Price of Anarchy for First Price Auction with Risk-Averse Bidders

Zhiqiang Zhuang
Tianjin University
Tianjin, China
zhuang@tju.edu.cn

Kewen Wang*
Griffith University
Brisbane, Australia
k.wang@griffith.edu.au

Zhe Wang
Griffith University
Brisbane, Australia
zhe.wang@griffith.edu.au

ABSTRACT

Inquiry into the price of anarchy (POA) for auctions is almost confined within the risk-neutral setting. Nonetheless, empirical and experimental studies suggest that real-world agents are more or less risk-averse rather than strictly risk-neutral. In this paper, we study the POA of first-price single-item auctions (FPA) with risk-averse bidders. For completeness, we consider both risk-averse and risk-neutral sellers. In the former, we establish that the POA is $1/2$ for both the symmetric FPA and FPA in general. In the latter, we show that the POA can be arbitrarily bad for the symmetric FPA and characterise the conditions for the POA to be constant. In response to a fairness issue in the case of risk-neutral sellers, we propose the notion of suboptimal social welfare. We subsequently derive POA bounds with respect to this new notion where the bounds are parameterised by two variables that capture the value range of the utility function.

KEYWORDS

Mechanism Design; Auction Theory; Price of Anarchy

ACM Reference Format:

Zhiqiang Zhuang, Kewen Wang, and Zhe Wang. 2023. Price of Anarchy for First Price Auction with Risk-Averse Bidders. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 7 pages.

1 INTRODUCTION

Practical and theoretical studies of auctions mostly assume *risk-neutral* agents who are indifferent to risk. Accordingly, their utility functions are *quasi-linear*, that is a bidder's utility is the difference between her valuation of the allocated item and her payment [10]. However, it has long been noticed that in reality agents often exhibit *risk-averse* behaviour, with a tendency to prefer a certain payment to an uncertain higher payment [2, 13, 20]. In the standard formulation of risk-averse agents, their utility functions are *concave* rather than quasi-linear. Risk aversion poses tremendous difficulties in the design and analysis of auctions. It remains a mystery how to make the best out of real-world risk attitudes and the risk of ignoring them. We aim to bring some clarity by exploring the *price of anarchy* (POA) for *first-price single-item auctions* (FPA) with risk-averse bidders.

The POA measures the inefficiency of an economic system due to the lack of coordination between self-interested participants [9]. In auctions, the POA is defined as the ratio between the social welfare in the worst equilibrium and the optimal social welfare. Following [18], we take optimal social welfare as the denominator when calculating the ratio. Hence the more inefficient an auction format is, the closer its POA is to zero. The ongoing inquiry into the POA for the risk-neutral setting have established POA bounds for many important auction formats as well as general frameworks for proving such bounds [7, 18, 19].

The derivation of POA does not necessarily involve the exact form of the equilibrium. Essentially, we can obtain an approximation guarantee for the equilibrium performance of an auction without knowing the equilibrium at all. The derivation effort is therefore especially beneficial for auctions with no known equilibrium or with an equilibrium that cannot be expressed explicitly. This is precisely the case for auctions in the risk-averse setting where little is known even for the most simple auction formats. For example, the equilibrium bidding strategy for a FPA is the solution to a differential equation that has no explicit expression [14]. Despite the apparent benefits, there has been almost no attempt to establish POA bounds for the risk-averse setting. We plan to fill this void by investigating the POA of FPA. While most investigations of risk aversion focus on the bidders, the sellers are by no means always neutral to risk. We try to exhaust all possible scenarios, considering both risk-averse and risk-neutral sellers in our investigation.

With a quasi-linear utility function, when calculating social welfare (i.e., the total utility of the bidders and the seller), the payment from the winning bidder to the seller results in a utility loss for the bidder that equals the utility gain for the seller. This means social welfare is simply the winning bidder's valuation (of the auctioned item). The efficiency of an auction, therefore, relies solely on how the auctioned item is allocated. An auction achieves optimal social welfare if it always allocates the item to the highest valued bidder. This is no longer the case in the risk-averse setting. The utility gain and loss, due to the payment transfer, cannot cancel out each other. Thus the amount of payment also matters, only a specific one gives rise to optimal social welfare. Our first technical result is a concrete expression of optimal social welfare in the risk-averse setting as a function of the bidders' valuations for both the cases of risk-averse and risk-neutral sellers. The function specifies the amount of payment for optimality. It also makes clear that optimal allocation coincides with that of the risk-neutral setting.

Although simple to ascertain, the result reveals the distinct feature of auctions in the risk-averse setting that both the allocation and payment play a part in achieving optimal social welfare. When the seller is risk-neutral, the optimality gives rise to an unsettling issue, though. The optimal payment can be negative or larger than the winning bidder's valuation which causes negative utility for the seller and respectively for the bidder. Thus either the bidder or the seller may suffer a loss for overall optimality. In response to the fairness issue, we propose the notion of *suboptimal social welfare*

*Corresponding author

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

which is the maximum achievable social welfare conditional on everyone having a non-negative utility.

Since we can view risk neutrality as a special case of risk aversion for which the extent of risk aversion is zero,[1] we are expecting worse POA for the risk-averse setting than the risk-neutral one. In the former, we have a larger pool of utility functions to pick the worst one. Firstly, with risk-averse bidders and sellers, we establish the POA of $1/2$ for symmetric FPA and FPA in general. Secondly, with risk-averse bidders and risk-neutral sellers, we show that the POA is zero even for the symmetric case. Subsequently, we look into the POA with respect to suboptimal social welfare. Unfortunately, without proper restriction, the POA is still arbitrary close to zero. We provide sufficient conditions for it to be constant in both the symmetric and the general setting. From the sufficient conditions we articulate lower bounds which are parameterised by two variables that capture the range of utility functions.

Our main strategy in establishing the POA and lower bound of it is to partition the space of joint valuation distributions and *Bayesian Nash Equilibriums* (BNEs) by the expected value of the winning bid. More specifically, we partition it into two subspaces one of which contains the valuation distributions and BNEs that induce an expected winning bid that is greater than one-half of the expected maximum valuation and another that contains the rest (i.e., induces an expected winning bid that is less than one-half of the expected maximum valuation). For the former subspace, a POA bound can be established from the winning bid alone. For the latter subspace, we adapt a common deviation strategy for proving POA bounds in the risk-neutral setting. The expected value of the winning bid induced in this subspace makes the adaptation possible.

2 PRELIMINARY

In this paper, we denote a vector as a boldface letter such as $\mathbf{v}, \mathbf{b}$ and a scalar as a lowercase letter possibly with subscripts such as v_i, b_i. Given a vector $\mathbf{b}$, b_i represents the ith component, $\mathbf{b}_{-i}$ the remaining components, and $(b_i', \mathbf{b}_{-i})$ the vector where the ith component is b_i' and the remaining components are $\mathbf{b}_{-i}$.

We adopt the standard independent private value model of an auction. There is a single and indivisible item to be auctioned by a seller to n bidders. Each bidder i has a private value v_i over the item which is drawn independently from a distribution F_i. The distributions $F_1, \ldots, F_n$ are common knowledge among the bidders. We assume the distributions have a bounded support. A valuation profile is a vector of n private values denoted as $\mathbf{v} = (v_1, \ldots, v_n)$ where v_i is the ith bidder's value.

Of the many ways to formulate risk-averse agents, we adopt the most common and least controversial one, that is each risk-averse agent has an identical and twice differentiable utility function $u : \mathbb{R} \to \mathbb{R}$ over wealth that satisfies $u(0) = 0$, $u'(x) > 0$ and $u''(x) < 0$ for all $x \in \mathbb{R}$. In what follows we abbreviate the last two properties as $u' > 0$ and $u'' < 0$. Implicitly in this formulation is the assumption that the agents have an equivalent monetary value for the auctioned item. So if a winning bidder has a valuation of v and makes a payment of b to the seller, then her utility is $u(v - b)$. The extent of risk aversion is determined by the concavity of the utility function u. Common measures include the Arrow-Pratt coefficient of absolute risk aversion defined as $A(x) = -u''(x)/u'(x)$ and the coefficient of relative risk aversion defined as $R(x) = -x \cdot u''(x)/u'(x)$. These measures may be constant, increasing or decreasing with respect to wealth. For example, a utility function is decreasing absolute risk-averse if $A(x)$ decreases as x increases. From here forward, all risk-averse agents have $u(\cdot)$ as their utility function and all bidders are risk-averse.

A bidding profile is a vector of n bids denoted as $\mathbf{b} = (b_1, \ldots, b_n)$. We let $u_i(\mathbf{b}; v_i)$ denotes bidder i's utility when her value is v_i and the bidding profile is $\mathbf{b}$. A strategy profile $\mathbf{s} = (s_1, \ldots, s_n)$ is a vector of functions where each function s_i maps a bidder's valuation to her bid. We use $\mathbf{s}(\mathbf{v})$ to denote the bidding profile resulting from the valuation profile $\mathbf{v}$. A strategy profile is a BNE if and only if reporting $s_i(v_i)$ maximizes i's expected utility for all i, that is

$$\mathbb{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{s}(\mathbf{v}); v_i)] \geq \mathbb{E}_{\mathbf{v}_{-i}}[u_i(b_i', \mathbf{s}_{-i}(\mathbf{v}); v_i)].$$

for every bidder i, every possible valuation v_i, and every possible deviating bid b_i'.

The *social welfare* of an auction is the sum of the bidders' and the seller's utility resulted in the auction. Given the format of the auction, social welfare is determined by the bidding and the valuation profile. In a FPA, the social welfare with respect to a bidding profile $\mathbf{b}$ and a valuation profile $\mathbf{v}$, denoted as $SW(\mathbf{b}; \mathbf{v})$, is such that

$$SW(\mathbf{b}; \mathbf{v}) = u(v - b) + b$$

when the seller is risk neutral, and

$$SW(\mathbf{b}; \mathbf{v}) = u(v - b) + u(b)$$

when the seller is risk-averse, where v is the valuation of the winning bidder and b her bid (i.e., $b = \max(\mathbf{b})$). The optimal social welfare with respect to a valuation profile $\mathbf{v}$, denoted as $OPT(\mathbf{v})$, is the maximum achievable social welfare when the valuation profile is $\mathbf{v}$.

3 OPTIMAL AND SUBOPTIMAL SOCIAL WELFARE

In this section, we derive a concrete expression of the optimal social welfare respectively for when bidders and sellers are risk-averse, and for when bidders are risk-averse and sellers are risk-neutral. We also introduce a notion of suboptimal social welfare that arises naturally when sellers are risk-neutral.

With a quasi-linear utility function, the payment deducted from the winning bidder's utility is compensated by the gain in the seller's utility, hence social welfare is simply the winning bidder's private value and the optimal one the highest private value possessed by the bidders. Consequently, optimal social welfare is solely determined by the allocation of the auctioned item. This is no longer true in the risk-averse setting in which the concavity of the utility function dictates that the amount of payment also plays a role. We refer to the allocation and payment that lead to optimal social welfare as *optimal allocation* and *optimal payment* respectively.

Starting with risk-averse bidders and risk-neutral sellers, an auction achieves optimal social welfare by allocating the auctioned item to the bidder with the highest valuation v and letting her pay b to the seller such that $u'(v - b) = 1$.

[1] Strictly speaking, our formulation of risk aversion does not permit risk neutrality. But it can be arbitrarily close to the risk-neutral setting.

Lemma 1. *In a FPA with risk-averse bidders and a risk-neutral seller,*

$$OPT(\mathbf{v}) = u(\max(\mathbf{v}) - b) + b$$

where $u'(\max(\mathbf{v}) - b) = 1$.

Proof. In a FPA, the social welfare is $u(v-b)+b$ where v and b are respectively the valuation and bid of the winning bidder. Since $\frac{\partial[u(v-b)+b]}{\partial v} = u'(v-b)$ and $u'(x) > 0$ for all $x \in \mathbb{R}$, $u(v-b)+b$ is an increasing function of v. Since $\frac{\partial[u(v-b)+b]}{\partial b} = 1 - u'(v-b)$, $\frac{\partial^2[u(v-b)+b]}{\partial b^2} = u''(v-b)$ and $u''(x) < 0$ for all $x \in \mathbb{R}$, $u(v-b)+b$ is a strictly concave function of b. The concavity implies that, given a v, $u(v-b)+b$ is at its global maximum when $1-u'(v-b) = 0$. Hence, given a valuation profile $\mathbf{v}$, the optimal social welfare is achieved by allocating the auctioned item to the bidder with the highest valuation and charging the bidder b where $u'(\max(\mathbf{v}) - b) = 1$. □

By the above lemma, no matter what $\max(\mathbf{v})$ is, the optimal payment b always adjusts the winning bidder's wealth $\max(\mathbf{v}) - b$ to an amount at which the marginal utility is 1. For the remaining of this paper, we reserve the letter x to denote this amount. That is, we write $OPT(\mathbf{v})$ as $u(x) + \max(\mathbf{v}) - x$ where $u'(x) = 1$. In this form, the optimal payment is $\max(\mathbf{v}) - x$. Clearly, it may be negative (i.e., $x > \max(\mathbf{v})$) or greater than the highest valuation (i.e., $x < 0$) which results in negative utility respectively for the seller and the winning bidder. For the sake of fairness, neither case is desirable. Moreover, since in a BNE, the winning bid is always between 0 and the winning bidder's valuation, for some utility functions (e.g., those with $x < 0$) none of the valuation distributions and BNEs will result in an optimal social welfare. Thus the optimal social welfare obtained in Lemma 1 is not a realistic goal to be aiming at. These motivate the notion of *suboptimal social welfare*, which we define as the highest achievable social welfare while guaranteeing non-negative utility for all agents. Formally this is the highest achievable social welfare with a payment in-between zero and the highest valuation, that is

$$SOPT(\mathbf{v}) = \begin{cases} u(\max(\mathbf{v})) & \text{if } x \geq \max(\mathbf{v}) \\ u(x) + \max(\mathbf{v}) - x & \text{if } 0 < x < \max(\mathbf{v}) \\ \max(\mathbf{v}) & \text{if } x \leq 0 \end{cases}$$

The rationale for $SOPT(\mathbf{v})$ is as follows: if $0 < x < \max(\mathbf{v})$, then the optimal payment already guarantees non-negative utility, thus $SOPT(\mathbf{v}) = OPT(\mathbf{v})$, and if $x \geq \max(\mathbf{v})$ ($x \leq 0$); then the optimal payment is less than 0 (resp. greater than $\max(\mathbf{v})$), thus setting the payment to 0 (resp. $\max(\mathbf{v})$) achieves the highest social welfare with non-negative utility.

Moving on to risk-averse bidders and sellers, we achieve optimal social welfare by allocating the item to the bidder with the highest valuation v and letting her pay $v/2$ to the seller. In other words, we achieve optimal social welfare when every bidder bids half of her value.

Lemma 2. *In a FPA with risk-averse bidders and a risk-averse seller,*

$$OPT(\mathbf{v}) = 2 \cdot u(\max(\mathbf{v})/2).$$

Proof. The social welfare is $u(v-b) + u(b)$ where v and b are respectively the valuation and bid of the winning bidder. Since $\frac{\partial[u(v-b)+u(b)]}{\partial v} = u'(v-b)$ and $u'(x) > 0$ for all $x \in \mathbb{R}$, $u(v-b)+u(b)$ is an increasing function of v. Since $\frac{\partial[u(v-b)+u(b)]}{\partial b} = u'(b) - u'(v-b)$, $\frac{\partial^2[u(v-b)+u(b)]}{\partial b^2} = u''(b)+u''(v-b)$, and $u''(x) < 0$ for all $x \in \mathbb{R}$, $u(v-b)+u(b)$ is a strictly concave function of b. The concavity implies, for any v, $u(v-b)+u(b)$ is at its global maximum when $u'(b) - u'(v-b) = 0$, that is $b = v/2$. Hence, given a valuation profile $\mathbf{v}$, the optimal social welfare is achieved by allocating the auctioned item to the bidder with the highest valuation and charges the bidder $\max(\mathbf{v})/2$. □

Noticeably, the optimal payment in the above lemma is independent of the exact form of the utility function. This contributes to better POA bounds with risk-averse sellers than with risk-neutral ones. One a side note this particular form of optimal social welfare also enforces an absolutely fair share of the overall utility between the wining bidder and the seller, as both parties end up with identical utility.

In the next two sections, we investigate how inefficient FPA is with risk-averse bidders. The inefficiency is measured by the extent to which the POA of the auction format is close to zero:

$$POA = \inf_{\mathcal{F},\mathbf{s},n,u} \frac{\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})]}{\mathbb{E}_{\mathbf{v}}[OPT(\mathbf{v})]}.$$

That is, POA is the minimum ratio of the expected social welfare and the optimal one ranging over all joint valuation distributions $\mathcal{F}$, all BNE $\mathbf{s}$, all choice n for the number of bidders and all forms of the utility function $u(\cdot)$ satisfying $u(0) = 0$, $u' > 0$ and $u'' < 0$. When $\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})]/\mathbb{E}_{\mathbf{v}}[OPT(\mathbf{v})]$ evaluates to $0/0$, we define it to be 1.

4 RISK-AVERSE SELLERS

In this section, we establish the POA of FPA with risk-averse bidders and sellers. As we will see, compared with the case of risk-averse bidders and risk-neutral sellers, the alignment of the seller's risk attitude with the bidders' lead to better POA bounds.

We start with the symmetric setting in which each bidder's valuation is drawn independently from a common distribution. According to [14], in this setting a FPA has a unique symmetric (pure strategy) BNE in which the bidding function is non-decreasing. Therefore the winning bidder in such an auction always has the highest valuation, just as in the outcome that achieves optimal social welfare. Since the allocation is optimal, we can concentrate on the inefficiency caused by non-optimal payment. We establish that the POA for a symmetric FPA is 1/2.

Theorem 3. *The price of anarchy of the symmetric first-price single-item auction format with risk-averse bidders and sellers is 1/2.*

Proof. Let the valuation and the equilibrium bid of the winning bidder be v and b respectively. Then $SW(\mathbf{s}(\mathbf{v});\mathbf{v}) = u(v-b) + u(b)$ and $OPT(\mathbf{v}) = 2 \cdot u(v/2)$. Since $0 \leq b \leq v$, $u(v-b)+u(b)$ is at its minimum when $b = v$ or $b = 0$. In either case, $u(v-b)+u(b) = u(v)$. Thus

$$SW(\mathbf{s}(\mathbf{v});\mathbf{v}) \geq u(v) > u(v/2) = 1/2 \cdot OPT(\mathbf{v}).$$

To show that the lower bound of $1/2$ is also the upper bound, we will construct an instance of symmetric FPA that achieves no more than $1/2$ of the corresponding optimal social welfare. Suppose there is only one bidder and the density function f of the distribution of valuations is such that $f(v) = 1$ for some $v > 0$. Thus in the BNE the bidder bids 0 and wins the auctioned item. This auction outcome results in the social welfare of $u(v)$. Since $u'' < 0$ and $u''(x)$ can be arbitrarily small for $v/2 < x < v$, $u(v)$ can be arbitrarily close to $u(v/2)$, in which case the social welfare is arbitrarily close to $1/2$ of the optimal one.

□

Theorem 3 echos with the dual determining factors in achieving optimal social welfare in the risk-averse setting. Symmetric FPA has an optimal allocation, but its non-optimal payment can take away as much as half of the achievable social welfare.

In the proof of the theorem, we establish the $1/2$ upper bound through the single bidder scenario. Another way is to consider an infinite number of bidders. Due to the unlimited competition, the BNE strategy is for each bidder to bid their true valuation which results in the social welfare of $u(\max(\mathbf{v}))$ and the ratio of $1/2$. Perhaps it is worth mentioning that, when the number of bidders is more than one but finite, the ratio is strictly greater than $1/2$. In these cases, the BNE strategy is to bid between 0 and their valuations which achieves the social welfare of $u(h) + u(l)$ for $\max(\mathbf{v})/2 \leq h < \max(\mathbf{v})$ and $0 < l \leq \max(\mathbf{v})/2$. While $u(h)$ can be arbitrary close to $u(\max(\mathbf{v})/2)$, $u(l)$ is strictly greater than 0, thus the ratio is at least $1/2 + u(l)/u(\max(\mathbf{v})/2)$. In other words, for cases that matter, the efficiency of symmetric FPA is better than what the POA suggests.

Moving on to the general setting in which bidders' valuations are drawn independently from possibly different distributions, we establish that the POA is also $1/2$.

Theorem 4. *The price of anarchy of the first-price single-item auction format with risk-averse bidders and sellers is* $1/2$.

Proof. Any pair of joint valuation distribution $\mathcal{F}$ and BNE $\mathbf{s}$ (for the distribution) induces a fixed expected winning bid $\mathbb{E}_{\mathbf{v}}[\max(\mathbf{s}(\mathbf{v}))]$ and a fixed expected maximum valuation $\mathbb{E}_{\mathbf{v}}[\max(\mathbf{v})]$. Thus we can partition the space of valuation distribution and BNE pairs into a subspace in which the induced expected winning bid and expected maximum valuation are such that $\mathbb{E}_{\mathbf{v}}[\max(\mathbf{s}(\mathbf{v}))] > \mathbb{E}_{\mathbf{v}}[\max(\mathbf{v})/2]$ and another subspace in which the induced expected values are such that $\mathbb{E}_{\mathbf{v}}[\max(\mathbf{s}(\mathbf{v}))] \leq \mathbb{E}_{\mathbf{v}}[\max(\mathbf{v})/2]$. For simplicity, we write $\max(\mathbf{v})$ and $\max(\mathbf{s}(\mathbf{v}))$ as v and b respectively.

For the former subspace, since $\mathbb{E}_{\mathbf{v}}[b] > \mathbb{E}_{\mathbf{v}}[v/2]$ we have

$$\begin{aligned}\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})] &= \mathbb{E}_{\mathbf{v}}[u(v'-b)+u(b)]\\ &\geq \mathbb{E}_{\mathbf{v}}[u(b)]\\ &> \mathbb{E}_{\mathbf{v}}[u(v/2)]\end{aligned}$$

where v' is the expected valuation of the winning bidder. The equality follows from the definition of social welfare in a FPA, where $u(v'-b)$ is the winning bidder's utility and $u(b)$ the seller's. The first inequality follows from the fact that a bidder cannot have negative utility in a BNE, that is $\mathbb{E}_{\mathbf{v}}[u(v'-b)] \geq 0$. Since the expected optimal social welfare is $\mathbb{E}_{\mathbf{v}}[2 \cdot u(v/2)]$, a lower bound for the POA for this subspace of valuation distributions and BNEs is $\mathbb{E}_{\mathbf{v}}[u(v/2)]/\mathbb{E}_{\mathbf{v}}[2 \cdot u(v/2)] = 1/2$.

For the latter subspace in which $\mathbb{E}_{\mathbf{v}}[b] \leq \mathbb{E}_{\mathbf{v}}[v/2]$, we adapt the traditional deviation strategy for proving POA bounds to the risk-averse setting. Let $x_i^*(\mathbf{v})$ be the indicator variable for whether or not bidder i is the one with the highest valuation.

Let $\mathbf{s}$ be a BNE in the subspace. Suppose bidder i bids half of her value instead of $s_i(v_i)$. If she wins, her utility is $u(v_i - v_i/2) = u(v_i/2) \geq u(v_i/2 - b)$. If she loses, it must be that $v_i/2 < b$ and she obtains utility $0 \geq u(v_i/2 - b)$. Thus in either case

$$u_i(v_i/2, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i) \geq u(v_i/2 - b)$$

which implies

$$u_i(v_i/2, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i) \geq u(v_i/2 - b) \cdot x_i^*(\mathbf{v}).$$

Summing the inequality over all bidders, we have

$$\sum_{i=1}^{n} u_i(v_i/2, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i) \geq u(v/2 - b). \tag{1}$$

Since $\mathbf{s}$ is a BNE,

$$\mathbb{E}_{\mathbf{v}_{-i}}[u_i(\mathbf{s}(\mathbf{v}); v_i)] \geq \mathbb{E}_{\mathbf{v}_{-i}}[u_i(v_i/2, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i)].$$

Taking expectation over v_i and summing up the n inequalities of the above form, we have

$$\sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{s}(\mathbf{v}); v_i)] \geq \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[u_i(v_i/2, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i)].$$

Therefore

$$\begin{aligned}\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})] &= \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{s}(\mathbf{v}); v_i)] + \mathbb{E}_{\mathbf{v}}[u(b)]\\ &\geq \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[u_i(v_i/2, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i)] + \mathbb{E}_{\mathbf{v}}[u(b)].\end{aligned}$$

Combining with inequality (1), the above inequality implies

$$\begin{aligned}E_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})] &\geq \mathbb{E}_{\mathbf{v}}[u(v/2 - b) + u(b)]\\ &= u(\mathbb{E}_{\mathbf{v}}[v/2 - b]) + u(\mathbb{E}_{\mathbf{v}}[b])\\ &= u(\mathbb{E}_{\mathbf{v}}[v/2] - \mathbb{E}_{\mathbf{v}}[b]) + u(\mathbb{E}_{\mathbf{v}}[b]).\end{aligned}$$

Since $\mathbb{E}_{\mathbf{v}}[b] \leq \mathbb{E}_{\mathbf{v}}[v/2]$, $u(\mathbb{E}_{\mathbf{v}}[v/2] - \mathbb{E}_{\mathbf{v}}[b]) + u(\mathbb{E}_{\mathbf{v}}[b])$ is at its minimum when $\mathbb{E}_{\mathbf{v}}[b] = 0$ or $\mathbb{E}_{\mathbf{v}}[b] = v/2$. Since, in either case $u(\mathbb{E}_{\mathbf{v}}[v/2] - \mathbb{E}_{\mathbf{v}}[b]) + u(\mathbb{E}_{\mathbf{v}}[b]) = \mathbb{E}_{\mathbf{v}}[u(v/2)]$, we have $u(\mathbb{E}_{\mathbf{v}}[v/2] - \mathbb{E}_{\mathbf{v}}[b]) + u(\mathbb{E}_{\mathbf{v}}[b]) \geq \mathbb{E}_{\mathbf{v}}[u(v/2)]$ which implies

$$\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})] \geq \mathbb{E}_{\mathbf{v}}[u(v/2)] = 1/2 \cdot \mathbb{E}_{\mathbf{v}}[OPT(\mathbf{v})].$$

We have established the lower bound of $1/2$ for both subspaces, thus $1/2$ is a lower bound for the space of all pairs of valuation distributions and BNEs. It is an immediate consequence of Theorem 3 that $1/2$ is also an upper bound.

□

It may look peculiar that the POA for symmetric FPA is the same as that for FPA in general. Only the payment can be non-optimal for the former whereas both the allocation and payment can be non-optimal for the latter. While non-optimal payment alone can take away half of the optimal social welfare, our results indicate that it does not get any worse when the allocation is also non-optimal.

The reason for the peculiar finding is that we defined the POA to be the worst ratio over all possible forms of the utility function $u(\cdot)$. In the symmetric setting, it is easy to see from the proof of Theorem 3 that the POA can be expressed as

$$\mathbb{E}_{\mathbf{v}}(u(v))/(2 \cdot \mathbb{E}_{\mathbf{v}}(u(v/2)))$$

where v is the highest valuation among the bidders. The ratio is at its minimum when the utility function is such that $u(v)$ and $u(v/2)$ are arbitrarily close, which gives rise to the POA of $1/2$. That is we have a better ratio with other utility functions. In the general setting, we established that a lower bound for the POA is

$$\mathbb{E}_{\mathbf{v}}(u(v/2))/(2 \cdot \mathbb{E}_{\mathbf{v}}(u(v/2))) = 1/2$$

regardless of the form of the utility function $u(\cdot)$. So if we defined the POA with respect to a fixed utility function, the symmetric setting would have a better POA.

The POA for FPA with risk-neutral agents is $1 - 1/e^2 \approx 0.8647$ [7] for which the only source of inefficiency is the non-optimal allocation. It is clear that in the risk-averse setting where inefficiency can also be caused by non-optimal payment, FPA has lower efficiency.

5 RISK-NEUTRAL SELLERS

In this section, we study the POA of FPA with risk-averse bidders and risk-neutral sellers. In this setting, a key factor in determining optimal social welfare is the value of x where $u'(x) = 1$. We first present lower bounds of social welfare for when $x \leq 0$, $0 \leq v \leq x$, and $0 \leq x \leq v$ where v is the winning bidder's valuation. These bounds follow immediately from properties of $u(\cdot)$.

Lemma 5. *If $u'(x) = 1$, then the following holds:*

(1) *if $x \leq 0$, then $u(v - b) + b \geq u(v)$,*
(2) *if $0 \leq v \leq x$, then $u(v - b) + b \geq v$,*
(3) *if $0 \leq x \leq v$, then $u(v - b) + b \geq \min(v, u(v))$*

where $v \geq 0$ and $0 \leq b \leq v$.

In summary, no matter the value of x, a lower bound of social welfare is either v or $u(v)$, whichever is smaller. Also, b denotes the bid of the winning bidder and as all bids in a BNE are between zero and the bidder's value so is the value of b.

It turns out that the difference between the risk attitudes of the bidders and sellers is a game-changer. The POA of FPA is zero even in the symmetric setting. To see this, note that $OPT(\mathbf{v}) = u(x) + (\max(\mathbf{v}) - x)$. Since $u(x) - x$ can be arbitrarily large, so is $OPT(\mathbf{v})$. The actual social welfare however is always finitely large. Thus the ratio between the expected social welfare and the optimal one can be arbitrarily close to zero.

This isn't the only problem, because, as discussed, achieving optimal social welfare may result in negative utility for either the bidder or the seller, which gives rise to a fairness issue. We, therefore, turn our attention to suboptimal social welfare and with respect to which we derive POA bounds.

Unfortunately, the POA is still zero. For a constant one, we have to restrict the range of utility functions. The following theorem gives sufficient and necessary conditions of a constant POA with respect to suboptimal social welfare for symmetric FPA.

Theorem 6. *The price of anarchy of the symmetric first-price single-item auction format with risk-averse bidders, risk-neutral sellers, and sub-optimal social welfare is a constant iff*

$$v/\beta \leq u(v) \leq \alpha \cdot v$$

for all $v > 0$, some $\alpha \geq 1$ and some $\beta \geq 1$ and the constant is at least $\min(1/\alpha, 1/(2 \cdot \beta))$.

Proof. Since $SOPT(\mathbf{v}) = u(x) - x + \max(\mathbf{v})$ when $0 < x < \max(\mathbf{v})$, to prevent $SOPT(\mathbf{v})$ from being infinitely large, $u(x) - x$ must be bounded from above, that is $u(x) \leq \alpha \cdot x$ for some $\alpha > 1$.

Let $x = 0$ and let there be only one bidder who has a valuation of $v > 0$. Then, in any BNE $\mathbf{s}$, the bidder bids 0. Thus $SW(\mathbf{s}(\mathbf{v}), \mathbf{v}) = u(v)$ and $SOPT(\mathbf{v}) = v$. Hence to prevent $\mathbb{E}_{\mathbf{v}}[u(v)]/\mathbb{E}_{\mathbf{v}}[v]$ from approaching 0, $u(v)$ must be bounded from below, that is $u(v) \geq v/\beta$ for some $\beta \geq 1$.

We have shown that $v/\beta \leq u(v) \leq \alpha \cdot v$ is necessary for a constant POA. In the reaming of the proof, we will show that it is also sufficient for the $\min(1/\alpha, 1/(2 \cdot \beta))$ lower bound. In the symmetric setting, the winning bidder is the one with the highest valuation. We let the valuation and the equilibrium bid be v and b respectively. We will show the ratio of the social welfare and the suboptimal one is at least $\min(1/\alpha, 1/(2 \cdot \beta))$ respectively for when $x \leq 0$, $x \geq v$ and $0 < x < v$.

Case 1, $x \leq 0$: Then $SOPT(\mathbf{v}) = v$ and $SW(\mathbf{s}(\mathbf{v}); \mathbf{v}) = u(v - b) + b > u(v) \geq v/\beta$ where the first inequality follows from part 1 of Lemma 5. So the ratio for this case is at least $(v/\beta)/v = 1/\beta$.

Case 2, $x \geq v$: Then $SOPT(\mathbf{v}) = u(v) \leq \alpha \cdot v$ and $SW(\mathbf{s}(\mathbf{v}); \mathbf{v}) = u(v - b) + b > v - b + b = v$ where the last inequality follows from part 2 of Lemma 5. So the ratio for this case is at least $v/(\alpha \cdot v) = 1/\alpha$.

Case 3, $0 < x \leq v$: Then either $v \leq x$ or $v > x$. If $v \leq x$, then this is identical to Case 2 in which the ratio is at least $1/\alpha$. So suppose $v > x$. Then it follows from part 3 of Lemma 5 that $SW(\mathbf{s}(\mathbf{v}); \mathbf{v}) = u(v - b) + b \geq \min(v, u(v))$.

Let $c \in \mathbb{R}$ be such that $u(c) = c$. If $v \leq c$, then, by the concavity of $u(\cdot)$, we have $u(v) > v$ which means $u(v - b) + b \geq v$. Since $SOPT(\mathbf{v}) = u(x) + v - x \leq (\alpha - 1) \cdot x - v \leq \alpha \cdot v$, the ratio is at least $v/(\alpha \cdot v) = 1/\alpha$. If $v > c$, then we have $u(v) < v$ which means $u(v - b) + b \geq u(v) \geq v/\beta$. Since $u(x) < u(v) < v$, $SOPT(\mathbf{v}) = u(x) + v - x \leq 2 \cdot v$. Thus the ratio is at least $(v/\beta)/(2 \cdot v) = 1/(2 \cdot \beta)$. □

Moving on to the general setting, we establish that if the utility function is such that $v/\beta \leq u(v) \leq \alpha \cdot v$ for some $\alpha \geq 1$ and $\beta \geq 1$, then a FPA achieves at least $\min(1/(2 \cdot \alpha), 1/(4 \cdot \beta))$ of the suboptimal social welfare.

Theorem 7. *The price of anarchy (with respect to sub-optimal social welfare) of the first-price single-item auction format with risk-averse bidders, risk-neutral sellers is at least $\min(1/(2 \cdot \alpha), 1/(4 \cdot \beta))$ if $v/\beta \leq u(v) \leq \alpha \cdot v$ for some $\alpha \geq 1$ and $\beta \geq 1$.*

Proof. As in the proof of Theorem 4 we write $\max(\mathbf{v})$ and $\max(\mathbf{s}(\mathbf{v}))$ as v and b respectively and partition the space of valuation distribution and BNE pairs into a subspace in which the expected winning bid and expected maximum valuation are such that $\mathbb{E}_{\mathbf{v}}[\max(\mathbf{s}(\mathbf{v}))] > \mathbb{E}_{\mathbf{v}}[\max(\mathbf{v})/2]$ and another subspace in which the expected values are such that $\mathbb{E}_{\mathbf{v}}[\max(\mathbf{s}(\mathbf{v}))] \leq \mathbb{E}_{\mathbf{v}}[\max(\mathbf{v})/2]$.

For the former subspace, since $\mathbb{E}_{\mathbf{v}}[b] > \mathbb{E}_{\mathbf{v}}[v/2]$ we have

$$\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})] = \mathbb{E}_{\mathbf{v}}[u(v' - b) + b] \geq \mathbb{E}_{\mathbf{v}}[b] > \mathbb{E}_{\mathbf{v}}[v/2]$$

where v' is the winning bidder's valuation. Depending on the value of x, $SOPT(\mathbf{v})$ is v, $u(v)$, or $u(x) + v - x$ and $SOPT(\mathbf{v}) \leq \alpha \cdot v$ for all three cases. Thus a lower bound for the POA is $(v/2)/(\alpha \cdot v) = 1/(2 \cdot \alpha)$.

For the latter subspace, by following the deviation strategy in the proof of Theorem 4, we have

$$\sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{s}(\mathbf{v}); v_i)] \geq \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[u_i(v_i/2, \mathbf{s}_{-i}(\mathbf{v}_{-i}); v_i)] \geq \mathbb{E}_{\mathbf{v}}[u(v/2 - b)]$$

for any valuation distribution and BNE $\mathbf{s}$ within the subspace. The risk neutrality of the seller does not play a part in deriving the above inequalities. What it does is in deriving the seller's expected utility which is $\mathbb{E}_{\mathbf{v}}[b]$. Therefore

$$\mathbb{E}_{\mathbf{v}}[SW(\mathbf{s}(\mathbf{v});\mathbf{v})] = \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{s}(\mathbf{v}); v_i)] + \mathbb{E}_{\mathbf{v}}[b] \geq \mathbb{E}_{\mathbf{v}}[u(v/2 - b) + b].$$

It follows from Lemma 5 that $u(v/2 - b) + b \geq \min(v/2, u(v/2))$. We will derive a lower bound for the ratio of social welfare and the suboptimal one respectively for the case of $x \leq 0$, $x \leq v$ and $x \geq v$.

Case 1, $x \leq 0$: then $SOPT(\mathbf{v}) = v$ and $v/2 > u(v/2)$ which means the ratio is at least

$$u(v/2)/v \geq (v/(2 \cdot \beta))/v = 1/(2 \cdot \beta).$$

Case 2, $x \geq v$: then $SOPT(\mathbf{v}) = u(v) \leq \alpha \cdot v$ and $v/2 < u(v/2)$ which means the ratio is at least

$$(v/2)/(\alpha \cdot v) = 1/(2 \cdot \alpha).$$

Case 3, $x \leq v$: Let c be such that $u(c) = c$, thus $c > u(x) > x$. If $v/2 \geq c$, then $v/2 \geq u(v/2)$ and $SOPT(\mathbf{v}) = u(x) + v - x \leq 2 \cdot v$ which means the ratio is at least

$$u(v/2)/(2 \cdot v) \geq (v/(2 \cdot \beta))/(2 \cdot v) = 1/(4 \cdot \beta).$$

If $v/2 \leq c < v$, then $v/2 \leq u(v/2)$ and $SOPT(\mathbf{v}) = u(x)+v-x \leq 2 \cdot v$ which means the ratio is at least

$$(v/2)/(2 \cdot v) = 1/4.$$

If $v \leq c$, then $v/2 \leq u(v/2)$ and $SOPT(\mathbf{v}) = u(x) + v - x \leq (\alpha - 1) \cdot x - v \leq \alpha \cdot v$ which means the ratio is at least

$$(v/2)/(\alpha \cdot v) = 1/(2 \cdot \alpha).$$

Thus we conclude that the POA for the latter subspace is at least $\min(1/(2 \cdot \alpha), 1/(4 \cdot \beta))$ which is also the lower bound of the POA for the space of all valuation distribution and BNE.

□

The lower bound of $\min(1/(2 \cdot \alpha), 1/(4 \cdot \beta))$ is one half of the one we derived in the symmetric setting where the only source of inefficiency is non-optimal payment. The results suggest that non-optimal allocation alone takes away 1/2 of the achievable social welfare. But we cannot be certain of it as the bound may not be tight.

6 RELATED WORK

For auctions with risk-averse agents, the main research efforts aim to optimise revenue rather than social welfare. [14, 15] characterised the optimal auction formats for revenue maximisation where [15] adopted a constant absolute risk-averse utility function and [14] a more general one. For classic auction formats, it is shown that FPA generally outperforms second-price auction (SPA) [3, 14, 16]. Since the optimal auctions articulated in [14, 15] are far from practical, many have attempted to fine-tune classic auctions to make the best out of risk aversion. [4–6] studied the best reserve price to set for FPA and SPA. Building on [5], [1] investigated the additional option of charging an entry fee.

Listed as an open question in [18], the POA for auctions in the risk-averse setting is mostly uncharted territory. The topic is better understood in the risk-neutral setting. A major achievement is the so-called *smoothness framework* [17, 19] which provides a general and modular approach to proving POA bounds. The key idea is that for many auctions, there exists a deviation strategy for the bidders that achieves a certain utility that is bounded by some faction of the optimal social welfare minus the payment. We also used the deviation approach in our proof of Theorem 4 and 7. However, deviation alone is in general not sufficient to derive POA bounds in the risk-averse setting. This is largely due to the demand for a specific amount of payment for optimality, which is beyond the consideration of the smoothness framework.

Despite the difficulty, [8] tried to adapt the smoothness framework to the risk-averse setting. Their main result is that, for an auction in the risk-neutral setting, if a POA of α is provable through the smoothness framework, then the auction has a POA of at least $\alpha/2$ in the risk-averse setting,[2] provided that the deviation does not incur negative utility and the utility function is normalised. Due to the normalisation, the result of [8] do not hold for general concave utility functions like ours. Moreover, [8] considered only risk-neutral sellers.

In the non-auction domain, [11, 12] established inefficiency bounds for risk-averse selfish routing. It remains to be seen if their approaches are transferable to auctions.

7 CONCLUSION

In this paper, we explored the POA of FPA with risk-averse bidders. Deriving POA bounds in the risk-averse setting is intrinsically more difficult than in the risk-neutral setting. The concavity of the utility function and its implications on optimal social welfare are not the only difficulties. According to our definition of POA, there lies the task of identifying the worst of such concave functions in terms of efficiency.

When the seller is also risk-averse, we established that the POA is 1/2 for both the symmetric FPA and FPA in general. When the seller is risk-neutral, we demonstrated that the POA is zero. Because of a fairness issue with optimal social welfare in this setting, we opt to derive POA bounds with respect to suboptimal social welfare. It is still zero, but we are able to articulate a parameterised lower bound with a simple expression.

[2] It is $2 \cdot \alpha$ in [8] which takes the optimal social welfare as the nominator in defining the POA.

Two new conceptual aspects of risk-aversion and POA arise while obtaining the technique results. The first one is the distinction between risk-averse and risk-neutral sellers. Our findings indicate the seller's risk attitude can have a significant impact of the efficiency of FPA. When the risk attitude of the bidder aligns with that of the seller, the auction has better efficiency. The second one is the "non-optimal allocation versus non-optimal payment" problem in determine the inefficiency of FPA. Both forms of non-optimality contribute to the inefficiency, the question is by how much it is attributed to one particular form.

There are several routes for further work. While aiming for generality, we made minimum assumptions over the utility functions, it holds the promise to investigate the POA with respect to some specific risk-averse utility functions that are well-studied. These include, for example, the constant relative risk aversion utility function $u(z) = z^{\alpha}$ and the constant absolute risk aversion utility function $u(z) = 1 - exp(-\alpha z)$ where α is the measure of risk aversion. We know more about the equilibrium with such utility functions [10]. Furthermore we think that a clear separation of the inefficiency caused by non-optimal allocation and non-optimal payment deserves further investigation. A rigorously way to go about it is work out the minimum ratio between the expected valuation of the winning bidder and the expected highest valuation. This will give us the exact amount of inefficecny caused by non-optimal allocation.

ACKNOWLEDGMENTS

This work was partially supported by National Natural Science Foundation of China (NSFC) (61976153).

REFERENCES

[1] Indranil Chakraborty. 2019. RESERVE PRICE VERSUS ENTRY FEE IN STANDARD AUCTIONS. *Economic Inquiry* 57, 1 (2019), 648–653.

[2] James C. Cox, Vernon L. Smith, and James M. Walker. 1985. Experimental Development of Sealed-Bid Auction Theory; Calibrating Controls for Risk Aversion. *The American Economic Review* 75, 2 (1985), 160–165.

[3] Charles A. Holt. 1980. Competitive Bidding for Contracts under Alternative Auction Procedures. *Journal of Political Economy* 88, 3 (1980), 433–445.

[4] Audrey Hu. 2011. How bidder's number affects optimal reserve price in first-price auctions under risk aversion. *Economics Letters* 113, 1 (2011), 29–31.

[5] Audrey Hu, Steven A Matthews, and Liang Zou. 2010. Risk aversion and optimal reserve prices in first- and second-price auctions. *Journal of Economic Theory* 145, 3 (2010), 1188–1202.

[6] Audrey Hu, Steven A Matthews, and Liang Zou. 2019. Low Reserve Prices in Auctions. *The Economic Journal* 129, 622 (2019), 2563–2580.

[7] Yaonan Jin and Pinyan Lu. 2022. First Price Auction is $1 - 1/e^2$ Efficient. In *proceedings of the 63rd Annual Symposium on Foundations of Computer Science*.

[8] Thomas Kesselheim and Bojana Kodric. 2018. Price of Anarchy for Mechanisms with Risk-Averse Agents. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming*. 155:1–155:14.

[9] Elias Koutsoupias and Christos Papadimitriou. 2009. Worst-case equilibria. *Computer Science Review* 3, 2 (2009), 65–69.

[10] Vijay Krishna. 2009. *Auction Theory* (second ed.). Academic Press.

[11] Thanasis Lianeas, Evdokia Nikolova, and Nicolas E. Stier-Moses. 2016. Asymptotically Tight Bounds for Inefficiency in Risk-Averse Selfish Routing. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*. 338–344.

[12] Thanasis Lianeas, Evdokia Nikolova, and Nicolas E. Stier-Moses. 2019. Risk-Averse Selfish Routing. *Mathematics of Operations Research* 44, 1 (2019), 38–57.

[13] 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.

[14] Eric Maskin and John Riley. 1984. Optimal Auctions with Risk Averse Buyers. *Econometrica* 52, 6 (1984), 1473–1518.

[15] Steven A Matthews. 1983. Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory* 30, 2 (1983), 370–400.

[16] John G. Riley and William F. Samuelson. 1981. Optimal Auctions. *The American Economic Review* 71, 3 (1981), 381–392.

[17] Tim Roughgarden. 2015. Intrinsic Robustness of the Price of Anarchy. *J. ACM* 62, 5 (2015).

[18] Tim Roughgarden, Vasilis Syrgkanis, and Éva Tardos. 2017. The Price of Anarchy in Auctions. *Journal of Artificial Intelligence Research* 59, 1 (2017), 59–101.

[19] Vasilis Syrgkanis and Eva Tardos. 2013. Composable and Efficient Mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. 211–220.

[20] Shoshana Vasserman and Mitchell Watt. 2021. Risk aversion and auction design: Theoretical and empirical evidence. *International Journal of Industrial Organization* 79 (2021), 102758.