Settling the Efficiency of the First-Price Auction

YAONAN JIN
Columbia University
and
PINYAN LU
Shanghai University of Finance and Economics

We survey the main result from [Jin and Lu 2022]: For the first-price auction, the tight Price of Anarchy is $1-1/e^2 \approx 0.8647$, which closes the gap between the best known bounds [0.7430, 0.8689].

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

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: First-Price Auction, Price of Anarchy

1. FIRST-PRICE EQUILIBRIA CAN BE COMPLICATED

In 1961, Nobel Laureate Vickrey initiated auction theory. At the center of his work [Vickrey 1961] was the solution concept of Nash equilibrium [Nash 1950] for auctions as non-cooperative games. This game-theoretic approach has shaped modern auction theory and has had tremendous influence also on other areas in mathematical economics [Holt and Roth 2004]. In particular, Vickrey systematically investigated the first-price auction (or its equivalent, the Dutch auction), the most common auction format in real business.

In the first-price auction, the auctioneer (seller) sells an indivisible item to n potential bidders (buyers). The rule is as simple as it gets: All bidders simultaneously submit bids to the auctioneer (each of which is unknown to the other bidders); the highest bidder wins the item, paying his/her own bid. As simple as the rule is, the bidders' optimal bidding strategies can be sophisticated. From one bidder's own perspective, a higher bid means a higher payment on winning but also a better chance of winning. Accordingly, this bidder's optimal bidding strategy depends on the competitive environment, which in turn is determined by the other bidders' bidding strategies. This scenario is precisely a *non-cooperative game* and the standard solution concept is *Bayesian Nash equilibrium.* To get a better sense, let us consider a warm-up example from Vickrey's original work.

Example 1.1. Consider two bidders: Alice and Bob have (independent) uniform random values $v_1, v_2 \sim U[0, 1]$ and respectively bid $b_1 = \frac{v_1}{2}$ and $b_2 = \frac{v_2}{2}$. The value distributions and bidding strategies determine the (independent) bid distributions of bidders. In this example, they are (independent) uniform random bids $b_1, b_2 \sim U[0, \frac{1}{2}]$, whose CDFs are $B_1(b) = B_2(b) = \min(2b, 1)$. By bidding $b \geq 0$, Alice wins with probability $\min(2b, 1)$ and gains a utility $(v_1 - b)$ conditioned on winning. Her expected utility is $(v_1 - b) \cdot \min(2b, 1)$, which is maximized when $b = \frac{v_1}{2}$. Thus,

Authors' addresses: `yj2552@columbia.edu`, `lu.pinyan@mail.shufe.edu.cn`

her current strategy $b_1 = \frac{v_1}{2}$ is optimal. By symmetry, Bob's current strategy $b_2 = \frac{v_2}{2}$ is also optimal. In sum, the above strategy profile $(b_1,\, b_2) = (\frac{v_1}{2},\, \frac{v_2}{2})$ is an equilibrium, in the sense that no bidder can gain a better utility by deviating from his/her current strategy.

For clarity, let us formalize the model. Each bidder $i \in [n]$ independently draws his/her value from a distribution $v_i \sim V_i$. Only with this knowledge and depending on his/her own strategy s_i, bidder i submits a (possibly) random bid $b_i = s_i(v_i)$. Then over the randomness of the other bidders' values $\boldsymbol{v}_{-i} = (v_k)_{k \neq i}$ and strategies $\boldsymbol{s}_{-i} = \{s_k\}_{k \neq i}$, bidder i wins with probability $\mathrm{x}_i(b_i) \in [0,\, 1]$ and gains an expected utility $u_i(v_i,\, b_i) = (v_i - b_i) \cdot \mathrm{x}_i(b_i)$.

Definition 1.2. A strategy profile $\boldsymbol{s} = \{s_i\}_{i \in [n]}$ is a Bayesian Nash Equilibrium when: For each bidder $i \in [n]$ and any value $v \in \mathrm{supp}(V_i)$, the current strategy $s_i(v)$ is optimal, i.e., $\mathbf{E}_{s_i}\big[u_i(v,\, s_i(v))\big] \geq u_i(v,\, b)$ for any deviation bid $b \geq 0$.

Example 1.1 is special in that bidders have identically distributed values. This *symmetric* setting is well understood: The first-price auction has a *unique* Bayesian Nash equilibrium [Chawla and Hartline 2013], which is *fully efficient* – the bidder with the highest value always wins the item. Instead, the current research trend focuses on the *asymmetric* setting, where bidders' values are distinguished by their distributions. Again, let us get a better sense through two concrete examples.

Example 1.3. Consider two bidders: Alice has a fixed value $v_1 \equiv 2$ and always bids $s_1(v_1) \equiv 1$. Bob has a uniform random value $v_2 \sim U[0,\, 1]$ and *truthfully* bids his value $s_2(v_2) = v_2$, i.e., the distribution $B_2(b) = \min(b,\, 1)$. By bidding $b \geq 0$, Alice gains an expected utility $(v_1 - b) \cdot \min(b,\, 1)$, for which her current strategy $s_1(v_1) \equiv 1$ is optimal. Bob cannot gain a positive utility since his value $v_2 \sim U[0,\, 1]$ is at most Alice's bid $s_1(v_1) \equiv 1$, so his current strategy $s_2(v_2) = v_2$ is also optimal. In sum, this strategy profile (s_1, s_2) is an equilibrium.

Unlike the symmetric setting, this auction game has *multiple* equilibria. E.g., it is easy to verify that the same strategy $s_1(v_1) \equiv 1$ for Alice and a different strategy $\tilde{s}_2(v_2) = \max(\frac{2v_2 - 1}{v_2},\, 0)$ for Bob (i.e., the distribution $\tilde{B}_2(b) = \frac{1}{2-b}$ for $b \in [0,\, 1]$) are another equilibrium.

Both equilibria $(s_1,\, s_2)$ and $(s_1,\, \tilde{s}_2)$ presented in Example 1.3 have two features: (i) The strategies $s_1(v_1)$, $s_2(v_2)$, $\tilde{s}_2(v_2)$ are *pure strategies* – Each of them has no randomness and is just a function of values. (ii) Both equilibria are *fully efficient* akin to the symmetric setting – Alice always has the highest value $\equiv 2$ and always wins the item. However, these are not always the case in the asymmetric setting, such as in the next example, which only slightly modifies Example 1.3 by changing the fixed value of Alice from 2 to 1.

Example 1.4. Consider two bidders: Alice has a fixed value $v_1 \equiv 1$ and her bid $s_1(v_1) \sim B_1$ follows the distribution $B_1(b) = \frac{1}{4b-2}\exp(\frac{4b-3}{2b-1})$ for $b \in [\frac{1}{2},\, \frac{3}{4}]$. Bob has a uniform random value $v_2 \sim U[0,\, 1]$ and his bid $s_2(v_2) = \max(\frac{4v_2 - 1}{4v_2},\, 0)$ follows the distribution $B_2(b) = \frac{1}{4-4b}$ for $b \in [0,\, \frac{3}{4}]$. By bidding $b \geq 0$, Alice gains an expected utility $(v_1 - b) \cdot B_2(b)$, which is maximized $= \frac{1}{4}$ anywhere between $b \in [0,\, \frac{3}{4}]$, so her current strategy $s_1(v_1) \sim B_1$ is optimal. By elementary algebra, we can check that Bob's current strategy $s_2(v_2) = \max(\frac{4v_2-1}{4v_2}, 0)$ also is optimal. In sum, this strategy profile (s_1, s_2) is an equilibrium.

In Example 1.4, Alice has a *mixed strategy* – a fixed value $v_1 \equiv 1$ but a random bid $s_1(v_1) \sim B_1$. Moreover, the equilibrium $(s_1,\, s_2)$ is *not* fully efficient. E.g., with a value $v_2 = \frac{3}{4}$, although not the highest value $< v_1 \equiv 1$, Bob bids $s_2(v_2) = \frac{2}{3}$ and wins with probability $B_1(\frac{2}{3}) = \frac{3}{2e} \approx 0.5518$. Indeed, Example 1.4 has infinite equilibria, among which the given one (s_1, s_2) has the relatively "simplest" format. But none of those equilibria is a pure equilibrium or is fully efficient, despite the fact that Example 1.4 is only a minor modification of Example 1.3.

From the above examples, we observe that Bayesian Nash equilibria can be very complicated and sensitive to the value distributions, even if the auction rules are quite simple. After extensive study for more than 60 years, the first-price auction and its equilibria remain the centerpiece of modern auction theory and have promoted a rich literature (e.g., see the survey [Roughgarden et al. 2017] and references therein). These efforts are justifiable as the study of the first-price auction and its equilibria is both theoretically challenging and practically important.

2. (IN)EFFICIENCY OF FIRST-PRICE EQUILIBRIA

Among various aspects of the first-price auction, *efficiency* at equilibria is of primary interest. In economics, efficiency measures to what extent a recourse can be allocated to the persons who value it the most and thus maximize the *social welfare*, especially in a competitive environment. As shown in Example 1.4, the first-price auction generally is not fully efficient at an equilibrium: The winner has the highest *bid* but possibly not the highest *value*; this crucially depends on both (i) the instance $\boldsymbol{V}$ itself and (ii) which Bayesian Nash equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V})$ it falls into. Earlier works in economics focus on (generalizing) the conditions for the value distribution $\boldsymbol{V} = \{V_i\}_{i\in[n]}$ that guarantee the full efficiency.

However, the quality of (in)efficiency should not be all-or-nothing. E.g., when the highest value is 1, a value-0.99 bidder versus a value-0.01 bidder is very different, although neither of them is fully efficient. Towards a quantitative analysis, Koutsoupias and Papadimitriou introduced a new measure on the efficiency degradation under selfish behaviors, the *Price of Anarchy* [Koutsoupias and Papadimitriou 1999] (which is an analog of the "approximation ratio" in theoretical computer science). For the first-price auction, denote by $\mathsf{OPT}(\boldsymbol{V})$ the expected optimal social welfare from an instance $\boldsymbol{V}$, and by $\mathsf{FPA}(\boldsymbol{V},\, \boldsymbol{s})$ the expected social welfare at an equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V})$, then the Price of Anarchy is defined as follows.

Definition 2.1. For the first-price auction, the Price of Anarchy is given by

$$\mathsf{PoA} \;:=\; \inf\left\{ \frac{\mathsf{FPA}(\boldsymbol{V},\, \boldsymbol{s})}{\mathsf{OPT}(\boldsymbol{V})} \;\middle|\; (\boldsymbol{V},\, \boldsymbol{s}) \in \mathbb{V} \times \mathbb{BNE} \text{ and } \mathsf{OPT}(\boldsymbol{V}) < +\infty \right\}.$$

Clearly, PoA's are bounded between [0, 1]; a larger ratio means a higher efficiency and the $= 1$ ratio means the full efficiency. Specifically for the first-price auction, Syrgkanis and Tardos proved the first nontrivial lower bound of $1 - 1/e \approx 0.6321$ [Syrgkanis and Tardos 2013]. Later, Hoy, Taggart and Wang derived an improved lower bound of ≈ 0.7430 [Hoy et al. 2018]. On the other hand, Hartline, Hoy and

Taggart showed a concrete instance of ratio ≈ 0.8689 [Hartline et al. 2014], which remains the best known upper bound.

Despite the prevalence of the first-price auction and much effort in studying its efficiency, there has been a persistent gap in the state of the art. In the authors' recent work [Jin and Lu 2022], this long-standing open problem is finally solved.

THEOREM 2.2. *For the first-price auction,* $\mathsf{PoA} = 1 - 1/e^2 \approx 0.8647$.

Remarkably, neither of the best known lower and upper bounds [0.7430, 0.8689] are tight; we close the gap by improving both bounds to $1 - 1/e^2 \approx 0.8647$. This tight bound is not only of theoretic interest but has further implications in real business as it is fairly close to 1. Namely, the efficiency degradation at any equilibrium is small, no worse than 13.53%, which might be acceptable given other merits of the first-price auction.

En route to the tight PoA, we obtain new perspectives, characterizations, and properties of equilibria in the first-price auction. These might be of independent interest and will hopefully find applications in future research. E.g., the authors' follow-up [Jin and Lu 2023] uses these to strengthen the above result by showing that the *Price of Stability* for the first-price auction is the same $= 1 - 1/e^2$.

3. COMPARISON WITH PREVIOUS TECHNIQUES

On the Price of Anarchy in auctions, the canonical method is the *smoothness techniques* proposed by Roughgarden [Roughgarden 2015] and developed by Syrgkanis and Tardos [Syrgkanis and Tardos 2013]. The past decade has seen an abundance of its applications and extensions (see the survey [Roughgarden et al. 2017]).

However, the smoothness framework has an intrinsic restriction – it focuses on the structure of an auction game but ignores the *independence*, i.e., both the independence of value distributions $\boldsymbol{V} = \{V_i\}_{i \in [n]}$ and the independence of strategies $\boldsymbol{s} = \{s_i\}_{i \in [n]}$. Hence, although powerful for proving *lower bounds*, the smoothness framework often suffers from certain bottlenecks for proving *tight bounds.*

For the first-price auction, the smoothness-based bound of $1 - 1/e \approx 0.6321$ by Syrgkanis and Tardos [Syrgkanis and Tardos 2013] is tight when *correlated* distributions are allowed, i.e., no improvement is possible without reasoning about the independence. Later, Hoy, Taggart, and Wang [Hoy et al. 2018] obtained a better lower bound of ≈ 0.7430, by partially combining the independence arguments into the smoothness framework.

In sum, towards tight bounds for the first-price auction and other auctions, the primary consideration is: What is the consequence of the *independence* of values $\boldsymbol{v}$ and strategies $\boldsymbol{s}$? Our answer is, the bids $\boldsymbol{s}(\boldsymbol{v}) = (s_i(v_i))_{i \in [n]}$ are also independent and follow a product distribution $\boldsymbol{B}(\boldsymbol{V},\, \boldsymbol{s}) = \{B_i\}_{i \in [n]}$. We thus switch to another representation of equilibria, the bid distributions $\boldsymbol{B}(\boldsymbol{V},\, \boldsymbol{s})$ resulted from equilibria (*equilibrium bid distributions*), over the original representation $(\boldsymbol{V},\, \boldsymbol{s}) \in \mathbb{V} \times \mathbb{BNE}$.

—The original representation has two drawbacks: (i) One value distribution $\boldsymbol{V} \in \mathbb{V}$ can have *multiple* or even *infinite* equilibria. (ii) One equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V})$ generally has a *non-analytic* solution; even an efficient algorithm for computing (approximate) equilibria from value distributions is unknown.

—The new representation addresses both drawbacks: (i) One equilibrium bid distribution $\boldsymbol{B}(\boldsymbol{V},\, \boldsymbol{s})$ backward determines a *unique* equilibrium $(\boldsymbol{V},\, \boldsymbol{s}) \in \mathbb{V} \times \mathbb{BNE}$. (ii) The reconstruction of $(\boldsymbol{V},\, \boldsymbol{s})$ turns out to have an *analytic* solution.

The first part of [Jin and Lu 2022] establishes a bijection between the two search spaces $\mathbb{V} \times \mathbb{BNE}$ and $\mathbb{B}_{\mathsf{EQ}} := \{\boldsymbol{B}(\boldsymbol{V},\, \boldsymbol{s}) \,|\, (\boldsymbol{V},\, \boldsymbol{s}) \in \mathbb{V} \times \mathbb{BNE}\}$: They are equivalent for searching the *worst-case instances* for the Price of Anarchy, but the new search space $\mathbb{B}_{\mathsf{EQ}}$ is technically easier to explore since it is *one* infinite set rather than the Cartesian product $\mathbb{V} \times \mathbb{BNE}$ of *two* infinite sets. Namely, we prove that:

PROPOSITION 3.1. *For the first-price auction, the* Price of Anarchy *is given by*

$$\mathsf{PoA} \;:=\; \inf \left\{ \frac{\mathsf{FPA}(\boldsymbol{B})}{\mathsf{OPT}(\boldsymbol{B})} \;\middle|\; \boldsymbol{B} \in \mathbb{B}_{\mathsf{EQ}} \text{ and } \mathsf{OPT}(\boldsymbol{B}) < +\infty \right\}.$$

This *representations switch* allows us to develop the first principle approach that directly characterizes the *worst-case instances*. More precisely, the second part of [Jin and Lu 2022] starts from the switched search space $\mathbb{B}_{\mathsf{EQ}}$, step-by-step narrows down it by showing stronger and stronger constraints, and eventually captures the worst-case instances together with the tight $\mathsf{PoA} = 1 - 1/e^2$.

Both parts “representations switch” and “worst-case instances characterization” devise new techniques. E.g., we provide the first *discretization scheme* of Bayesian Nash equilibria. Such a scheme is crucial for understanding *complexity* and *learnability* of equilibria, but is a notorious task: First, equilibria can be very sensitive to perturbations of value distributions $\boldsymbol{V} \to \tilde{\boldsymbol{V}}$ (cf. Examples 1.3 and 1.4). Second, even the existence of a perturbed equilibrium $\tilde{\boldsymbol{s}} \in \mathbb{BNE}(\tilde{\boldsymbol{V}})$ “close enough” to a given equilibrium $\boldsymbol{s} \in \mathbb{BNE}(\boldsymbol{V})$ is doubtful, since no algorithm for computing equilibria is available. Our scheme circumvents both issues and, roughly speaking, stems from a delicate usage of the representations switch.

We remark that our approach towards the tight $\mathsf{PoA} = 1 - 1/e^2$ is very general. We believe that the framework as well as the “en route” structural and technical byproducts can be adaptable to future research, including (i) the Price of Anarchy/Stability for other auctions, (ii) complexity and learnability of equilibria, and (iii) other relevant topics like Revenue Maximization in auctions (e.g., [Hartline and Roughgarden 2009; Hartline et al. 2014; Jin et al. 2019]).

REFERENCES

CHAWLA, S. AND HARTLINE, J. D. 2013. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, M. J. Kearns, R. P. McAfee, and É. Tardos, Eds. ACM, 181–196.

HARTLINE, J. D., HOY, D., AND TAGGART, S. 2014. Price of anarchy for auction revenue. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, M. Babaioff, V. Conitzer, and D. A. Easley, Eds. ACM, 693–710.

HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proceedings 10th ACM Conference on Electronic Commerce (EC-2009), Stanford, California, USA, July 6–10, 2009*, J. Chuang, L. Fortnow, and P. Pu, Eds. ACM, 225–234.

HOLT, C. A. AND ROTH, A. E. 2004. The nash equilibrium: A perspective. *Proceedings of the National Academy of Sciences 101,* 12, 3999–4002.

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, Los Angeles, CA, USA, June 25-29, 2018, I. Diakonikolas, D. Kempe, and M. Henzinger, Eds. ACM, 132–137.

Jin, Y. and Lu, P. 2022. First-price auction is $1 - 1/e^2$ efficient. To appear in *Proceedings of the 63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS), 2022*.

Jin, Y. and Lu, P. 2023. The price of stability for first-price auction. To appear in *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2023*.

Jin, Y., Lu, P., Qi, Q., Tang, Z. G., and Xiao, T. 2019. Tight approximation ratio of anonymous pricing. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019.* 674–685.

Koutsoupias, E. and Papadimitriou, C. H. 1999. Worst-case equilibria. In *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings*, C. Meinel and S. Tison, Eds. Lecture Notes in Computer Science, vol. 1563. Springer, 404–413.

Nash, J. F. 1950. Equilibrium points in n-person games. *Proceedings of the national academy of sciences 36,* 1, 48–49.

Roughgarden, T. 2015. Intrinsic robustness of the price of anarchy. *J. ACM 62,* 5, 32:1–32:42.

Roughgarden, T., Syrgkanis, V., and Tardos, É. 2017. The price of anarchy in auctions. *J. Artif. Intell. Res. 59*, 59–101.

Syrgkanis, V. and Tardos, É. 2013. Composable and efficient mechanisms. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, D. Boneh, T. Roughgarden, and J. Feigenbaum, Eds. ACM, 211–220.

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