Equilibrium Analysis in Markets with Asymmetric Utility Functions

Martin Bichler
Technical University of Munich
Munich, Germany
m.bichler@tum.de

Markus Ewert
Technical University of Munich
Munich, Germany
markus.ewert@tum.de

Axel Ockenfels
University of Cologne
Cologne, Germany
Max Planck Institute for Research on Collective Goods
Bonn, Germany
ockenfels@uni-koeln.de

ABSTRACT

Traditional auction theory assumes symmetric utility functions and equilibrium bidding strategies based on common and symmetric prior distributions. Although extensive literature addresses asymmetries in prior distributions, the assumption of symmetric utility functions persists. A long history of experimental research on auctions suggests that this assumption is hard to justify and that utility functions are not symmetric across bidders, probably driven by different behavioral motives such as risk-aversion, which leads to asymmetric bidding strategies. This observation is important for markets with human bidders and agentic markets with autonomous agents. Unfortunately, equilibrium analysis with asymmetric utility functions is technically challenging, and we are not aware of equilibrium predictions. We leverage recent advances in equilibrium learning to compute equilibrium in asymmetric auction models. First, we analyze asymmetries in isolated markets. Interestingly, we can show that in contrast to the canonical symmetric model, unilateral deviation from the symmetric risk-neutral equilibrium strategy leads to higher profit for the deviating bidder compared to the bidder who follows the symmetric equilibrium strategy. Second, we analyze agentic markets, where firms compete repeatedly and can parameterize agents to bid more or less aggressively. This leads to an interesting meta-game in which it is a Nash equilibrium for both users to select a risk-seeking agent in the first-price auction and a risk-averse agent in the all-pay auction.

KEYWORDS

Asymmetric Utility, Equilibrium Learning, Risk-aversion, Auction Theory

ACM Reference Format:
Martin Bichler, Markus Ewert, and Axel Ockenfels. 2025. Equilibrium Analysis in Markets with Asymmetric Utility Functions. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

In the standard symmetric independent private values model in auction theory, bidders are assumed to have a prior belief over their opponents' valuations (or costs), which is modeled with a symmetric prior distribution that is common knowledge among the bidders, and their utility functions and strategies are assumed to be symmetric [17]. These symmetry assumptions lead to unique equilibrium predictions for the wide-spread first-price sealed-bid auction with mild assumptions on the prior distributions [8, 19] and the all-pay auction [18] but are hard to justify from observations in the field or in the lab [16].

A significant theoretical literature deals with asymmetries in the prior distributions [22, 27], while the utility functions are still assumed to be symmetric. A common and symmetric prior is not an unreasonable assumption in many applications. For example, in industries with relatively standardized production processes and inputs, competitors have accurate estimates of cost structures in the industry [7]. However, decades of experimental research on auctions have provided overwhelming evidence that strategies are not symmetric and deviate from their prescribed risk-neutral equilibrium strategy in the lab where a common and symmetric prior distribution is given [16]. For example, bidders in first-price auctions with a smaller number (two or three) of competitors consistently bid higher (or lower) than the risk-neutral *Bayes-Nash Equilibrium* (BNE) strategy in the lab. There is also evidence for overbidding in the field [20]. In the all-pay auction, bidders underbid the risk-neutral BNE for low valuations and overbid it for high valuations. Müller and Schotter [24] refer to this pattern as *bifurcation*. Risk-aversion is among the most widely cited conjectures [6, 10, 12].[1] However, from years of laboratory experiments, it is also clear that the level of risk-aversion is not symmetric across bidders in the lab [16].

To the best of our knowledge, this is the first study that systematically analyzes the influence of asymmetric utility functions in auctions. We want to understand *competition with asymmetric utility functions* and ask the question of how a firm that faces risk-averse opponents should behave strategically if this firm wants to maximize its expected profit. The expected profit of a risk-neutral and quasi-linear bidder equals their expected utility. When we say that bidders are risk-averse, we refer to their utility model, which they use to derive their equilibrium strategy. However, ex-post we are interested in analyzing their expected profit.

We study two scenarios, in which we assume that the competing firms have different information levels about the risk-attitude expressed by the utility function of their opponents: First, we analyze

[1] Alternative utility models are based on regret [11], and, depending on the regret parameter, it leads to qualitatively similar bidding strategies than risk-aversion in first-price auctions.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

asymmetries in small and isolated markets as they can frequently be found in practice.[2] People compete in an auction once and they do not necessarily understand the level of risk-aversion of their opponent. A widespread phenomenon in such situations is that of *false consensus*. This phenomenon describes that humans assume that their personal decisions are over-represented and correct in the prevailing circumstances [21], meaning that each firm assumes that the opponent has the same utility function as themselves. If such bidders meet and derive an equilibrium strategy under this assumption, the result is not in equilibrium, but it might be a good description of what happens in such isolated, high-stakes auctions, where the level of risk-aversion is most likely not the same among bidders. We also analyze the equilibrium outcome in case bidders knew each other's utility functions. In both of these models, the risk-neutral firm that aims to maximize its expected payoff achieves a lower expected profit than the risk-averse opponent. In symmetric games such as rock-paper-scissors, an agent that deviates from the symmetric equilibrium strategy cannot improve their expected profit. Therefore, it is interesting that with several bidders who aim to maximize expected profit, the one who deviates from the symmetric and risk-neutral BNE achieves higher expected profit (not utility) than their opponents who stick to this risk-neutral equilibrium strategy. Note that the expected profit they both make with a deviating bidder is reduced compared to the auction if all bidders behaved risk-neutral. Equilibria in auctions with asymmetric utility functions were unknown so far.

Second, we analyze markets where competitors interact frequently, and interactions are automated using bidding bots. Display ad auctions are just one example of agentic markets where bidders interact frequently [9]. We do not aim to model the complex details of such markets, but abstract to a more stylized model where bidders can parameterize how aggressively they want their agents to bid in each auction. This is naturally done via the level of risk-aversion of the bidders. For example, more risk-averse bidders bid more aggressively in each auction with a higher likelihood of winning but lower profit in a first-price auction. If bidders compete frequently, they should aim for expected profit, and risk-aversion should have less of an impact. However, this does not mean that both bidders can be expected to use a risk-neutral bidding agent that just aims to maximize the expected profit in each round. Actually, the environment leads to a meta-game where each agent selects a risk parameter. Recent advances in equilibrium learning allow us to determine the expected payoff for each combination of utility functions and, with this, a payoff matrix for the meta-game. We find that it is a Nash equilibrium for all bidders to select a risk-seeking agent in the first-price auction. In contrast, bidders choose a risk-averse agent in equilibrium in the all-pay auction. Both equilibria lead to bidding strategies of the agents below those of the BNE in the stage game, i.e. a single auction in this repeated game. While our model abstracts from the details of real-world display ad auctions or other high-frequency agentic markets, it shows that there are incentives for bidders to deviate from the implementation of a myopic agent that maximizes expected profit in each stage game. In the literature on algorithmic collusion in auctions [3], the equilibrium of the stage game is used as a baseline to compare the outcomes of agent interaction against. Our analysis shows that in the Nash equilibrium of the meta-game, bidding below the BNE of a single auction is an equilibrium. This provides a rationale for why bidders in agentic auction markets might bid below the BNE [3].

Overall, our paper explores two new environments and the resulting equilibria in cases where (1) bidders meet once, and the utility functions are asymmetric, and (2) bidders meet frequently, and the utility functions of their agents can be parameterized. Such questions have not been analyzed in the past because finding equilibrium with asymmetric utility functions analytically usually leads to an intractable system of non-linear partial differential equations. The recent advances in equilibrium learning and the availability of equilibrium solvers [5] allow us to address such questions and provide answers to these new types of equilibrium problems.

[2]https://balticwind.eu/two-bidders-participating-in-the-lithuanian-offshore-wind-auction/, https://money.usnews.com/investing/news/articles/2024-03-26/two-bidders-in-talks-to-buy-german-retail-giant-galeria, https://latinfinance.com/daily-brief/2024/03/26/three-bidders-vie-for-peru-sewage-plant-contract/

2 THE MODEL

Bayesian Games. Auctions are modeled as Bayesian games, defined by a quintuple $G = (N, \mathcal{V}, \mathcal{A}, F, u)$. N players, indexed by $i = 1, \dots, N$, participate in the game. $\mathcal{V} = \mathcal{V}_1 \times \cdots \times \mathcal{V}_N$ is the set of possible *type profiles* describing the private information available to the agents when choosing their actions. These types are drawn from some prior probability distribution F that is assumed to be common knowledge among the bidders. Given knowledge of their types, player i must then choose an action b_i from the set $\mathcal{A}_i$ of available actions. u is a vector of individual utility functions $u_i : \mathcal{V}_i \times \mathcal{A} \to \mathbb{R}$ describing the outcomes of the game. Crucially, for each player, these utilities depend only on their own type but on all players' chosen actions. Risk-neutral players are modeled via a quasi-linear ex-post utility function $u_i^{rn} = x_i v_i - p(b_i)$, where x_i is the allocation of the item to bidder i and $p(\cdot)$ the payment resulting from bid b_i. The item is allocated to the highest bidding agent in both auction types. While in the all-pay auction, the payment always corresponds to the submitted bid, only the winning bidder pays its bid in the first-price auction.

In order to maximize their own utility, every player i needs to decide on a strategy $\beta_i : \mathcal{V}_i \to \mathcal{A}_i$ that will prescribe her action b_i for a given valuation input v_i. A BNE describes a strategy profile $\beta^\star = (\beta_1^\star, \dots, \beta_N^\star)$ in which no agent can improve her expected utility by unilaterally deviating. We will write $\tilde{u}_i(\beta_i, \beta_{-i}) = \mathbb{E}_{v\sim F}[u_i(v_i, \beta_i(v_i), \beta_{-i}(v_{-i}))]$ for the (ex-ante) expected utility. Then, $\beta^\star$ is an (ex-ante) BNE if for all players i and all possible strategies β_i it holds that

$$\tilde{u}_i(\beta_i^\star, \beta_{-i}^\star) \geq \tilde{u}_i(\beta_i, \beta_{-i}^\star) \quad \forall \beta_i, \quad \forall i \in N.$$

Risk-aversion is modeled via a concave transformation of the risk-neutral utility function [2, 26] and can be expressed in various ways.

A specific type of *Constant Relative Risk Averse* (CRRA) is described as

$$u_{i,\eta}(u_i^{rn}) = (u_i^{rn})^\eta, \tag{1}$$

with $\eta \in [0, 1]$. A parameter $\eta > 1$ describes a risk-seeking behavior. The unique BNE bid function of a risk-neutral bidder with a uniform prior distribution on the interval $[0, 1]$ is a simple linear function $\beta^*(v) = \frac{v}{2}$ [28]. Also, the BNE bid function with a CRRA utility is linear with a uniform prior as we show in the following proposition.

Proposition 2.1. *The symmetric BNE bid function $\beta(v_i)$ for a risk-averse bidder with a CRRA utility function $(u_i^{rn})^\eta$ in a first-price auction with two bidders, where each bidder's valuation is uniformly distributed on $[0, 100]$ is $\beta(v_i) = \frac{v}{1+\eta}$.*

While it is probable that the Proposition 2.1 has been established previously, we were unable to find a specific reference. All proofs can be found in the Appendix A. The proposition will help us in our initial analysis of bidders with asymmetric utilities.

Equilibrium Learning. In this paper, we leverage SODA [5], an equilibrium learning technique, to find equilibrium in asymmetric auction games. Its idea is to discretize the action and valuation spaces and represent the bid functions of the agents as distributional strategies, i.e., mixed extensions of pure BNE strategies [23]. This allows us to compute the Nash equilibrium in the discretized game using online learning methods. In our experiments, we rely on the Frank-Wolfe algorithm [13] with a discretization size of 100 points.

If SODA converges to a pure-strategy equilibrium, then it has to be a BNE. In addition, an equilibrium verifier certifies that a strategy profile found with SODA is an approximate equilibrium. Ex-ante guarantees for convergence are much harder to derive. However, Ahunbay and Bichler [1] showed that the symmetric *Bayesian Coarse Correlated Equilibrium* (BCCE) of an all-pay auction is a singleton. No-regret learners converge to a BCCE, and SODA satisfies the no-regret property. A similar result with mild assumptions on the strategies analyzed also holds for the first-price auction. It does not hold for asymmetric environments as in this paper, but the ex-post verification can always certify that a strategy profile is an approximate equilibrium.

3 STRATEGIC ANALYSIS

In our analysis, we discuss three different environments in which the levels of available information regarding the opponents' utility vary. In the first environment, we assume an isolated market where none of the agents has information about the utility of the competitor, and thus, they both play the symmetric equilibrium strategy having *false consensus*. The asymmetric utilities lead to different strategies. In this disequilibrium, the risk-averse opponent bids more aggressively than a risk-neutral firm by submitting higher bids. We study if deviations from the risk-neutral BNE by the opponent lead to lower expected profit or not. Throughout, we distinguish between *expected profit* based on risk-neutral quasi-linear utility functions and the (expected) *utility* of risk-averse or risk-seeking bidders. Of course, ex-post, after the auction, bidders will always prefer the outcome with the higher profit. In our second environment, we assume that the risk attitude is known and analyze an *asymmetric equilibrium* with asymmetric utility functions. Again, we ask how the competitors fare regarding expected profit.

Third, if bidders interact very often on agentic markets such as display ad auctions, they utilize programmable bidding bots to submit their bids. We analyze an environment where the users can parameterize their agents to bid more or less aggressively. Although the impact of risk-aversion is much reduced in repeated games of this sort, the risk parameter can now be seen as the action of an agent in a meta-game. Our analysis shows that the profit-maximizing, risk-neutral strategy is not an equilibrium in the repeated game, and, depending on the auction format, it is a Nash equilibrium for the agents to either use a risk-averse or -seeking agent.

3.1 False Consensus

In our first environment, we model an opponent as being risk-averse and a firm aiming to maximize the expected profit. Under the false consensus effect, the opponent believes that the firm is also risk-averse to the same degree, and the firm assumes that the opponent is risk-neutral. This disequilibrium analysis provides a useful baseline for the equilibrium analysis provided next, and it is an interesting analysis for real-world auction markets, as described in the introduction.

Let us first analyze the first-price auction. With a uniform prior distribution and two bidders, the symmetric risk-neutral BNE in the independent private values model is linear, and so is the symmetric BNE of two risk-averse bidders i with a CRRA utility function (Proposition 2.1). It is well-known that if both bidders played a symmetric risk-neutral BNE strategy, any deviation would necessarily lead to a loss in profit for the deviating party compared to the equilibrium player [17, Sec. 2.3]. A crucial assumption in this theoretical analysis is that not only the prior distribution but also the strategies of the agents are symmetric. However, this is not necessarily the case if bidders' utilities are asymmetric. Interestingly, with symmetric priors but asymmetric strategies, the firm sticking to the symmetric risk-neutral BNE might fare worse and have lower profit than the overbidding one in terms of profit, as we will show. The following proposition makes the case for asymmetric and linear strategies in a model with symmetric uniform priors. The firm plays the risk-neutral BNE, while the opponent bids more aggressively, possibly driven by risk-aversion.

Proposition 3.1. *Suppose opponent j bids a linear strategy $\beta_j(v) = (\frac{1}{2} + \zeta)v$ with $0 < \zeta < 0.23$ in a first-price auction with two bidders and a uniformly distributed prior $U([0, 100])$. Opponent j faces a firm i bidding the symmetric risk-neutral equilibrium bidding strategy $\beta_i(v) = \frac{v}{2}$. In this asymmetric strategy profile, the opponent j achieves a higher expected profit than the risk-neutral firm.*

Thus, the opponent firm j achieves a higher expected profit if it slightly overbids. At first sight, this result is counter-intuitive. How can one agent gain by deviating unilaterally from the BNE? However, note that in this disequilibrium, both bidders are worse off compared to what they get in the risk-neutral BNE.

A $\zeta = 0.23$ is higher than the slopes of the bidding functions that were observed by Bichler et al. [4] in a dataset from Ockenfels and Selten [25]. This means a risk-neutral firm competing against such lab subjects would have a lower expected profit with this level of risk-aversion. Suppose the risk-neutral firm now understands the level of risk-aversion of the opponent and aims to best respond to this opponent. It turns out that the firm cannot do much better with this knowledge than by playing its risk-neutral BNE in a first-price auction, as we show in the following proposition.

Proposition 3.2. *In a first-price auction, where firm i competes against opponent j that uses a fixed, linear strategy β_j and both firms draw their valuations from a uniform distribution $U([0, 100])$, the* *best response of i is given by* $\beta_i(v_i) = min\left(\frac{v_i}{200}, \bar{b}_j\right)$, *where* $\bar{b}_j$ *is the highest bid of firm* j.

This observation that the risk-averse bidder actually achieves more profit (not utility) than the risk-neutral firm carries over to the two-player all-pay auction, where the risk-neutral BNE strategy is $\beta_i(v_i) = \frac{1}{2}v_i^2$ [17].

Proposition 3.3. *Suppose a opponent* j *uses a convex bidding strategy* $\beta_j(v) = (\frac{1}{2} + \zeta)\frac{1}{100}v^2$ *for a* $0 < \zeta < 0.59$ *in a two-player all-pay auction with a uniform prior distribution* $U([0, 100])$. *Firm* i *follows the symmetric risk-neutral BNE strategy,* $\beta_i(v) = \frac{v^2}{200}$. *In this asymmetric strategy profile, bidder* j *achieves a higher expected profit than agent* i.

3.2 Asymmetric Equilibrium

Next, we analyze the asymmetric equilibrium strategies that arise in auctions where the two bidders understand the asymmetries in their utility functions, one is risk-neutral (the firm) and the other is risk-averse (the opponent). Let's assume that the prior distribution is uniform with $U([0, 100])$. In a real-world example, the firm might be a professional trader participating in used-car auctions repeatedly, while the opponent might be a private person interested in a specific car who can be expected to be risk-averse. We are not aware of analytical solutions for such equilibrium problems and draw on equilibrium learning to find the resulting equilibria.

First, we analyze the standard environment with a risk-neutral firm and a risk-averse opponent that uses a CRRA utility. We analyze two levels of risk-aversion. Figures 1 and 2 show the resulting BNE functions for both players in a first-price and an all-pay auction. While we had simple linear BNE strategies with a uniform prior in the symmetric model, the asymmetric equilibrium strategies are not linear anymore.

Again, as in the false consensus case, the risk-averse bidder is better off in terms of expected profit than the risk-neutral bidder in equilibrium. The expected profit of the risk-averse bidder ($\eta = 0.80$) in equilibrium is 16.0 (16.5), while that of the risk-neutral bidder is only 15.4 (14.6) in the first-price (all-pay) auction. In contrast, if both played a risk-neutral BNE, their profit is 16.6. So, even in equilibrium, the risk-averse bidder achieves a higher profit compared to the risk-neutral. Figure 3 shows how these strategies change if there are two risk-averse and a risk-neutral bidder in both auction types. Still, the risk-averse bidder achieves a higher expected profit than their opponent in equilibrium, and both lose in expected profit by deviating from the risk-neutral BNE. The same observation holds for the case with two risk-neutral bidders and one risk-averse bidder. The latter achieves a higher expected payoff in equilibrium than the other two. Furthermore, this insight is robust against small variations of the risk parameter.

3.3 Meta Game

Finally, consider markets where both bidders are professionals who compete against each other frequently. The market for display ads, where the same advertisers compete for the same types of users thousands of times, provides a motivation. Bidding is automated by programmable bidding bots and needs to be completed in milliseconds. Similarly, professional car traders might compete in online markets for used cars over years, and even here, the valuation of cars based on public properties and the bidding can be automated. Firms can use parameterized agents to bid more or less aggressively. Our model abstracts from the details of these markets to analyze the question whether the BNE of the stage game is also an equilibrium in a repeated game. The folk theorems in repeated game theory already indicate that this is not necessarily the case, and that an abundance of Nash equilibrium payoff profiles might emerge in repeated games [14]. We want to understand the specific equilibria that emerge in such repeated auction games and whether the equilibria of the resulting meta-game yield higher or lower profit for the bidders.

In the meta-game, a bidder could parameterize their bidding agents to be risk-averse, which induces overbidding compared to the risk-neutral strategy or risk-seeking leading to underbidding the risk-neutral strategy. The choice of utility function becomes the actions of the players in the meta-game. It is not necessarily optimal for the agents to use a risk-neutral agent in such games, even though they want to maximize their expected profit only.

In what follows, we report the expected profit of both bidders in a first-price and all-pay auction if they select a risk-neutral, risk-averse, or risk-seeking utility function with a risk parameter η in a standard CRRA utility function. The payoff Table 1 of this meta-game shows that in the first-price auction, an equilibrium is to not use a risk-neutral utility function but instead commit to a risk-seeking utility model with a large parameter. In Table 3, we compare different degrees of this utility and show that this observation is robust against small changes. Thus, choosing a risk-seeking utility with a large parameter is an equilibrium in the meta-game. In contrast, in Table 2, we varied the parameter of a risk-averse utility function and showed that a risk-averse strategy is never an equilibrium in the meta-game. The observation that high levels of risk-seeking behavior are an equilibrium in this game rests on the curvature of the resulting bid functions.

For the all-pay auction, it is always a dominant strategy for both agents to pick a risk-averse utility function (also compare Tables 4 and 5). They would both pick the highest level of risk-aversion available. This is due to the fact that a higher degree induces a highly bifurcated pattern of the bid function in equilibrium, i.e., the convex shape of the bid function increases with the degree of risk-aversion. Consequently, the bidders would bid very low for *low* valuations and, thus, increase their payoff. At the same time, they would overbid for *high* valuations. A large parameter shifts the threshold between high and low valuations to the upper bound of the valuation support, such that the effect becomes larger. Figure 4 illustrates this for exemplary parameters.

For illustrative purposes, we focused on auctions with two bidders only. However, conducted the same experiments for the setting with three bidders. The analysis confirmed that the findings of the two-player setting generalize to environments with larger numbers of bidders: risk-averse bidders constitute an equilibrium in the all-pay auction, while the risk-seeking behavior of all agents is an equilibrium in the meta-game of the first-price auction. A similar analysis was performed with a Gaussian prior, showing that the results are robust to changes in the distributional assumptions. Due to space restrictions, we cannot report the detailed payoff tables in this paper.

Auction Type	Firm i (η)	Firm j: Risk-neutral		Firm j: Risk-aversion		Firm j: Risk-seeking	
		$\mathbb{E}_{v\sim F}[u_i]$	$\mathbb{E}_{v\sim F}[u_j]$	$\mathbb{E}_{v\sim F}[u_i]$	$\mathbb{E}_{v\sim F}[u_j]$	$\mathbb{E}_{v\sim F}[u_i]$	$\mathbb{E}_{v\sim F}[u_j]$
First-price	Risk-neutral	16.749	16.749	16.371	16.513	16.928	16.747
	Risk-aversion (0.95)	16.513	16.371	16.305	16.305	16.815	16.484
	Risk-seeking (1.05)	16.747	16.928	16.484	16.815	16.984	16.984
All-Pay	Risk-neutral	16.509	16.509	16.041	16.647	16.969	16.456
	Risk-aversion (0.95)	16.647	16.041	16.134	16.134	17.039	16.019
	Risk-seeking (1.05)	16.456	16.969	16.019	17.039	16.912	16.912

Table 1: Expected payoffs of two competing firms in a first-price and all-pay auction where both agents use the same parameterizations of the utility models.

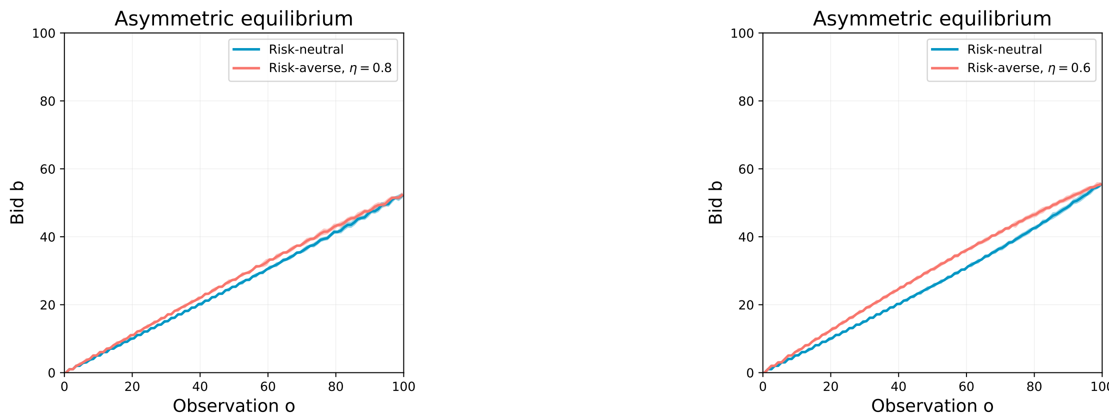

(a) $\eta = 0.80$

(b) $\eta = 0.60$

Figure 1: First-price auctions with risk-aversion

4 CONCLUSIONS

For the all-pay and the first-price auction in the symmetric independent private values model, there is a unique BNE prediction for risk-neutral bidders. It is well-known from decades of experimental research that assuming symmetric utility functions and symmetric equilibrium strategies is a too strong assumption. However, asymmetries in the utility functions led to intractable equilibrium problems, and we do not know of equilibrium predictions for such environments. Equilibrium learning has led to breakthroughs recently, allowing us to compute approximate BNEs even in asymmetric models. This allows us to explore new types of models.

Of course, if we deviate from the strong symmetry assumption, many environments are possible, differentiated by their risk attitude, and one could expect as many results as there are models. First, we explore isolated markets where bidders meet once and have different degrees of risk-aversion, a phenomenon that was often observed in the lab for small markets with only a few bidders. Interestingly, we show that in this model, the risk-averse bidder who deviates achieves higher expected profit compared to the agent who aims to maximize expected profit with a quasilinear utility function. This holds for environments with false consensus where bidders are in disequilibrium and for asymmetric equilibrium models. This is an interesting insight for high-stakes auctions as it shows how a risk-averse bidder has a substantial negative effect on a competitor that aims to maximize profit.

Second, we analyze agentic markets with repeated interaction, where the parameterization of the bidding agents results in a meta-game that has interesting and non-obvious equilibria. Across a wide range of risk parameters, we show that it is an equilibrium in the meta-game to choose an agent that is maximally risk-seeking in the first-price auction and one that is maximally risk-averse in the all-pay auction. Both equilibria of the meta-game lead to bidding strategies in the individual auctions that are less aggressive than the BNE of the stage game and thus increase the payoff of the bidders. The results carry over to auctions with more than two bidders and are robust to alternative prior distributions. Although this is a stylized model of an agentic market, the results suggest that one should expect bidding strategies of automated agents that lead to bids below the BNE strategy of a first-price or all-pay auction as soon as they understand the repeated nature of these games. This is important for agentic markets such as display ad auctions, where the same bidders compete for similar items very often.

Overall, this paper breaks new ground in analyzing auction markets with asymmetric utility functions, environments that have received little attention in multi-agent systems and auction theory.

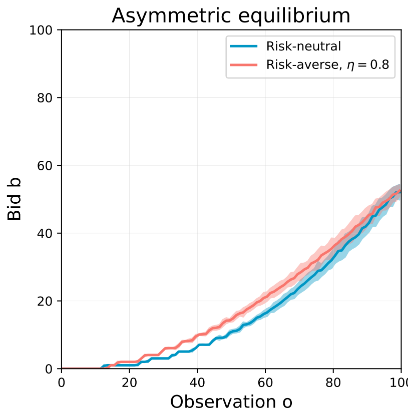

(a) $\eta = 0.80$

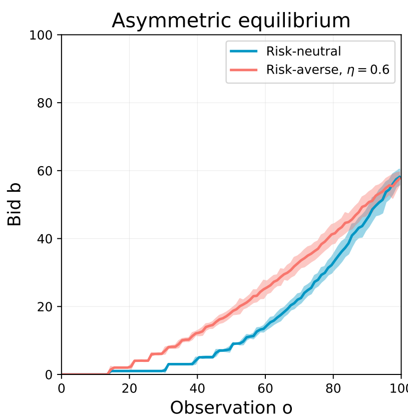

(b) $\eta = 0.60$

Figure 2: All-pay auctions with risk-aversion

ACKNOWLEDGMENTS

Bichler and Ewert gratefully acknowledge the financial support from the Deutsche Forschungsgemeinschaft (DFG) (grant no. BI 1057/1-10). Ockenfels gratefully acknowledges support by the German Science Foundation through Germany's Excellence Strategy (EXC 2126/1 390838866) and through the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 741409).

A PROOFS

A.1 Proof of Proposition 2.1

Proof. In a symmetric equilibrium, all bidders bid according to β. In a situation with two bidders, we need to show that if bidder 2 follows β, then it is also a best response for bidder 1 to follow β. Bidder 1 with true valuation v will make a bid b. Since β is invertible, we can find a valuation z such that $b = \beta(z)$. Bidder 1 wins the auction when the valuation of the other bidder is smaller than z. The probability of this event is $G(z) = z$. Bidder 1 chooses z to maximize their expected utility:

$$EU = G(z) \cdot u(v - \beta(z)), \tag{2}$$

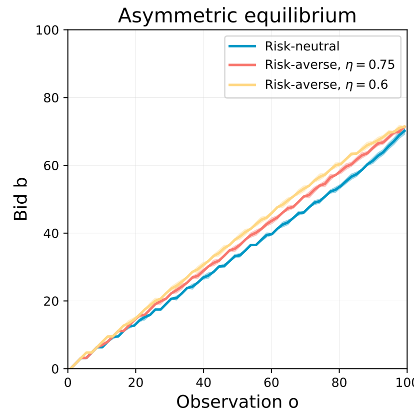

(a) First-price auction with three bidders

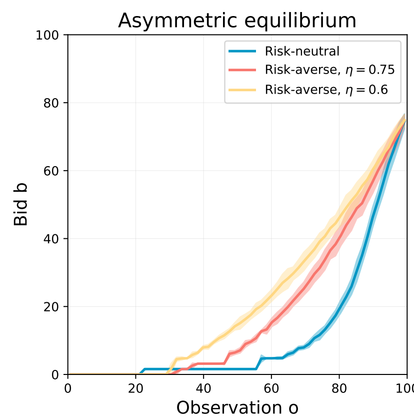

(b) All-pay auction with three bidders

Figure 3: Three bidders with risk-aversion.

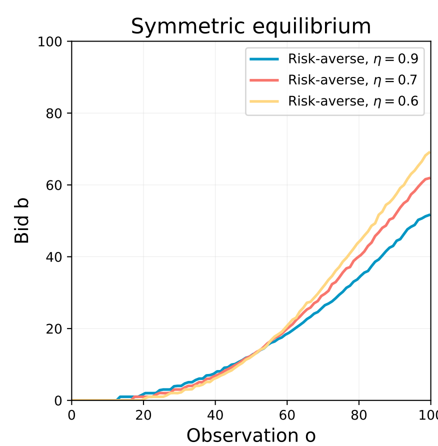

Figure 4: The symmetric BNE in an all-pay auction where bidders exhibit risk-aversion with different parameters.

which yields the first-order condition:

$$(v - \beta(z))^{\rho} \frac{1}{100} - \rho \cdot \frac{z}{100} \cdot (v - \beta(z))^{\rho - 1} \beta'(z) = 0. \tag{3}$$

In the symmetric equilibrium, we have $z = v$, and can derive an alternative form

$$(v - \beta(v))^{\rho-1} \cdot (\beta(v) + v \cdot (\rho\beta'(v) - 1)) = 0. \tag{4}$$

For $\rho \in (0, 1]$, there are several solutions to this differential equation. In equilibrium $\beta(0) = 0$, which yields the unique solution:

$$\beta^*(v) = \frac{v}{1+\rho}. \tag{5}$$

The second derivative is negative, which completes the proof. □

A.2 Proof of Proposition 3.1

Proof. The two agents choose their bids b_i and b_j according to some fixed strategies $\beta_i(v)$ and $\beta_j(v)$, where $v \sim F = U([0, 100])$. We assume that these strategies are linear and that j bids more aggressive than i, i.e. $\beta_i(v) = \lambda v$ and $\beta_j(v) = \lambda' v = (\lambda + \zeta)v$ with $\lambda \in (0, 100), \lambda + \zeta \in (0, 100)$, and $\zeta \geq 0$. After playing the auction, the agents observe their payoffs $p_i(v_i, b_i) = v_i - b_i$ and $p_j(v_j, b_j) = v_j - b_j$ if they won, otherwise $p_i = p_j = 0$. The expected payoffs for both players are:

$$\mathbb{E}_{v\sim F}(p_i(v_i, b_i)) = \int_0^{100} P(win)_i \cdot (t_i - \beta_i(t_i))dt_i \tag{6}$$

$$\mathbb{E}_{v\sim F}(p_j(v_j, b_j)) = \int_0^{v'} P(win)_j \cdot (t_j - \beta_j(t_j))dt_j + \int_{v'}^{100} (t_j - \beta_j(t_j))dt_j$$

Since j bids more aggressively than i, there exists a valuation v', where j wins with probability 1. Thus, we need to differentiate between the two cases in the definition of the corresponding expected payoff above. Given our assumptions, this valuation corresponds to $v' = 100\lambda/\lambda'$. The probability of winning $P(win)$ determines how likely it is that the submitted bid is higher than the bid of the opponent. It follows for agent i (and j analogously):

$$P(win)_i = P(\beta_j(v_j) < \beta_i(v_i)) = P(v_j < \beta_j^{-1}(\beta_i(v_i))) \tag{7}$$

$$= F(\beta_j^{-1}(\beta_i(v_i))) \overset{F\equiv U(0,100)}{=} \beta_j^{-1}(\beta_i(v_i))/100 \tag{8}$$

Consequently, the expected payoff of agent j is higher than the expected payoff of agent i if the following inequality holds:

$$\mathbb{E}_{v\sim F}(p_j(v_j, b_j)) \geq \mathbb{E}_{v\sim F}(p_i(v_i, b_i)) \tag{9}$$

The inverse bid functions of the players are given by $\beta_i^{-1}(b_i) = b_i/\lambda$ and $\beta_j^{-1}(b_j) = b_j/\lambda'$. Inserting the (inverse) bid functions, and winning probabilities leads to:

$$\frac{1}{100\lambda}\int_0^{v'} \lambda' \cdot t_j^2 - \lambda'^2 \cdot t_j^2 dt_j + \int_{v'}^{100} t_j - \lambda' \cdot t_j dt_j \geq \tag{10}$$

$$\frac{1}{100\lambda'}\int_0^{100} \lambda \cdot t_i^2 - \lambda^2 \cdot t_i^2 dt_i \tag{11}$$

We can integrate the term and substitute v', resulting in:

$$\frac{5000\lambda^2}{\lambda'} - \frac{5000\lambda^2}{3\lambda'^2} - \frac{10000\lambda}{3\lambda'} - 5000\lambda' + 5000 \geq 0 \tag{12}$$

The resulting equation can be used to determine for a given λ, at which ζ, the expected payoff of the more aggressive bidding agent is higher than for the other agent. Under the assumption that the firm i follows the risk-neutral BNE, the left-hand side of this equation is non-negative if $\zeta \leq 0.23$. □

A.3 Proof of Proposition 3.2

Proof. Assume that the strategy β_j of firm j is fixed and linear, and firm i has knowledge about it. *Ex-ante*, the expected utility of i is given by

$$\max_{b_i} \mathbb{E}_{v_j\sim F}(v_i - b_i)G_j(b_i), \tag{13}$$

where $G_j(\cdot)$ is the distribution of bids of firm j. In this setting, Guerre et al. [15] showed that $G_j(\cdot) = F(\cdot)$, with supports $supp(G_j) = [0, \beta_j(100)]$. We differentiate between two cases:

Case 1, $0 \leq b_i \leq \beta_j(100)$*:* The optimization problem is given by

$$\max_{b_i} \mathbb{E}_{v_j\sim F}(v_i - b_i)\frac{b_i}{\beta_j(100)} = \max_{b_i}(v_i - b_i)\frac{b_i}{\beta_j(100)}. \tag{14}$$

Solving for the first-order condition leads to

$$\frac{d}{db_i}(v_i - b_i)\frac{b_i}{\beta_j(1)} = (v_i - b_i)\frac{1}{\beta_j(1)} - \frac{b_i}{\beta_j(1)} \overset{!}{=} 0 \Rightarrow b_i = \frac{v_i}{2}, \tag{15}$$

which corresponds to the symmetric risk-neutral BNE.

Case 2, $\beta_j(100) \leq b_i$*:* The optimization problem is given by

$$\max_{b_i}(v_i - b_i) \cdot 1. \tag{16}$$

Thus, firm i does not overbid the highest bid of firm j. Taken together, the best response of i is given by

$$b_i = \begin{cases} \frac{v_i}{200}, & 0 \leq \frac{v_i}{200} \leq \beta_j(100) \\ \beta_j(100), & \beta_j(100) \leq \frac{v_i}{200}. \end{cases} \tag{17}$$

This strategy is also the best response in the *ex-ante* stage. □

A.4 Proof of Proposition 3.3

Proof. Assume two agents compete in an all-pay auction with a uniform prior distribution $F = U([0, 100])$. Both agents i and j choose their bids b_i and b_j according to some fixed strategies $\beta_i(v)$ and $\beta_j(v)$. We assume that both strategies are convex and agent j bids more aggressive than bidder i, i.e., $\beta_i(v) = \frac{\lambda}{100}v^2$ and $\beta_j(v) = \frac{\lambda'}{100}v^2 = \frac{(\lambda+\zeta)}{100}v^2$, where $\zeta > 0$ and $\lambda + \zeta < 1$. Both agents pay their bid regardless of whether they won the auction or not, and the winner obtains a payoff $p_k = v_k - b_k$, with valuation v_k and bid b_k of player k. Analogously to the proof of Proposition 3.1, we substitute the bid functions, the winning probabilities, and $v' = \sqrt{\lambda/\lambda'}100$ into the expected payoffs of the agent:

$$\mathbb{E}_{v\sim F}(p_j(v_j, b_j)) \geq \mathbb{E}_{v\sim F}(p_i(v_i, b_i)) \tag{18}$$

This leads to the following inequality:

$$\frac{10000\lambda}{3} - \frac{5000\lambda}{\lambda'} - \frac{10000\lambda'}{3} + \frac{10000\sqrt{\frac{\lambda'}{\lambda}}(\frac{\lambda}{\lambda'})^{3/2}}{3} - \tag{19}$$

$$\frac{10000\sqrt{\frac{\lambda}{\lambda'}}}{3} + 5000 \geq 0 \tag{20}$$

We can leverage this equation to determine for a given λ, at which ζ, the expected payoff of bidder j is higher than for agent i. For our setting, this is the case if $\zeta \in (0, 0.60)$. □

B FURTHER RESULTS

The following tables provide an in-depth overview of the results described in Section 3.3 above.

Table 2: Expected payoff of two competing firms that either exhibit risk-neutrality or risk-aversion in the first-price auction. Here the only Nash equilibrium is to adhere to the risk-neutral strategy.

Firm i	Firm j					
	Risk-neutral	Risk-aversion (η)				
	-	0.95	0.90	0.85	0.80	0.75
Risk-neutral	**(16.749, 16.749)**	(16.371, 16.513)	(16.036, 16.356)	(15.728, 16.203)	(15.413, 16.026)	(15.089, 15.825)
Risk-aversion (0.95)	(16.513, 16.371)	(16.305, 16.305)	(15.933, 16.088)	(15.628, 15.935)	(15.301, 15.749)	(14.978, 15.540)
Risk-aversion (0.90)	(16.356, 16.036)	(16.088, 15.933)	(15.837, 15.837)	(15.482, 15.618)	(15.157, 15.442)	(14.815, 15.239)
Risk-aversion (0.85)	(16.203, 15.728)	(15.935, 15.628)	(15.618, 15.482)	(15.336, 15.336)	(15.001, 15.124)	(14.665, 14.935)
Risk-aversion (0.80)	(16.026, 15.413)	(15.749, 15.301)	(15.442, 15.157)	(15.124, 15.001)	(14.833, 14.833)	(14.488, 14.611)
Risk-aversion (0.75)	(15.825, 15.089)	(15.540, 14.978)	(15.239, 14.815)	(14.935, 14.665)	(14.611, 14.488)	(14.293, 14.293)

Table 3: Expected payoff of two competing firms that either exhibit risk-neutrality or are risk-seeking in the first-price auction. Here the Nash equilibria are to play risk-seeking with a high parameterization.

Firm i	Firm j					
	Risk-neutral	Risk-seeking (η)				
	-	1.05	1.10	1.15	1.20	1.25
Risk-neutral	(16.749, 16.749)	(16.928, 16.747)	(17.190, 16.859)	(17.481, 16.958)	(17.745, 17.026)	(18.002, 17.110)
Risk-seeking (1.05)	(16.747, 16.928)	(16.984, 16.984)	(17.314, 17.154)	(17.612, 17.252)	(17.859, 17.343)	(18.115, 17.421)
Risk-seeking (1.10)	(16.859, 17.190)	(17.154, 17.314)	(17.391, 17.391)	(17.689, 17.532)	(17.986, 17.598)	(18.244, 17.695)
Risk-seeking (1.15)	(16.958, 17.481)	(17.252, 17.612)	(17.532, 17.689)	(17.766, 17.766)	(18.047, 17.889)	(18.339, 17.958)
Risk-seeking (1.20)	(17.026, 17.745)	(17.343, 17.859)	(17.598, 17.986)	(17.889, 18.047)	(18.131, 18.131)	(18.406, 18.224)
Risk-seeking (1.25)	(17.110, 18.002)	(17.421, 18.115)	(17.695, 18.244)	(17.958, 18.339)	(18.224, 18.406)	**(18.482, 18.482)**

Table 4: Expected payoff of two competing firms that either exhibit risk-neutrality or risk-aversion in the all-pay auction. Here the only Nash equilibrium is to use a risk-averse utility model with a parameter η that describes a high degree of risk-aversion.

Firm i	Firm j					
	Risk-neutral	Risk-aversion (η)				
	-	0.95	0.90	0.85	0.80	0.75
Risk-neutral	(16.509, 16.509)	(16.041, 16.647)	(15.642, 16.585)	(15.067, 16.604)	(14.569, 16.583)	(13.949, 16.493)
Risk-aversion (0.95)	(16.647, 16.041)	(16.134, 16.134)	(15.651, 16.192)	(15.137, 16.176)	(14.652, 16.202)	(14.025, 16.109)
Risk-aversion (0.9)	(16.585, 15.642)	(16.192, 15.651)	(15.732, 15.732)	(15.219, 15.733)	(14.687, 15.741)	(14.183, 15.710)
Risk-aversion (0.85)	(16.604, 15.067)	(16.176, 15.137)	(15.733, 15.219)	(15.304, 15.304)	(14.750, 15.254)	(14.216, 15.260)
Risk-aversion (0.80)	(16.583, 14.569)	(16.202, 14.652)	(15.741, 14.687)	(15.254, 14.750)	(14.811, 14.811)	(14.217, 14.709)
Risk-aversion (0.75)	(16.493, 13.949)	(16.109, 14.025)	(15.710, 14.183)	(15.260, 14.216)	(14.709, 14.217)	**(14.242, 14.242)**

Table 5: Expected payoff of two competing firms that either exhibit risk-neutrality or are risk-seeking in the all-pay auction. Here, the Nash equilibrium is to play the risk-neutral BNE.

Firm i	Firm j					
	Risk-neutral	Risk-seeking (η)				
	-	1.05	1.10	1.15	1.20	1.25
Risk-neutral	**(16.509, 16.509)**	(16.969, 16.456)	(17.364, 16.343)	(17.767, 16.268)	(18.070, 16.105)	(18.430, 16.008)
Risk-seeking (1.05)	(16.456, 16.969)	(16.912, 16.912)	(17.322, 16.810)	(17.705, 16.718)	(18.031, 16.587)	(18.404, 16.463)
Risk-seeking (1.10)	(16.343, 17.364)	(16.810, 17.322)	(17.208, 17.208)	(17.582, 17.147)	(17.959, 16.981)	(18.309, 16.853)
Risk-seeking (1.15)	(16.268, 17.767)	(16.718, 17.705)	(17.147, 17.582)	(17.512, 17.512)	(17.886, 17.399)	(18.259, 17.309)
Risk-seeking (1.20)	(16.105, 18.070)	(16.587, 18.031)	(16.981, 17.959)	(17.399, 17.886)	(17.765, 17.765)	(18.144, 17.668)
Risk-seeking (1.25)	(16.008, 18.430)	(16.463, 18.404)	(16.853, 18.309)	(17.309, 18.259)	(17.668, 18.144)	(18.054, 18.054)

REFERENCES

[1] Mete Şeref Ahunbay and Martin Bichler. 2025. On the Uniqueness of Bayesian Coarse Correlated Equilibria in Standard First-Price and All-Pay Auctions. *SIAM Symposium on Discrete Algorithms* (2025).

[2] Kenneth Joseph Arrow. 1965. *Aspects of the theory of risk-bearing*. Helsinki: Yrjö Jahnsonian Sä iö.

[3] Martino Banchio and Andrzej Skrzypacz. 2022. Artificial intelligence and auction design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 30–31.

[4] M. Bichler, M. Ewert, and A. Ockenfels. 2024. Estimating behavioral motives from experimental data. *TUM Working paper* 10, 1 (2024).

[5] Martin Bichler, Max Fichtl, and Matthias Oberlechner. 2024. Computing Bayes–Nash Equilibrium Strategies in Auction Games via Simultaneous Online Dual Averaging. *Operations Research* (2024). https://doi.org/10.1287/opre.2022.0287 arXiv:https://doi.org/10.1287/opre.2022.0287

[6] Martin Bichler, Kemal Guler, and Stefan Mayer. 2015. Split-Award Procurement Auctions—Can Bayesian Equilibrium Strategies Predict Human Bidding Behavior in Multi-Object Auctions? *Production and Operations Management* 24, 6 (2015), 1012–1027.

[7] Francesco Bof and Pietro Previtali. 2010. National models of public (e)-procurement in Europe. *Journal of e-Government Studies and Best Practices* 315295, 2010 (2010), 1–14.

[8] Shuchi Chawla and Jason D Hartline. 2013. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 181–196.

[9] Hana Choi, Carl F Mela, Santiago R Balseiro, and Adam Leary. 2020. Online display advertising markets: A literature review and future directions. *Information Systems Research* 31, 2 (2020), 556–575.

[10] James C. Cox, Vernon L. Smith, and James M. Walker. 1988. Theory and Individual Behavior of First-Price Auctions. *Journal of Risk and Uncertainty* 1, 1 (March 1988), 61–99. https://doi.org/10.1007/BF00055565

[11] Richard Engelbrecht-Wiggans and Elena Katok. 2007. Regret in auctions: Theory and evidence. *Economic Theory* 33, 1 (2007), 81–101.

[12] Gadi Fibich, Arieh Gavious, and Aner Sela. 2006. All-pay auctions with risk-averse players. *International Journal of Game Theory* 34, 4 (2006), 583–599.

[13] Marguerite Frank, Philip Wolfe, et al. 1956. An algorithm for quadratic programming. *Naval research logistics quarterly* 3, 1-2 (1956), 95–110.

[14] James W Friedman. 1971. A non-cooperative equilibrium for supergames. *The Review of Economic Studies* 38, 1 (1971), 1–12.

[15] Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. 2000. Optimal nonparametric estimation of first-price auctions. *Econometrica* 68, 3 (2000), 525–574.

[16] John H Kagel and Alvin E Roth. 2015. *The Handbook of Experimental Economics, Volume 2*. Princeton University Press.

[17] Vijay Krishna. 2009. *Auction theory*. Academic press.

[18] Vijay Krishna and John Morgan. 1997. An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory* 72, 2 (1997), 343–362.

[19] Bernard Lebrun. 2006. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior* 55, 1 (2006), 131–151.

[20] Jingfeng Lu and Isabelle Perrigne. 2008. Estimating risk aversion from ascending and sealed-bid auctions: The case of timber auction data. *Journal of Applied Econometrics* 23, 7 (2008), 871–896.

[21] Gary Marks and Norman Miller. 1987. Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin* 102, 1 (1987), 72–90.

[22] Eric Maskin and John Riley. 2000. Asymmetric auctions. *The review of economic studies* 67, 3 (2000), 413–438.

[23] Paul R. Milgrom and Robert J. Weber. 1985. Distributional strategies for games with incomplete information. *Mathematics of Operations Research* 10, 4 (1985), 619–632.

[24] Wieland Müller and Andrew Schotter. 2010. Workaholics and dropouts in organizations. *Journal of the European Economic Association* 8, 4 (2010), 717–743.

[25] Axel Ockenfels and Reinhard Selten. 2005. Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior* 51, 1 (2005), 155–170.

[26] John W Pratt. 1978. Risk aversion in the small and in the large. In *Uncertainty in Economics*. Elsevier, 59–79.

[27] Philip J Reny and Shmuel Zamir. 2004. On the existence of pure strategy monotone equilibria in asymmetric first-price auctions. *Econometrica* 72, 4 (2004), 1105–1125.

[28] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 1 (1961), 8–37.