

# Price of Anarchy for Mechanisms with Risk-Averse Agents

**Thomas Kesselheim**

University of Bonn, Institute of Computer Science, Bonn, Germany

thomas.kesselheim@uni-bonn.de

**Bojana Kodric**

MPI for Informatics and Saarland University, Saarbrücken, Germany

bojana@mpi-inf.mpg.de

---

## Abstract

We study the price of anarchy of mechanisms in the presence of risk-averse agents. Previous work has focused on agents with quasilinear utilities, possibly with a budget. Our model subsumes this as a special case but also captures that agents might be less sensitive to payments than in the risk-neutral model. We show that many positive price-of-anarchy results proved in the smoothness framework continue to hold in the more general risk-averse setting. A sufficient condition is that agents can never end up with negative quasilinear utility after playing an undominated strategy. This is true, e.g., for first-price and second-price auctions. For all-pay auctions, similar results do not hold: We show that there are Bayes-Nash equilibria with arbitrarily bad social welfare compared to the optimum.

**2012 ACM Subject Classification** Theory of computation → Algorithmic mechanism design, Theory of computation → Quality of equilibria

**Keywords and phrases** Mechanism Design, Price of Anarchy, Risk Aversion, Smoothness

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2018.155

**Related Version** A full version of the paper is available at [12], <https://arxiv.org/abs/1804.09468>.

**Funding** This work is supported by DFG through Cluster of Excellence MMCI.

## 1 Introduction

Many practical, “simple” auction mechanisms are not incentive compatible, making it beneficial for agents to behave strategically. A standard example is the first-price auction, in which one item is sold to one of  $n$  agents. Each of these agents is asked to report a valuation; the item is given to the agent reporting the highest value, who then has to pay what he/she reported. A common way to understand the effects of strategic behavior is to study resulting equilibria and to bound the *price of anarchy*. That is, one compares the social welfare that is achieved at the (worst) equilibrium of the induced game to the maximum possible welfare. Typical equilibrium concepts are Bayes-Nash equilibria and (coarse) correlated equilibria, which extend mixed Nash equilibria toward incomplete information or learning settings respectively. A key assumption in these analyses is that agents are *risk neutral*: Agents are assumed to maximize their expected quasilinear utility, which is defined to be the difference of the value associated to the outcome and payment imposed to the agent. So, an agent having a value of 1 for an item would be indifferent between getting this item with probability 10% for free and getting it all the time, paying 0.9.



© Thomas Kesselheim and Bojana Kodric;

licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).

Editors: Ioannis Chatzigiannakis, Christos Kaklamani, Dániel Marx, and Donald Sannella;

Article No. 155; pp. 155:1–155:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



However, there are many reasons to believe that agents are not risk neutral. For instance, in the above example the agent might favor the certain outcome to the uncertain one. Therefore, in this paper, we ask the question: *What “simple” auction mechanisms preserve good performance guarantees in the presence of risk-averse agents?*

The standard model of risk aversion in economics (see, e.g., [16]) is to apply a (weakly) concave function to the quasilinear term. That is, if agent  $i$ 's outcome is  $x_i$  and his payment is  $p_i$ , his utility is given as  $u_i(x_i, p_i) = h_i(v_i(x_i) - p_i)$ , where  $h_i: \mathbb{R} \rightarrow \mathbb{R}$  is a weakly concave, monotone function. That is, for  $y, y' \in \mathbb{R}$  and for all  $\lambda \in [0, 1]$ , it holds that  $h_i(\lambda y + (1 - \lambda)y') \geq \lambda h_i(y) + (1 - \lambda)h_i(y')$ . Agent  $i$  is risk neutral if and only if  $h_i$  is a linear function. If the function is strictly concave, this has the effect that, by Jensen's inequality, the utility for fixed  $x_i$  and  $p_i$  is higher than for a randomized  $x_i$  and  $p_i$  with the same expected  $v_i(x_i) - p_i$ .

We compare outcomes based on their social welfare, which is defined to be the sum of utilities of all involved parties including the auctioneer. That is, it is the sum of agents' utilities and their payments  $\text{SW}(\mathbf{x}, \mathbf{p}) = \sum_i u_i(x_i, p_i) + \sum_i p_i$ . In the quasilinear setting this definition of social welfare coincides with the sum of values  $\sum_i v_i(x_i)$ . With risk-averse utilities they usually differ. However, all our results bound the sum of values and therefore also hold for this benchmark.

We assume that the mechanisms are oblivious to the  $h_i$ -functions and work like in the quasilinear model. Only the individual agent's perception changes. This makes it necessary to normalize the  $h_i$ -functions because otherwise they could be on different scales, e.g., if  $h_1(y) = y$  and  $h_2(y) = 1000 \cdot y$ , which would be impossible for the mechanism to cope with without additional information. Therefore, we will assume that  $u_i(\mathbf{x}, p_i) = v_i(\mathbf{x})$  if  $p_i = 0$  and that  $u_i(\mathbf{x}, p_i) = 0$  if  $p_i = v_i$ . That is, for the two cases that  $p_i$  is either 0 or the full value, the utility matches exactly the quasilinear one. However, due to risk aversion, the agents might be less sensitive to payments.<sup>1</sup>

## 1.1 Our Contribution

We give bounds on the price of anarchy for Bayes-Nash and (coarse) correlated equilibria of mechanisms in the presence of risk-averse agents. Our positive results are stated within the smoothness framework, which was introduced by [24]. We use the version that is tailored to quasilinear utilities by [26], which we extend to mechanism settings with general utilities (for a formal definition see Section 4). Our main positive result states that the loss of performance compared to the quasilinear setting is bounded by a constant if a slightly stronger smoothness condition is fulfilled.

► **Main Result 1.** *Given a mechanism with price of anarchy  $\alpha$  in the quasilinear model provable via smoothness such that the deviation guarantees non-negative utility, then this mechanism has price of anarchy at most  $2\alpha$  in the risk-averse model.*

This result relies on the fact that the deviation action to establish smoothness guarantees agents non-negative utility. A sufficient condition is that all undominated strategies never have negative utility. First-price and second-price auctions satisfy this condition, we thus get constant price-of-anarchy bounds for both of these auction formats.

---

<sup>1</sup> We note here that this will not in turn allow the mechanism to arrive at huge utility gains, as compared to the quasilinear model, for example, by increasing payments arbitrarily. Indeed, Lemma 1 in Section 3 will show that the difference between the two optima is bounded by at most a multiplicative factor of 2.

In an all-pay auction every positive bid can lead to negative utility. Therefore, the positive result does not apply. As a matter of fact, this is not a coincidence because, as we show, equilibria can be arbitrarily bad.

► **Main Result 2.** *The single-item all-pay auction has unbounded price of anarchy for Bayes-Nash equilibria, even with only three agents.*

This means that although equilibria of first-price and all-pay auctions have very similar properties with quasilinear utilities, in the risk-averse setting they differ a lot. We feel that this to some extent matches the intuition that agents should be more reluctant to participate in an all-pay auction compared to a first-price auction.

In our construction for proving Main Result 2, we give a symmetric Bayes-Nash equilibrium for two agents. The equilibrium is designed in such a way that a third agent of much higher value would lose with some probability with every possible bid. Losing in an all-pay auction means that the agent has to pay without getting anything, resulting in negative utility. In the quasilinear setting, this negative contribution to the utility would be compensated by respective positive amounts when winning. For the risk-averse agent in our example, this is not true. Because of the risk of negative utility, he prefers to opt out of the auction entirely.

We also consider a different model of aversion to uncertainty, in which solution concepts are modified. Instead of evaluating a distribution over utilities in terms of their expectation, agents evaluate them based on the expectation minus a second-order term. We find that this model has entirely different consequences on the price of anarchy. For example, the all-pay auction has a constant price of anarchy in correlated and Bayes-Nash equilibria, whereas the second-price auction can have an unbounded price of anarchy in correlated equilibria.

## 1.2 Related Work

Studying the impact of risk-averseness is a regularly reoccurring theme in the literature. A proposal to distinguish between money and the utility of money, and to model risk aversion by a utility function that is concave first appeared in [1]. The expected utility theory, which basically states that the agent's behavior can be captured by a utility function and the agent behaves as a maximizer of the expectation of his utility, was postulated in [27]. This theory does not capture models that are standardly used in portfolio theory, "expectation minus variance" or "expectation minus standard deviation" [14], the latter of which we also consider in Section 7.

In the context of mechanisms, one usually models risk aversion by concave utility functions. One research direction in this area is to understand the effects of risk aversion on a given mechanism. For example, [6] studies symmetric equilibria in all-pay auctions with a homogenous population of risk-averse players. Due to symmetry and homogeneity, in this case, equilibria are fully efficient, that is, the price of anarchy is 1. In [18] a similar analysis for auctions with a buyout option is performed; [10] considers customers with heterogeneous risk attitudes in mechanisms for cloud resources. In [4] it is shown that for certain classes of mechanisms the correlated equilibrium is unique and has a specific structure depending on the respective valuations but independent of the actual utility function. One consequence of this result is that risk aversion does not influence the allocation outcome or the revenue.

Another direction is to design mechanisms for the risk-averse setting. For example, the optimal revenue is higher because buyers are less sensitive to payments. In a number of papers, mechanisms for revenue maximization are proposed [19, 17, 25, 11, 2, 7]. Furthermore, randomized mechanisms that are *truthful in expectation* lose their incentive properties if agents

are not risk neutral. Black-box transformations from truthful-in-expectation mechanisms into ones that fulfill stronger properties are given in [3] and [8].

Studying the effects of risk aversion also has a long history in game theory, where different models of agents' attitudes towards risk are analyzed. One major question is, for example, if equilibria still exist and if they can be computed [5, 9]. Price of anarchy analyses have so far only been carried out for congestion games. Tight bounds on the price of anarchy for atomic congestion games with affine cost functions under a range of risk-averse decision models are given in [23].

The smoothness framework was introduced by [24]. Among others, [26] tailored it to the quasilinear case of mechanisms. It is important to remark here that our approach is different from the one taken by [20]. They use the smoothness framework to prove generalized price of anarchy bounds for nonatomic congestion games in which players have biased utility functions. They assume that players are playing the “wrong game” and their point of comparison is the “true” optimal social welfare, meaning that the biases only determine the equilibria but do not affect the social welfare. We take the utility functions as they are, including the risk aversion, to evaluate social welfare in equilibria and also to determine the optimum, which makes our models incomparable.

For precise relation of von Neumann-Morgenstern preferences to mean-variance preferences, see for instance [15]. Mean-variance preferences were explored for congestion games in [21, 22], while [13] studies the bidding behavior in an all-pay auction depending on the level of variance-averseness.

## 2 Preliminaries

### 2.1 Setting

We consider the following setting: There is a set  $N$  of  $n$  players and  $\mathcal{X}$  is the set of possible outcomes. Each player  $i$  has a utility function  $u_i^{\theta_i}$ , which is parameterized by her type  $\theta_i \in \Theta_i$ . Given a type  $\theta_i$ , an outcome  $\mathbf{x} \in \mathcal{X}$ , and a payment  $p_i \geq 0$ , her utility is  $u_i^{\theta_i}(\mathbf{x}, p_i)$ . The traditionally most studied case are quasilinear utilities, in which types are valuation functions  $v_i \in \mathcal{V}_i$ ,  $v_i: \mathcal{X} \rightarrow \mathbb{R}$  and  $u_i^{v_i}(\mathbf{x}, p_i) = v_i(\mathbf{x}) - p_i$ . Throughout this paper, we will refer to quasilinear utilities by  $\hat{u}_i^{v_i}$ .

For fixed utility functions and types, the social welfare of an outcome  $\mathbf{x} \in \mathcal{X}$  and payments  $(p_i)_{i \in N}$  is defined as  $\text{SW}^\theta(\mathbf{x}, \mathbf{p}) := \sum_{i \in N} u_i^{\theta_i}(\mathbf{x}, p_i) + \sum_{i \in N} p_i$ . In the quasilinear case, this simplifies to  $\sum_{i \in N} v_i(\mathbf{x})$ . Unless noted otherwise, by  $\text{OPT}(\theta)$ , we will refer to the optimal social welfare under type profile  $\theta$ , i.e.,  $\max_{\mathbf{x}, \mathbf{p}} \text{SW}^\theta(\mathbf{x}, \mathbf{p})$ .

A *mechanism* is a triple  $(\mathcal{A}, X, P)$ , where for each player  $i$ , there is a set of actions  $\mathcal{A}_i$  and  $\mathcal{A} = \times_i \mathcal{A}_i$  is the set of action profiles,  $X: \mathcal{A} \rightarrow \mathcal{X}$  is an allocation function that maps actions to outcomes and  $P: \mathcal{A} \rightarrow \mathbb{R}_+^n$  is a payment function that maps actions to payments  $p_i$  for each player  $i$ . Given an action profile  $\mathbf{a} \in \mathcal{A}$ , we will use the short-hand notation  $u_i^{\theta_i}(\mathbf{a})$  to denote  $u_i^{\theta_i}(X(\mathbf{a}), p_i)$ .

### 2.2 Solution Concepts

In the setting of *complete information*, the type profile  $\theta$  is fixed. We consider (coarse) correlated equilibria, which generalize Nash equilibria and are the outcome of (no-regret) learning dynamics. A *correlated equilibrium (CE)* is a distribution  $\mathbf{a}$  over action profiles from  $\mathcal{A}$  such that for every player  $i$  and every strategy  $a_i$  in the support of  $\mathbf{a}$  and every action

$a'_i \in \mathcal{A}_i$ , player  $i$  does not benefit from switching to  $a'_i$  whenever he was playing  $a_i$ . Formally,

$$\mathbb{E}_{\mathbf{a}_{-i}|a_i}[u_i(\mathbf{a})] \geq \mathbb{E}_{\mathbf{a}_{-i}|a_i}[u_i(a'_i, \mathbf{a}_{-i})], \forall a'_i \in \mathcal{A}_i, \forall i .$$

In *incomplete information*, the type of each player is drawn from a distribution  $F_i$  over her type space  $\Theta_i$ . The distributions are common knowledge and the draws are independent among players. The solution concept we consider in this setting is the *Bayes-Nash Equilibrium*. Here, the strategy of each player is now a (possibly randomized) function  $s_i: \Theta_i \rightarrow \mathcal{A}_i$ . The equilibrium is a distribution over these functions  $s_i$  such that each player maximizes her expected utility conditional on her private information. Formally,

$$\mathbb{E}_{\theta_{-i}|\theta_i}[u_i^{\theta_i}(\mathbf{s}(\boldsymbol{\theta}))] \geq \mathbb{E}_{\theta_{-i}|\theta_i}[u_i^{\theta_i}(a_i, \mathbf{s}_{-i}(\boldsymbol{\theta}_{-i}))], \forall a_i \in \mathcal{A}_i, \forall \theta_i \in \Theta_i, \forall i .$$

The measure of efficiency is the expected social welfare over the types of the players: given a strategy profile  $\mathbf{s}: \times_i \Theta_i \rightarrow \times_i \mathcal{A}_i$ , we consider  $\mathbb{E}_{\boldsymbol{\theta}}[\text{SW}^{\boldsymbol{\theta}}(\mathbf{s}(\boldsymbol{\theta}))]$ . We compare the efficiency of our solution concept with respect to the expected optimal social welfare  $\mathbb{E}_{\boldsymbol{\theta}}[\text{OPT}(\boldsymbol{\theta})]$ .

The *price of anarchy (PoA)* with respect to an equilibrium concept is the worst possible ratio between the optimal expected welfare and the expected welfare at equilibrium, that is

$$\text{PoA} = \max_F \max_{D \in \text{EQ}(F)} \frac{\mathbb{E}_{\boldsymbol{\theta} \sim F}[\text{OPT}(\boldsymbol{\theta})]}{\mathbb{E}_{\boldsymbol{\theta} \sim F, \mathbf{a} \sim D}[\text{SW}^{\boldsymbol{\theta}}(\mathbf{a})]} ,$$

where by  $F = F_1 \times \dots \times F_n$  we denote the product distribution of the players' type distributions and by  $\text{EQ}(F)$  the set of all equilibria, which are probability distributions over action profiles.

We assume that players always have the possibility of not participating, hence any rational outcome has non-negative utility in expectation over the non-available information and the randomness of other players and the mechanism.


### 3 Modeling Risk Aversion

When modeling risk aversion, one wants to capture the fact that a random payoff (lottery)  $X$  is less preferred than a deterministic one of value  $\mathbf{E}[X]$ . The standard approach is, therefore, to apply a concave non-decreasing function  $h: \mathbb{R} \rightarrow \mathbb{R}$  to  $X$  and consider  $h(X)$  instead. By Jensen's inequality, we now know  $\mathbf{E}[h(X)] \leq h(\mathbf{E}[X])$ .

In the case of mechanism design, the utility of a risk-neutral agent is defined as the quasilinear utility  $v_i(\mathbf{x}) - p_i$ . That is, if an agent has a value of 1 for an item and has to pay 0.9 for it, then the resulting utility is 0.1. The expected utility is identical if the agent only gets the item with probability 0.1 for free. To capture the effect that the agent prefers the certain outcome to the uncertain one, we again apply a concave function  $h_i: \mathbb{R} \rightarrow \mathbb{R}$  to the quasilinear term  $v_i(\mathbf{x}) - p_i$ . We then consider utility functions  $u_i(\mathbf{x}, p_i) = h_i(v_i(\mathbf{x}) - p_i)$  in the setting described in Section 2. Note that the mechanisms we consider do not know the  $h_i$ -functions. They work as if all utility functions were quasilinear.

We want to compare outcomes based on their social welfare. We use the definition of social welfare being the sum of utilities of all involved parties including the auctioneer. That is,  $\text{SW}(\mathbf{x}, \mathbf{p}) = \sum_{i \in N} u_i(\mathbf{x}, p_i) + \sum_{i \in N} p_i$ . It is impossible for any mechanism to choose good outcomes for this benchmark if the  $h_i$ -function are arbitrary and unknown. Therefore, we assume that utility functions are normalized so that the utility matches the quasilinear one for  $p_i = 0$  and  $p_i = v_i(\mathbf{x})$  (see Figure 1). In more detail, we assume the following *normalized risk-averse utilities*:

1.  $u_i^{v_i}(\mathbf{x}, p_i) \geq u_i^{v_i}(\mathbf{x}, p'_i)$  if  $p_i \leq p'_i$  (monotonicity)
2.  $u_i^{v_i}(\mathbf{x}, p_i) = 0$  if  $p_i = v_i(\mathbf{x})$  (normalization at  $p_i = v_i(\mathbf{x})$ )



■ **Figure 1** Normalized risk-averse utility function (bold) and quasilinear utility function for a fixed allocation  $\mathbf{x}$  and varying payment  $p_i$ .

3.  $u_i^{v_i}(\mathbf{x}, p_i) = v_i(\mathbf{x})$  if  $p_i = 0$  (normalization at  $p_i = 0$ )
4.  $u_i^{v_i}(\mathbf{x}, p_i) \geq v_i(\mathbf{x}) - p_i$  if  $0 \leq p_i \leq v_i(\mathbf{x})$  and  $u_i^{v_i}(\mathbf{x}, p_i) \leq v_i(\mathbf{x}) - p_i$  otherwise (relaxed concavity)

Assumption 4 is a relaxed version of concavity that suffices our needs for the positive results. Our negative results always fulfill concavity.

As an effect of normalization, the optimal social welfare of the risk-averse setting can be bounded in terms of the optimal sum of values, which coincides with the social welfare for quasilinear utilities.

► **Lemma 1.** *Given valuation functions  $(v_i)_{i \in N}$  and normalized risk-averse utilities  $(u_i^{v_i})_{i \in N}$ , let  $OPT$  denote the optimal social welfare with respect to utilities  $(u_i^{v_i})_{i \in N}$  and  $\widehat{OPT}$  denote the optimal social welfare with respect to quasilinear utilities  $(\hat{u}_i^{v_i})_{i \in N}$ . Then,  $OPT \leq 2\widehat{OPT}$ .*

**Proof.** Let  $\mathbf{x}, \mathbf{p}$  denote the outcome and payment profile that maximizes the social welfare  $\sum_{i \in N} u_i^{v_i}(\mathbf{x}, p_i) + \sum_{i \in N} p_i$ . Consider a fixed player  $i$ . If  $0 \leq p_i \leq v_i(\mathbf{x})$ , then by monotonicity of  $u_i^{v_i}(\mathbf{x}, \cdot)$  and Assumption 3,  $u_i^{v_i}(\mathbf{x}, p_i) + p_i \leq u_i^{v_i}(\mathbf{x}, 0) + p_i \leq 2v_i(\mathbf{x})$ . If  $p_i > v_i(\mathbf{x})$ , then we know from Assumption 4 that  $u_i^{v_i}(\mathbf{x}, p_i) + p_i \leq v_i(\mathbf{x})$ . So, always,  $u_i^{v_i}(\mathbf{x}, p_i) + p_i \leq 2v_i(\mathbf{x})$ .

By taking the sum over all players, we get  $OPT = \sum_{i \in N} u_i^{v_i}(\mathbf{x}, p_i) + \sum_{i \in N} p_i \leq \sum_{i \in N} 2v_i(\mathbf{x}) \leq 2\widehat{OPT}$ . ◀

As a consequence, the optimal social welfare changes only within a factor of 2 by risk aversion and we may as well take  $\widehat{OPT}$  as our point of comparison. A VCG mechanism, for example, is still incentive compatible under risk-averse utilities but optimizes the wrong objective function. Lemma 1 shows that it is still a constant-factor approximation to optimal social welfare. However, in simple mechanisms, the agents' strategic behavior may or may not change drastically under risk aversion, depending on the mechanism. This way, equilibria and outcomes can possibly be very different.

## 4 Smoothness Beyond Quasilinear Utilities

Most of our positive results rely on the *smoothness* framework. It was introduced by [24] for general games. There are multiple adaptations to the quasilinear mechanism-design setting. We will use the one by [26]. As our utility functions will not be quasilinear, in this section we first observe that the framework can be extended to general utility functions. Note that throughout this section, the exact definition of  $OPT(\boldsymbol{\theta})$  is irrelevant. Therefore, it can be set to the optimal social welfare but also to weaker benchmarks depending on the setting.

► **Definition 2 (Smooth Mechanism).** A mechanism  $M$  is  $(\lambda, \mu)$ -smooth with respect to utility functions  $(u_i^{\theta_i})_{\theta_i \in \Theta_i, i \in N}$  for  $\lambda, \mu \geq 0$ , if for any type profile  $\boldsymbol{\theta} \in \times_i \Theta_i$  and for any

action profile  $\mathbf{a}$  there exists a randomized action  $a_i^*(\boldsymbol{\theta}, a_i)$  for each player  $i$ , such that  $\sum_i u_i^{\theta_i}(a_i^*(\boldsymbol{\theta}, a_i), \mathbf{a}_{-i}) \geq \lambda \text{OPT}(\boldsymbol{\theta}) - \mu \sum_i p_i(\mathbf{a})$ . We denote by  $u_i^{\theta_i}(\mathbf{a})$  the expected utility of a player if  $\mathbf{a}$  is a vector of randomized strategies.

Mechanism smoothness implies bounds on the price of anarchy. The following theorem and its proof are analogous to the theorems in [26], the proof is therefore deferred to the full version [12]. In cases where the deviation required by smoothness does not depend on  $a_i$ , the results extend to coarse correlated equilibria. The important point is that the respective bounds mostly do not depend on the assumption of quasilinearity.

► **Theorem 3.** *If a mechanism  $M$  is  $(\lambda, \mu)$ -smooth w.r.t. utility functions  $(u_i^{\theta_i})_{\theta_i \in \Theta_i, i \in N}$ , then any Correlated Equilibrium in the full information setting and any Bayes-Nash Equilibrium in the Bayesian setting achieves efficiency of at least  $\frac{\lambda}{\max\{1, \mu\}}$  of  $\text{OPT}(\boldsymbol{\theta})$  or of  $\mathbb{E}_{\boldsymbol{\theta}}[\text{OPT}(\boldsymbol{\theta})]$ , respectively.*

In the standard single-item setting, one item is auctioned among  $n$  players, with their valuations and actions (bids) both being real numbers. In the common auction formats, the item is given to the bidder with the highest bid.

In a *first-price auction*, the winner has to pay her bid; the other players do not pay anything. It is  $(1 - 1/e, 1)$ -smooth w.r.t. quasilinear utility functions. In an *all-pay auction*, all players have to pay their bid. It is  $(1/2, 1)$ -smooth w.r.t. quasilinear utility functions. These smoothness results were given by [26]. They also show that simultaneous and sequential compositions of smooth mechanisms are again smooth.

What is remarkable here is that first-price and all-pay auctions achieve nearly the same welfare guarantees. We will show that in the risk-averse setting this is not true. While the first-price auction almost preserves its constant price of anarchy, the all-pay auction has an unbounded price of anarchy, even with only three players.

## 5 Quasilinear Smoothness Often Implies Risk-Averse Smoothness (Main Result 1)

Our main positive result is that many price-of-anarchy guarantees that are proved via smoothness in the quasilinear setting transfer to the risk-averse one. First, we consider mechanisms that are  $(\lambda, \mu)$ -smooth with respect to quasilinear utility functions. We show that if the deviation strategy  $\mathbf{a}^*$  that is used to establish smoothness ensures non-negative utility, then the price-of-anarchy bound extends to risk-averse settings at a multiplicative constant loss.

► **Theorem 4.** *If a mechanism is  $(\lambda, \mu)$ -smooth w.r.t. quasilinear utility functions  $(\hat{u}_i^{v_i})_{i \in N, v_i \in V_i}$  and the actions in the support of the smoothness deviations satisfy  $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i$ , then any Correlated Equilibrium in the full information setting and any Bayes-Nash Equilibrium in the Bayesian setting achieves efficiency at least  $\frac{\lambda}{2 \cdot \max\{1, \mu\}}$  of the expected optimal even in the presence of risk averse bidders.*

Using Theorem 3, it suffices to prove the following lemma.

► **Lemma 5.** *If a mechanism is  $(\lambda, \mu)$ -smooth w.r.t. quasilinear utility functions  $(\hat{u}_i^{v_i})_{i \in N, v_i \in V_i}$  and the actions in the support of the smoothness deviations satisfy*

$$\hat{u}_i(a_i^*, \mathbf{a}_{-i}) \geq 0, \forall \mathbf{a}_{-i}, \forall i, \quad (1)$$

*then the mechanism is  $(\lambda/2, \mu)$ -smooth with respect to any normalized risk-averse utility functions  $(u_i^{v_i})_{i \in N, v_i \in V_i}$ .*

**Proof.** We start from an arbitrary action profile  $\mathbf{a}$  and want to satisfy Definition 2. Since there exist smoothness deviations s.t.  $\hat{u}_i(a_i^*, \mathbf{a}_{-i}) = v_i(\mathbf{x}(a_i^*, \mathbf{a}_{-i})) - p_i \geq 0, \forall \mathbf{a}_{-i}, \forall i$ , we know from property 4 of the risk aversion definition that  $u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \geq \hat{u}_i^{v_i}(a_i^*, \mathbf{a}_{-i})$ . Therefore,

$$\sum_i u_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \geq \sum_i \hat{u}_i^{v_i}(a_i^*, \mathbf{a}_{-i}) \geq \lambda \widehat{\text{OPT}} - \mu \sum_i p_i(\mathbf{a}) \geq \frac{\lambda}{2} \text{OPT} - \mu \sum_i p_i(\mathbf{a}),$$

where the last inequality follows from Lemma 1.  $\blacktriangleleft$

Note that in order for (1) to hold, it is sufficient if all undominated strategies guarantee non-negative quasilinear utility. For example, in a first-price auction, the only undominated bids are the ones from 0 to  $v_i$ . Regardless of the other players' bids, these can never result in negative utilities.

► **Corollary 6.** *Under normalized risk-averse utilities, the first-price auction has a constant price of anarchy for correlated and Bayes-Nash equilibria.*

We in addition note here that the first part of Property 4 of the normalization assumption,  $u_i^{v_i}(\mathbf{x}, p_i) \geq v_i(\mathbf{x}) - p_i, 0 \leq p_i \leq v_i(\mathbf{x})$ , is not crucial for obtaining a result similar to Theorem 4. Indeed, a relaxation of the form  $u_i^{v_i}(\mathbf{x}, p_i) \geq C \cdot (v_i(\mathbf{x}) - p_i), 0 \leq p_i \leq v_i(\mathbf{x})$  for  $0 < C < 1$ ,  $C$  constant, would incur a loss of at most a factor of  $C$  in the efficiency bound of Theorem 4. More details can be found in the full version [12].

For second-price auctions and their generalizations, for example, the just stated theorems do not suffice to prove guarantees on the quality of equilibria. One in addition needs a no-overbidding assumption and this is further taken care of in the framework of *weak smoothness*, also introduced in [26]. We defer all definitions and results that deal with weak smoothness, including the extension from quasilinear to general utility functions and risk-averse utilities yielding a constant loss as compared to the quasilinear case, to the full version [12].

We also consider the setting where players have hard *budget constraints*. Note that in this case the players' preferences are not quasilinear already in the risk neutral case. Informally, we show that if a mechanism is  $(\lambda, \mu)$ -smooth w.r.t. quasilinear utility functions, then the loss of performance in the budgeted setting is bounded by a constant, even in the presence of risk-averse agents. All details can be found in the full version [12].

## 6 Unbounded Price of Anarchy for All-Pay Auctions (Main Result 2)

From the previous section, we infer that the constant price-of-anarchy bounds for first-price and second-price auctions immediately extend to the risk-averse setting. This is not true for all-pay auctions; by definition there is no non-trivial bid that always ensures non-negative utility. Indeed, as we show in this section, the price of anarchy is unbounded in the presence of risk-averse players. Missing calculation details can be found in the full version [12].

► **Theorem 7.** *In an all-pay auction with risk-averse players, the PoA is unbounded.*

The general idea is to construct a Bayes-Nash equilibrium with two players that very rarely have high values and only then bid high values. We then add a third player who always has a high value. However, as the first two players bid high values occasionally, there is no possible bid that ensures he will surely win. This means, any bid has a small probability of not getting the item but having to pay. Risk-averse players are more inclined to avoid this kind of lotteries. In particular, making our third player risk-averse enough, he prefers the sure zero utility of not participating to any way of bidding that always comes with a small probability of negative utility.



**Proof of Theorem 7.** We consider two (mildly) risk-averse players who both have the same valuation distributions and a third (very) risk-averse player with a constant value. For a large number  $M > 5$ , the first two players have values  $v_1$  and  $v_2$  drawn independently from distributions with density functions of value  $2 \cdot (1 - (M - 1) \cdot \varepsilon)$  on the interval  $[1/2, 1)$  and value  $\varepsilon$  on the interval  $[1, M]$ , where  $\varepsilon = 1/M^2$ . The third player always has value  $1/3 \cdot \ln(M/2)$  for winning.

We will construct a symmetric pure Bayes-Nash equilibrium involving only the first two players. It will be designed such that for the third player it is a best response to always bid 0, i.e., to opt out of the mechanism and never win the item. So, the combination of these strategies will be a pure Bayes-Nash equilibrium for all three players.

Note that the social welfare of any equilibrium of this form is upper-bounded by the optimal social welfare that can be achieved by the first two bidders. By Lemma 1, it is bounded by

$$\mathbb{E}[\text{SW}] \leq 2 \cdot \mathbb{E}[\max\{v_1, v_2\}] \leq 2 \cdot \mathbb{E}[v_1 + v_2] = 2 \cdot (\mathbb{E}[v_1] + \mathbb{E}[v_2]) = 4 \cdot \mathbb{E}[v_1] \leq 4 .$$

Furthermore, the third player’s value  $v_3 = 1/3 \cdot \ln(M/2)$  is a lower bound to the optimal social welfare in the construction containing all three players. So, as pointwise  $\text{OPT}(v) \geq 1/3 \cdot \ln(M/2)$ , where  $v = (v_1, v_2, v_3) \in \mathcal{V}$  denotes the valuation profile, this implies that the price of anarchy can be arbitrarily high.

We define the utility functions by setting

$$u_i^{v_i}(b_i) = \begin{cases} \frac{h_i(v_i - b_i)}{h_i(v_i)} \cdot v_i & \text{if } b_i \text{ is the winning bid} \\ \frac{h_i(-b_i)}{h_i(v_i)} \cdot v_i & \text{otherwise} \end{cases} \quad (2)$$

For the first two players, we use  $h_i(x) := 1 - e^{-x}$ ,  $i \in \{1, 2\}$ , so in particular increasing and concave. For the third player, we set  $h_i(x) = x$  for  $x \geq 0$  and  $h_i(x) = C \cdot x$  for  $x < 0$ , where  $C = (16 \cdot \frac{1}{3} \cdot \ln M/2) \cdot M^2 \geq 1$ . Again this function is increasing and concave<sup>2</sup>. Note that the utility functions also satisfy normalizations at  $p_i = v_i(\mathbf{x})$  and at  $p_i = 0$ . We see that in this example risk aversion has the effect of heavily penalizing payments without winning the auction.

► **Claim 8.** *With the third player not participating, it is a symmetric pure Bayes-Nash equilibrium for the first two players to play according to bidding function  $\beta: \mathcal{V}_i \rightarrow \mathbb{R}_+$ ,  $i \in \{1, 2\}$ , such that*

$$\beta(x) = \int_{\frac{1}{2}}^x \frac{f(t)(e^t - 1)}{F(t) + (1 - F(t))e^t} dt , \quad (3)$$

where  $F$  denotes the cumulative distribution function of the value and  $f$  denotes its density.

**Proof.** We will argue that playing according to  $\beta$  is always the unique best response if the other player is playing according to  $\beta$ , too. Due to symmetry reasons, it is enough to argue about the first player.

Let us fix player 1’s value  $v_1 = x$  and consider the function  $g: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$  that is defined by  $g(y) = \mathbb{E}[u_1^x(b_1 = y, b_2 = \beta(v_2), b_3 = 0)]$ . We claim that  $g$  is indeed maximized at  $y = \beta(x)$ .

<sup>2</sup> Its slope is not an absolute constant. This is indeed necessary because the price of anarchy can be bounded in terms of the slopes of the  $h_i$ -functions as we show in the full version [12].

We have<sup>3</sup>

$$\begin{aligned} g(y) &= \Pr[\beta(v_2) \leq y] \cdot \frac{h_1(x-y)}{h_1(x)} \cdot x + (1 - \Pr[\beta(v_2) \leq y]) \cdot \frac{h_1(-y)}{h_1(x)} \cdot x \\ &= \frac{x}{h_1(x)} \left[ F(\beta^{-1}(y)) (h_1(x-y) - h_1(-y)) + h_1(-y) \right] \\ &= xe^y F(\beta^{-1}(y)) + \frac{x(1-e^y)}{1-e^{-x}} . \end{aligned}$$

The first derivative of this function is given by

$$g'(y) = xe^y F(\beta^{-1}(y)) + xe^y \frac{d}{dy} F(\beta^{-1}(y)) - \frac{x}{1-e^{-x}} e^y .$$

The inverse function theorem implies  $\frac{d}{dy} F(\beta^{-1}(y)) = \frac{f(\beta^{-1}(y))}{\beta'(\beta^{-1}(y))}$ . Furthermore, as  $\beta'(t) = \frac{f(t)(e^t-1)}{F(t)+(1-F(t))e^t}$ , we get for all  $t$  that  $\frac{f(t)}{\beta'(t)} = \frac{F(t)+(1-F(t))e^t}{e^t-1} = (1-F(t)) + \frac{1}{e^t-1}$ . This simplifies  $g'(y)$  to

$$g'(y) = xe^y + \frac{xe^y}{e^{\beta^{-1}(y)}-1} - \frac{xe^y}{1-e^{-x}} = \frac{xe^y}{(1-e^{-x})(e^{\beta^{-1}(y)}-1)} \left( 1 - e^{-x+\beta^{-1}(y)} \right) .$$

Notice that the factor  $\frac{xe^y}{(1-e^{-x})(e^{\beta^{-1}(y)}-1)}$  is always positive. Therefore, we observe that  $g'(y) = 0$  if and only if  $e^{-x+\beta^{-1}(y)} = 1$ , which is equivalent to  $y = \beta(x)$ . Furthermore,  $g'(y) > 0$  for  $y < \beta(x)$  and  $g'(y) < 0$  for  $y > \beta(x)$ . This means that  $y = \beta(x)$  has to be the (unique) global maximum of  $g(y)$ . ◀

► **Claim 9.** *If the first two players are bidding according to (3), then it is a best response for the third player to always bid 0.*

**Proof.** We now show that the very risk-averse third player with valuation  $1/3 \cdot \ln(M/2)$  will indeed bid 0 because every bid  $b'_3 > 0$  will cause negative expected utility.

We distinguish two cases. For values of  $b'_3 > \frac{1}{16}$ , we use that with a small probability one of the two remaining players has a valuation of at least  $M-1$ , which leads to negative utility. For  $b'_3 \leq \frac{1}{16}$  on the other hand, he loses so often that his expected utility is again negative.

Let us first assume that the third player bids  $b'_3$  with  $\frac{1}{16} < b'_3 \leq v_3$ . In this case, with probability more than  $\varepsilon$  one of the first two players has value of at least  $M-1$ . The bid of this player with  $v_i \geq M-1$  can be estimated as follows

$$\begin{aligned} \beta(v_i) &\geq \beta(M-1) \geq \int_{M/2}^{M-1} \frac{f(t)(e^t-1)}{1+(1-F(t))e^t} dt = \int_{M/2}^{M-1} \frac{\varepsilon(e^t-1)}{1+\varepsilon(M-t)e^t} dt \\ &\geq \frac{1}{2}(1-e^{-\frac{M}{2}}) \int_{M/2}^{M-1} \frac{\varepsilon e^t}{\varepsilon(M-t)e^t} dt = \frac{1}{2}(1-e^{-\frac{M}{2}}) \ln(M/2) > \frac{1}{3} \ln(M/2) , \end{aligned}$$

which means that by bidding  $b'_3$  the third player loses with probability of at least  $\varepsilon = 1/M^2$ . For the expected utility, this implies

$$\mathbb{E}[u_3(b'_3, \mathbf{b}_{-3})] \leq (1-\varepsilon)(v_3 - b'_3) + \varepsilon(-C \cdot b'_3) < \frac{1}{3} \ln \frac{M}{2} - \frac{1}{16} \cdot 16 \cdot \frac{1}{3} \ln \frac{M}{2} = 0 .$$

<sup>3</sup> Note that the first step assumes tie breaking in favor of player 1. This is irrelevant for the future steps as the involved probability distributions are continuous.

In the case where the third player bids  $b'_3$ ,  $0 < b'_3 \leq \frac{1}{16}$ , we need to be a bit more careful with estimating the winning probability. We first give a lower bound on the bidding function of the first player for  $v_1 < 1$

$$\beta(v_1) \geq \int_{1/2}^{v_1} \frac{2(1 - \frac{M-1}{M^2})(e^t - 1)}{2(t - \frac{1}{2})(1 - \frac{M-1}{M^2}) + (1 - 2(t - \frac{1}{2})(1 - \frac{M-1}{M^2})) \cdot e^t} dt > \frac{1}{4} \left( v_1 - \frac{1}{2} \right) .$$

Since for  $v_1 \geq 1$ ,  $\beta(v_1) > \frac{1}{16}$  with probability 1, this implies that with  $b'_3$ , the third player has a winning probability of at most

$$\Pr[\beta(v_1) \leq b'_3] \leq \Pr \left[ \frac{1}{4} \left( v_1 - \frac{1}{2} \right) \leq b'_3 \right] = \Pr \left[ v_1 \leq 4b'_3 + \frac{1}{2} \right] < 2 \cdot 4b'_3 .$$

Now, having in mind that  $C = (16v_3) \cdot M^2 \geq 32 \cdot v_3$ , the utility can be estimated as follows

$$\begin{aligned} \mathbb{E}[u_3(b'_3, \mathbf{b}_{-3})] &\leq \Pr[\beta(v_1) \leq b'_3] \cdot (v_3 - b'_3) - \Pr[\beta(v_1) > b'_3] \cdot 32 \cdot v_3 \cdot b'_3 \\ &< 8b'_3 \left( -v_3 - \frac{1}{16} \right) < 0 . \end{aligned}$$

So also in this case, the expected utility is negative. ◀

Combining the two claims, we have constructed a class of distributions and Bayes-Nash equilibria with unbounded price of anarchy. ◀

As a final remark, we note that the first two bidders occasionally bid high only due to risk aversion. In a symmetric Bayes-Nash equilibrium of the all-pay auction in the quasilinear setting, all bids are always bounded by the expected value of a player (see full version [12]). Therefore, such an equilibrium would not work as a point of departure.

## 7 Variance-Aversion Model

In this section, we consider a different model that tries to capture the effect that agents prefer certain outcomes to uncertain ones. It is inspired by similar models in game theory and penalizes variance of random variables. Rather than reflecting the aversion in the utility functions, it is modeled by adapting the solution concept.

In the usual definition of equilibria involving randomization, the utility of a randomized strategy profile is set to be the expectation over the pure strategies. The definition we consider here is modified by subtracting the respective standard deviation. For a player  $i$ , the utility of a randomized strategy profile  $\mathbf{a}$  is given as  $u_i^{v_i}(\mathbf{a}) = \mathbb{E}_{\mathbf{b} \sim \mathbf{a}}[\hat{u}_i^{v_i}(\mathbf{b})] - \gamma \sqrt{\text{Var}[\hat{u}_i^{v_i}(\mathbf{b})]}$ , so a player's utility for an action profile is his expected quasilinear utility for this profile minus the standard deviation multiplied by a parameter  $\gamma$  that determines the degree of variance-averseness,  $0 \leq \gamma \leq 1$ . As already mentioned,  $\hat{u}_i(\mathbf{a})$  denotes the quasilinear utility of player  $i$  for the action profile  $\mathbf{a}$ .

Bayes-Nash Equilibria and correlated equilibria can be defined the same way as before, always replacing expectations by the difference of expectation and standard deviation. The formal definition for  $s(\mathbf{v})$  being a Bayes-Nash equilibrium in this setting is that  $\forall i \in N$ ,  $\forall v_i \in \Theta_i$ ,  $a_i \in \mathcal{A}_i$ ,

$$\begin{aligned} \mathbb{E}_{\mathbf{v}_{-i}}[\hat{u}_i^{v_i}(s(\mathbf{v})) \mid v_i] - \gamma \sqrt{\text{Var}[\hat{u}_i^{v_i}(s(\mathbf{v})) \mid v_i]} \\ \geq \mathbb{E}_{\mathbf{v}_{-i}}[\hat{u}_i^{v_i}(a_i, s_{-i}(\mathbf{v}_{-i})) \mid v_i] - \gamma \sqrt{\text{Var}[\hat{u}_i^{v_i}(a_i, s_{-i}(\mathbf{v}_{-i})) \mid v_i]} . \end{aligned}$$

Note that we again evaluate social welfare as agents perceive it. That is, for a randomized strategy profile  $\mathbf{a}$ , we set  $SW^{\mathbf{v}}(\mathbf{a}) = \sum_i u_i^{v_i}(\mathbf{a}) + \sum_i p_i(\mathbf{a})$ .

Our first result shows that first-price and notably also all-pay auctions have a constant price of anarchy in this setting. Note that even though the proof looks a lot like smoothness proofs, it is not possible to phrase it within the smoothness framework, since here we are dealing with a different solution concept.

► **Theorem 10.** *Bayes-Nash Equilibria and Correlated Equilibria of the first-price and all-pay auction have a constant price of anarchy in this model.*

**Proof.** For simplicity, we will show the claim only for Bayes-Nash equilibria. The proof for correlated equilibria works the same way with minor modifications to the notation.

Assume  $\mathbf{b}$  is a Bayes-Nash equilibrium. We claim that  $\mathbb{E}_{\mathbf{v}} [SW^{\mathbf{v}}(\mathbf{b})] \geq \frac{1}{16} \cdot \mathbb{E}_{\mathbf{v}}[\text{OPT}]$ , where OPT denotes the value of social welfare in the allocation that maximizes it, i.e. maximized sum of utility and payments of the agents.

Consider a fixed player  $j$  and a fixed valuation  $v_j$ . Let  $q = \Pr[\max_{i \neq j} b_i \leq \frac{1}{4} \cdot v_j]$  denote the probability that no other player's bid exceeds  $\frac{1}{4} \cdot v_j$ .

Assume first that  $q \leq \frac{3}{4}$ . Then, because the total social welfare lower bounded by the payments  $\mathbb{E}_{\mathbf{v}_{-j}|v_j} [SW^{\mathbf{v}}(\mathbf{b})] \geq (1 - q) \frac{1}{4} v_j \geq \frac{1}{16} v_j$ .

On the other hand, if  $q \geq \frac{3}{4}$ , since  $\mathbb{E}_{\mathbf{v}_{-j}|v_j} [SW^{\mathbf{v}}(\mathbf{b})] \geq \mathbb{E}_{\mathbf{v}_{-j}|v_j} [u_j^{v_j}(\frac{v_j}{4}, b_{-j})]$ ,

$$\mathbb{E}_{\mathbf{v}_{-j}|v_j} [SW^{\mathbf{v}}(\mathbf{b})] \geq v_j q - \frac{1}{4} v_j - \gamma v_j \sqrt{q(1-q)} \geq \left( \frac{2 - \gamma\sqrt{3}}{4} \right) v_j \geq \frac{1}{16} v_j ,$$

where the first inequality is in fact an equality for the all-pay auction.

From here, by taking the expectation over  $v_j$  and by weighing the right hand side by the probability that OPT takes a particular agent, the theorem follows. ◀

This is contrasted by a correlated equilibrium with 0 social welfare in a setting with positive values. Indeed, for the special case of  $\lambda = 1$ , we see that the variance-averse model further differs from the risk-averse model described in previous sections.

► **Observation 1.** *The PoA for CE of second price auctions is unbounded if  $\gamma = 1$ .*

The proof can be found in the full version [12]. This is not only a difference between smoothness and weak smoothness. Our final result is a mechanism that is  $(\lambda, \mu)$ -smooth for constant  $\lambda$  and  $\mu$  but has unbounded price of anarchy.

► **Theorem 11.** *For any constant  $\gamma > 0$  there is a mechanism that is  $(\lambda, \mu)$ -smooth with respect to quasilinear utility functions for constant  $\lambda$  and  $\mu$  but has unbounded price of anarchy in the variance-aversion model.*

**Proof.** Consider a setting with two items and two players, who have unit-demand valuation functions such that  $\frac{1}{c} v_{i,1} \leq v_{i,2} \leq c v_{i,1}$  for constant  $c \geq 1$ . The players' possible actions are to either report one of the two items as preferred or to opt out entirely. Our mechanism first assigns player 1 her (claimed) favorite item, then assigns player 2 the remaining one unless she opts out. There are no payments. Obviously, this mechanism is  $(\frac{1}{c}, 0)$ -smooth because the allocation is within a  $\frac{1}{c}$ -factor of the optimal allocation by construction.

We will now construct a mixed Nash equilibrium of bad welfare. To this end, let  $v_{1,1} = v_{1,2} = \varepsilon$  for some small  $\varepsilon > 0$ . This makes player 1 indifferent between items 1 and 2. In particular, it is a best response to ask for item 1 with probability  $\frac{q-1}{q}$  and for item 2 with

probability  $\frac{1}{q}$ . We note at this point that in a Bayes-Nash equilibrium we could make this respective action the unique best response by having random types.

For player 2, we set  $v_{2,1} = c$ ,  $v_{2,2} = 1$ . She has the choice of participating or opting out. Opting out implies utility 0, whereas participating implies

$$u_2(\mathbf{a}) = \frac{c+q-1}{q} - \gamma \sqrt{\frac{(c-1)^2(q-1)}{q^2}} = \frac{(c-1)(1-\gamma\sqrt{q-1})}{q} + 1$$

Now, if we set  $q = c - 1$ , then  $u_2(\mathbf{a}) = 2 - \gamma\sqrt{c-2}$  which is negative for  $c > \frac{4}{\gamma^2} + 2$ . We further set  $c = \frac{4}{\gamma^2} + 3$ . That is, player 2 prefers to opt out. This outcome has social welfare  $\varepsilon$  whereas the optimal social welfare is  $c$ . ◀

Note that this last example shows that variance-averseness yields very strange preferences for lotteries. In our example, the variance-averse player prefers not to participate although any outcome in the (free) lottery has positive value.

---

## References

- 1 Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Commentaries of the Imperial Academy of Science of Saint Petersburg*, 1738. Republished *Econometrica* 22(1), 23–36 (1954).
- 2 Anand Bhalgat, Tanmoy Chakraborty, and Sanjeev Khanna. Mechanism design for a risk averse seller. In *Proc. 8th Intl. Workshop Internet & Network Economics (WINE)*, pages 198–211, 2012.
- 3 Shaddin Dughmi and Yuval Peres. Mechanisms for risk averse agents, without loss. *CoRR*, abs/1206.2957, 2012.
- 4 Paul Dütting, Thomas Kesselheim, and Éva Tardos. Mechanism with unique learnable equilibria. In *Proc. 15th Conf. Econom. Comput. (EC)*, pages 877–894, 2014.
- 5 Amos Fiat and Christos H. Papadimitriou. When the players are not expectation maximizers. In *Proc. 3rd Intl. Symp. Algorithmic Game Theory (SAGT)*, pages 1–14, 2010.
- 6 Gadi Fibich, Arieh Gavious, and Aner Sela. All-pay auctions with risk-averse players. *Int. J. Game Theory*, 34(4):583–599, 2006.
- 7 Hu Fu, Jason Hartline, and Darrell Hoy. Prior-independent auctions for risk-averse agents. In *Proc. 14th Conf. Electr. Commerce (EC)*, pages 471–488. ACM, 2013.
- 8 Martin Hoefer, Thomas Kesselheim, and Berthold Vöcking. Truthfulness and stochastic dominance with monetary transfers. *ACM Trans. Econom. Comput.*, 4(2):11:1–11:18, 2016.
- 9 Darrell Hoy. The concavity of atomic splittable congestion games with non-linear utility functions. In *EC 2012, Workshop on Risk Aversion in Algorithmic Game Theory and Mechanism Design*, 2012.
- 10 Darrell Hoy, Nicole Immorlica, and Brendan Lucier. On-demand or spot? selling the cloud to risk-averse customers. In *Proc. 12th Intl. Conf. Web and Internet Economics (WINE)*, pages 73–86, 2016.
- 11 Audrey Hu, Steven A Matthews, and Liang Zou. Risk aversion and optimal reserve prices in first-and second-price auctions. *J. Econom. Theory*, 145(3):1188–1202, 2010.
- 12 Thomas Kesselheim and Bojana Kodric. Price of anarchy for mechanisms with risk-averse agents. *CoRR*, abs/1804.09468, 2018. URL: <http://arxiv.org/abs/1804.09468>.
- 13 Bettina S Klose and Paul Schweinzer. Auctioning risk: The all-pay auction under mean-variance preferences. Technical report, University of Zurich, 2014. doi:10.5167/uzh-67043.
- 14 Harry M Markowitz. *Portfolio selection: efficient diversification of investments*, volume 16. Yale university press, 1968.

- 15 Harry M. Markowitz. Mean-variance approximations to expected utility. *Europ. J. Oper. Res.*, 234(2):346–355, 2014.
- 16 Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. *Microeconomic theory*, volume 1. Oxford university press New York, 1995.
- 17 Eric Maskin and John Riley. Optimal auctions with risk averse buyers. *Econometrica*, pages 1473–1518, 1984.
- 18 Timothy Mathews and Brett Katzman. The role of varying risk attitudes in an auction with a buyout option. *J. Econom. Theory*, 27(3):597–613, 2006.
- 19 Steven A Matthews. Selling to risk averse buyers with unobservable tastes. *J. Econom. Theory*, 30(2):370–400, 1983.
- 20 Reshef Meir and David C. Parkes. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. *SIGMETRICS Performance Evaluation Review*, 43(3):67–70, 2015.
- 21 Evdokia Nikolova and Nicolás E. Stier Moses. A mean-risk model for the traffic assignment problem with stochastic travel times. *Oper. Res.*, 62(2):366–382, 2014.
- 22 Evdokia Nikolova and Nicolás E. Stier Moses. The burden of risk aversion in mean-risk selfish routing. In *Proc. 16th Conf. Econom. Comput. (EC)*, pages 489–506, 2015.
- 23 Georgios Piliouras, Evdokia Nikolova, and Jeff S. Shamma. Risk sensitivity of price of anarchy under uncertainty. *ACM Trans. Econom. Comput.*, 5(1):5:1–5:27, 2016.
- 24 Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. 41st Symp. Theory of Computing (STOC)*, pages 513–522, 2009.
- 25 Mukund Sundararajan and Qiqi Yan. Robust mechanisms for risk-averse sellers. In *Proc. 11th Conf. Econom. Comput. (EC)*, pages 139–148, 2010.
- 26 Vasilis Syrgkanis and Éva Tardos. Composable and efficient mechanisms. In *Proc. 45th Symp. Theory of Computing (STOC)*, pages 211–220, 2013.
- 27 John von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton, New Jersey: Princeton University Press. XVIII, 625 p., 1944.