Efficiency of the First-Price Auction in the Autobidding World

Yuan Deng
Google Research
dengyuan@google.com

Jieming Mao
Google Research
maojm@google.com

Vahab Mirrokni
Google Research
mirrokni@google.com

Hanrui Zhang
Chinese University of Hong Kong
hanrui@cse.cuhk.edu.hk

Song Zuo
Google Research
szuo@google.com

Abstract

We study the price of anarchy of first-price auctions in the autobidding world, where bidders can be either utility maximizers (i.e., traditional bidders) or value maximizers (i.e., autobidders). We show that with autobidders only, the price of anarchy of first-price auctions is $1/2$, and with both kinds of bidders, the price of anarchy degrades to about 0.457 (the precise number is given by an optimization). These results complement the recent result by Jin and Lu [2022] showing that the price of anarchy of first-price auctions with traditional bidders is $1 - 1/e^2$. We further investigate a setting where the seller can utilize machine-learned advice to improve the efficiency of the auctions. There, we show that as the accuracy of the advice increases, the price of anarchy improves smoothly from about 0.457 to 1.

1 Introduction

Autobidding, the procedure of delegating the bidding tasks to automated agents to procure advertisement opportunities in online ad auctions, is becoming the prevalent bidding methods and contributing to more than 80% of total online advertising traffic [Dolan, 2020]. Autobidding significantly simplifies the interaction between the advertisers and the ad platform: instead of submitting bids for each auction separately, the advertisers only need to specify their high-level objectives and constraints and the automated agents bid on behalf of the advertisers. A popular autobidding strategy is value maximization subject to a target return-on-investment (ROI) constraint, in which the autobidding agents aim to maximize their value, such as the number of clicks/conversions, subject to a minimum admissible ratio on the return-on-investment constraint. This type of bidders has been referred as *value maximizers* by the recent line of research on autobidding [Deng et al., 2021, Balseiro et al., 2021a].

As the behavior model of automated bidders is very different from the classic (quasi-linear) *utility maximizers* maximizing the difference between value and payment, there has been a growing body of literature revisiting the effectiveness of the existing mechanisms through the lens of the autobidding world [Aggarwal et al., 2019, Deng et al., 2021, Balseiro et al., 2021a, Liaw et al., 2022, Babaioff et al., 2021]. In particular, it is well-known that running the second-price auction against utility maximizers results in socially optimal outcomes; however, Aggarwal et al. [2019] show that the price of anarchy (PoA) [Koutsoupias and Papadimitriou, 1999], the ratio between the worst welfare

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

https://doi.org/10.52202/079017-4420

in equilibrium and the socially optimal welfare, of running the second-price auction against value maximizers is $1/2$ (when the bidders do not adopt dominated strategies[1]).

In addition to the change in bidding methods, there has been an industry-wide change in auction formats in recent years: shifting from the second-price auction to the first-price auction, especially in the display ad markets, culminating in the move to the first-price auction from the Google Ad Exchange in September 2019 [Paes Leme et al., 2020]. The shift towards the first-price auction in combination with the increasing adoption of autobidding raises the question of understanding *the PoA of running the first-price auction in the autobidding world.*

1.1 Our Results

In this paper, we characterize the PoA of running the first-price auction in the fully autobidding world, where all bidders are value maximizers, and in the mixed autobidding world, where bidders could be either value maximizers or utility maximizers. Our results complement the very recent result from Jin and Lu [2022] showing that the PoA of running the first-price auction is $1 - 1/e^2$ in a world without autobidding, i.e., all bidders are utility maximizers. The results are summarized in Table 1 and we also include the known results for second-price auctions for comparison.

Mechanism	First-price auction	Second-price auction
Full autobidding	$1/2$	$1/2$ [Aggarwal et al., 2019]
Mixed autobidding	$\min_{t\in[0,1]} \frac{1+t\ln t}{2-t+t\ln t} \approx 0.457$	$1/2$ [Balseiro et al., 2021b]
No autobidding	$1 - 1/e^2 \approx 0.865$ [Jin and Lu, 2022]	1 [Vickrey, 1961]

Table 1: Summary of PoAs in second-price and first-price auctions under full autobidding, mixed autobidding, and no autobidding environments. We note that for the second-price auction, bidders are assumed to not adopt dominated strategies; otherwise the PoA can be 0.

Our results demonstrate that in a full autobidding world, the PoA of running the second-price auction and the first-price auction is the same, i.e, $1/2$; however, in the mixed autobidding world with both utility maximizers and value maximizers, the PoA of running the first-price auction is strictly less than the PoA of running the second-price auction. It is worth highlighting that our PoA bound for the first-price auction in the mixed autobidding world involves solving an optimization problem $\min_{t\in[0,1]} \frac{1+t\ln t}{2-t+t\ln t}$ and it turns out that this is the right bound as we are able to construct an instance matching this bound.

In addition to the above results, we further adopt the environment introduced by Balseiro et al. [2021b], where the seller can leverage machine-learned advice that approximates the buyers' values for designing mechanisms. Balseiro et al. [2021b] show that as the advice gets more and more accurate in predicting the buyers' values, setting reserves based on the advice in the second-price auction makes the PoA approaching 1. In our setting of using machine-learned advice in the first-price auction, we prove that setting reserves based on machine-learned advice is also effective in improving PoA in both full and mixed autobidding world, and the PoA approaches 1 as the quality of the machine-learned advice increases.

1.2 Related Works

For first-price auction with utility maximizers, the price of anarchy has been extensively studied [Roughgarden et al., 2017]. In the complete-information setting with additive values, there is no efficiency loss, i.e., the price of anarchy is 1 [Feldman et al., 2016]. For Bayesian settings, Feldman et al. [2013] first showed that the price of anarchy is at least $1/2$ for subadditive valuations and Syrgkanis and Tardos [2013] showed that the price of anarchy is at least $1 - 1/e$ for submodular valuations. These bound are shown to be tight for buyers with submodular valuations and subadditive valuations [Christodoulou et al., 2016]. When agents' values are independently distributed, the price of anarchy was later improved to ≈ 0.743 by Hoy et al. [2018] and Hartline et al. [2014] gave a concrete instance to establish an upper bound of ≈ 0.869. More recently, Jin and Lu [2022] show that

[1]When discussing the second-price auction, we assume the bidders do not adopt dominated strategies. In the case when the bidders could adopt dominated strategies, the PoA of running the second-price auction can be 0.

the price of anarchy for independently distributed values is exactly $1 - 1/e^2$. Paes Leme et al. [2020] provide models and analyses to explain why Display Ads market shifts to use first-price auctions.

Motivated by the emergence of autobidding in online advertising, there is a recent line of work studying auction design for value maximizers in an autobidding world. Aggarwal et al. [2019] design optimal bidding strategies in truthful auctions, show the existence of equilibrium, and prove welfare price of anarchy results. Deng et al. [2021] and Balseiro et al. [2021b] show how boosts and reserves can be used to improve the welfare efficiency guarantees given machine-learned advice approximating buyers' values. Balseiro et al. [2021a, 2022] give the characterization of the revenue-optimal auctions under various information structure. Mehta [2022] and Liaw et al. [2022] study how randomized and non-truthful mechanisms can help improve welfare efficiency for the case of two bidders. In particular, Liaw et al. [2022] show that the PoA of running the first-price auction is $1/2$ when restricting to equilibria of deterministic bids. In this work, we consider general equilibria with randomized bidding strategies. Deng et al. [2024] study the efficiency, as measured by the price of anarchy, of the generalized second-price auction in the autobidding world, and give a parametrized bound. Liaw et al. [2024] consider a setting with value maximizers only, which is similar to our "full autobidding" setting. The main difference is that they focus on the effects of budget constraints on top of ROI / ROS constraints. The high-level message there is that using the fractional optimum as the benchmark, first-price auctions become much less efficient with budget constraints, but this efficiency loss can be circumvented if we consider the (weaker) integral optimum as the benchmark. Gaitonde et al. [2023] take a somewhat different perspective, and give a bidding dynamics that does not necessarily converge, but on average achieves good welfare. Conitzer et al. [2022] study "pacing" in first-price auctions, where constrained bidders optimize their bids in multiple auctions by choosing a single multiplier, and bidding the product of this multiplier and their value in each auction. In contrast, we do not restrict the strategy space of bidders in this paper. Feng et al. [2024] study a meta-game in the context of autobidding, where human bidders strategically specify their constraints (e.g., the target ROI). They show that such strategic selection does not hurt the welfare too much in first-price auctions. For a more comprehensive overview, see, e.g. [Aggarwal et al., 2024].

2 Preliminaries

Ad auctions. Following prior work on autobidding, we consider a setting with n bidders $[n]$ and m auctions $[m]$, where each bidder participates in all auctions simultaneously. We generally use i to index bidders, and j to index auctions. Each bidder i has a value $v_{i,j} \geq 0$ for winning in each auction j. We consider first-price auctions — arguably the most common and well-known auction format — as the auction mechanism throughout the paper. That is, each bidder i submits a bid $b_{i,j}$ in each auction j. In each auction j, the bidder with the highest bid wins and is charged the winning bid as the payment. For the existence of equilibrium in the first-price auction, we assume ties are broken in favor of the bidder with the higher value within an auction, which is a standard assumption in the literature [Maskin and Riley, 2000].[2] If multiple bidders have the same value, then ties are broken in favor of the bidder with the smallest index.[3] We let $x_{i,j} = x_{i,j}(\{b_{i',j}\}_{i'})$ be the indicator variable indicating that i wins in auction j. Then, the value each bidder i receives in auction j is $\mathsf{val}_{i,j} = \mathsf{val}_{i,j}(\{b_{i',j}\}_{i'}) = x_{i,j}(\{b_{i',j}\}_{i'}) \cdot v_{i,j}$, and the payment i makes in auction j is $p_{i,j} = p_{i,j}(\{b_{i',j}\}_{i'}) = x_{i,j}(\{b_{i',j}\}_{i'}) \cdot b_{i,j}$. We remark that bidders may employ randomized bidding strategies, in which case $\{b_{i,j}\}$, $\{x_{i,j}\}$, $\{\mathsf{val}_{i,j}\}$ and $\{p_{i,j}\}$ are all random variables. Throughout the paper, we omit the dependence of $x_{i,j}$, $\mathsf{val}_{i,j}$ and $p_{i,j}$ on $\{b_{i',j}\}_{i'}$ whenever it is clear from the context. For brevity, we write $\boldsymbol{v}_i = \{v_{i,j}\}_j$, $\boldsymbol{b}_i = \{b_{i,j}\}_j$, etc.

Utility maximizers and value maximizers. We consider two types of bidders: traditional (quasi-linear) utility maximizers and (ROI-constrained) value maximizers. In the rest of the paper, we use B_u to denote the set of utility maximizers and B_v the set of value maximizers, where we always have $B_u \cap B_v = \emptyset$ and $B_u \cup B_v = [n]$. For each bidder i, let $\mathsf{val}_i = \sum_j \mathsf{val}_{i,j}$ be the total value

[2] Maskin and Riley [2000] assume ties are broken with a second-round second-price auction among the highest bidders to guarantee the existence of equilibrium in the presence of discrete values. This tie-breaking rule is equivalent to favoring the bidder with the higher value for standard utility maximizers, while the two tie-breaking rules can be different for value maximizers. We choose the one without the additional definition of second-round second-price auction for simplicity.

[3] Our positive results are oblivious to tiebreaking, i.e., they remain valid no matter what tiebreaking rule is used.

 https://doi.org/10.52202/079017-4420

i receives from all auctions, and $p_i = \sum_j p_{i,j}$ be the total payment i makes. Fixing other bidders' bidding strategies, a utility maximizer i bids in a way such that the total quasi-linear utility, i.e., $\mathsf{val}_i - p_i$, is maximized, without additional constraints. On the other hand, a value maximizer i bids in a way such that the total value val_i that i receives is maximized, subject to the ROI constraint that the ratio between the total value val_i and the total payment p_i is at least some predetermined threshold. Without loss of generality we assume this threshold is 1[4], i.e., we require the total payment p_i is at most the total value val_i for each value maximizer i. Formally, fixing other bidders' bids, a value maximizer i solves the following optimization problem to determine their bids $\{b_{i,j}\}_j$:

$$\begin{aligned} \text{maximize} \quad & \mathbb{E}\left[\mathsf{val}_i\right] \\ \text{subject to} \quad & \mathbb{E}\left[\mathsf{val}_i - p_i\right] \geq 0. \end{aligned}$$

Here, the expectations are taken over the randomness introduced by bidders bidding randomly. We note that when deciding how to bid, i can observe the distributions, but not the realizations, of the other bidders' bids.

Equilibria and the price of anarchy. The focus of this paper is on the efficiency of the first-price auction in equilibrium. That is, when each bidder i's bidding strategy $\boldsymbol{b}_i$ is optimal for that bidder's objective, given all other bidders' bidding strategies $\{\boldsymbol{b}_{i'}\}_{i' \neq i}$ (or $\boldsymbol{b}_{-i}$ for brevity). In other words, each bidder's strategy is a best response to other bidders' strategies. Note that since each bidder i cannot observe the realizations of other bidders' bids, all bidders' bids $\{\boldsymbol{b}_i\}_i$ must be independent. The way we measure efficiency is through the classical notion of the price of anarchy (PoA) [Koutsoupias and Papadimitriou, 1999], i.e., the ratio between the welfare resulting from the worst equilibrium possible, and the optimal welfare. Formally, given a problem instance of $(n, m, B_u, B_v, \{v_{i,j}\})$, we are interested in the instance-wise PoA defined as follows:

$$\mathsf{PoA}(n, m, B_u, B_v, \{v_{i,j}\}) = \inf_{\{\boldsymbol{b}_i\}_i \text{ form an equilibrium}} \frac{\sum_i \mathbb{E}[\mathsf{val}_i]}{\sum_j \max_i v_{i,j}}.$$

Our goal is to pin down the PoA of the first-price auction, which is the infimum of the instance-wise PoA taken over all instances of the problem. We assume $\sum_j \max_i v_{i,j} > 0$ in order for the PoA to be well-defined.

Equilibrium bidding behavior of value maximizers. In this paper, we do not put further restrictions on bidders' equilibrium bidding behavior, such as no-overbidding. This is in contrast to some prior work: In prior work, since overbidding is a dominated strategy for utility maximizers, the no-overbidding assumption is often imposed for analyzing the PoA of utility maximizers in second-price auctions and first-price auctions (otherwise, the PoA could be 0). However, in the model we consider, assuming no overbidding for value maximizers with ROI constraints may rule out their optimal strategies. More importantly, with the no overbidding assumption, the optimal strategy for value-maximizers would be to bid their values in first-price auctions since such a strategy maximizes their bids without violating the ROI constraints, leading to the trivial welfare-optimal outcome in a world with value-maximizers only. In other words, the welfare inefficiency in the autobidding world is mainly caused by the overbidding behavior from value maximizers, and such inefficiency can only be analyzed without assumptions such as no-overbidding (although such assumptions might make sense in other contexts).

3 PoA of the First-price Auction in a Full Autobidding World

In the literature [Aggarwal et al., 2019, Deng et al., 2021, Balseiro et al., 2021b], the analyses of PoA for the second-price auction in a full autobididng world mainly leverage the fact that the second-price auction is truthful for utility maximizers, a nice property implying that value maximizers adopt *uniform bidding* strategy in equilibrium [Aggarwal et al., 2019]. Here, a uniform bidding strategy for a value maximizer i is a strategy in which she bids $k_i \cdot v_{i,j}$ with a constant bid multiplier k_i in each auction j. However, as the first-price auction is not truthful for utility maximizers, a uniform bidding strategy is generally no longer a best response for a value maximizer. A value maximizer i may adopt

[4]See Appendix A for why this is without loss of generality, as well as further justifications for considering the liquid welfare in the autobidding world.

a *non-uniform* bidding strategy using different bid multiplier $k_{i,j}$ in different auction j, and the bid multipliers can even be *randomized*.

Liaw et al. [2022] show that the PoA of running the first-price auction in a full autobidding world is $1/2$ when restricting to equilibria of bidding in a deterministic way (i.e., restricted to pure Nash equilibria). In this section, we devise a novel argument to analyze equilibria with possibly randomized non-uniform bidding strategies to show that the PoA of the first-price auction in a full autobidding world is exactly $1/2$.

Theorem 1. *When all bidders are value maximizers, the PoA of the first-price auction in a full autobidding world is* $1/2$*. Formally,*

$$\inf_{n,m,\{v_{i,j}\}} \mathsf{PoA}(n, m, B_u = \emptyset, B_v = [n], \{v_{i,j}\}) = 1/2.$$

All missing proofs are deferred to Appendix B. We remark that the proof of Theorem 1 in fact establishes the following claim, which will be a building block in our analysis of PoA of the first-price auction in a mixed autobidding world with both value maximizers and utility maximizers.

Lemma 1. *For any* n*,* m*,* B_u*,* B_v*, and* $\{v_{i,j}\}$*, if* $\{\boldsymbol{b}_i\}_i$ *form an equilibrium, then*

$$\sum_{i \in B_v} \mathbb{E}[\mathsf{val}_i] + \sum_{i,j:\mathsf{rw}(j) \in B_v} \mathbb{E}[p_{i,j}] \geq \sum_{j:\mathsf{rw}(j) \in B_v} \max_i v_{i,j},$$

where $\mathsf{rw}(j) = \operatorname{argmax}_i v_{i,j}$ *denotes the "rightful winner" of each auction* j*, with ties broken in favor of the bidder with the smallest index.*

4 PoA of the First-price Auction in a Mixed Autobidding World

In this section, we proceed to the main result of the paper: characterizing the exact PoA of running the first-price auction in a mixed autobidding world with both utility maximizers and value maximizers participating in the auctions. The entire section is devoted to the proof of the following claim.

Theorem 2. *The PoA of the first-price auction in a mixed autobidding world is given by:*

$$\inf_{n,m,B_u,B_v,\{v_{i,j}\}} \mathsf{PoA}(n, m, B_u, B_v, \{v_{i,j}\}) = \min_{t \in [0,1]} \frac{1 + t \ln t}{2 - t + t \ln t} \approx 0.457.$$

Roadmap. In the rest of the section, we first establish a PoA lower bound. At first sight, one might be tempted to adapt the proof of Theorem 1, hoping to get the same lower bound of $1/2$ in the presence of utility maximizers. In doing so, one faces the following intuitive difficulty: for a utility maximizer i, the value i receives when best responding is not necessarily lower bounded by the value i receives when using the proxy bidding strategy $b'_{i,j} = v_{i,j}$. This is because aiming to maximize utility, i may bid in a way such that both the value i receives and the payment i makes are smaller, but their difference (which is the utility i gets) is larger. One simple fix is to consider instead the following proxy strategy: $b''_{i,j} = v_{i,j}/2$. Adapting the proof of Theorem 1, it is not hard to show this alternative proxy strategy gives a PoA lower bound of $1/4$. However, it is unlikely that such a fix would yield the tight ratio of about 0.457.

In order to overcome this obstacle, we present a novel local analysis (Lemma 3) for the first-price auction that handles auctions where the rightful winner is a utility maximizer, which roughly says that if the expected value that the rightful winner receives from such an auction is small, then the expected payment that other bidders make in the same auction must be large. This argument relies on a characterization of the structure of the best response of a utility maximizer. Given the local analysis, one can already get a lower bound of about 0.317 on the PoA by exploiting the tradeoff between value and payment.

To obtain the right ratio, we further combine the local analysis for utility maximizers with the bound for value maximizers (Lemma 1) discussed earlier. The idea is to consider 4 quantities separately: (A) the total value that all value maximizers receive, (B) the total value that all utility maximizers receive from auctions where they are the rightful winner, (C) the total payment made in auctions where the rightful winner is a value maximizer, and (D) the total payment made in auctions where the rightful winner is a utility maximizer, by bidders who are not the rightful winner.

 https://doi.org/10.52202/079017-4420

One can show that the welfare in equilibrium is lower bounded by a very specific combination of these 4 quantities (Lemma 4). Moreover, Lemma 1 provides a tradeoff between (A) and (C), and the local analysis (Lemma 3) provides a tradeoff between (B) and (D). Therefore, one can lower bound the PoA by solving an optimization problem involving a few number of parameters to get the worst-case tradeoffs. We show that this optimization gives the right ratio of about 0.457 (Lemma 5). Finally, we present a problem instance where there exists an equilibrium where the gap between the actual welfare and the optimal welfare is precisely the right ratio (Lemma 6). This, together with the matching lower bound, concludes the proof of Theorem 2.

4.1 A Local Analysis

In this subsection, we propose a local analysis of the first-price auction to analyze the tradeoff between welfare and payment in equillibrium when the rightful winner of the auction is a utility maximizer. Although the analysis is mainly designed for utility maximizers towards establishing the lower bound of PoA, it also works when the rightful winner of the auction is a value maximizer. For the remainder of this subsection, we fix n, m, $\{v_{i,j}\}$, and an equilibrium profile $\{\boldsymbol{b}_i\}_i$, and consider any fixed bidder i and auction j where i is the rightful winner. For brevity, we also let $v = v_{i,j}$.

Value-payment frontiers. We first try to understand the maximum expected value i can get in auction j when paying a particular amount in expectation. Let F be the CDF of the highest other bid, i.e., for any $x \geq 0$,

$$F(x) = \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} \leq x\right].$$

Note that in the first-price auction j, if bidder i bids t, the bidder i pays $t \cdot F(t)$ and earns value $v \cdot F(t)$. For any $x \geq 0$, consider a function $G : \mathbb{R}_+ \to [0, v]$ such that

$$G(x) = \sup_{t \geq 0:\ t \cdot F(t) \leq x} v \cdot F(t).$$

In other words, G maps a desired expected payment x to the maximum (or strictly speaking, supremum) value i can get by bidding deterministically, when paying at most the desired payment x in expectation. By bidding randomly, i can further achieve any value-payment pair that is a convex combination of points on G. In light of this, we further consider the *concave envelope H of G*, which captures the value-payment frontier when i can bid randomly. Note that for utility maximizers, although there always exists a deterministic best response (since all bids with a positive probability in a randomized bidding strategy must result in the same quasi-linear utility), the equilibrium bidding strategies may still be randomized for the sake of stability.

Characterizing the best response. Given the definition of the value-payment frontier H, one can further observe that when best responding, i's expected utility $\mathbb{E}[\mathsf{val}_{i,j} - p_{i,j}]$ depends only on H. In fact, intuitively, since H is concave, i's expected value and payment in auction j must correspond to the "rightmost" point on H where the derivative is larger than or equal to 1 (this is informal since (1) the point may not be unique and (2) the derivative may be undefined). Formally, when i is a utility maximizer, i's best response in j, $b_{i,j}$, has the following property.

Lemma 2. *For any $b \geq 0$ in the support of $b_{i,j}$, we have for any $w \geq b \cdot F(b)$, $H(w) - H(b \cdot F(b)) \leq w - b \cdot F(b)$.*

We remark that the same property also holds when i is a value maximizer, with minor modifications in the proof.

A tradeoff between value and payment in equilibrium. We are ready to present a tradeoff between the value i receives and the payment other bidders make in auction j. The intuition is that given Lemma 2, the curve H must be relatively flat to the right of $(p_{i,j}, H(p_{i,j}))$ (i.e., with slope at most 1), which means the CDF function F of the highest other bid must also be with a relatively small slope. On the other hand, the payment made by other bidders can be written as

$$b_{i,j} \cdot (1 - F(b_{i,j})) + \int_{b_{i,j}}^{\infty} (1 - F(t))\,\mathrm{d}t.$$

This means if $F(b_{i,j})$, the probability that i wins when bidding $b_{i,j}$, is small, then since F is with a small slope to the right of $b_{i,j}$, the payment must be relatively large. The following claim formally captures this tradeoff.

Lemma 3. *There exists some $x \in [0,1]$ such that*

$$\mathbb{E}[\mathsf{val}_{i,j}] = x \cdot v \quad \textit{and} \quad \sum_{i' \neq i} \mathbb{E}[p_{i',j}] \geq (1 - x + x \ln x) \cdot v.$$

We remark that Lemma 3 alone already implies a lower bound of about 0.317 on the PoA: summing the bound in Lemma 3 over i and j where $\mathsf{rw}(j) = i$ and exploiting the convexity of the mapping $z \mapsto 1 - z + z \ln z$, one can show there is some $x \in [0,1]$ such that

$$\sum_i \mathbb{E}[\mathsf{val}_i] \geq x \cdot \sum_j \max_i v_{i,j} \quad \text{and} \quad \sum_i \mathbb{E}[p_i] \geq (1 - x + x \ln x) \cdot \sum_j \max_i v_{i,j}.$$

Since $\sum_i \mathbb{E}[\mathsf{val}_i] \geq \sum_i \mathbb{E}[p_i]$, the worst case scenario is when $x = 1 - x + x \ln x$, which gives $x \approx 0.317$, a lower bound on the PoA.

4.2 Putting Everything Together

In this subsection, we combine Lemma 1 and Lemma 3 to obtain an improved (and in fact, tight) lower bound on the PoA of the first-price auction in the mixed autobidding world. Fixing n, m, B_u, B_v and $\{v_{i,j}\}$, we consider four key quantities in the proof. (A): the total value that all value maximizers receive; (B): the total value that all utility maximizers receive from auctions where they are the rightful winner; (C): the total payment made in auctions where the rightful winner is a value maximizer; and (D): the total payment made in auctions where the rightful winner is a utility maximizer, by bidders who are not the rightful winner. Formally,

$$A = \sum_{i \in B_v} \mathbb{E}[\mathsf{val}_i], \qquad B = \sum_{i \in B_u, j : \mathsf{rw}(j) = i} \mathbb{E}[\mathsf{val}_{i,j}]$$

$$C = \sum_{i,j : \mathsf{rw}(j) \in B_v} \mathbb{E}[p_{i,j}], \qquad D = \sum_{i,j : \mathsf{rw}(j) \in B_u \text{ and } \mathsf{rw}(j) \neq i} \mathbb{E}[p_{i,j}].$$

We first show that the welfare can be bounded by a combination of the above quantities.

Lemma 4. *For any bidding strategies $\{\mathbf{b}_i\}_i$ that form an equilibrium, we always have*

$$\sum_i \mathbb{E}[\mathsf{val}_i] \geq \max\{A, C + D\} + B.$$

We are now ready to prove our main PoA lower bound.

Lemma 5. *In a mixed autobidding world with both utility maximizers and value maximizers, the PoA of the first-price auction has the following lower bound:*

$$\inf_{n,m,B_u,B_v,\{v_{i,j}\}} \mathsf{PoA}(n, m, B_u, B_v, \{v_{i,j}\}) \geq \min_{t \in [0,1]} \frac{1 + t \ln t}{2 - t + t \ln t} \approx 0.457.$$

Remark 1. *Our analysis involves bounds (e.g., Lemma 3) of similar forms to those normally seen in the smoothness framework [Syrgkanis and Tardos, 2013] and its instantiations (e.g., [Hartline et al., 2014]), and it is conceivable that one might be able to fit part of our analysis in the smoothness framework. However, it is unlikely that doing so would substantially simplify or change the nature of the analysis. In particular, in order to instantiate the smoothness framework in the proper way that produces the tight bounds, presumably one would need to carry out essentially the same argument as in our analysis.*

4.3 A Matching Upper Bound

To conclude this section, we present a problem instance in which our PoA lower bound can be actually achieved, even with $n = 2$ bidders and $m = 2$ auctions.

Lemma 6. *When $n = m = 2$, there exists B_u, B_v and $\{v_{i,j}\}$ such that*

$$\mathsf{PoA}(n, m, B_u, B_v, \{v_{i,j}\}) \leq \min_{t \in [0,1]} \frac{1 + t \ln t}{2 - t + t \ln t} \approx 0.457.$$

Theorem 2 is a direct corollary of Lemma 5 and Lemma 6.

https://doi.org/10.52202/079017-4420

5 Improved Efficiency with ML Advice

In this section, we turn to a more practically-motivated setting, where the seller can rely on machine-learned advice to set reserves and thereby further improve the efficiency of the first-price auction. In particular, we show that as the advice gets more and more accurate, there is a way to set reserves such that the PoA of the first-price auction approaches 1.

Setup. Following prior work Balseiro et al. [2021b], we consider the setting in which for each auction j, the seller can access a signal s_j, which approximates the highest value $\max_i v_{i,j}$ in this auction. Formally, we assume there is a constant $\gamma \in [0,1]$ such that $s_j \in [\gamma \cdot \max_i v_{i,j}, \max_i v_{i,j}]$ for each auction j.[5] As Balseiro et al. [2021b] argue, the model is a reasonable simplification of the actual problem, which enables clean theoretical results while still making practical sense. In particular, the assumption that the signal never overestimates the value can be relaxed: if there is a (high-probability) bound on how much the signal might overestimate the maximum value, then one can scale everything down accordingly and retain the PoA bound in Theorem 3 (with high probability).

First-price auction with machine-learned reserves. The way we exploit machine-learned signals is to use them as reserve prices in the first-price auction. To be more specific, in each auction j, we set a reserve $r_j = s_j$, such that a bidder i can win only if i's bid $b_{i,j} \geq r_j$. In particular, if $\max_i b_{i,j} < r_j$, then no one wins in auction j. This is different from the first-price auction without reserves, where there is always exactly 1 winner.

We overload the notations and let $x_{i,j} = x_{i,j}(r_j, \{b_{i',j}\}_{i'})$ be the indicator variable that i wins in auction j, and define $\{\mathsf{val}_{i,j}\}$ and $\{p_{i,j}\}$ in a similar way such that all these variables depend on both the reserves $\{r_j\}$ and the bidding profile $\{b_{i',j}\}_{i'}$. For any n, m, B_u, B_v, and $\{v_{i,j}\}$, we consider the following parametrized instance-wise PoA when each $r_j = s_j \in [\gamma \cdot \max_i v_{i,j}, \max_i v_{i,j}]$:

$$\mathsf{PoA}_\gamma(n, m, B_u, B_v, \{v_{i,j}\}) = \inf_{\substack{\{b_{i,j}\} \text{ form an equilibrium} \\ \forall j, r_j \in [\gamma \cdot \max_i v_{i,j}, \max_i v_{i,j}]}} \frac{\sum_i \mathbb{E}[\mathsf{val}_i]}{\sum_j \max_i v_{i,j}}.$$

Our main result in the section is captured by the following theorem.

Theorem 3. *For any n, m, B_u, B_v, $\{v_{i,j}\}$ and $\gamma \in [0,1]$,*

$$\mathsf{PoA}_\gamma(n, m, B_u, B_v, \{v_{i,j}\}) \geq \min_{t \in [0,1]} \frac{1 + t\ln t - \gamma t(1 + \ln t)}{2 - t - \gamma + (1-\gamma)t\ln t}.$$

One way to interpret the lower bound is it provides a smooth transition between the case without machine-learned advice (or with totally unreliable advice, i.e., $\gamma = 0$), and the case where the advice is perfectly accurate (i.e., $\gamma = 1$) and the seller knows precisely the value of the rightful winner in each auction. In the former scenario corresponding to $\gamma = 0$, we recover the tight PoA bound stated in Theorem 2. In the latter scenario corresponding to $\gamma = 1$, as one would expect, the theorem gives a lower bound of 1, which means the auctions are perfectly efficient. More generally, for γ between 0 and 1, the lower bound in Theorem 3 increases monotonically as γ increases, since the fraction within the $\min$ increases in γ for any fixed $t \in [0,1]$:

$$\frac{1 + t\ln t - \gamma t(1 + \ln t)}{2 - t - \gamma + (1-\gamma)t\ln t} = \frac{(1-t)^2}{(2 - t + t\ln t)(1 + t\ln t) - (1 + t\ln t)^2\gamma} + \frac{t + t\ln t}{1 + t\ln t},$$

where $1 + t\ln t \in (0,1]$ and $2 - t + t\ln t \geq 1 \geq \gamma \geq \gamma \cdot (1 + t\ln t)$.

In order words, the efficiency of the first-price auction with machine-learned reserves improves as the reserves become more and more accurate. To better illustrate the rate at which the efficiency improves, we plot our lower bound against the accuracy γ in Figure 1. As the plot shows, our lower bound improves approximately linearly as the accuracy γ increases from 0 to 1.

[5] As we focus on the single-slot environment in this paper, for ease of presentation, our assumption here is weaker than Balseiro et al. [2021b] as they assume for each bidder i and each auction j, there is a signal $s_{i,j} \in [\gamma \cdot v_{i,j}, v_{i,j}]$, and setting $s_j := \max_i s_{i,j}$ gives the signals which can be used in our setting. Our result can be extended to the multi-slot environment by adopting the assumption from Balseiro et al. [2021b].

https://doi.org/10.52202/079017-4420

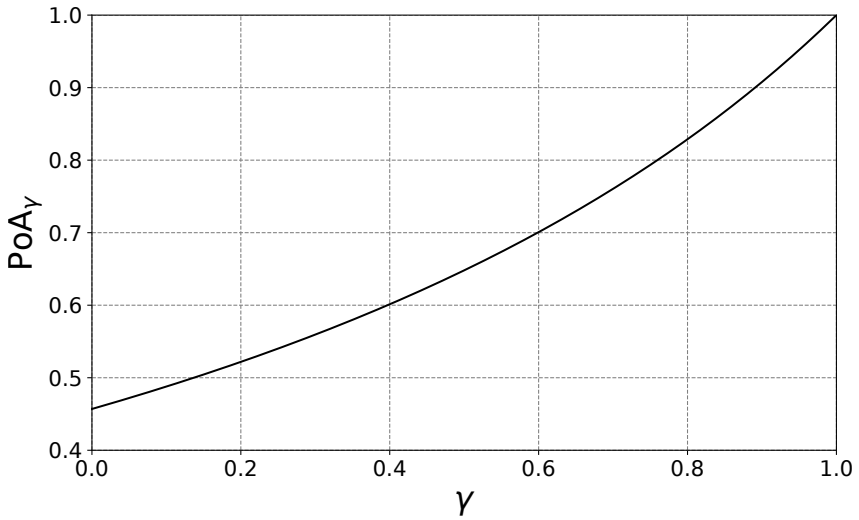

Figure 1: The efficiency of the first-price auction with machine-learned reserves when the accuracy of the reserves is $\gamma \in [0, 1]$.

To prove Theorem 3, the plan is to apply Lemma 4, which provides a lower bound on the welfare in any equilibrium, and consider tradeoffs between the 4 quantities involved in the lower bound. The key difference is that we obtain better tradeoffs with machine-learned reserves. In the rest of the section, we fix n, m, B_u, B_v, and $\{v_{i,j}\}$. With machine-learned reserves, the tradeoff given by Lemma 1 improves to the following (note the factor of $(1-\gamma)$ on the left hand side).

Lemma 7. *In auctions with machine-learned reserves of accuracy $\gamma \in [0, 1]$, for any $\{\boldsymbol{b}_i\}_i$ that form an equilibrium,*

$$(1-\gamma) \cdot \sum_{i \in B_v} \mathbb{E}[\mathsf{val}_i] + \sum_{i,j:\mathsf{rw}(j) \in B_v} \mathbb{E}[p_{i,j}] \geq \sum_{j:\mathsf{rw}(j) \in B_v} \max_i v_{i,j}.$$

We note that the above lemma also implies a lower bound of $1/(2-\gamma)$ on the PoA of the first-price auction with machine-learned reserves in the full autobidding world in which all bidders are value maximizers. When $\gamma \to 0$ (i.e., reserves are extremely unreliable), the bound degrades to $1/2$; when $\gamma \to 1$ (i.e., reserves are extremely accurate), the bound approaches 1, which means the auctions are efficient.

Corollary 1. *For any n, m, $\{v_{i,j}\}$ and $\gamma \in [0, 1]$, $\mathsf{PoA}_\gamma(n, m, B_u = \emptyset, B_v = [n], \{v_{i,j}\}) \geq 1/(2-\gamma)$.*

For utility maximizers, we adapt the local analysis, and obtain the following improved tradeoff between value and payment.

Lemma 8. *In auctions with machine-learned reserves of accuracy $\gamma \in [0, 1]$, for any $\{\boldsymbol{b}_i\}_i$ that form an equilibrium, and any bidder i and auction j where $\mathsf{rw}(j) = i$, there exists $x \in [0, 1]$ such that*

$$\mathbb{E}[\mathsf{val}_{i,j}] = x \cdot v_{i,j} \quad \textit{and} \quad \sum_{i' \neq i} \mathbb{E}[p_{i',j}] \geq (1 - x + (1-\gamma) x \ln x) \cdot v_{i,j}.$$

We can now prove Theorem 3 by plugging Lemma 7 and Lemma 8 into Lemma 4.

Acknowledgments and Disclosure of Funding

We thank anonymous reviewers for their helpful feedback.

References

Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. Autobidding with constraints. In *International Conference on Web and Internet Economics*, pages 17–30. Springer, 2019.

 https://doi.org/10.52202/079017-4420

Gagan Aggarwal, Ashwinkumar Badanidiyuru, Santiago R Balseiro, Kshipra Bhawalkar, Yuan Deng, Zhe Feng, Gagan Goel, Christopher Liaw, Haihao Lu, Mohammad Mahdian, et al. Auto-bidding and auctions in online advertising: A survey. *ACM SIGecom Exchanges*, 22(1):159–183, 2024.

Moshe Babaioff, Richard Cole, Jason Hartline, Nicole Immorlica, and Brendan Lucier. Non-quasi-linear agents in quasi-linear mechanisms. In *Proceedings of 12th Innovations in Theoretical Computer Science*, 2021.

Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, 2021a.

Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Optimal mechanisms for value maximizers with budget constraints via target clipping. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 2022. URL `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4081457`.

Santiago R. Balseiro, Yuan Deng, Jieming Mao, Vahab S. Mirrokni, and Song Zuo. Robust auction design in the auto-bidding world. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 17777–17788, 2021b. URL `https://proceedings.neurips.cc/paper/2021/hash/948f847055c6bf156997ce9fb59919be-Abstract.html`.

George Christodoulou, Annamária Kovács, Alkmini Sgouritsa, and Bo Tang. Tight bounds for the price of anarchy of simultaneous first-price auctions. *ACM Transactions on Economics and Computation (TEAC)*, 4(2):1–33, 2016.

Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Nicolas E Stier-Moses, Eric Sodomka, and Christopher A Wilkens. Pacing equilibrium in first price auction markets. *Management Science*, 68(12):8515–8535, 2022.

Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Towards efficient auctions in an auto-bidding world. In *Proceedings of The Web Conference 2021*, 2021.

Yuan Deng, Mohammad Mahdian, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. Efficiency of the generalized second-price auction for value maximizers. In *Proceedings of the ACM on Web Conference 2024*, pages 46–56, 2024.

Shelagh Dolan. US Programmatic Digital Display Ad Spending Will Grow Despite Pandemic-Related Recession. https://www.emarketer.com/content/us-programmatic-digital-display-ad-spending-will-grow-despite-pandemic-related-recession, 2020. Accessed: 9 Jul 2022.

Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 201–210, 2013.

Michal Feldman, Brendan Lucier, and Noam Nisan. Correlated and coarse equilibria of single-item auctions. In *International Conference on Web and Internet Economics*, pages 131–144. Springer, 2016.

Yiding Feng, Brendan Lucier, and Aleksandrs Slivkins. Strategic budget selection in a competitive autobidding world. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 213–224, 2024.

Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Budget pacing in repeated auctions: Regret and efficiency without convergence. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251, page 52. Schloss Dagstuhl–Leibniz-Zentrum f ür Informatik, 2023.

Jason Hartline, Darrell Hoy, and Sam Taggart. Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 693–710, 2014.

Darrell Hoy, Samuel Taggart, and Zihe Wang. A tighter welfare guarantee for first-price auctions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 132–137, 2018.

Yaonan Jin and Pinyan Lu. First price auction is 1–1/e2 efficient. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 179–187. IEEE Computer Society, 2022.

Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science*, pages 404–413. Springer, 1999.

Christopher Liaw, Aranyak Mehta, and Andres Perlroth. Efficiency of non-truthful auctions under auto-bidding, 2022. URL `https://arxiv.org/abs/2207.03630`.

Christopher Liaw, Aranyak Mehta, and Wennan Zhu. Efficiency of non-truthful auctions in auto-bidding with budget constraints. In *Proceedings of the ACM on Web Conference 2024*, pages 223–234, 2024.

Eric Maskin and John Riley. Equilibrium in sealed high bid auctions. *The Review of Economic Studies*, 67(3):439–454, 2000.

Aranyak Sitanshu Mehta. Auction design in an auto-bidding setting: Randomization improves efficiency beyond vcg. In *Proceedings of the ACM Web Conference 2022*, pages 173–181, 2022.

Renato Paes Leme, Balasubramanian Sivan, and Yifeng Teng. Why do competitive markets converge to first-price auctions? In *Proceedings of The Web Conference 2020*, pages 596–605, 2020.

Tim Roughgarden, Vasilis Syrgkanis, and Eva Tardos. The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101, 2017.

Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 211–220, 2013.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

A More on Liquid Welfare

Following prior work on autobidding [Aggarwal et al., 2019, Deng et al., 2021, Balseiro et al., 2021a, Liaw et al., 2022], we consider the liquid welfare as the objective throughout this paper. The liquid welfare is the maximum amount of payment that can be extracted from the bidders given a certain allocation. Formally, assuming each bidder i's target ROI ratio is τ_i, and each bidder i's value in each auction j is $v_{i,j}$, the liquid welfare under allocation $\{x_{i,j}\}$ is defined as follows:

$$\sum_{i,j} x_{i,j} \cdot v_{i,j}/\tau_i.$$

In particular, we let $\tau_i = 1$ for each utility maximizer i, which simply means i is individually rational. One can check the above definition is consistent with the classical social welfare when all bidders are utility maximizers, and it is also consistent with prior work on autobidding when all bidders are value maximizers. The benchmark we consider in this paper is the first-best liquid welfare, i.e.,

$$\sum_{j} \max_{i} v_{i,j}/\tau_i.$$

This can be achieved by allocating to the bidder i with the highest target-adjusted value $v_{i,j}/\tau_i$ in each auction j. Observe that the quantity that actually matters in the (first-best) liquid welfare is the target-adjusted value $v_{i,j}/\tau_i$. Therefore, without loss of generality, we can replace each bidder i's value in each auction j with the target-adjusted value $v_{i,j}/\tau_i$, and assume all bidders' target ROI ratios are 1.

The liquid welfare makes more sense than the classical social welfare in the presence of value maximizers, because a bidder with a high target ROI ratio and low values *behaves* in the same way as a bidder with a low target ROI ratio and high values, as long as their target-adjusted values are the same. There is no way for the mechanism to distinguish between the two types of bidders, and thus, it is infeasible to try to maximize the classical social welfare.

 https://doi.org/10.52202/079017-4420

B Omitted Proofs

Proof of Theorem 1. First we provide a simple hard instance, which implies an upper bound of $1/2$ on the PoA. Consider an instance with $n = 2$ bidders and $m = 2$ auctions, where both bidders are value maximizers. We set $v_{1,1} = 1$, $v_{2,2} = 1 - \varepsilon$, and $v_{1,2} = v_{2,1} = 0$, where $\varepsilon > 0$ is a small quantity to be fixed later. Consider the following deterministic bidding strategies: $b_{1,1} = b_{2,1} = 0$ (so bidder 1 wins in auction 1 by tiebreaking and pays 0), $b_{1,2} = 1$, and $b_{2,2} = 0$ (so bidder 1 wins in auction 2 and pays 1). First we verify this is in fact an equilibrium. Observe that the ROI constraints of both bidders are satisfied. Bidder 1 does not want to deviate from $\boldsymbol{b}_1$, since bidder 1 already gets the maximum value possible, i.e., 1. Bidder 2 would ideally want to bid higher (possibly randomly) in auction 2 in order to win with some probability and get positive value. However, given $b_{1,2} = 1$, bidder 2 would have to bid at least 1 in order to win, which means $p_{2,2} \geq 1$ whenever $x_{2,2} = 1$. This would violate bidder 2's ROI constraint. So both bidders' strategies are best responses, and therefore form an equilibrium. On the other hand, the optimal welfare is $v_{1,1} + v_{2,2} = 2 - \varepsilon$, and the welfare resulting from the above equilibrium is only $v_{1,1} = 1$, which means the PoA is at most $1/(2 - \varepsilon)$. Letting $\varepsilon \to 0$ yields an upper bound of $1/2$.

Now we focus on lower bounding the instance-wise PoA fixing any n, m, and $\{v_{i,j}\}$. Consider any equilibrium, induced by bidding strategies $\{\boldsymbol{b}_i\}_i$. Our task is to show that

$$2 \sum_i \mathbb{E}[\mathsf{val}_i] \geq \sum_j \max_i v_{i,j},$$

where the expectations, as well as all expectations in the rest of the proof, are taken over $\{\boldsymbol{b}_i\}_i$. For each auction j, let $\mathsf{rw}(j) = \operatorname{argmax}_i v_{i,j}$ be the "rightful winner" of auction j, where ties are broken in favor of the bidder with the smaller index. Now fix a bidder i, and consider the following proxy bidding strategy $\boldsymbol{b}'_i$ of bidder i: $b'_{i,j} = v_{i,j}$ in each auction j, i.e., the proxy bidding strategy $\boldsymbol{b}'_i$ is a strategy in which bidder i bids truthfully. Then, for each auction j where $i = \mathsf{rw}(j)$, the probability that i wins in auction j under the proxy bidding strategy $\boldsymbol{b}'_i$, fixing all other bidders' strategies $\boldsymbol{b}_{-i}$, is

$$\Pr_{\boldsymbol{b}_{-i}}[x_{i,j}(b'_{i,j}, \boldsymbol{b}_{-i,j}) = 1] = \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} \leq b'_{i,j}\right] = \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} \leq v_{i,j}\right].$$

Also, observe that i's ROI constraint is satisfied under the proxy bidding strategy $\boldsymbol{b}'_i$, since under the proxy bidding strategy $\boldsymbol{b}'_i$, the following holds deterministically:

$$\mathsf{val}_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) - p_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i}) = \sum_j x_{i,j}(b'_{i,j}, \boldsymbol{b}_{-i,j}) \cdot (v_{i,j} - b'_{i,j}) = \sum_j x_{i,j}(b'_{i,j}, \boldsymbol{b}_{-i,j}) \cdot (v_{i,j} - v_{i,j}) = 0.$$

Taking the expectation over $\boldsymbol{b}_{-i}$, we see that i's ROI constraint is satisfied (and is in fact binding). The above observations imply the following fact: There exists a bidding strategy (i.e., $\boldsymbol{b}'_i$), under which i's ROI constraint is satisfied, and the expected value i receives is

$$\mathbb{E}[\mathsf{val}_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})] \geq \sum_{j : i = \mathsf{rw}(j)} \Pr_{\boldsymbol{b}_{-i}}[x_{i,j}(b'_{i,j}, \boldsymbol{b}_{-i,j}) = 1] \cdot v_{i,j} = \sum_{j : i = \mathsf{rw}(j)} \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} \leq v_{i,j}\right] \cdot v_{i,j}.$$

Observe that since i is a value maximizer and $\boldsymbol{b}_i$ is a best response to $\boldsymbol{b}_{-i}$, we must have

$$\mathbb{E}[\mathsf{val}_i(\boldsymbol{b}_i, \boldsymbol{b}_{-i})] \geq \mathbb{E}[\mathsf{val}_i(\boldsymbol{b}'_i, \boldsymbol{b}_{-i})] \geq \sum_{j : i = \mathsf{rw}(j)} \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} \leq v_{i,j}\right] \cdot v_{i,j}.$$

On the other hand, fix a bidder i and consider the total payment collected in any auction j with $\mathsf{rw}(j) = i$, under $(\boldsymbol{b}_i, \boldsymbol{b}_{-i})$. By the properties of the first-price auction, we have

$$\begin{aligned}\sum_{i'} \mathbb{E}[p_{i',j}(b_{i,j}, \boldsymbol{b}_{-i,j})] &= \mathbb{E}\left[\max_{i'} b_{i',j}\right] \geq \mathbb{E}\left[\max_{i' \neq i} b_{i',j}\right] \\ &\geq \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} > v_{i,j}\right] \cdot v_{i,j} = \left(1 - \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} \leq v_{i,j}\right]\right) \cdot v_{i,j}.\end{aligned}$$

Summing over all auctions j with $\mathsf{rw}(j) = i$, we have

$$\sum_{i', j : i = \mathsf{rw}(j)} \mathbb{E}[p_{i',j}(\boldsymbol{b}_i, \boldsymbol{b}_{-i})] \geq \sum_{j : i = \mathsf{rw}(j)} \left(1 - \Pr_{\boldsymbol{b}_{-i}}\left[\max_{i' \neq i} b_{i',j} \leq v_{i,j}\right]\right) \cdot v_{i,j}.$$

Combining this with the bound on the total expected value i receives, we get

$$\mathbb{E}[\mathsf{val}_i(\boldsymbol{b}_i, \boldsymbol{b}_{-i})] + \sum_{i', j : i = \mathsf{rw}(j)} \mathbb{E}[p_{i',j}(\boldsymbol{b}_i, \boldsymbol{b}_{-i})] \geq \sum_{j : i = \mathsf{rw}(j)} v_{i,j}.$$

Further summing over i (and omitting dependence on $(\boldsymbol{b}_i, \boldsymbol{b}_{-i})$), the above gives

$$\sum_i \mathbb{E}[\mathsf{val}_i] + \sum_i \mathbb{E}[p_i] \geq \sum_j \max_i v_{i,j}.$$

Finally, observe that for any bidder i, the ROI constraint implies $\mathbb{E}[p_i] \leq \mathbb{E}[\mathsf{val}_i]$, and therefore we can conclude the proof with

$$2 \sum_i \mathbb{E}[\mathsf{val}_i] \geq \sum_i \mathbb{E}[\mathsf{val}_i] + \sum_i \mathbb{E}[p_i] \geq \sum_j \max_i v_{i,j}. \qquad \square$$

Proof of Lemma 2. Suppose towards a contradiction that there exists $w > b \cdot F(b)$ such that $H(w) - H(b \cdot F(b)) > w - b \cdot F(b)$, or equivalently, $H(w) - w > H(b \cdot F(b)) - b \cdot F(b)$. Let $b'_{i,j}$ be a possibly randomized bidding strategy that realizes the point $(w, H(w))$, i.e., when using $b'_{i,j}$, i receives value $H(w)$ and pays w in expectation in auction j. Then, the utility that i receives from j when using the new strategy $b'_{i,j}$ is $H(w) - w$. On the other hand, consider the deterministic bidding strategy $b''_{i,j} = b$. Since b is in the support of $b_{i,j}$, $b''_{i,j}$ must also be a best response. However, under $b''_{i,j}$, i receives utility $H(b \cdot F(b)) - b \cdot F(b)$, which is strictly worse than the utility i receives when using the alternative strategy $b'_{i,j}$. In other words, $b''_{i,j}$ cannot be a best response, a contradiction. $\square$

Proof of Lemma 3. Let $x = \mathbb{E}[F(b_{i,j})]$. Clearly $x \in [0, 1]$ and $\mathbb{E}[\mathsf{val}_{i,j}] = x \cdot v$. Our remaining task is to show

$$\sum_{i' \neq i} \mathbb{E}[p_{i',j}] \geq (1 - x + x \ln x) \cdot v.$$

Recall that

$$\sum_{i' \neq i} \mathbb{E}[p_{i',j}] = \mathop{\mathbb{E}}_{b_{i,j}} \left[b_{i,j} \cdot (1 - F(b_{i,j})) + \int_{b_{i,j}}^{\infty} (1 - F(t)) \, \mathrm{d}t \right].$$

As a result, in order to lower bound the payment, we want to pointwise upper bound F. Fix any b in the support of the best-response strategy $b_{i,j}$, we must have

$$H(b \cdot F(b)) = G(b \cdot F(b)) = v \cdot F(b).$$

By Lemma 2, for any $w \geq b \cdot F(b)$, we have $H(w) - H(b \cdot F(b)) \leq w - b \cdot F(b)$, which implies that $H(w) \leq H(b \cdot F(b)) + w - b \cdot F(b)$. Also, since H is the concave envelope of G, the same upper bound applies to G, and we have

$$G(w) \leq H(w) \leq H(b \cdot F(b)) + w - b \cdot F(b) = v \cdot F(b) + w - b \cdot F(b).$$

Due to the monotonicity of $F(x)$, we have that for any $t \geq b$, we have $t \cdot F(t) \geq b \cdot F(b)$. Therefore, given the relation between G and F, the above implies that for any $t \in [b, v)$ (note that this range might be empty),

$$v \cdot F(t) \leq G(t \cdot F(t)) \leq v \cdot F(b) + t \cdot F(t) - b \cdot F(b) \implies F(t) \leq \frac{(v - b) \cdot F(b)}{v - t}.$$

We next bound the payment conditioned on $b_{i,j} = b$. We consider two possible cases: $b \geq v$ and $b < v$. If $b \geq v$, we simply have

$$\begin{aligned} \sum_{i' \neq i} \mathbb{E}[p_{i',j} \mid b_{i,j} = b] &= b \cdot (1 - F(b)) + \int_b^{\infty} (1 - F(t)) \, \mathrm{d}t \\ &\geq b \cdot (1 - F(b)) \geq (1 - F(b)) \cdot v \geq (1 - F(b) + F(b) \ln F(b)) \cdot v. \end{aligned}$$

 https://doi.org/10.52202/079017-4420

In particular, the last inequality is because $F(b)\ln F(b) \le 0$ when $F(b) \in [0,1]$. If $b < v$, we must have $b = v - (v-b)\cdot 1 \le v - (v-b)\cdot F(b)$, and as a result,

$$\begin{aligned}
\sum_{i' \neq i} \mathbb{E}[p_{i',j} \mid b_{i,j} = b] &= b \cdot (1 - F(b)) + \int_b^\infty (1 - F(t))\,\mathrm{d}t \\
&\ge b \cdot (1 - F(b)) + \int_b^{v-(v-b)\cdot F(b)} \left(1 - \frac{(v-b)\cdot F(b)}{v-t}\right)\mathrm{d}t \\
&= (1 - F(b)) \cdot v + (v-b)\cdot F(b) \cdot \ln(v-t)\Big|_b^{v-(v-b)\cdot F(b)} \\
&= (1 - F(b)) \cdot v + (v-b)\cdot F(b) \cdot \ln F(b) \\
&\ge (1 - F(b) + F(b)\ln F(b)) \cdot v.
\end{aligned}$$

Combining the two cases and further taking the expectation over $b_{i,j}$, we get

$$\sum_{i' \neq i} \mathbb{E}[p_{i',j}] \ge \mathop{\mathbb{E}}_{b_{i,j}}\left[(1 - F(b_{i,j}) + F(b_{i,j})\ln F(b_{i,j}))\cdot v\right] \ge (1 - x + x\ln x)\cdot v,$$

where the last inequality is due to Jensen's inequality, the fact that $z \mapsto 1 - z + z\ln z$ is convex on $[0,1]$, and we chose $x = \mathbb{E}[F(b_{i,j})]$. □

Proof of Lemma 4. Observe that

$$\begin{aligned}
\sum_i \mathbb{E}[\mathsf{val}_i] &= \sum_{i \in B_v} \mathbb{E}[\mathsf{val}_i] + \sum_{i \in B_u} \mathbb{E}[\mathsf{val}_i] \\
&= \max\left\{\sum_{i\in B_v} \mathbb{E}[\mathsf{val}_i], \sum_{i \in B_v} \mathbb{E}[p_i]\right\} + \sum_{i \in B_u} \mathbb{E}[\mathsf{val}_i] \\
&= \max\left\{\sum_{i\in B_v} \mathbb{E}[\mathsf{val}_i], \sum_{i \in B_v} \mathbb{E}[p_i]\right\} + \sum_{i \in B_u, j:\mathsf{rw}(j)=i} \mathbb{E}[\mathsf{val}_{i,j}] + \sum_{i \in B_u, j:\mathsf{rw}(j)\neq i} \mathbb{E}[\mathsf{val}_{i,j}] \\
&= \max\left\{A, \sum_{i \in B_v} \mathbb{E}[p_i]\right\} + B + \sum_{i \in B_u, j:\mathsf{rw}(j)\neq i} \mathbb{E}[\mathsf{val}_{i,j}].
\end{aligned}$$

Note that for each $i \in B_u$ and j, we always have $\mathbb{E}[\mathsf{val}_{i,j}] \ge \mathbb{E}[p_{i,j}]$ since i is best responding and the utility i receives in j is at least 0. Therefore, we can further relax the above into the following:

$$\begin{aligned}
\sum_i \mathbb{E}[\mathsf{val}_i] &\ge \max\left\{A, \sum_{i \in B_v} \mathbb{E}[p_i]\right\} + B + \sum_{i \in B_u, j:\mathsf{rw}(j)\neq i} \mathbb{E}[p_{i,j}] \\
&\ge \max\left\{A, \sum_{i \in B_v} \mathbb{E}[p_i] + \sum_{i \in B_u, j:\mathsf{rw}(j)\neq i} \mathbb{E}[p_{i,j}]\right\} + B.
\end{aligned}$$

Rearranging the sums within the $\max$, we have that

$$\begin{aligned}
\sum_{i \in B_v} \mathbb{E}[p_i] + \sum_{i \in B_u, j:\mathsf{rw}(j)\neq i} \mathbb{E}[p_{i,j}] &= \sum_i \mathbb{E}[p_i] - \sum_{i \in B_u, j:\mathsf{rw}(j)=i} \mathbb{E}[p_{i,j}] \\
&= \sum_{i,j:\mathsf{rw}(j)\in B_v} \mathbb{E}[p_{i,j}] + \sum_{i,j:\mathsf{rw}(j)\in B_u \text{ and } \mathsf{rw}(j)\neq i} \mathbb{E}[p_{i,j}] \\
&= C + D.
\end{aligned}$$

Substituting in the $\max$ gives the bound to be proved. □

Proof of Lemma 5. Fix any n, m, B_u, B_v, and $\{v_{i,j}\}$. Consider any bidding strategies $\{\boldsymbol{b}_i\}_i$ that form an equilibrium. Let

$$V_1 = \sum_{j:\mathsf{rw}(j)\in B_v} \max_i v_{i,j}, \quad V_2 = \sum_{j:\mathsf{rw}(j)\in B_u} \max_i v_{i,j}.$$

By Lemma 4, we have

$$\frac{\sum_i \mathbb{E}[\mathsf{val}_i]}{\sum_j \max_i v_{i,j}} \geq \frac{\max\{A, C+D\} + B}{V_1 + V_2}.$$

By Lemma 1, we have $A + C \geq V_1$. So there must be some $x \in [0,1]$, such that $A \geq x \cdot V_1$ and $C \geq (1-x) \cdot V_1$. As for B and D, we apply Lemma 3. For each j where $\mathsf{rw}(j) \in B_u$, Lemma 3 states that there exists some $y_j \in [0,1]$, such that

$$\mathbb{E}[\mathsf{val}_{\mathsf{rw}(j),j}] = y_j \cdot \max_i v_{i,j}, \quad \sum_{i \neq \mathsf{rw}(j)} \mathbb{E}[p_{i,j}] \geq (1 - y_j + y_j \ln y_j) \cdot \max_i v_{i,j}.$$

Let $y = B/V_2$. Since $B = \sum_{j:\mathsf{rw}(j) \in B_u} \mathbb{E}[\mathsf{val}_{\mathsf{rw}(j),j}] \leq \sum_{j:\mathsf{rw}(j) \in B_u} \max_i v_{i,j} = V_2$, we have $y \in [0,1]$. In particular,

$$y = \frac{\sum_{j:\mathsf{rw}(j) \in B_u} \mathbb{E}[\mathsf{val}_{\mathsf{rw}(j),j}]}{\sum_{j:\mathsf{rw}(j) \in B_u} \max_i v_{i,j}} = \frac{\sum_{j:\mathsf{rw}(j) \in B_u} y_j \cdot \max_i v_{i,j}}{\sum_{j:\mathsf{rw}(j) \in B_u} \max_i v_{i,j}}.$$

Since $t \mapsto 1 - t + t \ln t$ is convex, by Jensen's inequality, we must have

$$\begin{aligned} D &= \sum_{j:\mathsf{rw}(j) \in B_u} \sum_{i \neq \mathsf{rw}(j)} \mathbb{E}[p_{i,j}] \geq \sum_{j:\mathsf{rw}(j) \in B_u} (1 - y_j + y_j \ln y_j) \cdot \max_i v_{i,j} \\ &\geq (1 - y + y \ln y) \sum_{j:\mathsf{rw}(j) \in B_u} \max_i v_{i,j} = (1 - y + y \ln y) \cdot V_2. \end{aligned}$$

Plugging in the above derived lower bounds for A, B, C, and D, we have

$$\frac{\sum_i \mathbb{E}[\mathsf{val}_i]}{\sum_j \max_i v_{i,j}} \geq \frac{\max\{A, C+D\} + B}{V_1 + V_2} \geq \frac{\max\{x \cdot V_1, (1-x) \cdot V_1 + (1 - y + y \ln y) \cdot V_2\} + y \cdot V_2}{V_1 + V_2},$$

where $V_1, V_2 \geq 0$, $V_1 + V_2 > 0$ and $x, y \in [0,1]$. In other words, we must have

$$\mathsf{PoA}(n, m, B_u, B_v, \{v_{i,j}\}) \geq \inf_{V_1, V_2 \geq 0, V_1 + V_2 > 0, x, y \in [0,1]} \frac{\max\{x \cdot V_1, (1-x) \cdot V_1 + (1 - y + y \ln y) \cdot V_2\} + y \cdot V_2}{V_1 + V_2}.$$

A careful analysis (see Lemma 9 in Appendix C) of the right-hand-side shows that it equals to

$$\min_{y \in [0,1]} \frac{1 + y \ln y}{2 - y + y \ln y} \approx 0.457,$$

which is achieved when $x = 1$ and $V_1 = (1 - y + y \ln y) \cdot V_2$. This concludes the proof. □

Proof of Lemma 6. Let $B_u = \{1\}$ and $B_v = \{2\}$. We choose $v_{1,1} = 1$ and $v_{1,2} = v_{2,1} = 0$, and leave $v_{2,2} > 0$ to be fixed later. The idea is to let bidder 2 win in auction 2 for free, and therefore bidder 2 gains enough slackness in the ROI constraint and can therefore compete with bidder 1 in auction 1. To be specific, in order for bidder 2 to win for free in auction 2, we choose $b_{1,2} = b_{2,2} = 0$ so that bidder 2 wins by tiebreaking and pays 0. We want bidder 1's value-payment frontier in auction 1 to have the following shape: the curve starts at $(0, t)$ for some $t \in [0,1]$ (to be fixed later), and then increases linearly with slope precisely 1, until it hits the upper bound, i.e., when the y-axis becomes $v_{1,1} = 1$. This ensures that bidder 1 bidding $b_{1,1} = 0$ deterministically is a best response. Translating this to bidder 2's bidding strategy, we would like that for any $b \geq 0$,

$$\Pr[b_{2,1} \leq b] = \min\left\{1, \frac{t}{1-b}\right\}.$$

One can show that the expected payment that bidder 2 makes in auction 1 is $1 - t + t \ln t$. In order to satisfy bidder 2's ROI constraint, we choose $v_{2,2} = 1 - t + t \ln t$. One may check that the above bidding strategies in fact form an equilibrium. Note that $\mathbb{E}[\mathsf{val}_1] = \Pr[b_{2,1} = 0] \cdot v_{1,1} = t$, and $\mathbb{E}[\mathsf{val}_2] = v_{2,2} = 1 - t + t \ln t$.

Considering the ratio between the optimal welfare and the welfare resulting from the above equilibrium, we have

$$\mathsf{PoA}(n, m, B_u, B_v, \{v_{i,j}\}) \leq \frac{\mathbb{E}[\mathsf{val}_1 + \mathsf{val}_2]}{v_{1,1} + v_{2,2}} = \frac{1 + t \ln t}{2 - t + t \ln t},$$

 https://doi.org/10.52202/079017-4420

for any $t \in [0,1]$. Minimizing over t gives

$$\mathsf{PoA}(n, m, B_u, B_v, \{v_{i,j}\}) \le \min_{t \in [0,1]} \frac{1 + t \ln t}{2 - t + t \ln t}. \qquad \square$$

Proof of Lemma 7. The proof is similar to that of Lemma 1 at a high level, except that the expected payment is higher, because in each auction j, the rightful winner $\mathsf{rw}(j)$ must bid at least r_j (which is a constant fraction of $v_{\mathsf{rw}(j),j}$) unless the highest competing bid is at least $v_{\mathsf{rw}(j),j}$ with probability 1.

Again fix any equilibrium strategies $\{\boldsymbol{b}_i\}$. By the same argument as in the proof of Lemma 1, for any value maximizer $i \in B_v$, we have

$$\mathbb{E}[\mathsf{val}_i] \ge \sum_{j : i = \mathsf{rw}(j)} v_{i,j} \cdot \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right].$$

Also, observe that for any value maximizer $i \in B_v$ and auction j where $\mathsf{rw}(j) = i$, at least one of the following properties must hold: Either (1) $b_{i,j} \ge r_{i,j}$ deterministically, or (2) $\Pr[\max_{i' \ne i} b_{i',j} \ge v_{i,j}] = 1$. This is because if neither of the two holds, then i for any $b < r_j$ in the support of $b_{i,j}$, i would be better off by moving probability mass from b to $v_{i,j}$, which preserves feasibility and strictly increases the value i receives from j. For each j where (1) holds, we have

$$\begin{aligned}
\sum_{i'} \mathbb{E}[p_{i',j}] &\ge v_{i,j} \cdot \left(1 - \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right]\right) + r_j \cdot \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right] \\
&\ge v_{i,j} \cdot \left(1 - \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right]\right) + \gamma \cdot v_{i,j} \cdot \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right] \\
&= v_{i,j} - (1 - \gamma) \cdot v_{i,j} \cdot \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right].
\end{aligned}$$

For each j where (2) holds, we immediately know that

$$\sum_{i'} \mathbb{E}[p_{i',j}] \ge v_{i,j} \cdot \Pr\left[\max_{i' \ne i} b_{i',j} \ge v_{i,j}\right] = v_{i,j} \ge v_{i,j} - (1 - \gamma) \cdot v_{i,j} \cdot \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right].$$

So combining the two cases and summing over j where $i = \mathsf{rw}(j)$,

$$\sum_{i', j : \mathsf{rw}(j) = i} \mathbb{E}[p_{i',j}] \ge \sum_{j : \mathsf{rw}(j) = i} \left(v_{i,j} - (1 - \gamma) \cdot v_{i,j} \cdot \Pr\left[\max_{i' \ne i} b_{i',j} \le v_{i,j}\right]\right).$$

Combining this with the lower bound on $\mathbb{E}[\mathsf{val}_i]$ above, we get

$$(1 - \gamma) \cdot \mathbb{E}[\mathsf{val}_i] + \sum_{i', j : \mathsf{rw}(j) = i} \mathbb{E}[p_{i',j}] \ge \sum_{j : \mathsf{rw}(j) = i} v_{i,j}.$$

Further summing over $i \in B_v$ gives the inequality to be proved. $\square$

Proof of Lemma 8. The first part of the reasoning is almost identical to the proof of Lemma 3: For any b in the support of $b_{i,j}$, Lemma 2 gives us a way to upper bound the CDF of the highest other bid, and subsequently to lower bound the payment made by other bidders. Without repeating the entire argument, we start directly from this lower bound. Let F be the CDF of the highest other bid, i.e., for any $t \ge 0$,

$$F(t) = \Pr\left[\max_{i' \ne i} b_{i',j} \le t\right].$$

When $b \ge v$, we have

$$\sum_{i' \ne i} \mathbb{E}[p_{i',j} \mid b_{i,j} = b] \ge (1 - F(b)) \cdot v \ge (1 - F(b) + (1 - \gamma) \cdot F(b) \ln F(b)) \cdot v_{i,j}.$$

When $b < v$, again we must have one of the following two: Either (1) $b \ge r_j \ge \gamma \cdot v_{i,j}$, or (2) $F(v_{i,j}) = 1$. When (1) holds, we have

$$\sum_{i' \ne i} \mathbb{E}[p_{i',j} \mid b_{i,j} = b] \ge (1 - F(b)) \cdot v_{i,j} + F(b) \ln F(b) \cdot (v_{i,j} - b) \ge (1 - F(b) + (1 - \gamma) F(b) \ln F(b)) \cdot v_{i,j}.$$

When (2) holds, clearly we have

$$\sum_{i' \neq i} \mathbb{E}[p_{i',j} \mid b_{i,j} = b] \geq v_{i,j} \cdot F(v_{i,j}) = v_{i,j} \geq (1 - F(b) + (1-\gamma)F(b) \ln F(b)) \cdot v_{i,j}.$$

Combining these cases and taking the expectation over $b_{i,j}$ (again letting $x = \mathbb{E}[F(b_{i,j})]$), we get

$$\sum_{i' \neq i} \mathbb{E}[p_{i',j}] \geq \mathbb{E}[(1 - F(b_{i,j}) + (1-\gamma) \cdot F(b_{i,j}) \ln F(b_{i,j})) \cdot v_{i,j}] \geq (1 - x + (1-\gamma) \cdot x \ln x) \cdot v_{i,j},$$

where the last inequality from Jensen's inequality since $z \mapsto 1 - z + (1-\gamma) \cdot z \ln z$ is convex on $[0,1]$ for any $\gamma \in [0,1]$. □

Proof of Theorem 3. Reproducing the argument in the proof of Lemma 5 (but with the improved tradeoffs given by Lemma 7 and Lemma 8), we can lower bound $\mathsf{PoA}_\gamma(n, m, B_u, B_v, \{v_{i,j}\})$ by

$$\inf_{V_1, V_2 \geq 0, V_1 + V_2 > 0, x, y \in [0,1]} \frac{\max\{x \cdot V_1, (1 - (1-\gamma)x) \cdot V_1 + (1 - y + (1-\gamma) y \ln y) \cdot V_2\} + y \cdot V_2}{V_1 + V_2}.$$

A careful analysis (see Lemma 10 in Appendix C) of the infimum shows that it equals to

$$\min_{y \in [0,1]} \frac{1 + y \ln y - \gamma y (1 + \ln y)}{2 - y - \gamma + (1-\gamma) y \ln y}.$$

□

C Analysis of the Infimum

Lemma 9.

$$\inf_{V_1, V_2 \geq 0, V_1 + V_2 > 0, x, y \in [0,1]} \frac{\max\{x \cdot V_1, (1-x) \cdot V_1 + (1 - y + y \ln y) \cdot V_2\} + y \cdot V_2}{V_1 + V_2} = \min_{y \in [0,1]} \frac{1 + y \ln y}{2 - y + y \ln y}.$$

Proof. Without loss of generality, suppose $V_1 + V_2 = 1$, so the domain is compact, and we only need to prove that

$$\inf_{V_1 + V_2 = 1, V_1, V_2, x, y \in [0,1]} \frac{\max\{x \cdot V_1, (1-x) \cdot V_1 + (1 - y + y \ln y) \cdot V_2\} + y \cdot V_2}{V_1 + V_2} = \min_{y \in [0,1]} \frac{1 + y \ln y}{2 - y + y \ln y}.$$

Since the term inside the infimum is continuous on the domain we can replace the $\inf$ with a $\min$. In other words, there exists V_1^*, V_2^*, x^* and y^* that achieve the minimum. If $V_2 = 0$, then clearly it is lower bounded by $1/2$ (which is clearly suboptimal). So we must have $V_2^* > 0$.

Now observe that the two terms within the inner $\max$ must be equal when the minimum is achieved. This is because: if $V_1^* \geq (1 - y^* + y^* \ln y^*) \cdot V_2^*$, then x^* must equalize the two terms (any other choice of x^* would be suboptimal). Otherwise, the inner $\max$ must be $(1-x^*) \cdot V_1^* + (1 - y^* + y^* \ln y^*) \cdot V_2^*$, which means x^* must be 1, in which case V_1^* does not appear at all. But then, increasing V_1^* and decreasing V_2^* simultaneously would strictly decrease the value, which contradicts the optimality of these choices. So we must have

$$x^* \cdot V_1^* = (1-x^*) \cdot V_1^* + (1 - y^* + y^* \ln y^*) \cdot V_2^* \implies V_1^* = \frac{(1 - y^* + y^* \ln y^*) \cdot V_2^*}{2x^* - 1} \text{ and } x^* \in (1/2, 1].$$

Plugging this in, the minimum simplifies to

$$\frac{x^* \cdot (1 + y^* + y^* \ln y^*) - y^*}{2x^* - y^* + y^* \ln y^*}.$$

Observe that the denominator of the above fraction is always positive whenever $x^* \in (1/2, 1]$ and $y^* \in [0,1]$. Moreover, depending on the value of y^*, the partial derivative of the above fraction with respect to x^* is either never 0 or always 0. In the former case, it must be the case that $x^* = 1$ (any other choice cannot be optimal); in the latter case, the value of x^* does not matter, so without loss of generality we can choose $x^* = 1$. Plugging this in, the minimum simplifies to

$$\frac{1 + y^* \ln y^*}{2 - y^* + y^* \ln y^*} = \min_{y \in [0,1]} \frac{1 + y \ln y}{2 - y + y \ln y}.$$

This concludes the proof. □

https://doi.org/10.52202/079017-4420

Lemma 10.

$$\inf_{V_1,V_2\geq 0, V_1+V_2>0, x,y\in[0,1]} \frac{\max\{x\cdot V_1, (1-(1-\gamma)x)\cdot V_1 + (1-y+(1-\gamma)y\ln y)\cdot V_2\} + y\cdot V_2}{V_1+V_2}$$
$$= \min_{y\in[0,1]} \frac{1+y\ln y-\gamma y(1+\ln y)}{2-y-\gamma+(1-\gamma)y\ln y}.$$

Proof. When $\gamma = 1$, it is easy to check the above minimum is 1. In the rest of the proof, we assume $\gamma \in [0, 1)$. We break this into two cases:

- When
$$V_1 \leq \gamma\cdot V_1 + (1-y+(1-\gamma)y\ln y)\cdot V_2 \implies V_1 \leq \frac{(1-y+(1-\gamma)y\ln y)\cdot V_2}{1-\gamma},$$
the minimum becomes
$$\min \frac{(1-(1-\gamma)x)\cdot V_1 + (1+(1-\gamma)y\ln y)\cdot V_2}{V_1+V_2},$$
and the minimum must be achieved when $x = 1$, in which case it further simplifies to
$$\min \frac{\gamma\cdot V_1 + (1+(1-\gamma)y\ln y)\cdot V_2}{V_1+V_2}.$$
Since $y\ln y \in [-1/e, 0]$ and $(1+(1-\gamma)y\ln y) > \gamma$, we want V_1 to be as large as possible, which means $V_1 = (1-y+(1-\gamma)y\ln y)\cdot V_2/(1-\gamma)$. Plugging this in, the minimum becomes
$$\min \frac{1+y\ln y-\gamma y(1+\ln y)}{2-y-\gamma+(1-\gamma)y\ln y}.$$
- When
$$V_1 \geq \gamma\cdot V_1 + (1-y+(1-\gamma)y\ln y)\cdot V_2 \implies V_1 \geq \frac{(1-y+(1-\gamma)y\ln y)\cdot V_2}{1-\gamma},$$
the minimum must be achieved when the two terms within the inner $\max$ are equalized, i.e.,
$$x\cdot V_1 = (1-(1-\gamma)x)\cdot V_1 + (1-y+(1-\gamma)y\ln y)\cdot V_2 \implies V_1 = \frac{(1-y+(1-\gamma)y\ln y)\cdot V_2}{(2-\gamma)x-1}.$$
This also means we can restricted the domain of x to $x \in (1/(2-\gamma), 1]$. Plugging this in, the minimum simplifies to
$$\min \frac{x(1+y-\gamma y+(1-\gamma)y\ln y)-y}{(2-\gamma)x-y+(1-\gamma)y\ln y}.$$
Observe that the denominator is always positive for x and y in the domain, so the derivative of the fraction with respect to x is either never 0 or always 0. In the former case, the minimum can only be achieved when $x = 1$; in the latter case, the value of x does not matter, so we can without loss of generality choose $x = 1$. Plugging $x = 1$ in, the minimum simplifies to
$$\min \frac{1+y\ln y-\gamma y(1+\ln y)}{2-y-\gamma+(1-\gamma)y\ln y}.$$

Combining the two cases concludes the proof. □

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction discuss the claims proved in the paper.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses settings where the results do not apply, and also presents negative results.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

https://doi.org/10.52202/079017-4420

Justification: The full model is defined and all results are proved.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [NA]

   Justification: No experimental results.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: No data or code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [NA]

   Justification: No experimental results.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [NA]

   Justification: No experimental results.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

 https://doi.org/10.52202/079017-4420

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [NA]

   Justification: No experimental results.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: In particular, we have done our best to preserve anonymity.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

    Justification: The paper is theoretical without immediate broader impacts.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: No data or models released.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [NA]

    Justification: No existing assets used.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

 https://doi.org/10.52202/079017-4420

Answer: [NA]

Justification: No new assets introduced.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No crowdsourcing or human subjects involved.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No human subjects involved.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.