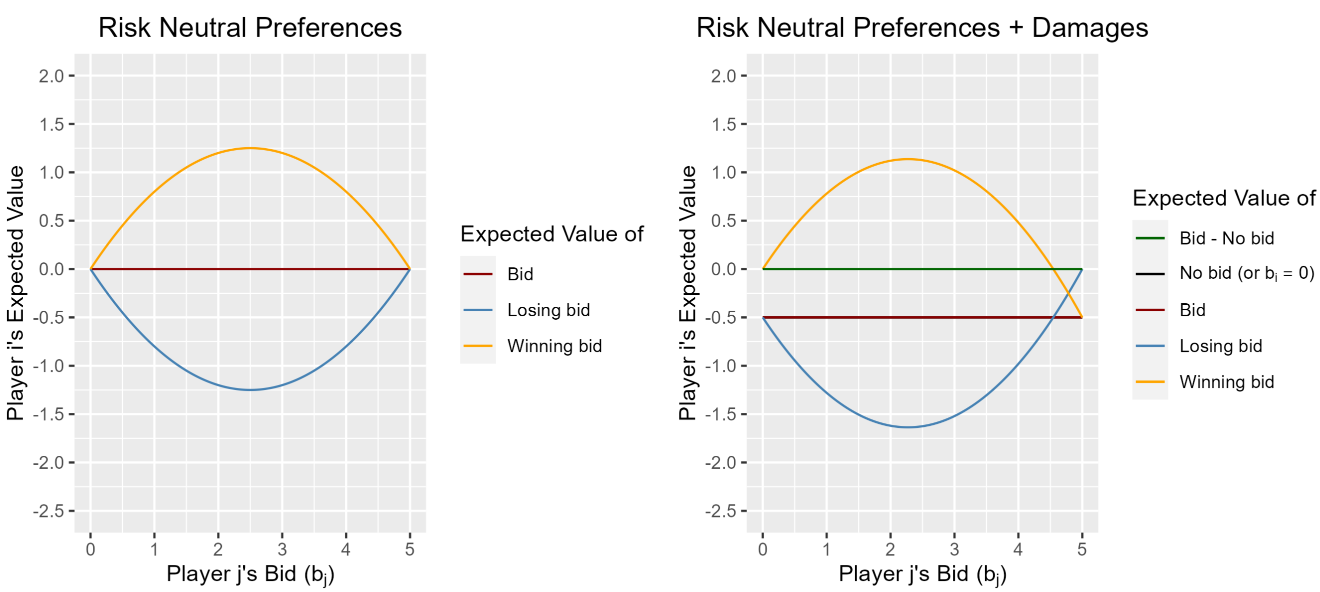

Figure 8: Expected Values against No-Damages NE Mixed Strategy, Risk-Neutral Preferences

Similarly, we can also explore the model with loss aversion alone ($\alpha = 0$) or risk aversion alone ($\lambda = 0$). Loss aversion on its own without risk aversion does not yield a bimodal probability distribution

[18]Note that, in the left panel, the expected value of no bid, a strictly positive bid and the difference between the two are all the same as in the left panel of Figure 7 at zero. In the right panel, the expected value of no bid and a strictly positive bid overlap, so that their difference is zero.

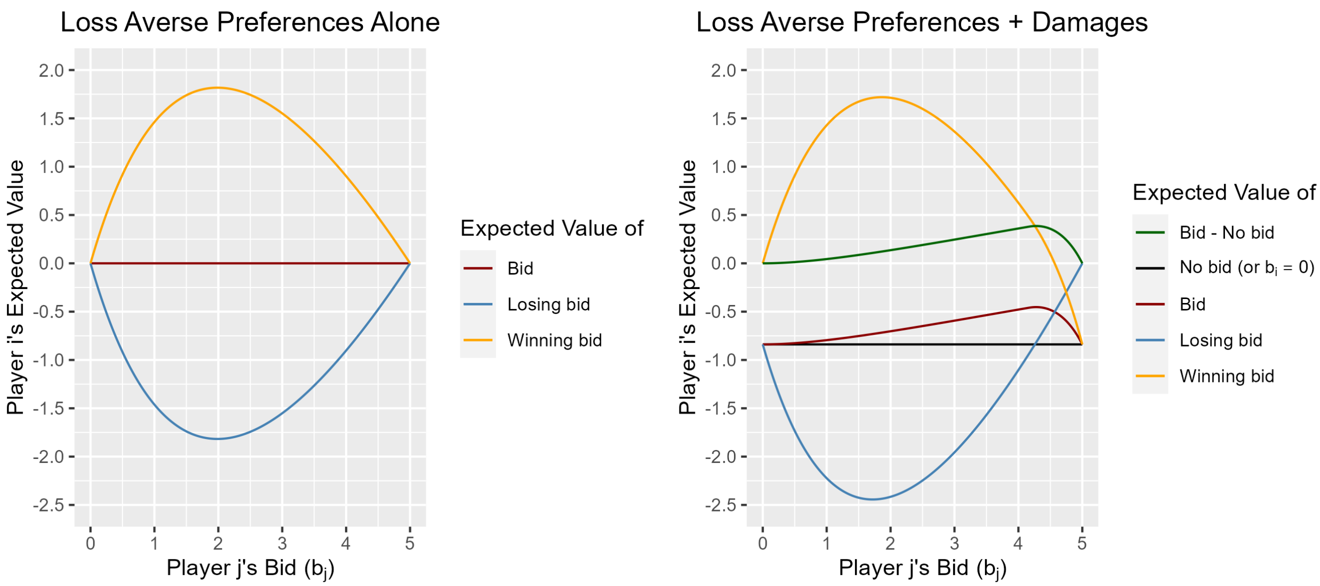

Figure 9: Expected Values against No-Damages NE Mixed Strategy, Loss-Averse Preferences

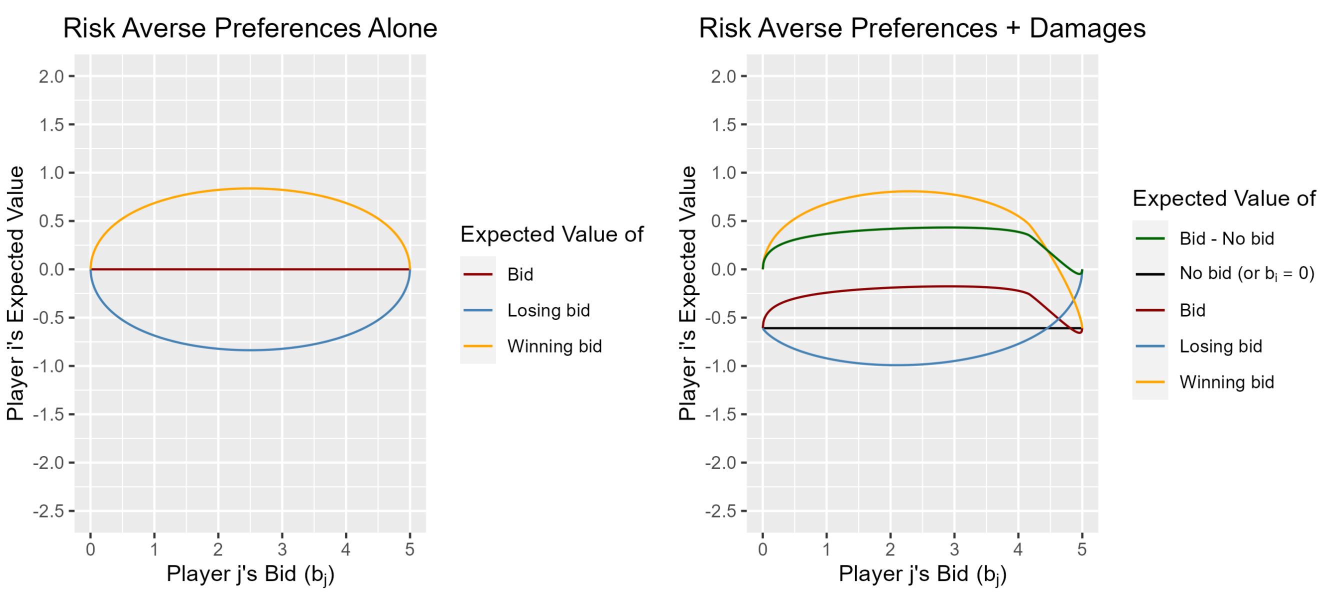

Figure 10: Expected Values against No-Damages NE Mixed Strategy, Risk-Averse Preferences

for the equilibrium mixed strategy as desired, but still implies that the mixed-strategy must change as seen in Figure 9 where the expected value of bidding (red line) is larger than the expected value of not bidding (black line). Risk aversion on its own without loss aversion does yield a bimodal distribution for the no-damages mixed strategy, but implies that the mixed-strategy must change even more than under loss aversion alone (See Figure 10, where the difference between expected value of bidding (red line) and no bid (black line) are further apart than in Figure 9).