An Improved Welfare Guarantee for First-Price Auctions

DARRELL HOY
Tremor Technologies
and
SAM TAGGART
Oberlin College
and
ZIHE WANG
Shanghai University of Finance and Economics

We highlight recent progress in worst-case analysis of welfare in first price auctions. It was shown in [Syrgkanis and Tardos 2013] that in any Bayes-Nash equilibrium of a first-price auction, the expected social welfare is at least a $(1 - 1/e) \approx .63$-fraction of optimal. This result uses *smoothness*, the standard technique for worst-case welfare analysis of games, and is tight if bidders' value distributions are permitted to be correlated. With independent distributions, however, the worst-known example, due to [Hartline et al. 2014], exhibits welfare that is a $\approx .89$-fraction of optimal. This gap has persisted in spite of the canonical nature of the first-price auction and the prevalence of the independence assumption. In [Hoy et al. 2018], we improve the worst-case lower bound on first-price auction welfare assuming independently distributed values from $(1 - 1/e)$ to $\approx .743$. Notably, the proof of this result eschews smoothness in favor of techniques which exploit independence. This note overviews the new approach, and discusses research directions opened up by the result.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Science**]: Economics

1. INTRODUCTION

The first-price auction is a ubiquitous resource allocation protocol. To sell a single item, the auctioneer solicits sealed bids from each of n interested bidders. The highest bidder wins the item, and pays their bid. Unfortunately, the simplicity of the format belies several difficulties. First, in a first-price auction, the highest-valued bidder need not always win in Bayes-Nash equilibrium [Vickrey 1961]. It is therefore necessary to analyze the structure of equilibria to quantify this welfare loss. Unfortunately, a second problem is that theorists have struggled to solve for equilibria in all but the simplest examples for decades.

In light of these obstacles, the first price auction is a canonical example of the success of worst-case analaysis, which allows theorists to reason broadly about welfare in equilibria without solving for them. In particular, [Syrgkanis and Tardos 2013] show that in any Bayes-Nash equilibrium of the first-price auction, the expected welfare is at least a $(1-1/e)$-fraction of optimal. Moreover, [Syrgkanis 2014] shows

Authors' addresses: `darrell.hoy@gmail.com`, `Sam.Taggart@oberlin.edu`, `wang.zihe@mail.shufe.edu.cn`

this result to be tight if one allows correlation in bidders' value distributions.

The first-price auction welfare guarantee of [Syrgkanis and Tardos 2013] is proven using *smoothness*, the standard technique for proving worst-case welfare guarantees in games. Smoothness arguments rely on minimal details of equilibrium, which allows smoothness-derived welfare guarantees to extend to a broad range of solution concepts, including mixed Nash equilibrium, Bayes-Nash equilibrium, and the outcomes of learning behavior [Syrgkanis and Tardos 2013; Hartline et al. 2015], and even to non-welfare objectives such as revenue [Hartline et al. 2014]. Moreover, smoothness is useful beyond first-price auctions, yielding guarantees for other natural ways of selling a single item such as all-pay auctions and a wide range of multi-item auctions, including greedy combinatorial auctions and smooth single-item auctions run simultaneously or in sequence [Syrgkanis and Tardos 2013]. A more comprehensive account of smoothness in auctions can be found in [Roughgarden et al. 2016].

As mentioned above, the smoothness-derived guarantee of $(1 - 1/e)$ for first-price auction in Bayes-Nash equilibrium is tight for correlated distributions. Under the more commonly-studied assumption of independent value distributions, however, the worst-known example is far from tight, with equilibrium capturing a $\approx .89$-fraction of the optimal welfare. Despite this apparent shortfall of the smoothness-derived bound, no better guarantee was known. Consequently, it was unclear whether the technical difficulties in solving for first-price auction equilibria prevented the discovery of instances with low welfare, or whether the detail-agnostic nature of smoothness was to blame. In [Hoy et al. 2018], we prove a stronger guarantee of $\approx .743$ for the first-price equilibrium under the assumption of independent values. This result is the first worst-case BNE welfare analysis which uses independence to surpass a smoothness-derived welfare bound.

In what follows, we offer a technical overview of the improved welfare guarantee of [Hoy et al. 2018], and discuss the approach's connections to the weaker $(1-1/e)$ bound. For expositional clarity, we will argue for a special case of the independently-distributed values setting, in which the highest-valued bidder is consistent across all realized value profiles. The general case extends the logic we will present via a conditioning argument. We refer the reader to the paper for futher details.

2. CENTRAL EXAMPLE AND NOTATION

Consider selling a single item to n bidders via a first-price auction. Each bidder i draws their value for the item from a prior with CDF F_i, which we assume to be independent and not necessarily identical across bidders. For simplicity of exposition, we will assume that bidder 1's value is a point mass at v_1, and that all other bidders' values are less than v_1 in every realization.

We study the first-price auction in Bayes-Nash equilibrium. Each bidder i realizes their value v_i and selects a sealed bid b_i. The highest bidder from the bid profile $\mathbf{b}$ wins the item and pays their bid. Each bidder i selects their bid b_i to maximize their expected utility $\tilde{u}_i(b_i) = \mathbb{E}_{v_{-i}}[(v_i - b_i)\tilde{x}_i(\mathbf{b}(\mathbf{v}))]$, where $\mathbf{b}(\mathbf{v})$ denotes the profile of bids under value profile $\mathbf{v}$, and $\tilde{x}_i(\mathbf{b}(\mathbf{v}))$ is an indicator of whether bidder i wins.

We study the objective of utilitarian social welfare, given by $\mathbb{E}_{\mathbf{v}}[\sum_i v_i\tilde{x}_i(\mathbf{b})]$. For our example, the optimal welfare is simply given by v_1, as bidder 1 consistently has the highest value. The expected welfare from equilibrium, meanwhile, can be broken into two terms: the welfare contribution from bidder 1, the *proper winner*, and the welfare contribution from other bidders, who are *improper winners*. If we let i^* denote the index of the auction's winner, then we have the following expression for equilibrium welfare:

$$\text{Equilibrium Welfare} = v_1 \Pr[i^* = 1] + \mathbb{E}[v_{i^*} \mid i^* \neq 1] \Pr[i^* \neq 1]. \tag{1}$$

In addition to greatly simplifying the analysis, the restricted form of our example parallels that of the worst-known examples. Both the tight example for the $(1-1/e)$ bound in the case of correlated values from [Syrgkanis 2014] and the worst-known example under independence, from [Hartline et al. 2014], involve a single highest bidder with a deterministic value. It is remains open whether or not such a structure will suffice to prove a tight bound under independent distributions.

3. STANDARD ANALYSIS

To highlight the connection between the standard $(1 - 1/e)$ guarantee for the first price auction and our improved bound, we first sketch a proof of the former guarantee. The analysis consists of two parts: we first derive a tradeoff between the expected utility of bidder 1 and the strength of the bid distribution from other players, and in particular the bids of improperly allocated bidders. Smoothness-based proofs achieve this by arguing about bidder 1's possible deviation bids, but we instead follow the slightly different approach of [Hartline et al. 2014]. Second, we note that players' bids lowerbound their values. Combining these two facts yields a tradeoff between the utility (and hence value) of the proper winner and the values of the improper winners, yielding the welfare guarantee.

Formally, let B_i be the CDF for player i's bid distribution (in aggregate over both their randomly drawn values and possible mixed strategies), and let $B_c(b) := \prod_{i\neq 1} B_i(b)$ denote the CDF of the competing bid distribution faced by player i. Note that B_c can be thought of as the interim allocation rule faced by player 1, i.e. $B_c(b) = \mathbb{E}_{\mathbf{b}_{-1}}[\tilde{x}_1(\mathbf{b})] := \tilde{x}_1(b)$. Further let u_1 denote the equilibrium utility for bidder 1.

We first make precise the intutive observation that if bidder 1's utility u_1 is low, they must be facing a strong competing bid distribution, and vice versa. We do so with the following lemma:

Lemma 3.1. *Let $\tau(z)$ denote the competing bid for player 1 with quantile z in B_c; i.e. $\tau(z) = B_c^{-1}(z)$. Then for all $z \in [0, 1]$, $\tau(z) \geq v_1 - u_1/z$.*

The proof follows from the fact that bidder 1 could bid $\tau(z)$ if they desired, and recieve probability of allocation z. Therefore, bidder 1's best response inequality implies $u_1 \geq (v_1 - \tau(z))z$. Rearranging yields the lemma.

Next, note that whenever a bidder i^* wins with a bid of b_{i^*}, it must be that their value is at least their bid, which we summarize as:

Lemma 3.2. *Let i^* denote the winner of the first price auction. Then $b_{i^*} \leq v_{i^*}$.*

To see how the $(1 - 1/e)$ guarantee follows from these two lemmas, note that when when player 1 loses, it is because their critical bid $\tau(z)$ realized from B_c was higher than their bid b_1. Moreover, this critical bid corresponds to the bid of some other player, who won the auction instead. Assuming player 1 bids b_1, we may therefore lowerbound the equilibrium welfare as:

$$\begin{aligned} v_1\tilde{x}_1(b_1) + \int_{\tilde{x}_1(b_1)}^1 \tau(z)\,dz &\geq v_1\tilde{x}_1(b_1) + \int_{\tilde{x}_1(b_1)}^1 v_1 - \frac{u_1}{z}\,dz \\ &= v_1 + u_1 \ln(\tilde{x}_1(b_1)) \\ &\geq v_1 + u_1 \ln\frac{u_1}{v_1} \\ &\geq (1-1/e)v_1. \end{aligned}$$

The first inequality follows from Lemma 3.1, the second from the fact that $\tilde{x}_1(b_1) \geq u_1/v_1$, and the third from the fact that $u_1 \in [0, v_1]$.

Note that this analysis not only did not use the independence of bidders' value distributions, but in fact *only used bidder 1's best response condition*, along with no overbidding for other bidders, and of course the semantics of the auction. Consequently, this analysis extends not just to Bayes-Nash equilibrium with correlated distributions, but to other solution concepts as well. Moreover, note that in our particular example, we have exactly accounted for the welfare contribution of bidder 1 in equilibrium, but have relied on the bids (and hence payments) of improperly allocated bidders as a lower bound on their values. This latter quantity is therefore the only natural target for more fine-grained analysis that utilizes independence. In other words, the above analysis perfectly captures the welfare contribution of the proper winner and the payments of the improper winners, but as we show below, underestimates the utiltiy contribution of the improper winners.

4. IMPROVED ANALYSIS: IMPROPER WINNERS

We now use the assumption that bidders' values are independnently distributed to improve the standard bound. As discussed, the proof of the $(1-1/e)$ bound relied on a loose bound on the values of improperly allocated bidders. We will use independence to improve this characterization.

Assume player 1 is bidding according to a mixed strategy, with CDF B_1. Then for any other bidder i, we may pick bid b_1 with quantile $q_1 = B_1(b_1)$ in player 1's bid distribution, and note that bidder i's best response condition is given by

$$(v_i - b_i)\tilde{x}_i(b_i) \geq (v_i - b_1)\tilde{x}_i(b_1), \tag{2}$$

where b_i is bidder i's equilibrium bid, and $\tilde{x}_i(b)$ is bidder i's probability of allocation should they bid b. Note that if bidders 1 and i were the only bidders in the auction, this would immediately imply an improved lower bound on v_i, as we would have $\tilde{x}_i(b_i) \leq 1$ and $\tilde{x}_i(b_1) = q_1$, yielding $(v_i - b_i) \geq (v_i - b_1)q_1$, which can be rearranged as

$$v_i \geq \frac{b_i - b_1 q_1}{1 - q_1}.$$

For $q_1 = 0$, this bound corresponds to the previous lower bound of b_i on v_i. For $q_1 > 0$, however, this yields strictly stronger bound, assuming $b_i \geq b_1$, as must be the case when bidder i is improperly allocated.

The above argument does not yet require independence, but fails with more than two bidders. In the presence of additional bidders, we employ a trick from [Kirkegaard 2009] which exploits independence. Note that for any bidder $i \in \{1, \ldots, n\}$, we may use independence to write the allocation probability $\tilde{x}_i(b)$ as $\prod_{j \neq i} B_j(b)$. Since player 1 prefers to bid b_1 instead of b_i, we have

$$(v_1 - b_1)\tilde{x}_1(b_1) \geq (v_1 - b_i)\tilde{x}_1(b_i) \tag{3}$$

Dividing (2) by (3) and cancelling factors of $B_j(b_1)$ and $B_j(b_i)$ for $j \notin \{1, i\}$ yields:

$$\frac{(v_i - b_i)B_1(b_i)}{(v_1 - b_i)B_i(b_i)} \geq \frac{(v_i - b_1)B_1(b_1)}{(v_1 - b_1)B_i(b_1)}. \tag{4}$$

As noted before, $B_1(b_1) = q_1$, and $B_1(b_i) \leq 1$. If we assume that $b_i \geq b_1$, then we also have $B_i(b_1) \leq B_i(b_i)$. We may use these facts to obtain the bound:

$$\frac{(v_i - b_i)}{(v_1 - b_i)} \geq \frac{(v_i - b_1)q_1}{(v_1 - b_1)},$$

which may be rearranged to prove the following lemma:

Lemma 4.1. *Let $B_1(b_1) = q_1$, and assume some other bidder i outbids bidder 1 with a bid of b_i. Then it must be that*

$$v_i \geq v_1 \frac{\frac{b_i}{v_1} - (1 - q_1)\frac{b_i}{v_1}\frac{b_1}{v_1} - q_1\frac{b_1}{v_1}}{1 - q_1 - \frac{b_1}{v_1} + q_1\frac{b_i}{v_1}}. \tag{5}$$

While this appears opaque, note that the ratios $\frac{b_i}{v_1}$ and $\frac{b_1}{v_1}$ are necessarily in $[0, 1]$, as no bidder will overbid, and as $v_i \leq v_1$. Note further that for $q_1 = 0$, this yields $v_i \geq b_i$, and for $q_1 = 1$, this yields $v_i \geq v_1$.

5. IMPROVED ANALYSIS: COMPLETE WELFARE

All that remains is to integrate the improved value lower bound of Lemma 4.1 into the analysis of the aggregate social welfare. The argument follows the pattern of the $(1 - 1/e)$ guarantee. Let $\underline{v}(v_1, q_1, b_1, b_i)$ denote the righthand side of (5), and let $b_1(q_1) = B_1^{-1}(q_1)$. We may lower bound the equilibrium welfare as:

$$\int_0^1 \left[v_1\tilde{x}_1(b_1(q_1)) + \int_{\tilde{x}_1(b_1)}^1 \underline{v}(v_1, q_1, b_1(q_1), \tau(z))\, dz \right] dq_1. \tag{6}$$

Moreover, we can observe that $\underline{v}(v_1, q_1, b_1, b_i)$ is monotone increasing in b_i, and hence we may use the lower bound of Lemma 3.1 in place of $\tau(z)$.

Evaluating the inside integral in (6) yields a welfare lower bound of

$$\int_0^1 v_1 - u_1(1 - q_1)\ln\left(1 + \tfrac{v_1 - b_1(q_1) - u_1}{(1-q_1)u_1}\right) dq_1.$$

We may further note that the integrand is increasing in $b_1(q_1)$, and hence take $b_1(q_1) = 0$ for all q_1 in the worst case to obtain:

$$\int_0^1 v_1 - u_1(1 - q_1)\ln\left(1 + \tfrac{v_1 - u_1}{(1-q_1)u_1}\right) dq_1.$$

Minimizing this function numerically over all choices of $u_1 \in [0, v_1]$ yields a worst-case welfare in equilibrium of $\approx .743v_1$.

6. CONCLUSION AND OPEN PROBLEMS

In this note we have offered a proof that the worst-case welfare of the first-price auction in Bayes-Nash equilibrium under the assumption that bidders' values are independently distributed is at least a $\approx .743$-fraction of optimal. We argued for a special case, where there was a single bidder with a deterministic value who would win with probability 1 under the welfare-optimal allocation. The argument followed a the same pattern as previous worst-case analyses of the first-price auction, but used independece to refine the lossiest step from previous analyses. The proof of the general theorem follows this same template, but requires a more complicated conditioning argument to handle uncertainty around the highest-valued bidder. The interested reader is invited to consult the full paper for details.

The key insight above was that standard analyses underestimate the utilities of improper winners. Because many of the standard analyses of smooth mechanisms can be framed like Section 3, there is hope that our approach may be more widely applicable. A natural first candidate would be the all-pay auction, which has a similar smoothness-based welfare guarantee in [Syrgkanis and Tardos 2013; Hartline et al. 2014]. The all-pay auction is well-understood in mixed strategies [Christodoulou et al. 2015], but lacks tight bounds beyond this setting. A more ambitious, but possibly fruitful direction would be to explore multi-item auctions. The $(1 - 1/e)$ guarantee from smoothness is tight for the simultaneous first-price auction with submodular valuations [Christodoulou et al. 2016], but there may be other natural auctions among the many where smoothness applies where our approach could yield progress.

Even for single-parameter settings with first-price payment semantics, open questions remain. First is the obvious task of proving tight bounds. The worst-known example has welfare which is a $\approx .89$-fraction of optimal, off from the $\approx .74$ guarantee. For matroid, combinatorial, or even k-unit auctions, the $(1-1/e)$ guarantee from smoothness extends naturally, but the approach outlined in this haves proved surprisingly challenging to extend, due to the delicate nature of the conditioning step omitted in the above analysis. Finally, [Hartline et al. 2014] show how to extend the analysis in Section 3 to the objective of revenue, under independent distributions. Can this bound be improved using the above techniques?

REFERENCES

Christodoulou, G., Kovács, A., Sgouritsa, A., and Tang, B. 2016. Tight bounds for the price of anarchy of simultaneous first-price auctions. *ACM Transactions on Economics and Computation (TEAC) 4,* 2, 9.

Christodoulou, G., Sgouritsa, A., and Tang, B. 2015. On the efficiency of all-pay mechanisms. In *Algorithms-ESA 2015*. Springer, 349–360.

Hartline, J., Hoy, D., and Taggart, S. 2014. Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 693–710.

Hartline, J., Syrgkanis, V., and Tardos, E. 2015. No-regret learning in bayesian games. In *Advances in Neural Information Processing Systems*. 3061–3069.

Hoy, D., Taggart, S., and Wang, Z. 2018. A tighter welfare guarantee for first-price auctions.

In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2018. ACM, New York, NY, USA, 132–137.

Kirkegaard, R. 2009. Asymmetric first price auctions. *Journal of Economic Theory 144,* 4 (July), 1617–1635.

Roughgarden, T., Syrgkanis, V., and Tardos, E. 2016. The price of anarchy in auctions. *arXiv preprint arXiv:1607.07684*.

Syrgkanis, V. 2014. *Efficiency of mechanisms in complex markets*. Cornell University.

Syrgkanis, V. and Tardos, E. 2013. Composable and efficient mechanisms. In *ACM Symposium on Theory of Computing*. 211–220.

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