Scale-robust Auctions*

Jason Hartline† Aleck Johnsen‡ Yingkai Li§

Abstract

We study auctions that are robust at any scale, i.e., they can be applied for selling both expensive and cheap items and achieve the best multiplicative approximations of the optimal revenue in the worst case. The optimal mechanism randomizes between selling at the second-price and a 2.45 multiple of the second-price.

Keywords— robustness, auction, multiplicative approximation.
JEL— D44, D82

1 Introduction

In many markets, it is customary to implement fixed proportional transaction fees regardless of the scale of the commodity. For instance, in real estate, agents typically charge a commission fee of around 6% regardless of the sale price of the house. Similarly, in digital application markets, Apple Store imposes a 30% service fee and Google Play charges a 15% service fee for each app purchase, irrespective of the app's price. Motivated by this feature of markets, we consider the design of auctions that are resilient to scale, i.e., that achieve a favorable revenue guarantee approximating the optimal revenue in a multiplicative manner.

We study a robust analysis framework in which the principal designs auctions that perform well at all scales (Hartline and Roughgarden, 2008). In this framework, the principal

*Work done in part while all authors are supported by NSF CCF 1618502. An extended abstract of this work has appeared in 61st Annual Symposium on Foundations of Computer Science (FOCS'20) under the title "Benchmark Design and Prior-independent Optimization". The authors thank Tan Gan, Shengwu Li, Eran Shmaya and Philipp Strack for helpful comments and suggestions.

†Northwestern University. Email: hartline@northwestern.edu.
‡Northwestern University. Email: aleckjohnsen@u.northwestern.edu.
§National University of Singapore. Email: yk.li@nus.edu.sg.

seeks an auction that is independent of the distribution over agents' values and, specifically, the scale of the distribution. The goal is to minimize the multiplicative approximation of the optimal mechanism in the worst case over a family of possible distributions.[1]

We study the single-item auction in a symmetric environment where the buyers' values are drawn independently and identically from a regular distribution.[2] For regular distributions, if the distribution is known by the principal, the second price auction with monopoly reserve is Bayesian optimal (Myerson, 1981). If the regular distribution is unknown, Bulow and Klemperer (1996) show that by adding an additional buyer, the seller can extract at least the optimal revenue (without the additional buyer) using the second-price auction. A corollary of this result is that with a fixed market size of n buyers, the second-price auction attains at least $1 - \frac{1}{n}$ fraction of the optimal revenue. Thus, in large markets where the number of buyers converges to infinity, the second-price auction is asymptotically optimal, while in small markets, the multiplicative gap between the optimal revenue and the second-price auction can be as bad as 2. Is the second-price auction, via this corollary of Bulow and Klemperer, the best scale-robust auction?

In this paper, we focus on the design of optimal scale-robust mechanisms in small markets. In particular, we focus on the extreme case proposed by Dhangwatnotai et al. (2015) where there are only two buyers. The restriction to small markets is consistent with our motivation of robust analysis. Unlike in large markets, where sellers can rely on abundant historical data to accurately estimate the valuation distributions of buyers, such data is insufficient in small markets. Therefore, a seller with limited information often find it natural to adopt the scale-robust approach for selling the goods. When there are only two buyers, Allouah and Besbes (2018) show that the second-price auction is indeed scale-robust optimal if the valuation distribution of the buyers satisfy the monotone hazard rate condition (MHR). However, Fu et al. (2015) show that the seller can improve her worst case approximation guarantee by randomly marking up the second-price if the valuation distribution only satisfies the regularity condition (which is weaker than MHR). The main intuition is that without MHR, the worst-case valuation distribution may be too heavy-tailed, and hence the seller benefits from randomization to hedge between that case that second-price auction is optimal and the case that the monopoly price is much higher than the second-price.

We identify the optimal scale-robust and dominant strategy incentive compatible (DSIC)

[1] This analysis framework is known as prior-independent approximation in the computer science literature following Hartline and Roughgarden (2008)

[2] A distribution is regular if its corresponding virtual value function is non-decreasing (Myerson, 1981).

mechanism for regular valuation distributions when there are two buyers,[3] which answers a major question left open from Dhangwatnotai et al. (2015), Fu et al. (2015), and Allouah and Besbes (2018). The optimal mechanism is a mixture between the second-price auction and the auction where these prices are marked up by a factor of about 2.45.[4] Note that our restriction to DSIC mechanisms is not without loss of generality when the seller can adopt more general and potentially non-truthful mechanisms (e.g., Caillaud and Robert, 2005; Feng et al., 2021). However, we aim to design auctions that are robust to the beliefs of all parties, and DSIC mechanisms provide max-min optimal revenue guarantees over worst-case beliefs of the buyers (Chung and Ely, 2007).

There are two main takeaways from our characterization of the optimal scale-robust mechanism. First, we show that in the scale-robust revenue maximization problem, it is optimal to exclude the low-value agent from winning the item, even if the higher-value agent does not receive it.[5] Intuitively, allocating the item to a low-value agent could be beneficial if their marginal revenue contribution is positive in a robust environment. However, as we illustrate in our paper, for any robust mechanism, the worst-case distribution assigns positive marginal revenue only to the highest value in its support. As a result, all lower value agents have a negative marginal revenue contribution, and any positive allocation to a lower-value agent reduces the expected revenue guarantee in our robust setting.

Another insight from our characterization is that achieving the optimal robust performance requires randomizing over only a single markup price for regular distributions, rather than a continuum of prices. This result stems from the balance between two opposing forces. When the markup strictly exceeds 1, no matter how small, the seller risks losing sales, leading to a non-continuous drop in revenue compared to a mechanism without markups. To justify this risk, the markup must be sufficiently high to generate substantial revenue when buyers' valuations are sufficiently dispersed. However, if the markup is too large, the probability that buyers are sufficiently apart diminishes, making the expected benefit again insufficient to offset the revenue loss from foregone trades. Our result identifies that there exists a unique intermediate markup price that optimally balances these effects, ensuring robust performance.

[3]We show that our proposed mechanism is optimal among mechanisms that satisfy a scale-invariance assumption. We conjecture that even mechanisms without scale-invariance cannot outperform our proposed mechanism. See Allouah and Besbes (2018) for a more detailed discussion on the scale-invariance assumption.

[4]An alternative view of the optimal mechanism is that the winning agent only receives the full item if his bid is sufficiently high compared to the second highest bid, and receives a "damaged" item, or equivalently a partial allocation of the item, if his bid is close to the second highest bid.

[5]This result may not hold in general for other robust optimization problems, such as consumer surplus maximization (see Hartline and Roughgarden, 2014), or revenue maximization with sample access (see Allouah et al., 2022).

The robust analysis framework in this paper is multiplicative approximation, i.e, the worst-case ratio between the performance of the Bayesian optimal mechanism which knows the distribution, and the performance of the designed mechanism. This robustness measure is not standard in the economic literature for mechanism design where max-min optimal (e.g., Bergemann and Schlag, 2011; Carroll, 2017; Carrasco et al., 2018; Carroll and Segal, 2019) or min-max regret (e.g., Bergemann and Schlag, 2011; Guo and Shmaya, 2019, 2022) are commonly adopted. To understand robustness to scale, neither of these prior frameworks can be applied as they give trivial solutions. In particular, the max-min optimal mechanism would focus on the smallest scale, which is where the performance is the lowest. Guarantees for the smallest scale would not translate to good performance at larger scales where there is much more to gain. On the other hand, the optimal min-max regret is achieved at large scales where there is the most to lose, and gives at small scales only the trivial guarantee that performance is non-negative. When the range of scales required in the robustness analysis is taken to the lower or upper limit, respectively, these frameworks provide only trivial guarantees.[6] In contrast, mechanisms with optimal worst-case approximation ratio provide the same good performance guarantee at all scales. Further comparison of robustness frameworks can be found in Appendix B.

1.1 Related Work

The scale-robust analysis framework gives a natural approach of identifying the robustly optimal mechanism. Previous literature has only identified optimal mechanisms in environments that are special cases of the fully general problem. Hartline and Roughgarden (2014) gave the optimal mechanism for revenue maximization in the sale of a single item to a single agent with value from a bounded support where the optimal mechanism posts a randomized price. For revenue maximization in the sale of an item to one of two agents with values drawn from an i.i.d. regular distribution, Dhangwatnotai et al. (2015) show that the second price auction is a 2-approximation. Fu et al. (2015) gave a randomized mechanism showing that this factor of 2 is not tight. Upper and lower bounds on this canonical problem were improved by Allouah and Besbes (2018) to be within $[1.80, 1.95]$ for DSIC mechanisms. The main result of our paper is to identify the optimal scale-robust mechanism for this environment with a factor of about 1.91.[7] This allows us not only to show that random mechanisms outperform deterministic ones in scale-robust auctions, as illustrated in Fu et al. (2015),

[6] Any mechanism is max-min optimal and min-max regret optimal since the optimal max-min value is 0 while the optimal min-max regret is unbounded.

[7] The Allouah and Besbes (2018) lower bound of 1.80 was proved under the same additional assumption of scale invariance as our lower bound of 1.91.

but also to identify the structure of the optimal randomization—specifically, that it assigns a positive winning probability only to the high value agent and randomizes over only one markup price that is strictly above 1. Finally, for this two agent problem with i.i.d. values from a distribution in the subset of regular distributions that further satisfy a monotone hazard rate condition, Allouah and Besbes (2018) show that the second-price auction is scale-robust optimal.

The restriction to DSIC mechanisms has the desirable property that agents' behaviors and the expected revenue in DSIC mechanisms do not rely on agents' information about each other, and the set of DSIC mechanisms is equivalent to the set of ex post implementable mechanisms (Bergemann and Morris, 2005). Without the restriction to DSIC mechanisms, Caillaud and Robert (2005) use an ascending auction in virtual value space to implement the Bayesian optimal mechanism. A critic for such implementation is that this mechanism takes the common knowledge assumption too literally and it is impractical for real-world applications. Feng and Hartline (2018) and Feng et al. (2021) show that there exist simple and practical non-incentive-compatible mechanisms that outperform the optimal DSIC mechanism, and further study of non-incentive-compatible mechanisms is still warranted in scale-robust analysis framework.

Our paper relates to the auction design literature with max-min optimal and min-max regret objectives when the principal is ignorant of the value distribution. For max-min optimization, Bergemann and Schlag (2011) and Carrasco et al. (2018) consider the design of the robustly optimal mechanism in the single-item, single-buyer setting. Bachrach and Talgam-Cohen (2022) extend the model to two i.i.d. buyers and the model with correlated valuations is considered in Che (2022). Both papers identify the second-price auction with random markups as the max-min optimal mechanism, where the distribution over markups rely on the expected value of each buyer. By contrast, the information about the expected value is not available to the principal in our model, and there exists a fixed distribution over markups that achieves the optimal approximation ratio.

For min-max regret optimization, the optimal distribution over prices for the single-item, single-buyer setting is characterized in Bergemann and Schlag (2008, 2011). Anunrojwong et al. (2022) show that a second-price auction with random reserve prices is robustly optimal when there are multiple agents even if the values of the agents can be correlated.

In contrast, by focusing on multiplicative approximations, our mechanism provides non-trivial and interesting insights on designing optimal robust mechanisms, e.g., randomization over a single markup price is optimal for robustness to scale. Moreover, compared to max-min optimal which is often too pessimistic, and the min-max regret which is often too optimistic, multiplicative approximation maintains a good balance between these two

situations. In Appendix B, we provide an illustration for why worst-case multiplicative approximation can be viewed as a measure that lies between the pessimistic and optimistic extremes.

2 Preliminaries

The principal sells a single item to $n = 2$ agents with private values $\boldsymbol{v} = (v_1, v_2)$. The agents have linear utilities, i.e., agent i's utility is $v_i\, x_i - p_i$ for allocation probability x_i and expected payment p_i. Agents' values are drawn independently and identically from a product distribution $\boldsymbol{F} = F \times F$ where F will denote the cumulative distribution function of each agent's value.

Mechanisms A mechanism M is defined by an ex post allocation and payment rule $\boldsymbol{x}^M$ and $\boldsymbol{p}^M$ which map the profile of values $\boldsymbol{v}$ to a profile of allocation probabilities and a profile of payments, respectively. We focus on mechanisms that are feasible, dominant strategy incentive compatible, and individually rational:

- For selling a single item, a mechanism is *feasible* if for all valuation profiles, the allocation probabilities sum to at most one, i.e., $\forall \boldsymbol{v},\ \sum_i x_i^M(\boldsymbol{v}) \leq 1$.
- A mechanism is *dominant strategy incentive compatible* if no agent i with value v_i prefers to misreport some value z: $\forall \boldsymbol{v}, i, z,\ v_i\, x_i^M(\boldsymbol{v}) - p_i^M(\boldsymbol{v}) \geq v_i\, x_i^M(z, \boldsymbol{v}_{-i}) - p_i^M(z, \boldsymbol{v}_{-i})$ where $(z, \boldsymbol{v}_{-i})$ denotes the valuation profile with v_i replaced with z.
- A mechanism is *individually rational* if truthful reporting always leads to non-negative utility: $\forall \boldsymbol{v}, i,\ v_i\, x_i^M(\boldsymbol{v}) - p_i^M(\boldsymbol{v}) \geq 0$.

Denote a family of feasible mechanisms by $\mathcal{M}$ and a mechanism in this family by M. The expected revenue of mechanism M when the value profile is $\boldsymbol{v}$ is denoted by $M(\boldsymbol{v})$. When evaluating the revenue of a mechanism in expectation over the distribution, we adopt the short-hand notation $M(F) = \mathbf{E}_{\boldsymbol{v}\sim F}[M(\boldsymbol{v})]$. Given a distribution F and a family of mechanisms $\mathcal{M}$, the optimal mechanism, denoted by OPT_F, maximizes the expected revenue of the principal:

$$\mathrm{OPT}_F = \operatorname*{argmax}_{M\in\mathcal{M}} M(F).$$

In our paper, we focus on robust mechanisms that are scale invariant.

Definition 1 (Scale Invariant)**.** *Given any incentive-compatible mechanism M with allocation rule $x^M(\boldsymbol{v})$, mechanism M is* scale invariant *if for each agent i, valuation profile $\boldsymbol{v}$ and any constant $\alpha > 0$, $x_i^M(\alpha \cdot \boldsymbol{v}) = x_i^M(\boldsymbol{v})$. Scale invariance further implies $M(a \cdot \boldsymbol{v}) = a \cdot M(\boldsymbol{v})$.*

Allouah and Besbes (2018) prove that the optimal scale-robust mechanism (defined later in Definition 2) among a broad family of mechanisms is scale invariant. They show that if $\lim_{\alpha \to 0} x_i(\alpha \cdot \boldsymbol{v})$ always exists for mechanisms in the family, then the optimal scale-robust mechanism is scale invariant. The latter assumption can be viewed as a simplicity condition on the allocation of the mechanism. Additionally, Allouah and Besbes conjecture that this weaker assumption is without loss; if true, the mechanism we identify as the optimal mechanism among scale-invariant mechanisms is also scale-robust optimal among all mechanisms.

Revenue Curves A mechanism's revenue can be easily and geometrically understood via the marginal revenue approach of Myerson (1981) and Bulow and Roberts (1989). For distribution F, the *quantile* q of an agent with value value v denotes how strong that agent is relative to the distribution F. Specifically, quantiles are defined by the mapping $Q_F(v) = \Pr_{z \sim F}\{z \geq v\}$. Denote the mapping back to value space by V_F, i.e., $V_F(q)$ is the value of the agent with quantile q. A single agent *price-posting revenue curve* $P(q)$ gives the revenue of posting a price such that the probability that the agent accepts the price is q. For an agent with value distribution F, price $V_F(q)$ is accepted with probability q, and its expected revenue is $P(q) = q \cdot V_F(q)$. A single agent *revenue curve* $R_F(q)$ gives the optimal revenue from selling to a single agent under the constraint that ex ante sale probability is q. By Bulow and Roberts (1989), the revenue curve R is always concave, and it coincides with the concave hull of the price-posting revenue curve P. In this paper, we focus on the family of regular distributions. Let $\mathcal{F}^{\text{Reg}}$ be the family of i.i.d. regular value distribution.

Assumption 1 (Regularity)**.** *A distribution F is regular if the price-posting revenue curve P is concave.*[8]

An immediate implication for regular distribution is that the price-posting revenue curve coincides with the revenue curve, i.e., $P = R$. The optimal mechanism for a single agent posts the *monopoly price* $V_F(\bar{q})$ which corresponds to the monopoly quantile $\bar{q} = \operatorname{argmax}_q R_F(q)$. In multi-agent settings, the expected revenue of any multi-agent mechanism M equals its expected surplus of marginal revenue.

[8] An equivalent definition for regularity is that the virtual value function $\phi(v) = v - \frac{1-F(v)}{f(v)}$ is non-decreasing in v.

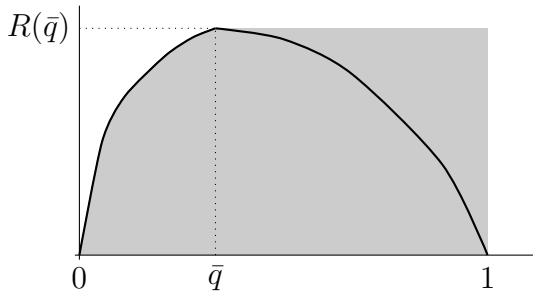


Figure 1: The solid black curve is the revenue curve $R(q)$ for the single-agent setting. The gray area is the area under the smallest monotone concave upper bound of the revenue curve, which is half of the optimal revenue.

Lemma 1 (Myerson, 1981)**.** *Given any incentive-compatible mechanism M with allocation rule $\boldsymbol{x}^M(\boldsymbol{v})$, the expected revenue of mechanism M for agents with regular distribution $\boldsymbol{F}$ is equal to its expected surplus of marginal revenue, i.e.,*

$$M(\boldsymbol{F}) = \sum\nolimits_i \mathbf{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[p_i^M(\boldsymbol{v})\right] = \sum\nolimits_i \mathbf{E}_{\boldsymbol{v}\sim\boldsymbol{F}}\left[R'_F(Q_F(v_i)) \cdot x_i^M(\boldsymbol{v})\right].$$

Corollary 1 (Myerson, 1981)**.** *For i.i.d., regular, single-item environments, the optimal mechanism* OPT_F *is the second-price auction with anonymous reserve equal to the monopoly price.*

The following lemma from Dhangwatnotai et al. (2015) follows from Lemma 1 and gives a geometric understanding of the expected revenues of second-price auction and optimal mechanism in two-agent settings. The geometry is illustrated in Figure 1. Intuitively, the revenue from each agent can be interpreted as being generated by a mechanism that is determined by the reported value of the other agent. For example, in a second-price auction, each agent effectively faces a distribution over posted prices, where the price always equals the reported value of the other agent. This perspective is particularly useful for analysis, as demonstrated in Lemmas 2 and 8.

Lemma 2 (Dhangwatnotai et al., 2015)**.** *In i.i.d. two-agent single-item environments,*

- *the expected revenue of second-price auction is twice the area under the revenue curve;*
- *the expected revenue of the optimal mechanism is twice the area under the smallest monotone concave upper bound of the revenue curve.*

Robust Objectives In this paper, we consider the model where the principal is ignorant of the true distribution over values. Instead, the principal knows that the true distribution

belongs to a family $\mathcal{F}$ and designs a mechanism that minimizes the worst case approximation ratio to the optimal revenue for distributions within $\mathcal{F}$. This scale-robust analysis framework is also referred to as the prior-independent mechanism design (Hartline and Roughgarden, 2008).

Definition 2 (Robust Framework)**.** *The* scale-robust analysis framework *is given by a family of mechanisms $\mathcal{M}$ and a family of distributions $\mathcal{F}$ and solves the program*

$$\beta \triangleq \min_{M \in \mathcal{M}} \max_{F \in \mathcal{F}} \frac{\mathrm{OPT}_F(F)}{M(F)}. \tag{β}$$

3 Optimal Scale-Robust Mechanisms

We solve for the optimal mechanism that is robust to scale for the revenue objective with the restriction to

- the family of i.i.d. regular value distribution $\mathcal{F}^{\mathrm{Reg}}$; and
- the family of feasible, incentive compatible, individually rational, and scale-invariant mechanisms $\mathcal{M}^{\mathrm{SI}}$.

The following discussion motivates these restrictions in the robust design problem.

- *Restrictions on Distributions.* In general, without assumptions of symmetric or regular value distributions, no mechanism achieves good scale-robust performance. Analyzing robust mechanism design without these assumptions does not allow us to distinguish good mechanisms from bad ones, as all mechanisms perform poorly in such settings. Moreover, nearly all papers on the scale-robust analysis framework focus on i.i.d. agents, and almost all studies on revenue maximization within this framework restrict attention to regular distributions.
- *Restrictions on Mechanisms.* The restriction to feasible and individually rational mechanisms is necessary to have a sensible optimization problem. The restriction to incentive compatibility is standard in almost all papers on the scale-robust analysis framework, with the exception of Feng and Hartline (2018), which shows that this restriction is not without loss of optimality. However, maintaining incentive compatibility ensures robustness to agents' beliefs as well. Additionally, the assumption of scale invariance is conjectured to be without loss and simplifies the structure of the robust mechanisms.

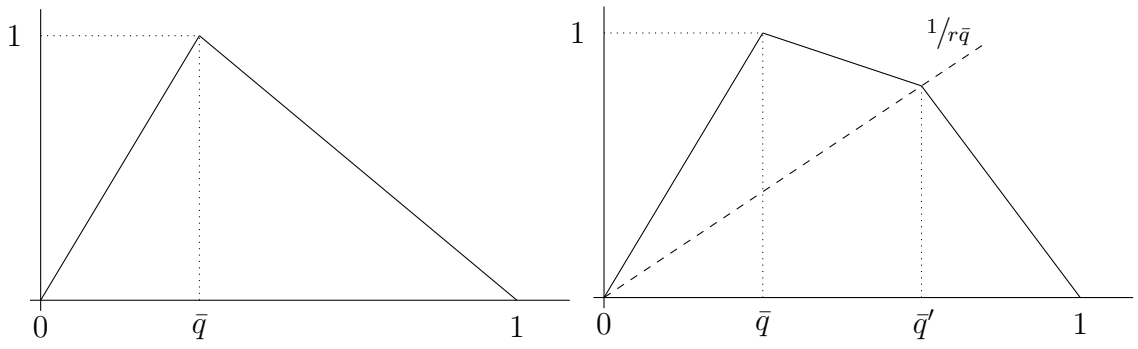


Figure 2: The left hand side is the revenue curve for triangle distribution $\mathrm{Tri}_{\bar{q}}$ and the right hand side is the revenue curve for quadrilateral distribution $\mathrm{Qr}_{\bar{q},\bar{q}',r}$. The definition of quadrilateral distribution $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ will be formally introduced later in Section 3.2.

Within the family of scale-invariant mechanisms, the following family of (stochastic) markup mechanisms is (essentially, in $n = 2$ agent environments) the restriction to the family of lookahead mechanisms (Ronen, 2001) that are scale invariant. Notice that the second-price auction is the 1-markup mechanism M_1. Our main result will show that the optimal scale-robust mechanism is a stochastic markup mechanism, and we identify the optimal distribution over markups for regular distributions.

Definition 3 (Markup Mechanism)**.** *For any parameter $r \geq 1$, the* r-markup mechanism M_r *identifies the agent with the highest-value (and ties broken uniformly at random) and offers this agent r times the second-highest value. A* stochastic markup mechanism *draws r from a given distribution on $[1, \infty)$. The family of stochastic markup mechanisms is $\mathcal{M}^{\mathrm{SMKUP}}$.*

In the analysis, given our restriction to scale-invariant mechanisms, it will be sufficient to consider distributions that are normalized so that the single-agent optimal revenue is $\max_q R(q) = 1$. An important family of distributions with revenue normalized to 1 are normalized triangle distributions, which have revenue curves shaped like triangles (Figure 2). Essentially, we will show that these distributions represent the worst-case scenarios that must be considered when designing scale-robust mechanisms.

Definition 4 (Triangle Distribution)**.** *A* normalized triangle distribution with monopoly quantile $\bar{q}$*, denoted* $\mathrm{Tri}_{\bar{q}}$*, is defined by the quantile function*

$$Q_{\mathrm{Tri}_{\bar{q}}}(v) = \begin{cases} \frac{1}{1+v(1-\bar{q})} & v \leq 1/\bar{q} \\ 0 & \textit{otherwise.} \end{cases}$$

The triangulation *of a normalized distribution with monopoly quantile $\bar{q}$ is* $\mathrm{Tri}_{\bar{q}}$*. The family of normalized triangle distributions is* $\mathcal{F}^{\mathrm{Tri}} = \{\mathrm{Tri}_{\bar{q}} : \bar{q} \in [0,1]\}$.

Intuitively, for any monopoly quantile $\bar{q}$, normalized triangle distributions is the distribution that is first order stochastically dominated by any other distribution with monopoly quantile $\bar{q}$. That is, in the single-agent problem, normalized triangle distributions minimize the expected revenue of any given mechanism while maintaining the optimal revenue and monopoly quantile unchanged.

Theorem 1. *For i.i.d., regular, two-agent and single-item environments, the optimal scale-invariant, incentive-compatible mechanism for optimization program* (β) *is* M_{α^*, r^*} *which randomizes over the second-price auction* M_1 *with probability* α^* *and* r^**-markup mechanism* M_{r^*} *with probability* $1 - \alpha^*$*, where* $\alpha^* \approx 0.806$ *and* $r^* \approx 2.447$*. The worst-case regular distribution for this mechanism is triangle distribution* $\mathrm{Tri}_{\bar{q}^*}$ *with* $\bar{q}^* \approx 0.093$ *and its approximation ratio is* $\beta \approx 1.907$.

In the two sections below we prove this theorem with the following main steps.

1. We characterize the optimal scale-robust mechanism under the restriction to stochastic markup mechanisms and triangle distributions. Specifically, we analyze the Nash equilibrium of a zero-sum game between nature and the mechanism designer, where nature selects from triangle distributions and the mechanism designer chooses stochastic markup mechanisms. This equilibrium, characterized in Theorem 2, coincides with the solution provided in Theorem 1.

2. We show that the stochastic markup mechanism and the triangle distribution in Theorem 1 are mutual best responses within the broader families of scale-invariant mechanisms and regular distributions. Specifically, Lemma 6 establishes that stochastic markup mechanisms are the mechanism designer's best response among general scale-invariant mechanisms when facing triangle distributions. Meanwhile, Lemma 7 shows that triangle distributions serve as worst-case distributions among all regular distributions for a family of stochastic markup mechanisms, including the one characterized in Theorem 1. These steps pose a major challenge in the paper, requiring innovative reduction techniques that build upon the concept of revenue curves.

Combining these results gives the theorem.

3.1 Stochastic Markup Mechanisms versus Triangle Distributions

In this section we characterize the solution to the scale-robust analysis framework restricted to stochastic markup mechanisms and triangle distributions. We first define a general family of truncated distributions, which will be important subsequently in the proof. Recall that for scale-invariant mechanisms, it is without loss to normalize the distributions to have monopoly revenue one.

Definition 5 (Truncated Distribution)**.** *A distribution is* truncated *if the highest-point in its support is the monopoly price (typically a point mass). The* truncation *of a distribution is the distribution that replaces every point above the monopoly price with the monopoly price. The family of truncated distributions is denoted* $\mathcal{F}^{\text{Trunc}}$.

Revenue of Various Mechanisms The three lemmas below give formulae for the revenue of the optimal mechanism, the second-price auction, and non-trivial markup mechanisms for triangle distributions. The formula for revenue of markup mechanisms is discontinuous at $r = 1$. Thus, in our discussion we will distinguish between the second-price auction M_1 and non-trivial markup mechanism M_r for $r > 1$.

Lemma 3. *For i.i.d., normalized truncated, two-agent, single-item environments, the optimal mechanism posts the monopoly price and obtains revenue* $2-\bar{q}$ *where* $\bar{q}$ *is the probability that an agent's value equals the monopoly price.*

Proof. The smallest monotone concave function that upper bounds the revenue curve is a trapezoid; its area is $\bar{q}/2 + 1 - \bar{q}$. The optimal revenue from two agents, by Lemma 2, is twice this area, i.e., $2 - \bar{q}$.[9] □

Lemma 4. *The revenue of the second-price auction* M_1 *for distribution* $\text{Tri}_{\bar{q}}$ *is 1, i.e.,* $M_1(\text{Tri}_{\bar{q}}) = 1$.

Proof. By Lemma 2, the revenue is twice the area under the revenue curve. That area is $1/2$; thus, the revenue is 1. □

Lemma 5. *The revenue of the* r*-markup mechanisms* M_r *on triangle distribution* $\text{Tri}_{\bar{q}}$*, for* $r \in (1, \infty)$ *and* $\bar{q} \in [0, 1)$*, is*

$$M_r(\text{Tri}_{\bar{q}}) = \frac{2r}{(1-\bar{q})(r-1)} \left(\frac{1-\bar{q}}{1-\bar{q}+\bar{q}r} + \frac{\ln\left(\frac{r}{1-\bar{q}+\bar{q}r}\right)}{1-r} \right).$$

[9] Equivalently, it is also easy to verify that for truncated distributions, the optimal mechanism is to post a price of $\frac{1}{\bar{q}}$ to both agents, which yields an expected revenue of $2 - \bar{q}$.

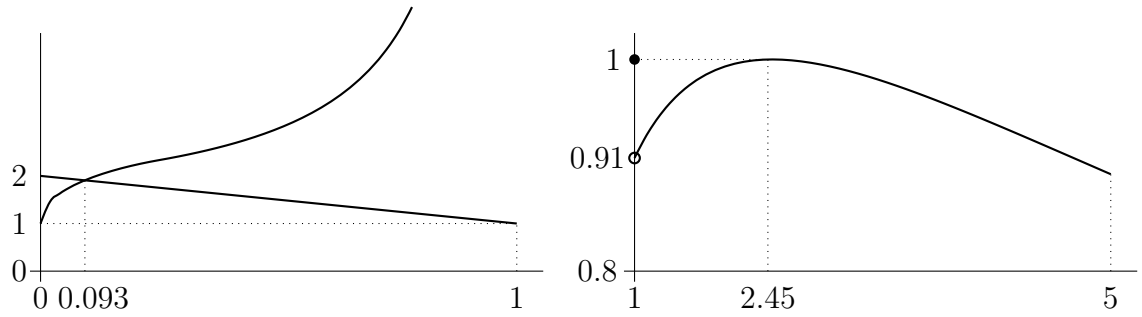


Figure 3: The figure on the left plots, as a function of $\bar{q}$, the approximation ratio $\mathrm{APX}_1(\bar{q})$ of the second-price auction M_1 against triangle distribution $\mathrm{Tri}_{\bar{q}}$ (straight line), and the approximation ratio $\mathrm{APX}_*(\bar{q})$ of the optimal non-trivial markup mechanism against triangle distribution $\mathrm{Tri}_{\bar{q}}$ (curved line). These functions cross at $\bar{q}^* = 0.0931057$. The figure on the right plots the revenue of the r markup mechanism M_r on triangle distribution $\mathrm{Tri}_{\bar{q}^*}$ as a function of markup r, i.e., $M_r(\mathrm{Tri}_{\bar{q}^*})$. Notice that, by choice of $\bar{q}^*$, the optimal non-trivial markup mechanism has the same revenue as the second-price auction.

The proof of Lemma 5 is straightforward and given in Appendix A. These lemmas allow us to numerically compute the expected revenues and approximation ratios of stochastic markup mechanisms given triangular distributions, which are illustrated in Figure 3.

Optimal Stochastic Markup Mechanism The following theorem characterizes the optimal stochastic markup mechanism that is robust to scale against triangle distributions. The parameters of this optimal mechanism are the solution to an algebraic expression (cf. Lemma 5) that we are unable to solve analytically. Our proof will instead combine numeric calculations of select points in parameter space with theoretical analysis to rule out most of the parameter space. For the remaining parameter space, we can show that the expression is well-behaved and, thus, numeric calculation can identify near optimal parameters. Discussion of this hybrid numerical and theoretical analysis can be found in Appendix A.

Theorem 2. *For i.i.d., triangle distribution, two-agent, single-item environments, the optimal stochastic markup mechanism for optimization program* (β) *is* M_{α^*,r^*} *which randomizes over the second-price auction* M_1 *with probability* α^* *and* r^**-markup mechanism* M_{r^*} *with probability* $1-\alpha^*$*, where* $\alpha^* \approx 0.806$ *and* $r^* \approx 2.447$*. The worst-case distribution for this mechanism is the triangle distribution* $\mathrm{Tri}_{\bar{q}^*}$ *with* $\bar{q}^* \approx 0.093$ *and its approximation ratio is* $\beta \approx 1.907$*.*

Intuitively, the optimization program (β) can be viewed as a zero-sum game between

the designer and an adversary, where the designer chooses a mechanism M, the adversary chooses a worst-case distribution F (and its induced revenue curve), and the payoff of the designer is the approximation ratio $\mathrm{OPT}_F(F)/M(F)$ (see Definition 2). The optimal solution to the optimization program (β) is essentially a Nash equilibrium strategy between the designer and the adversary in this zero-sum game.

The high level approach of this proof is to identify the triangle $\mathrm{Tri}_{\bar{q}^*}$ for which the designer is indifferent between the second price auction M_1 and the optimal (non-trivial) markup mechanism, denoted M_{r^*}. For such a distribution $\mathrm{Tri}_{\bar{q}^*}$, the designer is also indifferent (in minimizing the approximation ratio) between any mixture over M_1 (with probability α) and M_{r^*} (with probability $1-\alpha$), and all other r-markup mechanisms for $r \notin \{1, r^*\}$ are inferior (see Figure 3). We then identify the α^* for which the adversary's best response (in maximizing the approximation ratio) to M_{α^*,r^*} is the distribution $\mathrm{Tri}_{\bar{q}^*}$. This solution of M_{α^*,r^*} and $\mathrm{Tri}_{\bar{q}^*}$ is a Nash equilibrium between the designer and adversary and, thus, it solves the optimization problem. The parameters can be numerically identified as $\alpha^* \approx 0.80564048$, $r^* \approx 2.4469452$, $\bar{q}^* \approx 0.0931057$, and the approximation ratio is $\beta \approx 1.9068943$.

3.2 Mutual Best-response of Stochastic Markup Mechanisms and Triangle Distributions

In this section we show that stochastic markup mechanisms are a best response (for the designer) to truncated distributions and that truncated distributions are a best response (for the adversary) to stochastic markup mechanisms. Moreover, we show that among truncated distributions, triangle distributions are the best for the adversary. Triangle distributions are known to be worst case for other questions of interest in mechanism design, e.g., approximation by anonymous reserves and anonymous pricings (Alaei et al., 2018). The proof that triangle distributions are worst-case for two-agent revenue maximization under the scale-robust analysis framework is significantly more involved than these previous results.

3.2.1 Best Response of Stochastic Markup Mechanisms

Lemma 6. *For i.i.d., two-agent, single-item environments and any scale-invariant incentive-compatible mechanism M, there is a stochastic markup mechanism M' with (weakly) higher revenue (and weakly lower approximation ratio) on every truncated distribution F. I.e., $M'(F) \geq M(F)$.*

Proof. In a stochastic markup mechanism the price of the higher agent is a stochastic multiplicative factor $r \geq 1$ of the value of the lower agent (with ties broken randomly). To

prove this theorem we must argue that (a) if the agents are not tied, then revenue improves if the lower agent loses, (b) if the agents are tied, then revenue is unaffected by random tie-breaking, and (c) any such scale-invariant mechanism looks to the higher-valued agent like a stochastic posted pricing with price that is a multiplicative factor (at least one) of the lower-valued agent's value.

To see (a), note that the revenue of the mechanism is equal to its virtual surplus (Lemma 1) and for triangle distributions only the highest value in the support of the distribution has positive virtual value. Thus, any mechanism that sells to a strictly-lower-valued agent can be improved by not selling to such an agent.

To see (b), note that for any i.i.d. distribution the revenue of any mechanism is invariant to randomly permuting the identities of the agents. Thus, we can assume random tie-breaking.

To see (c), recall that the family of incentive-compatible single-agent mechanisms is equivalent to the family of random price postings. Once we have ruled out selling to the lower-valued agent, the mechanism is a single-agent mechanism for the higher-valued agent (with price at least the lower-valued agent's value. By the assumption that the mechanism is scale invariant, the distribution of prices offered to the higher-valued agent must be multiplicative scalings of the lower-valued agent's value. □

3.2.2 Best Response of Triangle Distributions

Next we will give a sequence of results that culminate in the observation that for any regular distribution and any stochastic markup mechanism with probability α at least $2/3$ on the second-price auction (which includes the optimal mechanism from Theorem 2) either the triangulation of the distribution or the point mass Tri_1 has (weakly) higher approximation ratio. As the notation indicates, the point mass distribution Tri_1 is a triangle distribution.

Lemma 7. *For i.i.d., two-agent, single-item environments and any regular distribution F and any stochastic markup mechanism M that places probability $\alpha \in [2/3, 1]$ on the second-price auction, either the triangulation of the distribution F^{Tri} or the point mass Tri_1 has (weakly) higher approximation ratio. I.e.,* $\max\left\{\frac{\text{OPT}_{F^{\text{Tri}}}(F^{\text{Tri}})}{M(F^{\text{Tri}})}, \frac{\text{OPT}_{\text{Tri}_1}(\text{Tri}_1)}{M(\text{Tri}_1)}\right\} \geq \frac{\text{OPT}_F(F)}{M(F)}$.

To prove this lemma we give a sequence of results showing that for any regular distribution, a corresponding truncated distribution is only worse; for any truncated distribution and a fixed stochastic markup mechanism (that mixes over M_1 and some M_r), a corresponding quadrilateral distribution (based on r) is only worse; and for any quadrilateral distribution, a corresponding triangle distribution (independent of r) is only worse. The theorem follows from combining these results. The first step assumes that the probability

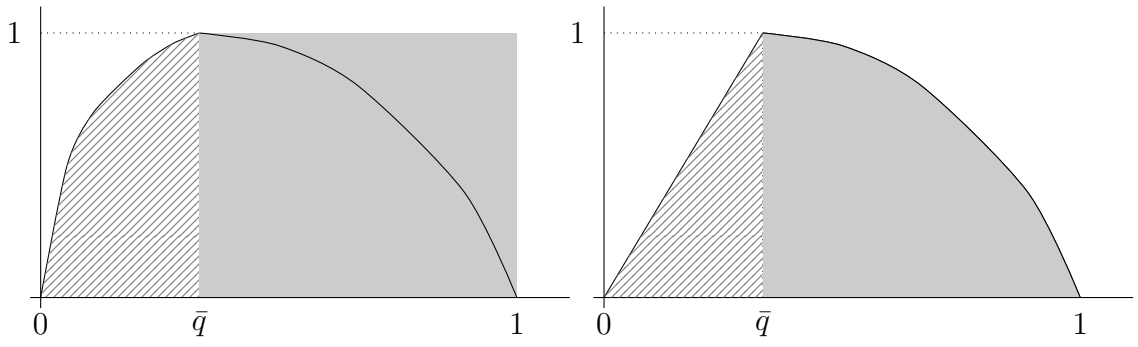


Figure 4: The illustration of the revenue decomposition of Lemma 8 for M on distribution F and truncation F' for the optimal mechanism and second-price auction. The thin black line on the left and right figures are the revenue curves corresponding to F and F', respectively. The dashed area on the left represents $\mathrm{OPT}_+ = \mathrm{SPA}_+$ and the gray area on the left represents $\mathrm{OPT}_- = \mathrm{OPT}'_-$. The dashed area on the right represents $\mathrm{OPT}'_+ = \mathrm{SPA}'_+$ and the gray area on the right represents $\mathrm{SPA}'_- = \mathrm{SPA}_-$.

that the stochastic markup mechanism places on the second price auction is $\alpha \in [1/2, 1]$; the last step further assumes that $\alpha \in [2/3, 1]$.

Best response of truncated distributions To begin, the following lemma shows that the best response of the adversary to a relevant stochastic markup mechanism is a truncated distribution. Recall that by Fu et al. (2015) the optimal scale-robust mechanism is strictly better than a 2-approximation. On the other hand, any stochastic markup mechanism that places probability α on the second-price auction M_1 has approximation ratio at least $1/\alpha$. Specifically, on the (degenerate) distribution that places all probability mass on 1, a.k.a. Tri_1, the approximation factor of such a stochastic markup mechanism is exactly $1/\alpha$. We conclude that all relevant stochastic markup mechanisms place probability $\alpha > 1/2$ on the second-price auction. Thus, this lemma applies to all relevant mechanisms.

Lemma 8. *For i.i.d., two-agent, single-item environments, any regular distribution F, and any stochastic markup mechanism M that places probability $\alpha \in [1/2, 1]$ on the second-price auction; either the truncation of the distribution F' or the point mass distribution* Tri_1 *has (weakly) higher approximation ratio. I.e.,* $\max\left\{\frac{\mathrm{OPT}_{F'}(F')}{M(F')}, \frac{\mathrm{OPT}_{\mathrm{Tri}_1}(\mathrm{Tri}_1)}{M(\mathrm{Tri}_1)}\right\} \geq \frac{\mathrm{OPT}_F(F)}{M(F)}$.

Proof. It can be assumed that the approximation of stochastic markup mechanism M on distribution F is at least $1/\alpha$ (where α denotes the probability that M places on the second-price auction). Notice that the revenue M on the point mass on 1 (a truncated distribution) is α and the optimal revenue on this distribution is 1. If the approximation factor

$\mathrm{OPT}_F(F)/M(F)$ is less than $1/\alpha$ then the point mass on 1 (a truncated distribution) achieves a higher approximation than F and the lemma follows. For the remainder of the proof, assume that the approximation factor of mechanism M on distribution F is more than $1/\alpha$.

View the stochastic markup mechanism M as a distribution over two mechanisms: the second-price auction M_1 with probability α, and M_*, a distribution over non-trivial markup mechanisms M_r with $r > 1$, with probability $1 - \alpha$. The optimal mechanism is OPT_F. Decompose the revenue from distribution F across these three mechanisms as follows. Denote the monopoly quantile of F by $\bar{q}$. See Figure 4.

- OPT_+ and OPT_- give the expected revenue of the optimal mechanism from each agent when their opponent has value above and below the monopoly price $V_F(\bar{q})$.
- $\mathrm{SPA}_+ = \mathrm{OPT}_+$ and SPA_- give the expected revenue of the second-price auction M_1 from each agent when their opponent has value above and below the monopoly price.
- MKUP_+ and MKUP_- give the expected revenue of the stochastic markup mechanism M_* when the realized prices are (strictly) above and (weakly) below the monopoly price.

Consider truncating the distribution F at the monopoly quantile $\bar{q}$ to obtain $F' \in \mathcal{F}^{\mathrm{Trunc}}$. Define analogous quantities (with identities):

- $\mathrm{OPT}'_+ < \mathrm{OPT}_+$ and $\mathrm{OPT}'_- = \mathrm{OPT}_-$.

  Identities follow from the geometric analysis of Lemma 2.
- $\mathrm{SPA}'_+ = \mathrm{OPT}'_+$ and $\mathrm{SPA}'_- = \mathrm{SPA}_-$.

  Identities follow from the geometric analysis of Lemma 2.
- $\mathrm{MKUP}'_+ = 0$ and $\mathrm{MKUP}'_- = \mathrm{MKUP}_-$.

  Values above the monopoly price are not supported by the truncated distribution, so the revenue from those prices is zero. On the other hand, prices (weakly) below the monopoly price are bought with the exact same probability as the cumulative distribution function F' and F are the same for these prices.

The remainder of the proof follows a straightforward calculation. Write the approximation ratio of M on distribution F (using the given identities) and rearrange:

$$\begin{aligned}\frac{\mathrm{OPT}_F(F)}{M(F)} &= \frac{\mathrm{OPT}_+ + \mathrm{OPT}_-}{\alpha\left(\mathrm{OPT}_+ + \mathrm{SPA}_-\right) + (1-\alpha)\left(\mathrm{MKUP}_+ + \mathrm{MKUP}_-\right)} \\ &= \frac{\mathrm{OPT}_+ + [\mathrm{OPT}_-]}{\alpha\ \mathrm{OPT}_+ + [\alpha\,\mathrm{SPA}_- + (1-\alpha)\left(\mathrm{MKUP}_+ + \mathrm{MKUP}_-\right)]}\end{aligned}$$

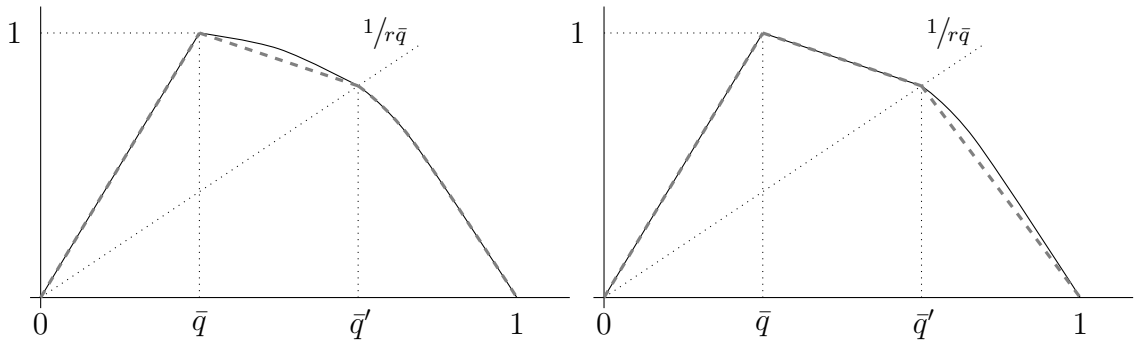


Figure 5: The main two steps of Lemma 10 are illustrated. In the first step (right-hand side), the revenue curves of distributions F^{Trunc} (thin, solid, black) and $F^{\dagger}$ (thick, dashed, gray) are depicted. In the second step, the revenue curves of the distributions $F^{\dagger}$ (thin, solid, black) and F^{Qr} (thick, dashed, gray) are depicted. In both cases the revenue of the r-markup mechanism is is higher on the thin, solid, black curve than the thick, dashed, gray curve.

Since the approximation ratio on F is at least $1/\alpha$, the ratio of the first term in the numerator and denominator is at most the ratio of the remaining terms [in brackets]:

$$\frac{1}{\alpha} = \frac{\text{OPT}_+}{\alpha\ \text{OPT}_+} \leq \frac{[\text{OPT}_-]}{[\alpha\,\text{SPA}_- + (1-\alpha)\,(\text{MKUP}_+ + \text{MKUP}_-)]}$$

Now write the approximation ratio of M on truncation F' (using the given identities) and bound:

$$\begin{aligned}
\frac{\text{OPT}_{F'}(F')}{M(F')} &= \frac{\text{OPT}'_+ + [\text{OPT}_-]}{\alpha\ \text{OPT}'_+ + [\alpha\ \text{SPA}_- + (1-\alpha)\ \text{MKUP}_-]} \\
&\geq \frac{\text{OPT}'_+ + [\text{OPT}_-]}{\alpha\ \text{OPT}'_+ + [\alpha\ \text{SPA}_- + (1-\alpha)\,(\text{MKUP}_+ + \text{MKUP}_-)]} \\
&\geq \frac{\text{OPT}_+ + [\text{OPT}_-]}{\alpha\ \text{OPT}_+ + [\alpha\ \text{SPA}_- + (1-\alpha)\,(\text{MKUP}_+ + \text{MKUP}_-)]} \\
&= \frac{\text{OPT}_F(F)}{M(F)}.
\end{aligned}$$

The calculation shows that, for any distribution F, the truncated distribution F' increases the approximation factor of the stochastic markup mechanism. Thus, the worst-case distribution is truncated. □

Best response of quadrilateral distributions The next step is to show that, among truncated distributions, the worst-case distribution for stochastic markup mechanisms are those with quadrilateral-shaped revenue curves, i.e., ones that are piecewise linear with

three pieces (see Figure 2). Recall that for a truncated distribution at monopoly quantile $\bar{q}$, the upper bound of the support is a point mass on $1/\bar{q}$.

Definition 6 (Quadrilateral Distribution)**.** *A* normalized quadrilateral distribution *with parameters* $\bar{q}, \bar{q}'$ *and* r *with* $r \geq 1$ *and* $\frac{\bar{q}r}{\bar{q}r+(1-\bar{q})} \leq \bar{q}' \leq \min\{r\bar{q}, 1\}$*, denoted by* $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ *is defined by quantile function as:*

$$Q_{\mathrm{Qr}_{\bar{q},\bar{q}',r}}(v) = \begin{cases} \frac{\bar{q}'}{\bar{q}'+vr\bar{q}(1-\bar{q}')} & v < 1/r\bar{q} \\ \frac{\bar{q}'\bar{q}(r-1)}{vr\bar{q}(\bar{q}'-\bar{q})+(r\bar{q}-\bar{q}')} & 1/r\bar{q} \leq v \leq 1/\bar{q} \\ 0 & 1/\bar{q} < v \end{cases}$$

The following lemma summarizes an analysis from Allouah and Besbes (2018) and is useful in bounding the revenue from markup mechanisms.

Lemma 9 (Allouah and Besbes, 2018)**.** *Consider the* r*-markup mechanism, two i.i.d. regular agents with value distribution* F*, quantile* $\bar{q}'$ *corresponding to the monopoly price divided by* r*, and the distribution* $\tilde{F}$ *that corresponds to* F *ironed on* $[\bar{q}', 1]$*: the virtual surplus from quantiles* $[\bar{q}', 1]$ *is higher for* F *than for* $\tilde{F}$*.*

Proof. The proof of this lemma is technical and non-trivial. It is given in the proof of Proposition 4 of Allouah and Besbes (2018). □

The next lemma reduces the worst case distribution from the family of truncated distributions to the family of quadrilateral distributions. The reduction is illustrated in Figure 5, by showing that ironing the revenue curves sequentially within $[\bar{q}, \bar{q}']$ and $[\bar{q}', 1]$ decreases the revenue of the stochastic markup mechanism. The optimal revenue is not affected because it is obtained using a reserve price corresponding to the monopoly quantile $\bar{q}$ and it is agnostic to the shape of the revenue curve for $q > \bar{q}$.

Lemma 10. *For i.i.d., two-agent, single-item environments, any truncated distribution* F^{Trunc}*, and any stochastic markup mechanism* $M_{\alpha,r}$ *with probability* α *on the second-price auction* M_1 *and probability* $1-\alpha$ *on non-trivial markup mechanism* M_r*; there is a quadrilateral distribution* F^{Qr} *with the same optimal revenue and (weakly) lower revenue in* $M_{\alpha,r}$*. I.e.,* $\mathrm{OPT}_{F^{\mathrm{Qr}}}(F^{\mathrm{Qr}}) = \mathrm{OPT}_{F^{\mathrm{Trunc}}}(F^{\mathrm{Trunc}})$ *and* $M_{\alpha,r}(F^{\mathrm{Qr}}) \leq M_{\alpha,r}(F^{\mathrm{Trunc}})$*.*

Proof. On any normalized truncated distribution with monopoly quantile $\bar{q}$, the optimal revenue is $2-\bar{q}$ (Lemma 3). Thus, to prove the lemma it is sufficient to show that for any truncated distribution $F^{\mathrm{Trunc}} \in \mathcal{F}^{\mathrm{Trunc}}$ with monopoly quantile $\bar{q}$ there is a normalized quadrilateral distribution $F^{\mathrm{Qr}} \in \mathcal{F}^{\mathrm{Qr}} \subset \mathcal{F}^{\mathrm{Trunc}}$ with monopoly quantile $\bar{q}$ and lower revenue in $M_{\alpha,r}$.

The quadrilateral distribution F^{Qr} is obtained by ironing F^{Trunc} on $[\bar{q}, \bar{q}']$ and $[\bar{q}', 1]$ where quantile $\bar{q}'$ satisfies $V_{F^{\text{Trunc}}}(\bar{q}) = r\, V_{F^{\text{Trunc}}}(\bar{q}')$. We consider an intermediary distribution $F^{\dagger}$ that is F^{Trunc} ironed only on $[\bar{q}, \bar{q}']$. See Figure 5. The proof approach is to show that $M_{\alpha,r}(F^{\text{Trunc}}) > M_{\alpha,r}(F^{\dagger}) > M_{\alpha,r}(F^{\text{Qr}})$.

As $M_{\alpha,r}$ is a convex combination of the second-price auction M_1 and the r-markup mechanism M_r. It suffices to show the inequalities above hold for both auctions. In fact, the result holds for the second-price auction from the geometric analysis of revenue of Lemma 2. The revenue of the second-price auction for two i.i.d. agents is twice the area under the revenue curve. As the revenue curve has strictly smaller area from F^{Trunc} to $F^{\dagger}$ to F^{Qr}, we have $M_1(F^{\text{Trunc}}) > M_1(F^{\dagger}) > M_1(F^{\text{Qr}})$. Below, we analyze the r-markup mechanism M_r.

The following price-based analysis shows that $M_r(F^{\text{Trunc}}) > M_r(F^{\dagger})$:

- The revenue from quantiles in $[0, \bar{q}]$ is unchanged.

  These quantiles are offered prices from quantiles in $[\bar{q}', 1]$. The values of quantiles $[0, \bar{q}]$ and $[\bar{q}', 1]$ are the same for both distributions; thus, the revenue is unchanged.

- The revenue from quantiles in $[\bar{q}, \bar{q}']$ decreases.

  These quantiles are offered prices from quantiles in $[\bar{q}', 1]$. For the distribution $F^{\dagger}$ relative to F^{Trunc}: Values are lower at any quantile $q \in [\bar{q}, \bar{q}']$; the distribution of prices (from quantiles in $[\bar{q}', 1]$) is the same. Thus, revenue is lower.

- The revenue from quantiles in $[\bar{q}', 1]$ is unchanged.

  These quantiles are in $[\bar{q}', 1]$ and are offered prices from quantiles in $[\bar{q}', 1]$. The distributions are the same for these quantiles; thus, the revenue is unchanged.

The following virtual-surplus-based analysis shows that $M_r(F^{\dagger}) > M_r(F^{\text{Qr}})$:

- The virtual surplus of quantiles in $[0, \bar{q}]$ is unchanged.

  These quantiles have the same virtual values under the two distributions and the same probability of winning, i.e., $1 - \bar{q}'$ (when the other agent's quantile is in $[\bar{q}', 1]$.

- The virtual surplus of quantiles in $[\bar{q}, \bar{q}']$ is decreased.

  Their prices come from quantiles in $[\bar{q}', 1]$ which are decreased; thus, their probabilities of winning are increased. Their virtual values are negative, so these increased probabilities of winning result in decreased virtual surplus.

- The virtual surplus of quantiles in $[\bar{q}', 1]$ is decreased.

  This result is given by Lemma 9. □

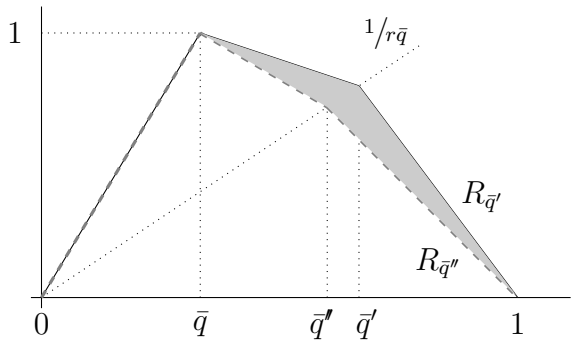


Figure 6: Illustrating the proof of Lemma 11, the difference of revenue for second price auction M_1 on revenue curves $R_{\bar{q}'}$ and $R_{\bar{q}''}$, which respectively correspond to quadrilateral distributions $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ and $\mathrm{Qr}_{\bar{q},\bar{q}'',r}$, is equal to twice of the gray area, which is at least $\bar{q}'' - \bar{q}'$. Moreover, the difference of revenue for the r-markup mechanism M_r on revenue curves $R_{\bar{q}'}$ and $R_{\bar{q}''}$ is at most $2(\bar{q}'' - \bar{q}')$.

Best response of triangle distributions We complete the proof of Lemma 7 by showing that triangle distributions lead to lower revenue than quadrilateral distributions. The intuition of the proof is illustrated in Figure 6. For any $r > 1$ and any stochastic markup mechanism $M_{\alpha,r}$ with probability $\alpha \in [2/3, 1]$, consider a family of quadrilateral distributions $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ parameterized by $\bar{q}'$. The optimal revenue is again not affected by $\bar{q}'$ while the revenue of $M_{\alpha,r}$ is monotone increasing in $\bar{q}'$. Thus the approximation ratio of $M_{\alpha,r}$ is maximized by minimal $\bar{q}'$ for which the degenerate quadrilateral $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ is a triangle.

Lemma 11. *For i.i.d., two-agent, single-item environments, normalized quadrilateral distribution* $\mathrm{Qr}_{\bar{q},\bar{q}',r}$*, and stochastic markup mechanism* $M_{\alpha,r}$ *with probability* $\alpha \in [2/3, 1]$ *on the second-price auction* M_1 *and probability* $1 - \alpha$ *on non-trivial markup mechanism* M_r*; the triangle distribution* $\mathrm{Tri}_{\bar{q}}$ *has the same optimal revenue and (weakly) lower revenue in* $M_{\alpha,r}$*. I.e.,* $\mathrm{OPT}_{\mathrm{Tri}_{\bar{q}}}(\mathrm{Tri}_{\bar{q}}) = \mathrm{OPT}_{\mathrm{Qr}_{\bar{q},\bar{q}',r}}(\mathrm{Qr}_{\bar{q},\bar{q}',r})$ *and* $M_{\alpha,r}(\mathrm{Tri}_{\bar{q}}) \le M_{\alpha,r}(\mathrm{Qr}_{\bar{q},\bar{q}',r})$.

Proof. By Lemma 3, the optimal revenues for quadrilateral distribution $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ and triangle distribution $\mathrm{Tri}_{\bar{q}}$ are the same (and equal to $2 - \bar{q}$). To show that the revenue of $M_{\alpha,r}$ is worse on $\mathrm{Tri}_{\bar{q}}$ than $\mathrm{Qr}_{\bar{q},\bar{q}',r}$, it suffices to show that the revenue on $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ is monotonically increasing in $\bar{q}'$. Specifically the minimum revenue is when the quadrilateral distribution is degenerately equal to the triangle distribution.

The proof strategy is to lower bound the partial derivative with respect to $\bar{q}'$ of the revenues of r-markup mechanism and the second-price auction for quadrilateral distributions

$\mathrm{Qr}_{\bar{q},\bar{q}',r}$ as

$$\frac{\partial M_r(\mathrm{Qr}_{\bar{q},\bar{q}',r})}{\partial \bar{q}'} \geq -2, \tag{1}$$

$$\frac{\partial M_1(\mathrm{Qr}_{\bar{q},\bar{q}',r})}{\partial \bar{q}'} \geq 1. \tag{2}$$

Thus, for mechanism $M_{\alpha,r}$ with $\alpha \geq 2/3$, we have

$$\frac{\partial M_{\alpha,r}(\mathrm{Qr}_{\bar{q},\bar{q}',r})}{\partial \bar{q}'} \geq \alpha - 2(1-\alpha) \geq 0$$

and revenue is minimized with the smallest choice of $\bar{q}'$ for which quadrilateral distribution $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ is degenerately a triangle distribution. It remains to prove bounds (1) and (2).

For simplicity, since in this section the only parameter we change in distribution $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ is $\bar{q}'$, we introduce the notation $P_{\bar{q}'}(v)$ to denote the revenue of posting price v, and $V_{\bar{q}'}(q)$ to denote the price v given quantile q when the distribution is $\mathrm{Qr}_{\bar{q},\bar{q}',r}$. The proof is illustrated in Figure 6.

We now prove bound (1). For any pair of quadrilateral distributions $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ and $\mathrm{Qr}_{\bar{q},\bar{q}'',r}$ with $\bar{q}'' \geq \bar{q}'$, we analyze the difference in revenue for posting price $r \cdot v_{(2)}$.

$$\begin{aligned}
&M_r(\mathrm{Qr}_{\bar{q},\bar{q}'',r}) - M_r(\mathrm{Qr}_{\bar{q},\bar{q}',r}) \\
&\quad = 2\int_{\bar{q}''}^{1} P_{\bar{q}''}(r \cdot V_{\bar{q}''}(q))\,dq - 2\int_{\bar{q}'}^{1} P_{\bar{q}'}(r \cdot V_{\bar{q}'}(q))\,dq \\
&\quad \geq 2\int_{\bar{q}''}^{1} P_{\bar{q}''}(r \cdot V_{\bar{q}''}(q))\,dq - 2\int_{\bar{q}''}^{1} P_{\bar{q}'}(r \cdot V_{\bar{q}'}(q))\,dq - 2(\bar{q}'' - \bar{q}') \\
&\quad \geq 2\int_{\bar{q}''}^{1} P_{\bar{q}'}(r \cdot V_{\bar{q}''}(q))\,dq - 2\int_{\bar{q}''}^{1} P_{\bar{q}'}(r \cdot V_{\bar{q}'}(q))\,dq - 2(\bar{q}'' - \bar{q}') \\
&\quad \geq -2(\bar{q}'' - \bar{q}').
\end{aligned}$$

The first equality is constructed as follows: Both agents face a random price that is r times the value of the other agent who has quantile q drawn from $U[0,1]$. The revenue from this price is given by, e.g., $P_{\bar{q}'}(r \cdot V_{\bar{q}'}(q))$ which is 0 when $q \leq \bar{q}'$. The first inequality holds because $P_{\bar{q}'}(r \cdot V_{\bar{q}'}(q)) \leq 1$ for any quantile q. The second inequality holds since the revenue from revenue curve $P_{\bar{q}''}$ is weakly higher than from revenue curve $P_{\bar{q}'}$ for any value v. The third inequality holds because (a) the prices of the first integral are higher than the prices of the second integral, i.e., $V_{\bar{q}''}(q) \geq V_{\bar{q}'}(q)$ for every q, and (b) because these prices are below the monopoly price for distribution $\mathrm{Qr}_{\bar{q},\bar{q}'',r}$ and so higher prices give higher revenue.

Therefore, we have

$$\frac{\partial M_r(\mathrm{Qr}_{\bar{q},\bar{q}',r})}{\partial \bar{q}'} = \lim_{\bar{q}''\to\bar{q}'} \frac{M_r(\mathrm{Qr}_{\bar{q},\bar{q}'',r}) - M_r(\mathrm{Qr}_{\bar{q},\bar{q}',r})}{\bar{q}'' - \bar{q}'} \geq -2.$$

We now prove bound (2). The revenue of the second price auction for two i.i.d. agents is twice the area under the revenue curve (Lemma 2). For quadrilateral distribution $\mathrm{Qr}_{\bar{q},\bar{q}',r}$ this revenue is calculated as:

$$\begin{aligned}
M_1(\mathrm{Qr}_{\bar{q},\bar{q}',r}) &= 2\int_0^1 R_{\bar{q}'}(q)\,dq \\
&= 2\int_0^{\bar{q}} R_{\bar{q}'}(q)\,dq + 2\int_{\bar{q}}^{\bar{q}'} R_{\bar{q}'}(q)\,dq + 2\int_{\bar{q}'}^1 R_{\bar{q}'}(q)\,dq \\
&= \bar{q} + (\bar{q}' - \bar{q})(1 + \frac{\bar{q}'}{r\cdot\bar{q}}) + (1-\bar{q}')\frac{\bar{q}'}{r\cdot\bar{q}} \\
&= \bar{q}' + (1-\bar{q})\frac{\bar{q}'}{r\cdot\bar{q}}.
\end{aligned}$$

Therefore, we have

$$\frac{\partial M_1(\mathrm{Qr}_{\bar{q},\bar{q}',r})}{\partial \bar{q}'} = 1 + \frac{1-\bar{q}}{r\cdot\bar{q}} \geq 1. \qquad \square$$

4 Conclusions

This paper introduces a framework for designing scale-robust auctions, ensuring optimal multiplicative revenue approximation across different valuation scales. We identify the optimal mechanism within this framework, which randomizes between the second-price auction and an auction that marks up the second-highest bid by a factor of approximately 2.45. This mechanism outperforms existing prior-independent approaches and provides a robust solution for small-market settings where distributional knowledge is limited. This characterization of the optimal mechanism provides insights into how auctioneers can design robust mechanisms without reliance on detailed distributional knowledge.

Future research can explore generalizations of scale-robust mechanisms beyond the two-bidder setting, as well as applications to multi-unit and combinatorial auctions. In particular, for environments with more than two bidders, an important open question is whether randomly marking up the second-highest bid suffices for robust optimality or if more complex pricing strategies, incorporating the third- or fourth-highest bids, are necessary. Additionally, understanding how the optimal random markup changes with respect to the tail of the distribution is an interesting direction, given that for MHR distributions, a fixed price

suffices, while for heavier-tailed regular distributions, random markups are necessary for optimality. Another promising avenue is to examine how scale-robust mechanisms perform in dynamic or repeated auction environments.

References

Alaei, S., Hartline, J., Niazadeh, R., Pountourakis, E., and Yuan, Y. (2018). Optimal auctions vs. anonymous pricing. *Games and Economic Behavior*.

Allouah, A., Bahamou, A., and Besbes, O. (2022). Pricing with samples. *Operations Research*, 70(2):1088–1104.

Allouah, A. and Besbes, O. (2018). Prior-independent optimal auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 503–503. ACM.

Anunrojwong, J., Balseiro, S., and Besbes, O. (2022). On the robustness of second-price auctions in prior-independent mechanism design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 151–152.

Bachrach, N. and Talgam-Cohen, I. (2022). Distributional robustness: From pricing to auctions. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, page 150.

Bergemann, D. and Morris, S. (2005). Robust mechanism design. *Econometrica*, pages 1771–1813.

Bergemann, D. and Schlag, K. (2008). Pricing without priors. *Journal of the European Economic Association*, 6(2-3):560–569.

Bergemann, D. and Schlag, K. (2011). Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *The American Economic Review*, 86(1):180–194.

Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90.

Caillaud, B. and Robert, J. (2005). Implementation of the revenue-maximizing auction by an ignorant seller. *Review of Economic Design*, 9:127–143.

Carrasco, V., Luz, V. F., Kos, N., Messner, M., Monteiro, P., and Moreira, H. (2018). Optimal selling mechanisms under moment conditions. *Journal of Economic Theory*, 177:245–279.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488.

Carroll, G. and Segal, I. (2019). Robustly optimal auctions with unknown resale opportunities. *The Review of Economic Studies*, 86(4):1527–1555.

Che, E. (2022). Robustly optimal auction design under mean constraints. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 153–181.

Chung, K.-S. and Ely, J. C. (2007). Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2015). Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333.

Feng, Y. and Hartline, J. D. (2018). An end-to-end argument in mechanism design (prior-independent auctions for budgeted agents). In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 404–415. IEEE.

Feng, Y., Hartline, J. D., and Li, Y. (2021). Revelation gap for pricing from samples. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1438–1451.

Fu, H., Immorlica, N., Lucier, B., and Strack, P. (2015). Randomization beats second price as a prior-independent auction. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 323–323. ACM.

Guo, Y. and Shmaya, E. (2019). Robust monopoly regulation. *R&R at American Economic Review*.

Guo, Y. and Shmaya, E. (2022). Regret-minimizing project choice. *accepted at Econometrica*.

Hartline, J. D. and Roughgarden, T. (2008). Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM Symposium on Theory of Computing*, pages 75–84.

Hartline, J. D. and Roughgarden, T. (2014). Optimal platform design. *arXiv preprint arXiv:1412.8518*.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6:58–73.

Ronen, A. (2001). On approximating optimal auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 11–17. ACM.

A Missing Proofs from Section 3

Proof of Theorem 2. As discussed in Section 3.1, we first identify the triangle distribution $\bar{q}^*$ and the r^* for which M_1 and M_{r^*} obtain the same ratio. Denote the approximation ratio for the second-price auction M_1 as $\text{APX}_1(\bar{q}) = 2-\bar{q}$ (the ratio of Lemma 3 to Lemma 4), which is continuous in $\bar{q}$. Denote the approximation ratio of the optimal markup mechanism against distribution $\text{Tri}_{\bar{q}}$ by $\text{APX}_*(\bar{q}) = \sup_{r>1} \frac{\text{OPT}_{\text{Tri}_{\bar{q}}}(\text{Tri}_{\bar{q}})}{M_r(\text{Tri}_{\bar{q}})}$. By Lemma 5, the approximation ratio $\text{APX}_*(\bar{q})$ is continuous in $\bar{q}$ as well. It is easy to verify that $\text{APX}_1(0) = 2 > \text{APX}_*(0) = 1$ while $\text{APX}_1(1) = 1 < \text{APX}_*(1) = \infty$. By continuity, there exists a $\bar{q}^*$ where these two functions cross, i.e., $\text{APX}_*(\bar{q}^*) = \text{APX}_1(\bar{q}^*)$. See Figure 3. By numerical calculation, $\bar{q}^* \approx 0.0931057$, and

$$r^* = \operatorname*{argmax}_{r>1} \frac{\text{OPT}_{\text{Tri}_{\bar{q}^*}}(\text{Tri}_{\bar{q}^*})}{M_r(\text{Tri}_{\bar{q}^*})} \approx 2.4469452.$$

The details of all numerical calculations are provided in the remainder of this section.

Now, fixing r^*, we search for α^* for which the adversary maximizes the approximation ratio of mechanism M_{α^*,r^*} by selecting triangle distribution $\text{Tri}_{\bar{q}^*}$. Denote by $\bar{q}_r(\alpha)$ the monopoly quantile as a function of α for the triangle distribution that maximizes the approximation ratio of mechanism $M_{\alpha,r}$, i.e.,

$$\bar{q}_r(\alpha) = \operatorname*{argmax}_{\bar{q}} \frac{\text{OPT}_{\text{Tri}_{\bar{q}}}(\text{Tri}_{\bar{q}})}{M_{\alpha,r}(\text{Tri}_{\bar{q}})}.$$

By numerical calculation, for any $r \in [2.445, 2.449]$, $\bar{q}_r(0.81) < \bar{q}^* < \bar{q}_r(0.8)$. Continuity of $\bar{q}_r(\cdot)$ for $r \in [2.445, 2.449]$ and $\alpha \in [0.8, 0.81]$ (formally proved in Appendix A.3), then implies that there exists α^* such that $\bar{q}_{r^*}(\alpha^*) = \bar{q}^*$. By numerical calculation, $\alpha^* \approx 0.80564048$. □

To identify the optimal mechanism for triangle distributions, we evaluate the ratio of revenues of markup mechanisms on triangle distributions to the optimal revenue. For distribution $\text{Tri}_{\bar{q}}$ the optimal revenue is $2 - \bar{q}$ (Lemma 3). The revenue for r-markup mechanism is calculated by Lemma 5. In this appendix, we drive the formula of Lemma 5 and show that it has bounded partial derivatives in both markup r and monopoly quantile $\bar{q}$. We then describe the details of the hybrid numerical and analytical argument of Theorem 2. Finally we give the proof of continuity of the adversary's best response distribution to the probability the mechanism places on the second-price auction.

A.1 Derivation and smoothness of Lemma 5

Proof of Lemma 5. Denote the quantile corresponding to the price $r\,V_{\mathrm{Tri}_{\bar{q}}}(q)$ for markup $r > 1$ as

$$\hat{Q}(q,r) = Q_{\mathrm{Tri}_{\bar{q}}}(r\,V_{\mathrm{Tri}_{\bar{q}}}(q)) = \begin{cases} \frac{q}{r-qr+q} & \text{if } r\,V_{\mathrm{Tri}_{\bar{q}}}(q) \le 1/\bar{q}, \\ 0 & \text{otherwise.} \end{cases}$$

When the quantile of the second highest agent is smaller than $\hat{Q}(\bar{q}, 1/r)$, the price $r \cdot v_{(2)}$ is higher than the support of the valuation distribution. Therefore, the revenue of posting price $r \cdot v_{(2)}$ to the highest bidder is

$$\begin{aligned} M_r(\mathrm{Tri}_{\bar{q}}) &= 2r \int_{\hat{Q}(\bar{q},1/r)}^{1} V_{\mathrm{Tri}_{\bar{q}}}(q)\hat{Q}(q,r)\,dq \\ &= 2r \int_{\hat{Q}(\bar{q},1/r)}^{1} \frac{1-q}{1-\bar{q}} \cdot \frac{1}{r-qr+q}\,dq = \frac{2r}{1-\bar{q}} \left[\frac{q}{r-1} + \frac{\ln(r-qr+q)}{(r-1)^2}\right]_{\frac{\bar{q}}{1/r-\bar{q}/r+\bar{q}}}^{1} \\ &= \frac{2r}{(1-\bar{q})(r-1)} \left(\frac{1-\bar{q}}{1-\bar{q}+\bar{q}r} - \frac{\ln\left(\frac{r}{1-\bar{q}+\bar{q}r}\right)}{r-1}\right), \end{aligned}$$

where the second equality holds just by the definition of the distribution. □

Consider the revenue of r-markup mechanism on the triangle distribution $\mathrm{Tri}_{\bar{q}}$ as a function of $r \in (1,\infty)$ and $\bar{q} \in [0,1]$. The formula for this revenue is given by Lemma 5. The following two claims show that the ratio of revenues has bounded partial derivative with respect to both $r \in (1,\infty)$ and $\bar{q} \in [0,1]$ and, thus, numerical evaluation of the revenue at selected parameters allows large regions of parameter space to be ruled out.

Claim 1. *For any distribution F and any constants $1 \le r_1 \le r_2$, we have $M_{r_1}(F) \ge {r_1}/{r_2}\, M_{r_2}(F)$.*

Claim 2. *For any mechanism M_r with $r \ge 1$, and any constants $0 \le \bar{q}_1 \le \bar{q}_2 < 1$, we have ${(1-\bar{q}_2)}/{(1-\bar{q}_1)}\, M_r(\mathrm{Tri}_{\bar{q}_2}) \le M_r(\mathrm{Tri}_{\bar{q}_1}) \le 2(\bar{q}_2 - \bar{q}_1) + M_r(\mathrm{Tri}_{\bar{q}_2})$.*

Proof of Claim 1. For any realized valuation profile, if the item is sold in mechanism M_{r_2}, then the item is sold in mechanism M_{r_1} since the price posted to the highest agent is smaller in mechanism M_{r_1}. Moreover, when the item is sold in mechanism M_{r_1}, the payment from agent with highest value is at least ${r_1}/{r_2}$ fraction of the payment in mechanism M_{r_2}. Taking expectation over the valuation profiles, we have $M_{r_1}(F) \ge {r_1}/{r_2} \cdot M_{r_2}(F)$. □

Proof of Claim 2. Consider $\hat{Q}(\cdot,\cdot)$ as defined in the proof of Lemma 5, above. By directly

comparing the revenue from two distributions,

$$\begin{aligned} M_r(\mathrm{Tri}_{\bar{q}_1}) &= 2r \int_{\hat{Q}(\bar{q}_1, 1/r)}^{1} V_{\mathrm{Tri}_{\bar{q}_1}}(q)\, \hat{Q}(q,r)\, dq \\ &\leq 2(-\hat{Q}(\bar{q}_1, 1/r) + \hat{Q}(\bar{q}_2, 1/r)) + 2r \int_{\hat{Q}(\bar{q}_2, 1/r)}^{1} V_{\mathrm{Tri}_{\bar{q}_1}}(q)\, \hat{Q}(q,r)\, dq \\ &\leq 2(\bar{q}_2 - \bar{q}_1) + 2r \int_{\hat{Q}(\bar{q}_2, 1/r)}^{1} V_{\mathrm{Tri}_{\bar{q}_2}}(q)\, \hat{Q}(q,r)\, dq \\ &= 2(\bar{q}_2 - \bar{q}_1) + M_r(\mathrm{Tri}_{\bar{q}_2}). \end{aligned}$$

The first equality holds because the quantile of $V_{\mathrm{Tri}_{\bar{q}_1}}(q) \cdot r$ is 0 for $q < \hat{Q}(\bar{q}_1, 1/r)$. The first inequality holds because $r \cdot V_{\mathrm{Tri}_{\bar{q}_1}}(q)\hat{Q}(q,r) \leq 1$ for any quantile q. The second inequality holds because $V_{\mathrm{Tri}_{\bar{q}_1}}(q) \leq V_{\mathrm{Tri}_{\bar{q}_2}}(q)$ for $\bar{q}_1 \leq \bar{q}_2$ and $q \geq \bar{q}_2$ by the definition of distributions $\mathrm{Tri}_{\bar{q}_1}$ and $\mathrm{Tri}_{\bar{q}_2}$, and $\hat{Q}(\bar{q}_2, 1/r) - \hat{Q}(\bar{q}_1, 1/r) \leq \bar{q}_2 - \bar{q}_1$. Moreover, we have

$$\begin{aligned} M_r(\mathrm{Tri}_{\bar{q}_1}) &= 2r \int_{\hat{Q}(\bar{q}_1, 1/r)}^{1} V_{\mathrm{Tri}_{\bar{q}_1}}(q)\hat{Q}(q,r)\, dq \\ &\geq 2r \int_{\hat{Q}(\bar{q}_2, 1/r)}^{1} V_{\mathrm{Tri}_{\bar{q}_1}}(q)\hat{Q}(q,r)\, dq \\ &\geq \frac{2r(1-\bar{q}_2)}{1-\bar{q}_1} \int_{\hat{Q}(\bar{q}_2, 1/r)}^{1} V_{\mathrm{Tri}_{\bar{q}_2}}(q)\hat{Q}(q,r)\, dq \\ &= \frac{1-\bar{q}_2}{1-\bar{q}_1} \cdot M_r(\mathrm{Tri}_{\bar{q}_2}), \end{aligned}$$

where the first inequality holds because $\bar{q}_1 \leq \bar{q}_2$ and function $\hat{Q}(q,r)$ is increasing in q. The second inequality holds because $V_{\mathrm{Tri}_{\bar{q}_1}}(q) \geq {}^{(1-\bar{q}_2)}/_{(1-\bar{q}_1)} \cdot V_{\mathrm{Tri}_{\bar{q}_2}}(q)$. □

A.2 Numerical and Analytical Arguments of Theorem 2

The proof of Theorem 2 is based on a hybrid numerical and analytical argument. We can numerically calculate the revenue of a mechanism M_r on a distribution $\mathrm{Tri}_{\bar{q}}$ via Lemma 5 and then we can argue, via Claim 2 and Claim 1, that nearby mechanisms and distributions have similar revenue. This approach will both allow us to argue about the structure of the solution and to identify the mechanism M_{α^*, r^*} and distribution of the solution $\mathrm{Tri}_{\bar{q}^*}$. Our subsequent discussion gives the details of these hybrid arguments.

We first approximate $\bar{q}^*$ by showing that $\bar{q}^* \in [0.09310569, 0.09310571]$. The parameters for this range are found by discretizing the space and finding the optimal choice of $\bar{q}^*$. Note that the optimal choice of $\bar{q}^*$ satisfies $M_1(\mathrm{Tri}_{\bar{q}^*}) = M_{r(\bar{q}^*)}(\mathrm{Tri}_{\bar{q}^*})$. Therefore, it is sufficient for us to show that for any quantile $\bar{q} \notin [0.09310569, 0.09310571]$, either $M_1(\mathrm{Tri}_{\bar{q}}) > M_{r(\bar{q})}(\mathrm{Tri}_{\bar{q}})$

or $M_1(\mathrm{Tri}_{\bar{q}}) < M_{r(\bar{q})}(\mathrm{Tri}_{\bar{q}})$.

First we show for any $\bar{q} \in [0, 0.09310569]$, $M_1(\mathrm{Tri}_{\bar{q}}) < M_{r(\bar{q})}(\mathrm{Tri}_{\bar{q}})$. Here we discretize the space $[0, 0.09310569]$ into Q_d with precision $\epsilon = 10^{-9}$. By numerically calculation using Lemma 5, we have

$$\min_{\bar{q} \in Q_d} M_{2.446946}(\mathrm{Tri}_{\bar{q}}) = M_{2.446946}(\mathrm{Tri}_{0.09310569}) \geq 1 + 10^{-8}$$

and for any $\bar{q} \in [0, 0.09310569]$, letting $\bar{q}_d$ be the largest quantile in Q_d smaller than or equal to $\bar{q}$, the minimum revenue for mechanism $M_{2.446946}$ is

$$M_{2.446946}(\mathrm{Tri}_{\bar{q}}) \geq \frac{1 - \bar{q}_d - \epsilon}{1 - \bar{q}_d} \cdot M_{2.446946}(\mathrm{Tri}_{\bar{q}_d}) \geq 1 + 8 \times 10^{-9} > M_1(\mathrm{Tri}_{\bar{q}}),$$

where the first inequality holds by Claim 2 and the second inequality holds because $\bar{q}_d \leq 0.1$.

Then we show for any $\bar{q} \in [0.09310571, 1]$, $M_1(\mathrm{Tri}_{\bar{q}}) > M_{r(\bar{q})}(\mathrm{Tri}_{\bar{q}})$. We discretize the space $[0.09310571, 1]$ into $\hat{Q}_d$ with precision $\hat{\epsilon} = 10^{-9}$. First note that $M_r(\mathrm{Tri}_{\bar{q}}) < 1$ for any $\bar{q} \geq 0.093$ and $r \geq 11$, since the expected probability the highest type got allocated is less than $\frac{1}{2}$, and hence the expected virtual value for mechanism M_r with distribution $\mathrm{Tri}_{\bar{q}}$ is less than 1. By Lemma 3, the revenue in this case is less than 1. With bounded range for optimal ratio r, we discretize the space $(1, 11]$ into R_d with precision $\epsilon_r = 10^{-9}$. By numerically calculation using Lemma 5, we have

$$\max_{\bar{q} \in \hat{Q}_d, r \in R_d} M_r(\mathrm{Tri}_{\bar{q}}) = M_{2.446945061}(\mathrm{Tri}_{0.09310571}) \leq 1 - 3 \times 10^{-8}$$

and for any $\bar{q} \in [0.09310571, 1]$ and any $r \in (1, 11]$, letting $\bar{q}_d$ be the largest quantile in $\hat{Q}_d$ smaller than or equal to $\bar{q}$ and r_d be the smallest number in R_d larger than or equal to r, the maximum revenue for distribution $\mathrm{Tri}_{\bar{q}}$ is

$$\max_{r \in (1,11]} M_r(\mathrm{Tri}_{\bar{q}}) \leq \frac{r_d}{r_d - \epsilon_r} \cdot (2\hat{\epsilon} + M_{r_d}(\mathrm{Tri}_{\bar{q}_d})) \leq 1 - 10^{-8} < M_1(\mathrm{Tri}_{\bar{q}}),$$

where the first inequality holds by Claim 1 and 2, and the second inequality holds because $r_d > 1$. Combining the numerical calculation, we have that $\bar{q}^* \approx 0.0931057$.

Note that both mechanism M_1 and M_{r^*} are the best responses for distribution $\mathrm{Tri}_{\bar{q}^*}$, achieving revenue 1, and hence the optimal approximation ratio is

$$\beta = \frac{\mathrm{OPT}_{\mathrm{Tri}_{\bar{q}^*}}(\mathrm{Tri}_{\bar{q}^*})}{M_{\alpha^*, r^*}(\mathrm{Tri}_{\bar{q}^*})} = 2 - \bar{q}^* \approx 1.9068943.$$

Next we show that by choosing ratio $r^* \approx 2.4469452$ and probability $\alpha^* \approx 0.80564048$,

the approximation ratio of mechanism M_{α^*,r^*} approximates β. Here we discretize the quantile space $[0,1]$ into Q'_d with precision $\epsilon' = 10^{-9}$, using the formula in Lemma 3 and Lemma 5, the triangle distribution that maximizes the approximation ratio for mechanism M_{α^*,r^*} is $\mathrm{Tri}_{0.093105694}$ with approximation ratio at most 1.9068943044. For any $\bar{q} \in [0, \frac{1}{2}]$, letting $\bar{q}_d$ be the largest quantile in Q'_d smaller than or equal to $\bar{q}$, the minimum revenue for mechanism M_{α^*,r^*} is

$$\begin{aligned} M_{\alpha^*,r^*}(\mathrm{Tri}_{\bar{q}}) &\geq \frac{1-\bar{q}_d-\epsilon'}{1-\bar{q}_d} \cdot M_{\alpha^*,r^*}(\mathrm{Tri}_{\bar{q}_d}) \\ &\geq \frac{1}{1.906894309}\,\mathrm{OPT}_{\bar{q}_d}(\mathrm{Tri}_{\bar{q}_d}) \geq \frac{1}{1.906894309}\,\mathrm{OPT}_{\bar{q}}(\mathrm{Tri}_{\bar{q}}), \end{aligned}$$

where the second inequality holds because $\bar{q}_d \leq \frac{1}{2}$ and the last inequality holds because $\bar{q}_d \leq \bar{q}$. For any $\bar{q} \in [\frac{1}{2}, 1]$, the minimum revenue for mechanism M_{α^*,r^*} is

$$M_{\alpha^*,r^*}(\mathrm{Tri}_{\bar{q}}) \geq \alpha^* \cdot M_1(\mathrm{Tri}_{\bar{q}}) \geq 0.8 \geq \frac{1}{1.875}\,\mathrm{OPT}_{\bar{q}}(\mathrm{Tri}_{\bar{q}}),$$

since for any $\bar{q} \in [\frac{1}{2}, 1]$, $M_1(\mathrm{Tri}_{\bar{q}_d}) = 1$ and $\mathrm{OPT}_{\bar{q}}(\mathrm{Tri}_{\bar{q}}) = 2 - \bar{q} \leq 1.5$. Therefore, $r^* \approx 2.4469452$ and probability $\alpha^* \approx 0.80564048$ are the desirable parameters, with error at most 2×10^{-8} in approximation ratio. By our characterization, the error solely comes from numerical calculation, finishing the numerical analysis for Theorem 2.

A.3 Continuity of Distribution in Probability of Second-price Auction

Recall the function $\bar{q}_r(\alpha)$ which gives the adversary's best-response triangle distribution the mechanism $M_{\alpha,r}$. The continuity of the function $\bar{q}_r(\alpha)$ is used to prove the existence of equilibrium between the randomized markup mechanism and the triangle distribution in Theorem 2. The following claim proves the continuity of the function $\bar{q}_r(\alpha)$, by numerically bounding the second derivative of the revenue ratio of the stochastic markup mechanism $M_{\alpha,r}$ on distribution $\mathrm{Tri}_{\bar{q}}$ with respect to α, the probability that the markup mechanism runs the second-price auction.

Claim 3. *Given any $r \in [2.445, 2.449]$, function $\bar{q}_r(\alpha)$ is continuous in α for $\alpha \in [0.8, 0.81]$.*

Proof of Claim 3. By Lemma 5 and Lemma 1, the approximation ratio of mechanism $M_{\alpha,r}$

for triangle distribution $\text{Tri}_{\bar{q}}$ is

$$\begin{aligned}\text{APX}(\alpha, r, \bar{q}) &= \frac{\text{OPT}_{\text{Tri}_{\bar{q}}}(\text{Tri}_{\bar{q}})}{\alpha \cdot M_1(\text{Tri}_{\bar{q}}) + (1-\alpha) M_r(\text{Tri}_{\bar{q}})} \\ &= \frac{2-\bar{q}}{\alpha + \frac{2r(1-\alpha)}{(1-\bar{q})(r-1)} \left(\frac{1-\bar{q}}{1-\bar{q}+\bar{q}r} + \frac{\ln\left(\frac{r}{1-\bar{q}+\bar{q}r}\right)}{1-r} \right)}\end{aligned}$$

The approximation ratio is a continuous function of $\alpha, \bar{q}$. Therefore, to show that fixing r, function $\bar{q}_r(\alpha)$ is continuous in α, it is sufficient to show that there is a unique $\bar{q}$ that maximizes $\text{APX}(\alpha, r, \bar{q})$ for $r \in [2.445, 2.449]$ and $\alpha \in [0.8, 0.81]$, or equivalently, we show that there is a unique $\bar{q}$ that minimizes $1/\text{APX}(\alpha, r, \bar{q})$. By Claim 1 and 2, we can discretize the quantile space and numerically verify that distributions with monopoly quantiles $\bar{q} \notin [0.093, 0.094]$ are suboptimal. Therefore, we prove the uniqueness of the maximizer by showing that the second order derivative of $1/\text{APX}(\alpha, r, \bar{q})$ is strictly positive for $\bar{q} \in [0.093, 0.094]$.

$$\begin{aligned}\frac{\partial^2 \frac{1}{\text{APX}(\alpha,r,\bar{q})}}{(\partial \bar{q})^2} =& \frac{4(1-\alpha)r \left(-\frac{r-1}{(1-\bar{q}+\bar{q}r)^2} + \frac{1}{(1-\bar{q})(1-\bar{q}+\bar{q}r)} - \frac{\log(\frac{r}{1-\bar{q}+\bar{q}r})}{(r-1)(1-\bar{q})^2} \right)}{(r-1)(2-\bar{q})^2} \\ &+ \frac{2(1-\alpha)r \left(-\frac{(r-1)^2}{(1-\bar{q}+\bar{q}r)^3} - \frac{r-1}{(1-\bar{q})(1-\bar{q}+\bar{q}r)^2} + \frac{2}{(1-\bar{q})^2(1-\bar{q}+\bar{q}r)} + \frac{2\log(\frac{r}{1-\bar{q}+\bar{q}r})}{(r-1)(1-\bar{q})^3} \right)}{(r-1)(2-\bar{q})} \\ &+ \frac{4(1-\alpha)r \left(-\frac{1}{1-\bar{q}+\bar{q}r} - \frac{\log(\frac{r}{1-\bar{q}+\bar{q}r})}{(r-1)(1-\bar{q})} \right) + 2\alpha(r-1)}{(r-1)(2-\bar{q})^3}\end{aligned}$$

By substituting the upper and lower bounds of $\alpha, r, \bar{q}$, we know that

$$\frac{\partial^2 \frac{1}{\text{APX}(\alpha,r,\bar{q})}}{(\partial \bar{q})^2} > 0.7$$

for $r \in [2.445, 2.449], \alpha \in [0.8, 0.81]$ and $\bar{q} \in [0.093, 0.094]$, which concludes the uniqueness of the maximizer and the continuity of function $\bar{q}_r(\alpha)$. □

B Discussion of Robustness Paradigms

This paper focused on the robustness paradigm of worst case multiplicative-approximation ratio. This section provides an informal illustration of and comparison between it and other prevalant robustness paradigms. Specifically, we illustrate the ideas in a *robust monopoly*

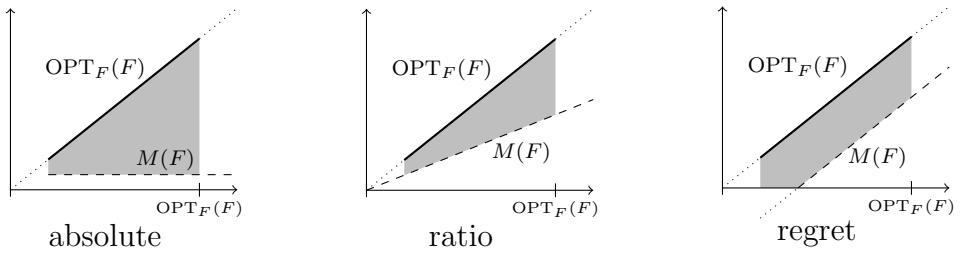


Figure 7: Comparison of three measures of robustness. The horizontal axis indexes the prior distributions F with respect to which we aim to be robust and is ordered by the performance of the optimal mechanism $\mathrm{OPT}_F(F)$. The vertical axis is the absolute performance. $M(F)$ is the expected performance of mechanism M given distribution F, and OPT_F is the Bayesian optimal mechanism with the knowledge about distribution F. Any mechanism M with performance curve within the shaded gray area is robustly optimal.

pricing problem where a monopoly seller aims to sell a single item to a buyer. The seller is uncertain about the distribution over values of the buyer except the fact that the distribution has support within $[1, H]$. This problem is considered in Bergemann and Schlag (2008) for min-max regret and in Hartline and Roughgarden (2014) for multiplicative approximation. Through this example, we will show that while the absolute max-min optimal focuses attention at small scales and min-max regret focuses at large scales, the multiplicative-approximation ratio focuses puts equal focus on all scales. These frameworks are illustrated in Figure 7 and the example is summarized in Table 1.

The absolute max-min framework is $\max_M \min_F M(F)$. For the max-min objective, the principal designs mechanisms that target the absolute worst case performance. Therefore, any mechanism that provides a performance guarantee between the optimal and the max-min value for all instances is admissible for the principal, i.e., any mechanism with performance curve within the gray area is max-min optimal for the principal. In particular, it is possible that the max-min optimal mechanism only provides the max-min value for all problem instances. However, on good instances, i.e., where $\mathrm{OPT}_F(F)$ is large, the gap between the optimal performance and the performance of the max-min robust mechanism can be very large. For the max-min objective in the robust monopoly pricing problem, characterizing the optimal mechanism is trivial, i.e., the max-min optimal mechanism is to sell the item with price 1, which obtains max-min revenue of 1 regardless of the value of the buyer. Now we evaluate this mechanism using other robust paradigms. It is easy to verify that if the actual distribution is a point mass at value H, the optimal revenue is H and the multiplicative approximation ratio is H and the regret is $H - 1$. Thus, the max-min optimal mechanism can have very poor performance under other robust paradigms.

	min revenue (large is better)	max approximation (small is better)	max regret (small is better)
max-min optimal mechanism	1	H	$H-1$
ratio optimal mechanism	$\frac{1}{1+\ln H}$	$1+\ln H$	$H-\frac{H}{1+\ln H}$
regret optimal mechanism	0	∞	$\frac{H}{e}$

Table 1: Comparisons of robust paradigms.

The min-max regret framework is $\min_M \max_F \mathrm{OPT}_F(F) - M(F)$. The min-max regret is often achieved at instances where there is the most to lose. The principal essentially targets the best case performance, and any mechanism that provides a performance guarantee that suffers at most an additive γ loss for all instances is regret optimal, where γ is the min-max regret. In this case, if $\mathrm{OPT}_F(F)$ is small, perhaps even smaller than γ, it is possible that the min-max regret optimal mechanism does not provide any non-trivial performance guarantee. Let us again consider the robust monopoly pricing problem for minimizing worst case regret and suppose that $H \geq e$. Bergemann and Schlag (2008) show that the min-max regret optimal mechanism is to post a randomized price p with cumulative distribution

$$G(p) = \begin{cases} 0 & p \in [1, \frac{H}{e}), \\ 1 + \ln \frac{p}{H} & p \in [\frac{H}{e}, H], \end{cases}$$

which guarantees min-max regret of $\frac{H}{e}$. Note that if the distribution over values has support less than $\frac{H}{e}$, the item is not sold with probability 1 and for any such distribution, the expected revenue given by this robust mechanism is 0. Thus, min-max regret provides a trivial guarantee when the optimal revenue is small. In particular, for the min-max regret optimal mechanism, the minimum revenue is 0 and the maximum multiplicative approximation ratio is infinity.

In contrast, the multiplicative approximation framework considered in this paper ensures that the robust mechanisms provide comparable performance to the Bayesian optimal for any instance. In particular, in the monopoly pricing example, Hartline and Roughgarden (2014) show that to minimize the multiplicative approximation ratio, the seller can post a price p with distribution $G(p) = \frac{1+\ln p}{1+\ln H}$ for any $p \in [1, H]$. The multiplicative approximation ratio is at most $1 + \ln H$ for all possible distributions. Moreover, the minimum revenue for the ratio optimal mechanism is $\frac{1}{1+\ln H}$ and the maximum regret is $H - \frac{H}{1+\ln H}$. As illustrated in Table 1, the multiplicative approximation framework provides a balanced performance between the extreme robust paradigms of absolute max-min and min-max regret.