# Simple Mechanisms for Agents with Non-linear Utilities*

Yiding Feng† Jason D. Hartline‡ Yingkai Li§

**Abstract**

We propose a simple class of pricing-based mechanisms for agents with non-linear utilities, such as risk aversion and budget constraints. We show that these mechanisms are max-min robustly optimal when the principal cannot observe the detailed non-linear utility functions and type distributions. Furthermore, we provide a general framework that extends the approximation guarantees of pricing-based mechanisms from linear to non-linear utilities in Bayesian environments. This extension allows us to show that many economic conclusions from auction theory based on linear utility models (e.g., Bulow and Roberts, 1989) approximately extend to non-linear utility models.

**Keywords**: simple mechanisms, non-linear utilities, pricing, marginal revenue maximization, approximation, robustness.
**JEL**: D44, D82

## 1 Introduction

Many classical papers on mechanism design, and in particular auction theory, focus on agents with linear utilities, i.e., the agents' utilities are linear functions of both their values for the goods, and their payments to the principal. The simplifying assumption of linear utilities allows the derivation of simple closed-form characterizations and nice economic interpretations of optimal mechanisms (e.g., Myerson, 1981; Bulow and Roberts, 1989). However, agents in practice deviate significantly

*This work is support by NSF CCF AF #1618502. This manuscript combines "Optimal Auctions vs. Anonymous Pricing: Beyond Linear Utility" as an abstract that appeared at the 20th ACM Conference on Economics and Computation (EC'19), and "Simple Mechanisms for Non-linear Agents" as an extended abstract that appeared at the 34th ACM-SIAM Symposium on Discrete Algorithms (SODA'23). The authors thank Yi-Chun Chen, Ben Golub, Nima Haghpanah, Deniz Kattwinkel, Ce Li, Harry Pei, Larry Samuelson, Xianwen Shi, Marciano Siniscalchi, Thomas Wiseman and the audiences at Simons Laufer Mathematical Sciences Institute (SLMath), the International Conference on Game Theory at Stony Brook and the World Congress of the Game Theory for helpful suggestions.

†Hong Kong University of Science and Technology. Email: ydfeng@ust.hk
‡Northwestern University. Email: hartline@northwestern.edu
§National University of Singapore. Email: yk.li@nus.edu.sg

from linear utilities, for example, through private budget constraints or risk aversion over probabilistic outcomes. Notably, optimal mechanisms for agents with non-linear utilities are in general not simple and therefore difficult to understand precisely. For example, the optimal mechanism for a single-item single-agent environment with a private budget constraint admits no closed-form characterization (Che and Gale, 2000).[1]

The dramatic differences for designing optimal mechanisms for agents with non-linear utilities raise the concern about the robustness of the economic conclusions we observed when assuming linear utilities. In this paper, by analyzing the performances of simple mechanisms, we identify conditions under which the economic conclusions derived from linear utility models approximately extend to general (i.e., non-linear) utility models.

Our definition of simple mechanisms are motivated by Bulow and Roberts (1989) who, as later interpreted by Alaei et al. (2013), show that for linear agents,[2] every mechanism can be interpreted as a *pricing-based mechanism*, i.e., mechanism where the menu offered to each agent is equivalent to a distribution over posted prices. As illustrated by Bulow and Roberts (1989), for agents with linear utilitiess, there is a unique mapping from values to quantiles based on the cumulative distribution function. Pricing-based mechanisms can thus be implemented in quantile space by specifying allocations and payments for these quantiles. Furthermore, an agent's quantile can be interpreted as the expected quantity sold when a price equal to the corresponding value is posted. Using this interpretation, Bulow and Roberts (1989) construct revenue curves that map quantities sold (quantiles) to expected revenue and define the marginal revenue for serving quantile $q$ as the derivative of the revenue curve evaluated at quantile $q$. They show that, for linear utility agents, the multi-agent revenue-optimal mechanism can be implemented as *marginal revenue maximization*, which always allocates the item to the agent with the quantile corresponding to the highest marginal revenue.

Pricing-based mechanisms for linear agents can be generalized to non-linear agents by considering *per-unit* prices, i.e., each agent who is given a per-unit price $p$ can purchase any lottery with winning probability $q \in [0, 1]$ and pay a deterministic price of $p \cdot q$. The implementation of this generalization is straightforward in single-agent settings.

However, when there are multiple non-linear agents, implementing the pricing-based mechanisms in quantile space analogous to linear agents is non-trivial since even the definition of quantiles for non-linear agents is not clear. This is because the type space of a non-linear agent can be multi-dimensional. For instance, for a risk averse agent, her private type could include both her value for the item and her risk attitude, and there does not exist a consistent ranking of types that

[1]Che and Gale (2000) provide a characterization showing that the optimal mechanism must be the solution of a differential equation. However, solving the differential equation given arbitrary type distribution is generally intractable.

[2]In this paper, we write "agents with linear utilities" as "linear agents" for short, and "agents with non-linear utilities" as "non-linear agents".

captures her demand for all possible menus. In our paper, we provide a natural definition that randomly maps the potentially multi-dimensional type of a non-linear agent to a single-dimensional quantile based on her demand function given per-unit prices. Such definition allows us to provide a construction that converts any pricing-based mechanism for linear agent to a pricing-based mechanism for non-linear agents with the same payoff guarantee for the designer (Theorem 1). Moreover, we extend the analysis from (Bulow and Roberts, 1989) to show that the marginal revenue maximization mechanism maximizes the expected revenue across all pricing-based mechanisms for non-linear agents.

In many practical applications involving non-linear agents, the principal often has limited knowledge regarding the preferences of the agents such as the agents' exact utility functions or type distributions. For instance, when agents are risk-averse, it is challenging for the principal to exactly know the risk attitudes of all agents. A paradigm for understanding the design of good mechanisms in those practical applications is the robustness framework (e.g., Carroll, 2015, 2017; Brooks and Du, 2021). We show that when the principal only knows the expected demand curve (for instance, based on the estimation of the expected demands of the agents using historical purchase data), the marginal revenue maximization mechanism is robustly optimal for non-linear agents. That is, the marginal revenue maximization mechanism maximizes the minimum expected revenue where the minimum is taken over all possible utility functions and type distributions that are consistent with the expected demands (Theorem 2). This robustness result suggests that when the details of non-linear utilities are unknown, the principal should refrain from customizing lotteries in pursuit of higher revenue. Instead, treating agents akin to those with linear utilities and applying the optimal pricing-based mechanism – marginal revenue maximization – is advisable.

Nonetheless, there exist applications where the principal can have more refined prior information. For those applications, not all mechanisms can be interpreted as pricing-based mechanisms and, in fact, pricing-based mechanisms are generally not optimal even in single-agent settings for non-linear agents (e.g., agents with budget constraints). To understand the additional benefit the principal can possibly obtain by figuring out the details of the environments and adopting sophisticated mechanisms targeting those details, we apply the method of approximations to quantify the maximum revenue loss of pricing-based mechanisms for non-linear agents.

The approximation guarantees of simple mechanisms have been widely studied for linear agents. Those approximation results allow qualitative conclusions about drivers of good economic outcomes. Consider the examples of selling $m$ units of identical items to multiple heterogeneous buyers. Example 1: Bulow and Roberts (1989) connects auction theory with classical microeconomics by showing that the revenue optimal mechanism can be interpreted as the marginal revenue maximization mechanism. Example 2: Yan (2011) shows that sequential posted pricing guarantees at least $1-\frac{1}{\sqrt{2\pi m}}$ of the optimal auction revenue, which indicates that simultaneity and competition are not salient features for revenue maximization, especially in large markets where the number of units $m$ is sufficiently large. Example 3: Jin et al. (2021) show that the multiplicative loss of posting an anonymous price can be as large as $\log m$ even when the agents' valuation distributions are regular. This implies that agents' identities and the seller's ability to third-degree price discriminate among the agents are crucial for revenue maximization, especially in large markets. See the survey of Hartline (2012) for detailed discussion of the method of approximation in economics.

We introduce a reduction framework that approximately extends the approximation guarantees of pricing-based mechanisms from linear agents to non-linear agents, and show that those mechanisms are approximately optimal for large families of non-linear agents. Specifically, given a pricing-based mechanism that guarantees a $\beta$-approximation (i.e., achieves at least $1/\beta$ fraction to the ex ante relaxation[3]) for linear agents and given non-linear agents that are $\zeta$-*resemblant*[4] of linear agents, our reduction framework shows that the analogous pricing-based mechanism for non-linear agents guarantees a $\beta\zeta$-approximation to the optimal (Theorem 3).[5] This reduction framework is general – it can be applied to any downward-closed feasibility constraint (e.g., single-item, multi-unit, matroid), many common objectives (e.g., revenue, welfare, or their convex combination).

Our reduction framework for non-linear agents allows us to observe not only that the main drivers of good mechanisms are similar for non-linear agents, but also that non-linear utility itself, in many of its forms, is not a main concern that necessitates specialized mechanism designs. As an illustration, agents with public budgets and regular valuation distributions are 1-resemblant of linear agents for revenue maximization, which implies, e.g., for $m$-unit environments of Example 2, that sequential posted pricing guarantees at least $1 - \frac{1}{\sqrt{2\pi m}}$ of the optimal revenue for such non-linear agents, and simultaneous implementation and competition among agents are not salient features for revenue maximization. Such conclusions are unlikely to be derived through the classical approach of analyzing the optimal mechanisms since the optimal mechanisms for non-linear agents often depend on the details of the utility structures, and are too complex to admit simple interpretations.

## 1.1 Discussion of Results on $\zeta$-Resemblance

To establish the approximations of pricing-based mechanisms for non-linear agents, we characterize broad families of non-linear agents that are $\zeta$-resemblant for small constant factors $\zeta$ (e.g.,

[3]This is an upper bound of the optimal, which will be formally defined in Section 5.

[4]We measure the resemblance of agents in terms of the (topological) closeness of the single-agent revenue curves, as defined in Bulow and Roberts (1989). We provide the details in the next subsection.

[5]Technically, our main results require that the agents utilities satisfy the von Neumann-Morgenstern expected utility representation (Morgenstern and von Neumann, 1953). However, the reduction framework for pricing-based mechanisms, such as sequential posted pricing, can also extend to agents that do not satisfy the expected utility representation (e.g., agents with endogenous valuation, Gershkov et al., 2021; Akbarpour et al., 2023). See Section 6 for a detailed discussion of non-expected utilities and other extensions.

agents with independent private budgets, or agents with risk aversion) and families that are not (e.g., agents whose budgets and values are correlated). For non-linear agents that are $\zeta$-resemblant, pricing-based mechanisms are approximately optimal if those mechanisms are approximately optimal for linear agents; thus, non-linearity of utility can be viewed as a detail that can be omitted from the model without significantly altering the main take-aways. On the other hand, with utility models that are not $\zeta$-resemblant for modest $\zeta$, non-linearity is a crucial feature that needs specific study for identifying forms of mechanisms lead to good economic outcomes.

The $\zeta$-resemblance property that governs our reduction is based on a single-agent price-posting approximation with ex ante supply constraints (cf. Bulow and Roberts, 1989). In particular, we say a non-linear agent is $\zeta$-resemblant if given any ex ante supply constraint $q \in [0, 1]$, the optimal payoff of the principal is at most $\zeta$ times the expected payoff when restricting the mechanisms to posting (randomized) per-unit prices. For linear agents, $\zeta = 1$, i.e., the optimal mechanisms under any ex ante supply constraint can be implemented as posting per-unit prices. For non-linear agents, a smaller parameter $\zeta$ results in a better approximation of posted pricing in single-agent problems, making non-linear agents more resemblant of linear agents.

In our paper, we focus on the designer's objective of revenue maximization and show that for canonical non-linear utility models, parameter $\zeta$ is bounded by small constants. We also discuss the extension for welfare maximization and their convex combinations in Section 6.1.

**Risk Averse Utility.** It is standard to model risk averse utility as a concave function that maps the agents' wealth to utility. For revenue maximization problems, risk aversion introduces a non-linearity into the incentive constraints of the agents, which in most cases makes mechanism design analytically intractable. We show that when the valuation distribution satisfies the monotone hazard rate property, without any assumption on the risk preference such as constant absolute risk aversion, a risk-averse agent is $e$-resemblant. Combined with our reduction framework, this implies that the additional multiplicative loss in approximation ratios for pricing-based mechanisms is at most $e$ when agents are risk averse. Note that our bound is worst-case, and as illustrated in Section 5.5, for common valuation distributions such as uniform distribution, the revenue losses of both marginal revenue maximization and sequential posted pricing are almost negligible when both the number of agents and number of items for sale are large.

**Budgeted Utility.** We consider a setting where agents have hard budget constraints for paying the principal. That is, each agent's utility is linear in payments if her payment is at most her budget, and is minus infinity if her payment exceeds the budget. We show that if the budgets are independent from the values, for regular valuation distributions,[6] a public budget agent is $1$-resemblant and a private budget agent is $3$-resemblant. This implies that simple mechanisms such as sequen-

[6]A valuation distribution is regular if the virtual value $v - \frac{1-F(v)}{f(v)}$ is weakly increasing in $v$.

tial posted pricing are approximately optimal for agents with independent budget distributions. In contrast, when the values are correlated with budgets, the principal can second-degree price discriminate against the agents using their budget constraints, thereby achieving almost full surplus extraction in some cases. In those cases, posted pricing mechanisms are not approximately optimal even in single-agent settings and such budgeted agents are not constant resemblant to linear agents.

## 1.2 Related Work

Mechanism design for non-linear agents is widely studied in the literature. Most relevantly, Baisa (2017) considers mechanism design without linear utilities and proposes a novel random allocation mechanism based on the reported demand curves. However, the mechanism implemented in Baisa (2017) offers a uniform price to all agents, which may be far from optimal as it does not price discriminate the agents based on their prior heterogeneity (cf. Jin et al., 2021). Kazumura et al. (2020) characterize the set of dominant strategy incentive compatible mechanisms in environments without linear utilities and provide a general procedure for optimizing over this class. However, depending on the utility models, obtaining a closed-form characterization using their procedure remains challenging. The main contribution of our paper is to propose a simple class of dominant strategy incentive compatible mechanisms, i.e., pricing-based mechanisms. We show that these mechanisms are robustly optimal when the detailed utility models are unknown and approximately optimal even compared to Bayesian incentive compatible mechanisms when the detailed utility models are known.[7]

The robust optimality of simple mechanisms in complex environments has also been established in Carroll (2015, 2017); Brooks and Du (2021); Hartline et al. (2020); Guo and Shmaya (2023); etc. The main intuition in that line of work is that simple mechanisms are robustly optimal when the principal is ambiguity averse and is not fully aware of the environment. Our robustness result shows that, for auctions with non-linear utilities, the simple mechanism of marginal revenue maximization is robustly optimal. In addition, our approximation results complement that line of work by showing that the principal may still wish to adopt simple mechanisms in practice since their performances are close to optimal. See Hartline and Lucier (2015) for a more detailed discussion on the benefits of adopting simple approximation mechanisms.

Frameworks for reducing approximation for non-linear agents to approximation for linear agents has also been studied in Alaei et al. (2013). Their framework converts marginal revenue mechanisms for linear agents into mechanisms for non-linear agents. Unlike our approach, which uses single-agent price-posting mechanisms as a building block, their framework involves single-

[7]Kushnir and Liu (2019) establish the equivalence between dominant strategy incentive compatible and Bayesian incentive compatible mechanisms under the assumptions of increasing differences over distributions and a convex-valued property. These assumptions are not made in our paper.

agent ex ante optimal mechanisms, which in general are hard to characterize (e.g., for risk-averse utilities). Due to this difference, the framework and its implementation in Alaei et al. (2013) are more complex, and the resulting mechanisms under their framework are not dominant strategy incentive compatible in general for non-linear agents.

In this work, as applications of our general framework, we focus on two non-linear models, agents with risk-averse attitudes and agents with budget constraints.

Most results for agents with risk-averse utilities consider the comparative performance of the first- and second-price auctions (cf. Holt Jr, 1980; Che and Gale, 2006). Matthews (1983) and Maskin and Riley (1984), however, characterize the optimal mechanisms for symmetric agents for constant absolute risk aversion and more general risk-averse models.

For symmetric agents with *public* budgets in single-item environments, Laffont and Robert (1996) and Maskin (2000) respectively study the revenue-maximization and welfare-maximization problems. Their results are generalized to i.i.d. values but asymmetric public budgets in Boulatov and Severinov (2021) and to symmetric agents with uniformly distributed private budgets in Pai and Vohra (2014). Che and Gale (2000) consider the single agent problem with *private* budget and valuation distribution that satisfies declining marginal revenues, and characterize the optimal mechanism by a differential equation. In environments without unit-demand constraints, a price-posting mechanism is optimal for selling a divisible good to a continuum of agents with private budgets if their valuations are regular with decreasing density (Richter, 2019) and is a 2-approximation in general (Abrams, 2006). For more general settings, no closed-form characterizations are known, and algorithmic solutions are provided in Alaei et al. (2012); Devanur and Weinberg (2017).

It is well known that simple mechanisms are approximately optimal for linear agents. For single item auction with linear agents, Alaei et al. (2019) show that the ratio between the optimal mechanism and anonymous pricing is at most $e$ under the regularity assumption, and later the ratio is improved to 2.62 by Jin et al. (2019). Chawla et al. (2010) show that the approximation ratio for order oblivious posted pricing is $2$, and Yan (2011) shows that the approximation ratio for sequential posted pricing is $e/(e-1)$.

For non-linear agents, given matroid environments, Chawla et al. (2011) show that a simple lottery mechanism is a constant approximation to the optimal ex post individually rational mechanism for agents with monotone-hazard-rate valuations and private budgets. In contrast, our approximation results are with respect to the optimal mechanism under interim individually rationality which can be arbitrarily larger than the benchmark from Chawla et al. (2011).

The single-agent analysis for quantifying $\zeta$-resemblance for non-linear agents is non-trivial in general. For risk-averse agents, our analysis builds on the welfare-revenue gap for linear agents developed in Hartline and Roughgarden (2009). For private-budget agents, our analysis builds on the type decomposition trick developed in Abrams (2006) and the single-sample approximation technique developed in Dhangwatnotai et al. (2015). The details of our technical innovations will be discussed in Appendix A.

# 2 Model

We consider the problem of selling $m$ units of indivisible and identical items to $n$ non-linear agents where each agent has unit demand for the items. In this paper, since all agents have unit demand, it is without loss to consider allocations to each agent as a lottery over $\{0, 1\}$.

**Agent's Utility Models.** There is a set of agents $N$ where $|N| = n$. An agent's *utility model* is defined as $(\mathcal{T}, \bar{F}, u)$ where $\mathcal{T}, \bar{F}$, and $u$ are the type space, type distribution and utility function. The outcome for an agent is the distribution over the pair $(x, p)$, where allocation $x \in \{0, 1\}$ and payment $p \in \mathbb{R}_+$. The utility function $u$ of the player is a mapping from her private type and the realized outcome to her von Neumann-Morgenstern expected utility for the outcome. Throughout this paper, we use the subscript $i$ to denote agent $i$ in multi-agent settings, and we will drop the subscripts without ambiguity when we discuss single-agent problems.

There are several specific utility models we are interested in this paper.

- **Linear utility:** The agent's private type is her value $v$ of the good. Given allocation $x$ and payment $p$, her utility is

$$u(v, x, p) = v \cdot x - p.$$

- **Non-linear utility:** For a non-linear agent with type $t$, her utility for allocation $x$ and payment $p$ is $u(t, x, p)$. In our general reduction framework for non-linear utilities, we do not impose any specific structural properties on the utility function $u$. However, there are several important special cases of non-linear utilities to which we apply our general framework.

  - **Risk-averse utility:** The agent's type is a pair $t = (v, \varphi)$ where $v$ is the private value and $\varphi$ is an increasing and concave function that represents the agent's private risk preference. We assume that $\varphi(0) = 0$. Given allocation $x$ and payment $p$, her utility is

$$u((v, \varphi), x, p) = \varphi(vx - p).$$

  - **Private-budget utility:** The agent's type is a pair $t = (v, w)$ where $v$ is the private value and $w$ is the private budget constraint. Given allocation $x$ and payment $p$, her utility is

$$u((v, w), x, p) = \begin{cases} vx - p & p \le w, \\ -\infty & p > w. \end{cases}$$

We use $F$ to denote the marginal distribution over values for agents with linear utilities, risk-averse utilities, or private-budget utilities.

**Mechanisms.** By the revelation principle (Myerson, 1981), we focus on revelation mechanisms $\{(x_i, p_i)\}_{i \in N}$, where agents simultaneously submit bids $\{b_i\}_{i \in N}$ from their type spaces to the mechanism, and each agent $i$ gets allocation $x_i(\{b_i\}_{i \in N})$ with payment $p_i(\{b_i\}_{i \in N})$. An allocation is feasible if $\sum_i x_i \le m$ where $m$ is the number of items.[8]

We consider mechanisms that satisfy *Bayesian incentive compatibility* (BIC), i.e., no agent can gain strictly higher expected utility than reporting her private type truthfully if all other agents are reporting their private types truthfully, and *interim individual rationality* (IIR), i.e., the expected utility is non-negative for all agents and all private types if all agents are reporting their private types truthfully mechanisms. For later discussion, we also define *dominant strategy incentive compatible* (DSIC) for a mechanism if no agent can gain strictly higher expected utility than reporting her private type truthfully, regardless of other agents' report. Note that DSIC coincides with BIC in single-agent environments for linear utilities.

**Pricings and Demands.** We study mechanisms based on simple per-unit posted pricing. The definition of the pricing-based mechanisms for general multi-agent environments will be introduced formally in Section 3. In this section, we will focus on pricing for single-agent environments.

**Definition 2.1.** *A* per-unit price $p$ *denotes the menu* $\{(z, z \cdot p) : z \in [0, 1]\}$. *The agent who chooses an option* $(z, z \cdot p)$ *receives an item with probability* $z$*, and pays prices* $z \cdot p$ *regardless of the allocation.*[9]

Given per-unit price $p$, the demand of the agent with type $t$ is

$$d^t(p) = \operatorname*{argmax}_{z \in [0,1]} \; z \cdot u(t, 1, z \cdot p) + (1 - z) \cdot u(t, 0, z \cdot p).$$

For example, a budgeted agent with value $v$ and budget $w$ given *per-unit price* $p$ will purchase the lottery $d = \min\{1, {}^{w}\!/\!{}_{p}\}$ if $v \ge p$, and purchase the lottery $d = 0$ otherwise. Throughout the

[8]In this paper, in order to simplify the exposition, we focus on the feasibility constraint where any allocation is feasible as long as the total allocation does not exceed the supply of the principal. Our results can be extended to multi-unit auctions with more general feasibility constraints. See Appendix B.6 for detailed discussions.

[9]We focus on all-pay format in our definition of per-unit posted pricing mechanisms. The payment format does not matter for linear agents, but it affects the incentives for non-linear agents. Our general reduction framework can also be extended to other payment formats such as winner-pays-it.

paper, we will assume that given any price $p \geq 0$, the utility of the agent is upper semi-continuous in chosen lottery $z$. This implies that the demand maximizing the utility can always be attained. Additionally, we assume that for any type $t \in \mathcal{T}$, $1 \in d^t(0)$ and there exists $p > 0$ such that $d^t(p) = 0$. We denote the aggregated demand correspondence as $D(p) = \mathbf{E}_{t \sim \bar{F}}[d^t(p)]$.

**Ordinary Goods Assumption.** In this paper, our focus is on the sale of ordinary goods, where the demand of the agent is non-increasing in the offered per-unit price. This assumption excludes Giffen goods or Veblen goods and is satisfied by many common utility models, such as risk-averse utility or private-budget utility.

**Assumption 1.** *The item is an* ordinary good*, i.e., for any type $t$, the demand $d^t(p)$ is upper hemi-continuous and weakly decreasing in $p$ in strong set order.*

In the special case where the optimal demand is unique, the ordinary good assumption is equivalent to the assumption that the demand function is weakly decreasing in per-unit price $p$.

**Principal's Objective.** We consider general payoff maximization for the principal, where welfare maximization, revenue maximization and their convex combinations are special cases of payoff maximization.

Specifically, we assume that the principal is an expected utility maximizer. The principal's payoff function is a mapping from the type-outcome pair of each agent to a real value, and it is additive and separable across different agents. That is, there exists payoff functions $\{\Psi_i\}_{i \in N}$ such that given allocations $\{x_i\}_{i \in N}$ and payments $\{p_i\}_{i \in N}$ for type profile $\{t_i\}_{i \in N}$, the payoff of the principal is

$$\Psi(\{t_i\}_{i \in N}, \{x_i\}_{i \in N}, \{p_i\}_{i \in N}) = \sum_i \Psi_i(t_i, x_i, p_i).$$

A special case of the payoff function is the revenue, where $\Psi_i(t_i, x_i, p_i) = p_i$ for any agent $i$.[10]

# 3 Implementation of Pricing-based Mechanisms

In Bayesian mechanism design, the taxation principle suggests that it is without loss to focus on *menu* mechanisms: Fixing any agent, the mechanism offers a menu of outcomes (i.e., her allocation and payment) to the agent, where the menu depends on other agents' bids.

For non-linear agents, the menu offered in the Bayesian optimal mechanism often requires complicated price discrimination schemes, and characterizing the optimal mechanism is challenging even in single-agent environments (c.f., Maskin and Riley, 1984; Maskin, 2000; Che and Gale,

[10] From now on, "payoff" refers to the principal's payoff.

2000). For example, to maximize the revenue from a single private-budget agent, the revenue optimal mechanism involves a complex menu of lotteries, and its menu size can be exponential in the size of the support of the budget distribution (Devanur and Weinberg, 2017).

Among all menu mechanisms, there is a subclass of simple mechanisms closely related to price posting – *pricing-based mechanisms*. The subclass of pricing-based mechanisms considers mechanisms where the menu (offered by the mechanism to each agent) is equivalent to posting a per-unit price. For linear agents, it is well known that there exist pricing-based mechanisms that are optimal in multi-agent environments (Myerson, 1981; Riley and Zeckhauser, 1983; Alaei et al., 2019). In this section, we provide an implementation of pricing-based mechanisms for non-linear agents. The performances of pricing-based mechanisms compared to the optimal in various non-linear utility environments are discussed later in Section 4 and Section 5.

## 3.1 Price-posting Payoff Curves

In this section, we provide an analysis of the expected payoff of the principal in single-agent environments. We first define the market clearing price for selling the item to an agent with ex ante supply constraint $q \in [0, 1]$.

**Definition 3.1.** *For any* $q \in [0, 1]$*, the* market clearing price $p^q$ *for the ex ante supply constraint* $q$ *is the highest per-unit price at which there exists a tie-breaking rule for demands such that the aggregated demand is* $q$*, i.e.,* $D(p^q) = q$.[11]

The existence of the market clearing price is implied by the assumption that the demand correspondence is upper hemi-continuous. Moreover, the ordinary goods assumption implies that the market clearing price is weakly decreasing in ex ante supply constraint $q$.

We define the *price-posting payoff curves* as the expected payoff in the single-agent problems by posting per-unit prices, as a function of ex ante supply constraints, and we refer to them as price-posting revenue curves when we focus on revenue as the payoff function. Moreover, since a price-posting payoff curve is not generally concave, we can apply the ironing technique from Myerson (1981) to get the concave hull of the price-posting payoff curves, which captures the optimal payoff the principal can get by posting random prices given the ex ante supply constraint.

**Definition 3.2.** *The* price-posting payoff curve $P(q)$ *for any* $q \in [0, 1]$*, is defined as the expected payoff from posting market clearing price* $p^q$*. The* ironed price-posting payoff curve $\bar{P}$ *is the concave hull of the price-posting payoff curve* $P$*.*

[11] In the case where the demand of the agent is not unique, we may need to consider randomized tie-breaking rules for the demand correspondence in order to satisfy the expected demand constraint.

Since the price-posting payoff curve is a single-dimensional function, the ironed price-posting payoff curve $\bar{P}$ can be obtained by randomizing over at most two per-unit prices for each ex ante constraint $q$.

Note that in single-agent environments, the expected payoff of the principal is uniquely captured by the price-posting payoff curves regardless of the details of the agent's utility models if the principal is restricted to posting (randomized) per-unit prices. Later in Sections 3.2 and 3.3, we will show that this observation extends for general pricing-based mechanisms when there are multiple agents.

## 3.2 Pricing-based Mechanisms for Linear Agents

We first introduce the pricing-based mechanisms for linear agents when there are multiple agents. We will introduce the concept of *quantiles* and then describe the pricing-based mechanisms for linear agents in quantile space.

**Quantiles for Linear Agents.** For an agent with linear utility, the quantile of value $v$ given value distribution $F$ is defined as

$$q(v) \in \left[\Pr_{z\sim F}[z \geq v], \Pr_{z\sim F}[z > v]\right],$$

where quantile $q(v)$ is drawn uniformly randomly if the above interval is not degenerate. For distributions without point masses, $q(v) = 1 - F(v)$. Conversely, the value corresponding to quantile $\hat{q}$ is $v(\hat{q}) \triangleq \sup_z\{q(z) \geq \hat{q}\}$. Note that for any distribution $F$ over value $v$, quantile $q$ is uniformly distributed between $[0, 1]$.

Based on this mapping for linear agents, any mechanism in the value space can be equivalently converted to a mechanism in the quantile space, and vice versa. Moreover, for a linear agent, the market clearing price $p^q$ coincides with the value corresponding to quantile $q$, i.e., $p^q = v(q)$. Without ambiguity, we also refer to $P(q)$ as the price-posting payoff for quantile $q$.

**Pricing-based Mechanisms for Linear Agents.** For linear agents, every deterministic, dominant strategy incentive compatible (DSIC) and interim individual rational (IIR) mechanism (e.g., the Bayesian optimal mechanism) can be implemented as a pricing-based mechanism (Alaei et al., 2013).[12] Based on our definition of quantiles for linear agents, the pricing-based mechanism for linear agents in quantile space is defined as follows:

- each agent $i$'s reported value $v_i$ is mapped to quantile $q_i$ based on $F_i$;

[12] Note that for linear agents, any BIC and IIR mechanism can be converted to a DSIC and IIR mechanism (Gershkov et al., 2013), and any randomized mechanism can be converted to a deterministic mechanism through purification (Chen et al., 2019) when type distributions are continuous. Moreover, the revenue optimal mechanism for linear agents is deterministic and DSIC.

- there is a threshold function $Q_i(q_{-i})$ for each agent $i$ such that agent $i$ is offered with a per-unit price $p^{Q_i(q_{-i})}$ where $q_{-i} \triangleq \{q_j\}_{j\in N\setminus\{i\}}$. Agent $i$ receives an item if and only if $q_i \leq Q_i(q_{-i})$.[13]

Essentially, $Q_i(q_{-i})$ is the threshold on allocation in the quantile space, and the pricing-based mechanism is uniquely determined by this profile of threshold functions.

Since the distribution over quantiles is uniform on $[0, 1]$ regardless of the valuation distribution, the expected payoff of the pricing-based mechanism $\mathcal{M}$ with threshold functions $\{Q_i\}_{i\in N}$ is uniquely pinned down by the price-posting payoff curves $\{P_i\}_{i\in N}$ as

$$\mathcal{M}(\{P_i\}_{i\in N}) = \mathbf{E}_{\{q_i\}_{i\in N}\sim U[0,1]^n}\left[\sum_{i\in N} P_i(Q_i(q_{-i}))\right].$$

In the next section, we will use this formula to derive the performance of pricing-based mechanisms for non-linear agents.

## 3.3 Pricing-based Mechanisms for Non-linear Agents

In this section, we provide a quantile-space implementation of pricing-based mechanisms for non-linear agents. For this purpose, we need to first provide a formal definition of quantiles for non-linear agents. For linear agents, the definition of quantiles is intuitive since higher values are mapped to lower quantiles according to the cumulative distribution function of the valuation distribution. In contrast, the construction of quantiles for non-linear agents is not obvious since the private types of non-linear agents can be multi-dimensional and there may not exist a natural order on types that is consistent for all mechanisms. For example, it is unclear how to rank a more risk averse agent with high value against a less risk averse agent with low value, or an agent with high value and low budget against another agent with low value and high budget. In this section, we introduce a definition of quantiles for non-linear agents based on their demands given market clearing prices under different ex ante supply constraints.

**Quantiles for Non-linear Agents.** For a non-linear agent with type $t$, define function $H^t(q) = d^t(p^q)$ where $d^t(p^q)$ is the demand of type $t$ under market clearing price $p^q$ that sells the item to the agent with probability $q$. Under the ordinary goods assumption, we can focus on tie-breaking rules such that the demand $d^t(p^q)$ is weakly increasing in $q$, which implies that function $H^t(q)$ is weakly increasing in $q$ for all type $t$.[14] Moreover, $H^t(0) = 0$ and $H^t(1) = 1$, and thus $H^t$ can be

[13] It is possible that in pricing-based mechanisms for linear agents, based on the tie breaking rules, agent $i$ receives an item only if $q_i < Q_i(q_{-i})$ for some $q_{-i}$. In this case, by slightly abusing the notation, we allow $Q_i(q_{-i})$ to take values such as $z^+$ for $z \in [0, 1]$, and we say $q_i \leq z^+$ if and only if $q_i < z$.

[14] In the case where the demands of the agent need to be randomized to satisfy the expected demand constraint, we

viewed as a distribution supported on $[0, 1]$ with $H^t(q)$ being the cumulative distribution function of quantile for any type $t$.

**Definition 3.3** (quantiles)**.** *For any non-linear agent with type $t$ and demand function $d^t(\cdot)$, the randomized quantile $q$ for type $t$ is drawn from the distribution $H^t$ with $H^t(q) \triangleq d^t(p^q)$.*

Our definition of quantiles for non-linear agents maps the potentially multi-dimensional types of non-linear agents into single-dimensional quantiles based on the demand functions for per-unit pricings. Moreover, given any market clearing price, types with higher demands are mapped to lower quantiles with higher probabilities, and hence are more likely to win an item in pricing-based mechanisms where items are allocated based on the thresholds in the quantile space. Note that the definition of quantile for non-linear agents encompasses the definition for linear agents. In particular, for a linear agent with continuous value distribution and realized value $v(q)$, the distribution $H^{v(q)}$ is a point-mass distribution at $q$ and $H^{v(q)}(\cdot)$ is a step function at $q$.

This definition of quantiles results in quantiles, for all distributions over non-linear preferences, being uniformly distributed. Specifically, for any non-linear agent, the cumulative distribution function of quantile $q$ is

$$\mathbf{Pr}_{t\sim\bar{F},q\sim H^t}[q \le z] = \mathbf{E}_{t\sim\bar{F}}\big[H^t(z)\big] = \mathbf{E}_{t\sim\bar{F}}\big[d^t(p^z)\big] = z$$

where the last inequality holds due to the definition of the market clearing price $p^z$.

**Pricing-based Mechanisms for Non-linear Agents.** Given the definition of quantiles for non-linear agents, we define the pricing-based mechanisms for non-linear agents in Algorithm 1, which first maps the types of all agents into quantiles based on Definition 3.3, and allocates items to agents based on thresholds in the quantile space.

**Theorem 1** (Implementation)**.** *For any non-linear agents with price-posting payoff curves $\{P_i\}_{i\in N}$, the pricing-based mechanism $\mathcal{M}$ with threshold functions $\{Q_i\}_{i\in N}$ defined in Algorithm 1 is DSIC and IIR for non-linear agents, and its expected payoff is*

$$\mathcal{M}(\{P_i\}_{i\in N}) = \mathbf{E}_{\{q_i\}_{i\in N}\sim U[0,1]^n}\left[\sum_{i\in N} P_i(Q_i(q_{-i}))\right].$$

consider a randomization over tie-breaking rules such that the demand function $d^t(p^q)$ is weakly increasing for all realized selections. Moreover, for different selections of demand function $d^t(p^q)$, the defined function $H^t(q)$ may be different for any given $t$ and $q$. However, the results in our paper hold for all selection rules satisfying the required weak monotonicity.

[15]To simplify the exposition, we only state the payment rules without the consideration of the tie-breaking rules in the pricing-based mechanism for linear agents. The details for constructing consistent tie-breaking rules are explained in Appendix C.

**Algorithm 1:** Pricing-based Mechanisms for Non-linear Agents

**Input:** Non-linear agents $\{(\mathcal{T}_i, \bar{F}_i, u_i)\}_{i \in N}$; a pricing-based mechanism $\mathcal{M}$ for linear agents with threshold function $\{Q_i\}_{i \in N}$ on quantiles.

1 For each agent $i$ with private type $t_i$, map the type to a random quantile $q_i$ according to cumulative distribution function $H_i^{t_i}(\cdot)$.

2 For each agent $i$, calculate quantile threshold as $\hat{q}_i = Q_i(q_{-i})$.

3 For each agent $i$, her allocation is $x_i = 1$ if $q_i \leq \hat{q}_i$ and $x_i = 0$ otherwise. The payment of agent $i$ is $p_i = p^{\hat{q}_i} \cdot d(t_i, p^{\hat{q}_i})$ regardless of the allocation.[15]

Theorem 1 implies that for any set of linear agents and non-linear agents with the same price-posting payoff curves $\{P_i\}_{i \in N}$, given any pricing-based mechanism $\mathcal{M}$ for linear agents, there exists a pricing-based mechanism for non-linear agents with the same allocation rule in the quantile space that guarantees the same expected payoff. Therefore, we can use $\mathcal{M}(\{P_i\}_{i \in N})$ to represent the expected payoff of pricing-based mechanism $\mathcal{M}$ given price-posting payoff curves $\{P_i\}_{i \in N}$ without the consideration of the detailed utility models of the agents. The proof of Theorem 1 relies on the following alternative interpretation of Algorithm 1.

**Price-posting Equivalence Interpretation.** For agent $i$ with type $t_i$, conditioning on threshold $\hat{q}_i = Q_i(q_{-i})$ which only depends on other agents' types, her expected allocation over the randomness of her quantile $q_i$ is

$$\mathbf{E}[\mathbf{1}[q_i \leq \hat{q}_i] \mid \hat{q}_i, t_i] = H_i^{t_i}(\hat{q}_i) = d(t_i, p^{\hat{q}_i})$$

and her payment is deterministically $p^{\hat{q}_i} \cdot d(t_i, p^{\hat{q}_i})$. Therefore, from the perspective of agent $i$, the constructed mechanism is equivalent to offering a market-clearing price $p^{\hat{q}_i}$ where the threshold quantile $\hat{q}_i$ only depends on other agents' types.

*Proof of Theorem 1.* Both DSIC and IIR of mechanism $\mathcal{M}$ are immediately guaranteed from the price-posting equivalence interpretation. Moreover, for any non-linear agent $i$, conditional on the event that her threshold quantile is $\hat{q}_i = Q_i(q_{-i})$, the expected payoff from agent $i$ in this event is $P_i(\hat{q}_i)$. Therefore, by linearity of expectation, the expected payoff of mechanism $\mathcal{M}$ is

$$\mathcal{M}(\{P_i\}_{i \in N}) = \mathbf{E}_{q \sim U[0,1]^n}\left[\sum_{i \in N} P_i(Q_i(q_{-i}))\right] \tag{1}$$

since the marginal distribution over quantile $q_i$ is drawn from a uniform distribution $U[0,1]$ for all agent $i$. $\square$

### 3.4 Marginal Payoff Maximization

The implementation of pricing-based mechanisms in quantile space is inspired by the marginal revenue maximization mechanism proposed in Bulow and Roberts (1989) for linear agents. Bulow and Roberts (1989) show that the marginal revenue maximization mechanism is revenue-optimal for linear agents. This result can be easily extended to general payoff maximization for linear agents. Specifically, Alaei et al. (2013) have shown that for linear agents, it is without loss of optimality to focus on pricing-based mechanisms. The expected payoff of a pricing-based mechanism $\mathcal{M}$ with threshold functions $Q_i(q_{-i})$ given price-posting payoff curves $\{P_i\}_{i\in N}$ can be reformulated in terms of marginal revenue as

$$\mathcal{M}(\{P_i\}_{i\in N}) = \mathbf{E}_{q\sim U[0,1]^n}\left[\sum_{i\in N} P_i(Q_i(q_{-i}))\right] = \sum_{i\in N}\mathbf{E}_{q\sim U[0,1]^n}[P_i(Q_i(q_{-i}))]$$
$$= \sum_{i\in N}\mathbf{E}_{q\sim U[0,1]^n}[P_i'(q_i)\cdot \hat{x}_i(q_i, q_{-i})] = \mathbf{E}_{q\sim U[0,1]^n}\left[\sum_{i\in N} P_i'(q_i)\cdot \hat{x}_i(q_i, q_{-i})\right]$$

where $\hat{x}_i(\cdot)$ is the allocation rule in quantile space induced by the threshold function $Q_i(q_{-i})$. That is, $\hat{x}_i(q_i, q_{-i}) = 1$ if and only if $q_i \le Q_i(q_{-i})$. The second and the last equality holds by the linearity of expectation, and the third equality holds since by the definition of $\hat{x}_i(\cdot)$, we have

$$P_i(Q_i(q_{-i})) = \int_0^1 P_i'(q_i)\cdot \hat{x}_i(q_i, q_{-i})\mathrm{d}q_i = \mathbf{E}_{q_i\sim U[0,1]}[P_i'(q_i)\cdot \hat{x}_i(q_i, q_{-i})]\,.$$

As interpreted by Bulow and Roberts (1989), $P_i'(q_i)$ can be viewed as the marginal payoffs for selling an item to agent $i$ with quantile $q_i$. To maximize the expected payoff, the marginal payoff maximization mechanism (MPM) allocates the items to agents with the highest marginal payoffs $P_i'(q_i)$. In the case where $P_i$ is concave for any agent $i$ and hence $P_i'(q_i)$ is monotone in $q_i$, the resulting mechanism is a pricing-based mechanism and maximizes the expected payoff of the principal. In the case that $P_i(q_i)$ is not concave in $q_i$, we can apply the ironing trick in Myerson (1981) by considering allocations that maximize the ironed marginal payoffs $\sum_{i\in N}\bar{P}_i'(q_i)\cdot \hat{x}_i(q_i, q_{-i})$.

For non-linear agents, since restricting attention to pricing-based mechanisms may not be without loss of generality, the marginal payoff maximization mechanism may not be optimal among all possible mechanisms. Nonetheless, by adopting the same argument as in linear agents, the marginal payoff maximization mechanism maximizes the principal's expected payoff among the subclass of pricing-based mechanisms regardless of the agents' utility models. Moreover, in Section 5, we will show that marginal payoff maximization is approximately optimal for broad classes of non-linear utilities.

Finally, note that the expected payoff of the marginal payoff maximization mechanism only depends on the price-posting payoff curves, not the detailed utility models. Therefore, we denote the payoff of marginal payoff maximization $(\mathrm{MPM})$ for agents with price-posting payoff curves $\{P_i\}_{i\in N}$ as $\mathrm{MPM}(\{P_i\}_{i\in N})$.

## 3.5 Illustrations of the Implementations

The description of pricing-based mechanisms in the quantile space is beneficial for understanding the economics behind payoff maximization, and will be useful for our subsequent analysis that shows its optimality or approximate optimality. In this section, we will first illustrate how simple pricing-based mechanisms are implemented for non-linear agents in the value space using examples.

**Marginal payoff maximization.** We first consider the case where agents are risk averse. We show that any risk-averse agent, when facing a per-unit price, will behave the same as a linear agent. The proof of this claim is provided in Appendix A.1.

**Claim 1.** *For any risk-averse agent with type $t = (v, \varphi)$ and for any per-unit price $p$, the expected utility of the agent is maximized at allocation $z \in \{0, 1\}$.*

Therefore, the price-posting payoff curves of risk-averse agents coincide with the price-posting payoff curves of linear agents with the same value distributions. Moreover, the pricing-based mechanisms for risk-averse agents also coincide with linear agents in the value space. For example, for i.i.d. linear agents with regular valuation distributions, the marginal revenue maximization mechanism can be implemented as the second price auction with monopoly price as the anonymous reserve (Myerson, 1981). Our observations above imply that for risk-averse agents with i.i.d. value distributions and arbitrary distribution on the risk aversion function, the marginal revenue maximization mechanism can also be implemented as the second price auction with monopoly price as the anonymous reserve.

The implementation of marginal revenue maximization is trickier for other non-linear utilities such as budgeted utilities. The main reason is that a budgeted agent behaves differently from a linear agent when facing per-unit prices. For example, consider two public-budget agents with values drawn from i.i.d. distributions that are uniform in $[0, 1]$. All agents have a public budget of $\frac{1}{2}$. In this example, it is easy to compute that for any agent $i$ with value $v_i$, the demand function is

$$d(v_i, p^q) = \begin{cases} 0 & v_i < \min\left\{\frac{1}{1+2q}, 1-q\right\} \\ \min\left\{1, \frac{1}{2}+q\right\} & v_i \geq \min\left\{\frac{1}{1+2q}, 1-q\right\}. \end{cases}$$

Moreover, the cutoff quantile for the marginal revenue being $0$ is $\frac{1}{2}$. Note that according to the demand function, any value $v_i < \frac{1}{2}$ has a deterministic quantile $1 - v_i > \frac{1}{2}$. Therefore, in the marginal revenue maximization mechanism, if both agents have values below $\frac{1}{2}$, the item is not sold and the revenue is 0. If only one agent has a value above $\frac{1}{2}$, the higher value agent receives the item with probability 1 and price 1. If both agents have values above $\frac{1}{2}$, and we assume that $v_1 \geq v_2 \geq \frac{1}{2}$ and tie is broken in favor of agent 1, then agent 1 will receive an item with probability $\frac{5}{8} + \frac{1}{2}(v_2 + v_2^2)$ and agent 2 will receive an item with the remaining probability. This allocation probability is calculated based on the mapping to randomized quantiles defined in Algorithm 1. In the last case, agent 1 always pays a deterministic price of $\frac{1}{2}$, and agent 2 pays a deterministic price of $\frac{1}{2}$ only with probability $1 - \frac{1}{2v_2}$.

**Sequential posted pricing.** Another simple pricing-based mechanism that is widely adopted in practice is sequential posted pricing (SPP). Intuitively, sequential posted pricing mechanisms approach the agents according to a given order and sell the items to agents using per-unit prices until the supply runs out. Next, we provide a formal definition of sequential posted pricing in quantile space.

**Definition 3.4.** *A* sequential posted pricing mechanism *is parameterized by* $(\{\hat{o}_i\}_{i\in N}, \{\hat{q}_i\}_{i\in N})$ *where* $\{\hat{o}_i\}_{i\in N}$ *denotes an order of the agents and* $\{\hat{q}_i\}_{i\in N}$ *denotes the cutoffs in quantile space. Agent* $i$ *receives an item if his randomized quantile is at most* $\hat{q}_i$ *and the items are not sold out to previous agents, i.e., the threshold functions* $\{Q_i\}_{i\in N}$ *satisfy that* $Q_i(q_{-i}) = \hat{q}_i \cdot \prod_{j:\hat{o}_j<\hat{o}_i} \mathbf{1}[q_j > \hat{q}_j]$.

Note that the payoff of the sequential posted pricing mechanism is also uniquely determined by the price-posting payoff curves $\{P_i\}_{i\in N}$ of the agents. Therefore, we denote $\mathrm{SPP}(\{P_i\}_{i\in N})$ as the optimal payoff among the class of sequential posted pricing mechanisms.

To implement the sequential posted pricing mechanism in the value space, it can be viewed as sequentially offering market clearing price $p^{\hat{q}_i}$ to each agent $i$ according to order $\{\hat{o}_i\}_{i\in N}$ if items are not sold out to previous agents. In the sequential posted pricing mechanism, each agent's optimal demand may be a non-degenerate lottery given per-unit prices, and we assume that the lottery is realized immediately after each agent's purchase decision.

# 4 Robust Optimality

In this section, we focus on a robust setting wherein the seller possesses only the knowledge of the agents' aggregated demand functions, lacking detailed information regarding the non-linear utility function or the type distribution. We show that the marginal revenue maximization mechanism (see Section 3.4) is robustly optimal in such environments.

Given any aggregated demand function $D(p)$ that is a non-increasing function,[16] let $\mathcal{P}(D)$ be

[16]We focus on non-increasing demand functions since it is without loss under the ordinary good assumption.

the set of feasible agent utility models $(\mathcal{T}, \bar{F}, u)$ such that for any $p \geq 0$,

$$D(p) = \mathbf{E}_{t\sim\bar{F}}\left[d^t(p;u)\right]$$

where $d^t(p;u)$ is the demand function of type $t \in \mathcal{T}$ induced by utility function $u$. Note that the aggregated demand function uniquely pins down the price-posting revenue curve $P$ regardless of the utility models.

For any mechanism $\mathcal{M}$, let $\mathcal{M}(\left\{(\mathcal{T}_i, \bar{F}_i, u_i)\right\}_{i\in N})$ be the expected revenue of mechanism $\mathcal{M}$ given a profile of utility models $\left\{(\mathcal{T}_i, \bar{F}_i, u_i)\right\}_{i\in N}$. In this section, we assume that the seller can only observe the aggregated demand function $D_i(p)$ for all agents $i$ without knowing the detailed utility models $(\mathcal{T}_i, \bar{F}_i, u_i)$. The justification is that the detailed utility models are often hard to estimate based on the historical purchase data from consumers for practical applications. However, it is relatively easier to estimate the demand functions based on posted pricing (Bajari et al., 2015).

When the seller is ignorant of the detailed utility models, the robust approach considers the objective of maximizing the expected revenue over the worst-case utility models that are consistent with the observed aggregated demand functions. That is, the seller chooses a mechanism $\mathcal{M}$ to maximize

$$\min_{(\mathcal{T}_i,\bar{F}_i,u_i)\in\mathcal{P}(D_i),\forall i\in N} \mathcal{M}(\left\{(\mathcal{T}_i, \bar{F}_i, u_i)\right\}_{i\in N}).$$

**Theorem 2.** *For any profile of non-increasing aggregated demand functions $\{D_i\}_{i\in N}$, if the seller does not observe the detailed profile of utility models, the marginal revenue maximization mechanism is the max-min revenue optimal mechanism, i.e.,*

$$\mathrm{MPM} \in \operatorname*{argmax}_{\mathcal{M}} \min_{(\mathcal{T}_i,\bar{F}_i,u_i)\in\mathcal{P}(D_i),\forall i\in N} \mathcal{M}(\left\{(\mathcal{T}_i, \bar{F}_i, u_i)\right\}_{i\in N}).$$

The interpretation of the robustness result is that if the seller is uncertain about the detailed utility models, they should avoid resorting to complex lottery mechanisms for second-degree price discrimination.

*Proof.* Note that in the robust environments, the seller does not know the detailed utility model to compute the demands based on the report types. Fortunately, the implementation of pricing-based mechanisms for non-linear agents in Algorithm 1 only requires the knowledge of the aggregated demand functions, and the seller can directly elicit the demand functions $d_i^{t_i}(p)$ as a function of $p$ from the agents in order to compute the randomized quantile $q_i$ for them. Using the same argument as in Theorem 1, it is a weakly dominant strategy for all agents to truthfully report their demand functions, and the expected revenue of the marginal revenue maximization mechanism is

$$\mathrm{MPM}(\{P_i\}_{i\in N}) = \max_{\text{pricing-based mech. } \mathcal{M}^\dagger} \mathcal{M}^\dagger(\{P_i\}_{i\in N})$$

regardless of the underlying utility model. Moreover, given any non-decreasing demand function $D_i$, there exists a linear utility model $(\mathcal{T}_i^*, \bar{F}_i^*, u_i^*)$ with $\mathcal{T}_i^* = \mathbb{R}$, $u_i^*(t, x, p) = t \cdot x - p$ such that $(\mathcal{T}_i^*, \bar{F}_i^*, u_i^*) \in \mathcal{P}(D_i)$. Moreover, for agents with linear utilities, the revenue-optimal mechanism is marginal revenue maximization even when the seller knows the detailed utility model. Therefore, for any mechanism $\mathcal{M}$, we have

$$\min_{(\mathcal{T}_i, \bar{F}_i, u_i)\in\mathcal{P}(D_i), \forall i\in N} \mathcal{M}(\{(\mathcal{T}_i, \bar{F}_i, u_i)\}_{i\in N}) \le \mathcal{M}(\{(\mathcal{T}_i^*, \bar{F}_i^*, u_i^*)\}_{i\in N})$$
$$\le \max_{\text{pricing-based mech. } \mathcal{M}^\dagger} \mathcal{M}^\dagger(\{P_i\}_{i\in N}) = \mathrm{MPM}(\{P_i\}_{i\in N}).$$

That is, the marginal revenue maximization mechanism is robustly optimal for non-linear agents.

$\square$

**Remark:** As illustrated in the proof of Theorem 2, the robust optimality of the marginal revenue maximization mechanism holds as long as the ambiguity set includes the linear utility model. This observation implies that our results remain valid even when the principal possesses more refined knowledge about the agents' utility models—for example, knowing that agents are risk-averse or face budget constraints. Based on this, we can immediately derive two corollaries.

- Suppose the principal knows that the agents are (weakly) risk-averse and knows their valuation distributions, but is uncertain about the specific risk preferences of the agents. Our result implies that marginal revenue maximization—which, in symmetric environments with risk-averse agents, can be implemented via a second-price auction with monopoly reserves—is robustly optimal in the presence of such uncertainty.

- Suppose the principal knows that the agents are risk-neutral but face budget constraints, while being uncertain about the joint distributions over values and budgets—except for their aggregated demand functions. Our result implies that marginal revenue maximization remains robustly optimal in this case.

# 5 Reduction Framework and Approximation Guarantees

For non-linear agents, the pricing-based mechanisms proposed in Section 3 in general are not optimal for payoff maximization if the principal has precise information regarding the environments.

In this section, we show that those mechanisms are approximately optimal for broad classes of non-linear utilities by providing a reduction framework that extends the approximation bounds of pricing-based mechanisms for linear agents to non-linear agents. An important application of our general reduction framework is that the robustly optimal mechanism, marginal payoff maximization, is approximately optimal against the optimal detail-dependent mechanism.

Specifically, our reduction framework (Theorem 3) will show that when non-linear agents have small resemblances, by focusing on pricing-based mechanisms that are robust to the details of the non-linear model, the payoff loss from this reduction is not large.

## 5.1 Optimal Payoff Curves and Resemblance

In this section, we analyze the single-agent problem with ex ante supply constraints. We define the *optimal payoff curves* as follows. The main difference between the optimal payoff curves and the price-posting payoff curves (Definition 3.2) is that the former does not restrict attention to mechanisms that post per-unit prices.

**Definition 5.1.** *The* optimal payoff curve $R(q)$ *is a mapping from any ex ante supply constraint* $q$ *to the optimal ex ante payoff for the single agent problem that in expectation sells the item with probability* $q$*, i.e.,*

$$\begin{aligned} R(q) = \max_{\pi \text{ is BIC, IIR}} \quad & \mathbf{E}_{(x,p)\sim\pi}[\mathbf{E}_{t\sim\bar{F}}[\Psi(t, x(t), p(t))]] \\ \textit{s.t.} \quad & \mathbf{E}_{(x,p)\sim\pi}[\mathbf{E}_{t\sim\bar{F}}[x(t)]] = q. \end{aligned}$$

**Lemma 1.** *The optimal payoff curve is concave.*

This lemma is immediate since in this case the space of incentive compatible mechanisms is closed under convex combination, and the payoff function is linear in the space of mechanisms. Next we review the relation between the optimal revenue curves and the concave hull of the price-posting revenue curves for linear agents.

**Lemma 2** (Bulow and Roberts, 1989)**.** *The optimal revenue curve* $R$ *of a linear agent is equal to her ironed price-posting revenue curve* $\bar{P}$*.*

In general, for non-linear agents, the optimal payoff (e.g., revenue) curves and the ironed price-posting payoff curves are not equivalent, and the mechanism maximizing the ex ante payoff is more complicated and extracts strictly higher payoff than the optimal price posting mechanism or randomizations over price posting mechanisms. We quantify the extent to which a non-linear agent resembles a linear agent based on the maximum payoff gap between posted pricing and optimal mechanisms in a single-agent environment with ex ante supply constraints.

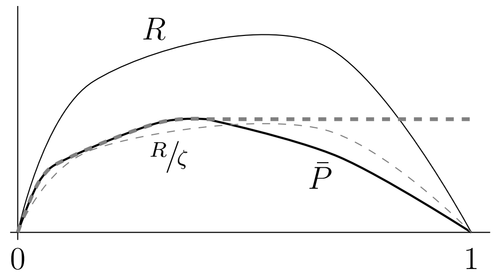


Figure 1: The thin dashed curve is the optimal payoff curve scaled by a factor of $\zeta$, and the thick dashed curve is the monotone extension of the ironed price-posting payoff curve.[17] The definition of $\zeta$ resemblance is satisfied if the thick dashed curve is always above the thin dashed line.

**$\zeta$-resemblance.** We introduce *$\zeta$-resemblance* of an agent to measure how closely a non-linear agent resembles a linear agent whose ironed price-posting payoff curve and optimal payoff curve are equal. This definition is illustrated in Figure 1. Note that based on our definition, Lemma 2 implies that a linear agent is 1-resemblant.

**Definition 5.2** ($\zeta$-resemblance)**.** *An agent's ironed price-posting payoff curve $\bar{P}$ is* $\zeta$-resemblant *to her optimal payoff curve $R$, if for all $q^\dagger \in [0, 1]$, there exists $q \leq q^\dagger$ such that $\bar{P}(q) \geq {}^1/_\zeta \cdot R(q^\dagger)$. Such an agent is* $\zeta$-resemblant*.*

Our definition implies that for a $\zeta$-resemblant agent, given any ex ante supply constraint $q$, there exists a distribution over per-unit prices with expected sale probability at most $q$ that achieves at least $\frac{1}{\zeta}$ fraction of the optimal payoff with this ex ante constraint.[18] Since posted pricing mechanisms are optimal mechanisms for a linear agent, smaller $\zeta$ implies that non-linear agents resemble linear agents better.

## 5.2 Reduction Framework

The reduction framework in this section relies on the following technique of ex ante relaxation for bounding the optimal payoff of the principal.

**Ex Ante Relaxation.** In general for non-linear agents, it is challenging to characterize the optimal mechanism and quantify the optimal payoff of the principal. Thus, we adopt the benchmark of *ex ante relaxation* (Yan, 2011), which serves as an upper bound of the optimal payoff. For selling $m$ units of identical items, a sequence of ex ante quantiles $\{q_i\}_{i \in N}$ is ex ante feasible if $\sum_i q_i \leq m$.

[17] We consider the monotone extension of the ironed price-posting payoff curve, as we allow for underselling when analyzing the optimal revenue from price-posting mechanisms under ex ante supply constraints. Our condition is weaker than requiring that $\bar{P}(q) \geq {}^1/_\zeta \cdot R(q)$ for all $q \in [0, 1]$.

[18] This is equivalent to an environment of selling items to a continuum of agents with a supply constraint, where $\zeta$-resemblance measures the worst-case gap (over all possible supplies) between optimal mechanisms and posted pricing mechanisms for non-linear agents.

We denote the set of ex ante feasible quantiles by EAF. The optimal ex ante payoff given the optimal payoff curves $\{R_i\}_{i \in N}$ is

$$\mathrm{EAR}(\{R_i\}_{i \in N}) = \max_{\{q_i\}_{i \in N} \in \mathrm{EAF}} \sum\nolimits_{i \in N} R_i(q_i).$$

Note that $\mathrm{EAR}(\{R_i\}_{i \in N})$ is an upper bound on the optimal payoff since any feasible mechanism must satisfy the ex ante feasibility constraints, and the ex ante payoff optimizes the principal's expected payoff under such constraints. Moreover, the ex ante payoff is uniquely determined by the optimal payoff curves regardless of the detailed utility models of the agents.

**Reduction of Approximation Guarantees.** Now we present the following theorem: a reduction framework that extends the approximation guarantees of every pricing-based mechanism for linear agents to non-linear agents.

**Theorem 3** (Reduction Framework)**.** *For any pricing-based mechanism $\mathcal{M}$ and any set of non-linear agents with price-posting payoff curves $\{P_i\}_{i \in N}$ and optimal payoff curves $\{R_i\}_{i \in N}$, if*

- *all non-linear agents are $\zeta$-resemblant; and*
- *mechanism $\mathcal{M}$ is a $\gamma$-approximation to the ex ante relaxation for linear agents with the same price-posting payoff curves, i.e., $\mathcal{M}(\{P_i\}_{i \in N}) \geq {}^{1}/_{\gamma} \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i \in N})$,*

*then mechanism $\mathcal{M}$ for non-linear agents (from Algorithm 1) is a $\gamma\,\zeta$-approximation to the ex ante relaxation, i.e., $\mathcal{M}(\{P_i\}_{i \in N}) \geq {}^{1}/_{\gamma\zeta} \cdot \mathrm{EAR}(\{R_i\}_{i \in N})$.*

Before the proof of Theorem 3, we first explain the conditions in Theorem 3, and in particular why we refer to condition $\mathcal{M}(\{P_i\}_{i \in N}) \geq {}^{1}/_{\gamma} \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i \in N})$ as the approximation guarantees for linear agents.

Note that for both pricing-based mechanism $\mathcal{M}$ and the ex ante relaxation, the expected payoff of the principal only depends on the price-posting payoff curves and optimal payoff curves respectively, not on the details of the utility models. Therefore, for any set $\mathcal{A}$ of non-linear agents with price-posting payoff curves $\{P_i\}_{i \in N}$, we can define a *linear agents analog* as a set $\mathcal{A}^L$ of agents with the same price-posting payoff curves.[19] By Theorem 1, the expected payoffs of mechanism $\mathcal{M}$ are the same for both non-linear agents $\mathcal{A}$ and the linear agents analog $\mathcal{A}^L$, which equals

[19]For payoff functions such as the revenue, the linear agents analog is well defined since the price-posting revenue curve $P(q)$ of a linear agent uniquely pins down her valuation distribution as $v(q) = \frac{P(q)}{q}$. For general payoff function, given the price-posting payoff curves $\{P_i\}_{i \in N}$ of the non-linear agents, there may not exist type distributions for linear agents such that their price-posting payoff curves coincide with $\{P_i\}_{i \in N}$. However, both the payoffs for pricing-based mechanisms and the ex ante relaxation are well defined given the payoff curves, and Theorem 3 does not rely on the existence of type distributions for linear agents that correspond to the given price-posting payoff curves. Hence, we can refer to the linear agents analog even without the existence of the underlying type distributions.

$\mathcal{M}(\{P_i\}_{i\in N})$ defined in Equation (1). Moreover, by Lemma 2, the optimal payoff curves for the linear agents analog are $\{\bar{P}_i\}_{i\in N}$, and hence the optimal ex ante payoff for the linear agents analog is $\mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$. Therefore, the condition that $\mathcal{M}(\{P_i\}_{i\in N}) \geq 1/\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$ can be interpreted as the approximation ratio of pricing-based mechanism $\mathcal{M}$ for the linear agents analog with respect to its ex ante relaxation benchmark.[20]

Another condition in Theorem 3 is the requirement that all non-linear agents are $\zeta$-resemblant. This condition is also necessary for pricing-based mechanisms to be approximately optimal for non-linear agents. Without this assumption, even in single-agent environments with supply constraints, there exists mechanisms with more complicated price discrimination schemes that significantly outperform simple pricing-based mechanisms.

*Proof of Theorem 3.* Given the assumption that pricing-based mechanism $\mathcal{M}$ is approximately optimal for linear agents, i.e., $\mathcal{M}(\{P_i\}_{i\in N}) \geq 1/\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$, it is sufficient to show that $\mathrm{EAR}(\{\bar{P}_i\}_{i\in N}) \geq 1/\zeta \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$ when all non-linear agents are $\zeta$-resemblant.

Let $\{q_i^\dagger\}_{i\in N} \in \mathrm{EAF}$ be the profile of optimal ex ante quantiles for optimal payoff curves $\{R_i\}_{i\in N}$. Since the ironed price-posting payoff curves $\{\bar{P}_i\}_{i\in N}$ are $\zeta$-resemblant to the optimal payoff curves $\{R_i\}_{i\in N}$, there exists a sequence of quantiles $\{q_i\}_{i\in N}$ such that for any agent $i$, $q_i \leq q_i^\dagger$ and $\bar{P}(q_i) \geq 1/\zeta \cdot R(q_i^\dagger)$. Note that $\{q_i\}_{i\in N}$ is also ex ante feasible. Therefore,

$$\mathrm{EAR}(\{R_i\}_{i\in N}) = \sum_{i\in N} R_i(q_i^\dagger) \leq \zeta \cdot \sum_{i\in N} \bar{P}_i(q_i) \leq \zeta \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N}). \qquad \square$$

Note that our reduction framework (Theorem 3) is useful for practical applications only if both conditions in the statement are satisfied for commonly adopted pricing-based mechanisms and non-linear utilities. The validity of these assumptions is provided in the following subsections.

## 5.3 Approximations of Pricing-based Mechanisms

In this section, we will mainly focus on the application of two simple mechanisms, marginal payoff maximization and sequential posted pricing. The definitions of these two mechanisms have been provided in Section 3.

**Marginal Payoff Maximization.** For linear agents, Bulow and Roberts (1989); Alaei et al. (2013) have shown that marginal payoff maximization is optimal for payoff maximization. If this mechanism is compared to the ex ante relaxation, Yan (2011) shows that the worst-case approximation factor is at most $1/(1 - \frac{1}{\sqrt{2\pi m}})$ for selling $m$ units of identical items. Combining this with Theorem 3, we obtain the following corollary.

[20]This condition is necessary for our reduction framework since linear utility is a special case of non-linear utilities.

**Corollary 1.** *For selling $m$ units of identical items to any set of non-linear agents with price-posting payoff curves $\{P_i\}_{i \in N}$ that are $\zeta$-resemblant to their optimal payoff curves $\{R_i\}_{i \in N}$, marginal payoff maximization is a $\zeta/(1-\frac{1}{\sqrt{2\pi m}})$-approximation to the ex ante relaxation for non-linear agents, i.e.,* $\mathrm{MPM}(\{P_i\}_{i \in N}) \geq \frac{1}{\zeta} \cdot (1-\frac{1}{\sqrt{2\pi m}}) \cdot \mathrm{EAR}(\{R_i\}_{i \in N})$.

**Sequential Posted Pricing.** Yan (2011) shows that the worst-case approximation factor of sequential posted pricing to the ex ante payoff is also at most $1/(1-\frac{1}{\sqrt{2\pi m}})$ for selling $m$ units to linear agents. Combining this with Theorem 3, we obtain the following corollary.

**Corollary 2.** *For selling $m$ units of identical items to any set of non-linear agents with price-posting payoff curves $\{P_i\}_{i \in N}$ that are $\zeta$-resemblant to their optimal payoff curves $\{R_i\}_{i \in N}$, sequential posted pricing is a $\zeta/(1-\frac{1}{\sqrt{2\pi m}})$-approximation to the ex ante relaxation for non-linear agents, i.e.,* $\mathrm{SPP}(\{P_i\}_{i \in N}) \geq \frac{1}{\zeta} \cdot (1-\frac{1}{\sqrt{2\pi m}}) \cdot \mathrm{EAR}(\{R_i\}_{i \in N})$.

**Discussions of Approximations.** Note that in the approximation factor, the term $1-\frac{1}{\sqrt{2\pi m}}$ converges to 1 in the large market where $m$ increases to infinity. The term $\zeta$ measures how close the non-linear agents resemble linear agents, and the approximation factor becomes better when non-linear agents become more resemblant of linear agents. Corollaries 1 and 2 implies that the approximation ratio will be close to 1 in large markets for non-linear agents that closely resemble linear agents. This implies that sequential posted pricings are approximately optimal for those non-linear agents. The interpretation is that competition and simultaneous implementation are not salient features for payoff maximization even when agents have non-linear utilities.

Moreover, even in small markets where $m$ is small, both marginal payoff maximization and sequential posted pricing remain a constant approximation to the optimal revenue. This holds for arbitrary non-linear agents that are $\zeta$-resemblant, and there is no additional requirement for the type distributions or detailed utility models. In particular, the results apply even when half of the agents are risk averse and the other half have private budget constraints. Although those two mechanisms still entail a constant loss in expected payoffs, the interpretation of the results is not to take the constant factors too literally. First, as we will illustrate later in Section 5.5, the actual performance of those mechanisms can be much better than the worst case bound given real instances. Secondly, the main economic insights should be derived through relative comparisons of approximation factors among different simple mechanisms. For instance, if the principal uses an anonymous pricing mechanism for revenue maximization when agents are asymmetric, the worst case approximation factor for linear agents is unbounded when the number of items is large. However, the two pricing-based mechanisms we propose are approximately optimal even for non-linear utilities as long as they are $\zeta$-resemblant for small $\zeta$.

## 5.4 Resemblance for Revenue Maximization

In this section, we instantiate our reduction framework by showing that the resemblance property is indeed satisfied for canonical models of non-linear utility. We focus on the objective of revenue maximization in this section, and show that both risk-averse agents and private-budget agents resemble linear agents under mild assumptions of the valuation distributions. In Appendix A.2, we show that those assumptions are necessary for pricing-based mechanisms to be approximately optimal even in single-agent settings without ex ante supply constraints. The extensions to welfare and other objective functions are discussed in Section 6.1

Since we will focus on the single-agent analysis in this section, we will drop the subscript representing the agent in all notations.

### 5.4.1 Risk-averse Agent

For a risk-averse agent, her private type is represented by a pair $t = (v, \varphi)$ where $v$ is the private value for the item and $\varphi$ represents the agent's risk preference. Note that we allow arbitrary correlations between the values and risk preferences, and we do not impose structures on the risk preferences such as constant absolute risk aversion (CARA) utilities.

Recall that $F$ is the marginal distribution over values. We assume that distribution $F$ satisfies the monotone hazard rate (MHR) condition. This condition is satisfied by many common distributions in practice such as uniform, exponential, and Gaussian distributions.

**Definition 5.3.** *A distribution $F$ with density function $f$ satisfies* monotone hazard rate (MHR) *condition if the hazard rate function $h(v) \triangleq \frac{f(v)}{1-F(v)}$ is non-decreasing in $v$.*

Essentially, Definition 5.3 implies that the tail probabilities in the valuation distribution are small, i.e., smaller than the exponential distribution (see, e.g., Allouah and Besbes, 2018). That is, the purpose of this assumption is to rule out fat-tail distributions.

**Proposition 1.** *A risk-averse agent is $e$-resemblant if her marginal value distribution satisfies MHR condition.*

The formal proof of the proposition can be found in Appendix A.1. At a high level, the proof consists of a two-step argument. First, we argue that for every mechanism, its expected revenue is upper bounded by the expected allocated value due to the individual rationality constraint. One of the main technical challenges for analyzing risk-averse agents is that the revenue-optimal mechanism is hard to characterize even for a single agent under specific structures such as CARA. Thanks to the first step, we bypass this challenge since it now suffices to compare the optimal expected allocated value achieved among all mechanisms with the optimal expected revenue among posted pricing mechanisms, in which we prove the gap is at most $e$ under the MHR condition.

### 5.4.2 Budgeted Agent

For a private-budgeted agent, her private type is represented by a pair $t = (v, w)$ where $v$ is the private value for the item and $w$ represents the agent's private budget constraint. In contrast to the previous resemblance analysis for risk-averse agents, where arbitrary correlations between values and risk preferences are allowed, here we make the assumption that values and budgets are independent.

Intuitively, when values and budgets are correlated, especially if values are negatively correlated with budgets, imposing a high payment with a small probability for low-value types would lead to high-value agents with low budgets being unable to select this option due to budget constraints. As a result, the information rent of the agents becomes almost negligible, and the optimal mechanism can extract revenue close to the full surplus in such cases.[21] However, such extreme price discrimination and full surplus extraction due to correlation structures are arguably impractical for real-world applications. Thus, in this section, we focus on cases where the budget distribution is independent of the valuation distribution for private-budget agents.

In addition to the independence assumption, we also impose the standard regularity condition on the marginal value distribution $F$ over values. This condition is strictly weaker than the MHR condition and can even include long-tail distributions such as Pareto distributions with a shape parameter at least 1.

**Definition 5.4.** *A distribution $F$ with density function $f$ is* regular *if the virtual value function $\phi(v) \triangleq v - \frac{1-F(v)}{f(v)}$ is non-decreasing in $v$.*

**Proposition 2.** *A private-budget agent is 3-resemblant if her value and budget are independently distributed, and her marginal value distribution is regular.*

The formal proof of the proposition can be found in Appendix A.4. Moreover, in Appendix A.3, we show that the approximation factor can be improved if the distribution over budgets is a point mass distribution, i.e., the budget of the agent is public. In particular, an agent with a public budget and a regular valuation distribution is 1-resemblant. That is, posting a per-unit price is revenue optimal for such a non-linear utility agent.

## 5.5 Examples for Approximation Guarantees

In this section, we use numerical evaluations to illustrate the approximations of sequential posted pricing mechanisms on more realistic instances. As discussed earlier, this also implies upper bounds on the approximations of marginal payoff maximization mechanisms. We will focus on

[21] We formalize this intuition in Example A.1.

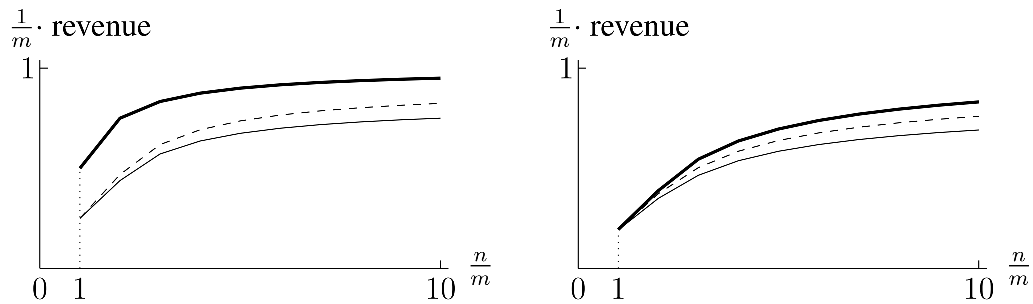


(a) risk-averse agents (b) private-budget agents

Figure 2: The above figure plot the revenue (normalized by the number of items $m$) as a function of $\frac{n}{m}$ for both risk-averse agents and private-budget agents. The expected revenue for sequential posted pricing is computed when $m = 5$ (thin solid line) and $m = 20$ (thin dashed line). An upper bound of the optimal revenue through ex ante relaxation is illustrated by the thick solid line. Note that this normalized upper bound does not depend on the value of $m$.

selling $m$ units of identical items to $n > m$ i.i.d. agents with the objective of revenue maximization. We only provide a simple illustration in this section, and the details for numerical computations are provided in Appendix B.7.

**Risk-averse Utility.** We consider agents with risk-averse utilities. Recall that a risk-averse agent's type is a pair $t = (v, \varphi)$. In particular, we focus on a simple example where each agent has a constant absolute risk aversion (CARA) utility function, i.e.,

$$\varphi(z) = \frac{1}{a}\left(1 - \exp\left(-az\right)\right)$$

for risk parameter $a > 0$, where $\exp(\cdot)$ is the exponential function, and a private value drawn from a uniform distribution on $[0, 1]$.

**Private-budget Utility.** We now consider agents with private-budget utilities where the values are independent of the budgets. This model is a generalization of the single-item setting studied in Che and Gale (2000) for a single agent and in Pai and Vohra (2014) for multiple agents. In both scenarios, the optimal mechanisms are complicated. For numerical analysis, we consider budget constrained agents whose value and budget are drawn independently from uniform distribution on $[0, 1]$.

**Numerical Evaluations.** Note that in general computing the optimal revenue for both risk-averse agents and private-budget agents is challenging. Therefore, in the comparison of revenues in Figure 2, instead of plotting the optimal revenue, we will plot the upper bound on the optimal revenue through the technique of ex ante relaxation. As illustrated in Figure 2, the gap between the expected revenue from sequential posted pricing and the upper bound of the optimal revenue is small, especially when both $n$ and $m$ are large. For example, the multiplicative revenue gap when $n = 200$ and $m = 20$ is less than $1.153$ for risk-averse agents and less than $1.095$ for private-budget agents.

**Discussions of Examples.** In the preceding numerical calculations, we showed that in i.i.d. environments, the revenue of sequential posted pricing mechanisms are approximately optimal for both risk-averse agents and private-budget agents. Since we replaced the optimal revenue with an upper bound derived through the ex ante relaxation, the actual gap between the optimal revenue and sequential posted pricing is even smaller compared to the illustration in Figure 2. The numerical analyses complement our theoretical results by showing that in practical scenarios, the actual performance of these simple mechanisms can significantly outperform the worst-case bounds, especially in large markets where both $n$ and $m$ are large.

# 6 Conclusions and Extensions

This paper develops a general framework for extending mechanism design results from linear to non-linear agents. By introducing a novel notion of random quantiles derived from the agents' demand functions, we provide an implementation of pricing-based mechanisms for non-linear agents in quantile space, which achieves the same payoff as in the linear case with the same aggregate demands. This implementation allows us to establish that the marginal revenue maximization mechanism, which is revenue-optimal for linear agents, remains robustly optimal for non-linear agents when the principal faces uncertainty about utility models and type distributions.

Building on this implementation, we also propose a reduction framework that extends the approximation guarantees of pricing-based mechanisms from linear to non-linear agents. Central to this framework is a novel resemblance property, which quantifies the gap between the optimal mechanism and posted pricing for single-agent problems with non-linear utilities and ex ante supply constraints. As applications of the framework, we derive approximation bounds for various mechanisms under different non-linear utility models (e.g., risk-averse and budgeted utility) in revenue-maximization settings.

We next discuss several important extensions of our framework.

## 6.1 Resemblance for General Payoff Functions

In this section, we show that our framework applies broadly beyond revenue maximization, e.g., welfare maximization or the convex combination of welfare and revenue maximization.

For the objective of welfare maximization, the VCG mechanism achieves the optimal welfare for linear agents. Since the behaviors of risk-averse agents are the same as linear agents when facing per-unit prices (Claim 1), the VCG mechanism also maximizes the welfare for those agents.

However, for general non-linear agents, VCG style mechanisms may not exist and the optimal mechanisms for welfare maximization may be very complex. For example, for agents with public budget constraints, Maskin (2000) shows that the welfare-optimal mechanism takes the form of an all-pay auction for symmetric agents, and Feng and Hartline (2018) further show that this mechanism cannot be implemented as a DSIC mechanism. However, by showing that budgeted agents resemble linear agents for the welfare, there exist simple pricing-based mechanisms, which are DSIC for agents with budget constraints, that are approximately optimal for welfare maximization even when agents have private budget constraints and when agents are asymmetric in ex ante. The proof of the following proposition is provided in Appendix B.1.

**Proposition 3.** *An agent with a private budget has the price-posting welfare curve $P$ that is $2$-resemblant to her optimal welfare curve $R$ if the budget is drawn independently from the valuation.*

For more complex payoff functions such as the convex combination of welfare and revenue maximization, we can extend our previous results to show that if an agent resembles linear agents for both welfare maximization and revenue maximization, this agent resembles linear agents for any convex combination of the two objectives. This observation relies on the following lemma, with proof provided in Appendix B.2.

**Lemma 3.** *If an agent is $\zeta$-resemblant for objective $1$ and $\zeta'$-resemblant for objective $2$ with non-negative values, then this agent is $(\zeta + \zeta')$-resemblant for any convex combination of the two objectives.*

## 6.2 Heterogeneous Utility Models

Our resemblant definitions are monotonic, formalized in the subsequent lemma. With this observation, our framework can be applied to environments with heterogeneous utility functions. For example, suppose some of the agents have private budget constraints and some of the agents are risk averse. If each agent $i \in N$ is $\zeta_i$-resemblant, sequential posted pricing for these agents is a $\frac{e}{e-1} \cdot \max_i\{\zeta_i\}$-approximation to the optimal ex ante relaxation. This observation relies on the following lemma, which can be derived directly based on the definition of resemblance.

**Lemma 4.** *For any $\zeta' \geq \zeta \geq 1$, $\zeta$-resemblant implies $\zeta'$-resemblant.*

## 6.3 Reduction Framework for Non-expected Utilities

In this paper, we have focused on designing simple mechanisms for agents with expected utility representations. For agents with non-expected utilities, given a distribution over per-unit prices, the demand of the agent may be different from her expected demand when she faces the realizations of the per-unit prices in the given distribution. Therefore, our implementation in Section 3 fails to construct an incentive-compatible mechanism for agents with non-expected utilities due to the inherent randomness of the per-unit prices based on the quantiles of other agents.

We show that for a special class of pricing-based mechanisms called *posted pricing mechanisms* (defined in Section 5.3), the payoff guarantees for linear agents can be approximately extended to agents with non-expected utilities. The details of the reduction are analogous to Theorem 3 and hence are deferred to Appendix B.3.

Note that the ordinary goods assumption is not required in the reduction framework for sequential posted pricing mechanisms. Consequently, these mechanisms can be effectively applied for the sale of Giffen goods or Veblen goods, providing approximately optimal revenue guarantees, irrespective of whether the non-linear agents have expected utility representations.

The crucial feature in sequential posted pricing mechanisms that allows for such generalization is that before each agent purchases a lottery from the seller, she is informed about the realizations of the per-unit prices she faces. Therefore, her demand choice in this multi-agent mechanism is consistent with the demand in her single-agent price-posting payoff curve without any assumptions, and thus the resulting mechanism is incentive compatible for any non-expected utilities and guarantees identical payoffs compared to the linear agent analog.

## 6.4 Anonymous Pricing

A desirable property for the multi-agent setting is anonymity. This requires that the per-unit prices posted to all agents are the same. However, the approximation guarantees of anonymous pricing for linear agents do not extend to non-linear agents with general payoff functions such as welfare maximization. The main reason is that given any quantile $q$, the market clearing price $p^q$ usually depends on the details of the utility function as well as the payoff function instead of just depending on the price-posting payoff curve. For example, for the objective of welfare maximization, although anonymous pricing guarantees 2-approximation for linear agents (Lucier, 2017), it may lead to huge welfare loss for non-linear agents. See Example B.2 in Appendix B.4 for an illustration.

In this section, we will focus on revenue maximization, and show that the approximation guarantees of anonymous pricing for linear agents extend to non-linear agents for revenue. The main reason why anonymous pricing extends for revenue maximization is because given any quantile $q$, the market clearing price $p^q$ is uniquely pinned down by the price-posting revenue curve $P(q)$ as $p^q = \frac{P(q)}{q}$. This definition does not depend on the details of the utility function.

In general multi-unit auctions, anonymous pricing is not a constant approximation to the optimal revenue even for linear agents (Jin et al., 2021). In the special case of single-item auctions, for linear agents, Alaei et al. (2019); Jin et al. (2019) showed that the central assumption for constant approximation of anonymous pricing is concavity of the price-posting revenue curves. Next, we provide a general reduction framework for anonymous pricing for non-linear agents in single-agent environments. Note that $\mathrm{AP}(\{P_i\}_{i\in N})$ is the optimal revenue from anonymous pricing when the price-posting revenue curves are $\{P_i\}_{i\in N}$. The proof of the following proposition is provided in Appendix B.4

**Proposition 4.** *Fix any set of (non-linear) agents with price-posting revenue curves $\{P_i\}_{i\in N}$ that are $\zeta$-resemblant to their optimal revenue curves $\{R_i\}_{i\in N}$. If the price-posting revenue curves are concave, then anonymous pricing is a $\zeta e$-approximation to the ex ante relaxation on the optimal revenue curves, i.e.,* $\mathrm{AP}(\{P_i\}_{i\in N}) \geq {}^1\!/\!_{\zeta e} \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$.

As instantiations of the reduction framework in Proposition 4, we can show that agents are $1$-resemblant and have concave price posting revenue curves when they have public budgets and regular valuation distributions, and they are $e$-resemblant when they are risk averse and their valuation distributions satisfy the monotone hazard rate condition.

## 6.5 Dominant Strategy Implementation

All mechanisms implemented in our paper are dominant strategy incentive compatible mechanisms. In contrast to linear agents, where any Bayesian incentive compatible mechanism can be implemented in dominant strategies for single item auctions (Gershkov et al., 2013), it is not without loss to consider dominant strategy incentive compatible mechanisms for non-linear agents (e.g., Feng and Hartline, 2018; Fu et al., 2018). Our results have implications for the line of work focusing on the design of strategically simple mechanisms (e.g., Chung and Ely, 2007; Börgers and Li, 2019). A consequence of our results is that for a broad family of non-linear agents, dominant strategy incentive compatible mechanisms are approximately optimal for any convex combination of welfare and revenue as the objective function.

# A Missing Proofs for Resemblance of Revenue Maximization

## A.1 Risk Aversion

We first show that given per-unit prices, a risk averse agent behave in the same way as a linear agent.

*Proof of Claim 1.* Given any per-unit price $p$, agent's expected utility with demand $z$ is

$$z \cdot \varphi(v - z \cdot p) + (1 - z) \cdot \varphi(-z \cdot p) \leq \varphi(z \cdot (v - p))$$

by Jensen's inequality since function $\varphi$ is concave, and the equality holds when $z \in \{0, 1\}$. Moreover, since $\varphi$ is monotonically increasing and $z \cdot (v - p)$ is maximized when $z$ is either $0$ or $1$, the expected utility is maximized at allocation $z \in \{0, 1\}$. $\square$

Next we prove the bound on resemblance for risk-averse agents.

**Proposition 1.** *A risk-averse agent is $e$-resemblant if her marginal value distribution satisfies MHR condition.*

Our proof relies on the following three auxiliary lemmas.

**Lemma 5.** *For every risk-averse agent and every mechanism that is BIC and IIR, the expected revenue is at most the expected allocated value.*

*Proof.* Let $x$ and $p$ be the (randomized) allocation and payment rule of the BIC, IIR mechanism. For an agent with risk-averse type $t = (v, \varphi)$, her expected utility is $\mathbf{E}[\varphi(vx(v) - p(v))]$. Since the mechanism is individual rational and $\varphi$ is concave,

$$\varphi\left(\mathbf{E}[vx(v) - p(v)]\right) \geq \mathbf{E}[\varphi(vx(v) - p(v))] \geq 0$$

Combining with the assumption that $\varphi$ is increasing with $\varphi(0) = 0$, we obtain $\mathbf{E}[vx(v)] \geq \mathbf{E}[p(v)]$ as desired. $\square$

**Lemma 6.** *For any MHR distribution $F$ and any threshold $\hat{v} \in \mathbb{R}_+$, the conditional distribution $F_{\geq \hat{v}}$ (conditioning on random variable $v \sim F$ is at least $\hat{v}$) is also MHR.*

*Proof.* It suffices the verify the monotonicity of the hazard rate function $h_{\geq \hat{v}}(\cdot)$ for conditional distribution $F_{\geq \hat{v}}$. Let $f$, $f_{\geq \hat{v}}$ be the density function of distribution $F$ and $F_{\geq \hat{v}}$, respectively. Note that

$$h_{\geq \hat{v}}(v) = \frac{f_{\geq \hat{v}}(v)}{1 - F_{\geq \hat{v}}(v)} = \frac{\frac{f(v)}{(1-F(\hat{v}))}}{\frac{1-F(v)}{(1-F(\hat{v}))}} = \frac{f(v)}{1 - F(v)} = h(v)$$

where $h(\cdot)$ is the hazard rate function for distribution $F$, the first and last equality holds by definition, and the second equality holds due to the definition of conditional distribution. Finally, invoking the monotonicity of hazard rate function $h(\cdot)$ finishes the proof. $\square$

**Lemma 7** (Hartline and Roughgarden, 2009)**.** *For any MHR value distribution $F$, it is guaranteed that* $\max_{p\in\mathbb{R}_+} p(1-F(p)) \geq \frac{1}{e}\cdot \mathbf{E}_{v\in F}[v]$.

As a sanity check, Lemma 7 suggests that for a linear agent with value drawn from MHR distribution $F$, the optimal revenue $\max_{p\in\mathbb{R}_+} p(1-F(p))$ is at least an $\frac{1}{e}$-fraction of the expected full surplus $\mathbf{E}_{v\in F}[v]$.

Now we are ready to prove Proposition 1.

*Proof of Proposition 1.* Fix an arbitrary $\hat{q} \in [0,1]$. Let $F$ be the marginal value distribution. To simplify the presentation, we assume $F$ is a continuous distribution. The analysis for general distribution can be extended straightforwardly. With slight abuse of notation, we denote $v(q) \triangleq F^{-1}(1-q)$ for each quantile $q \in [0,1]$. Invoking Lemma 5, the expected revenue of the $\hat{q}$ ex ante revenue-optimal mechanism is at most $\mathbf{E}_{q\sim U[0,1]}[v(q)\cdot \mathbf{1}[q\leq \hat{q}]] = \hat{q}\cdot \mathbf{E}_{q'\sim U[0,\hat{q}]}[v(q')]$. Note that the random variable $v(q')$ where $q' \sim U[0,\hat{q}]$ can be interpreted as the random variable drawn from conditional distribution $F_{\geq v\hat{q}}$. Invoking Lemmas 6 and 7, it ensures that there exists a price $p \geq v(\hat{q})$ such that its expected revenue is at least $\mathbf{E}_{q\sim U[0,1]}[v(q)\cdot \mathbf{1}[q\leq \hat{q}]]$. Since $p \geq v(\hat{q})$, the expected allocation probability is at most $\hat{q}$. Therefore, $\max_{q\leq\hat{q}} P(q) \geq \frac{1}{e}\cdot R(\hat{q})$ as desired. $\square$

## A.2 Necessity of Assumptions

**Example A.1** (Necessity of the independence between the value and budget distributions)**.** *Fix a large constant $h$. Consider a single agent with value $v$ drawn from $[1,h]$ with density function $\frac{h}{h-1}\frac{1}{v^2}$, and budget $w = 2h - v$, i.e., her value and budget are fully correlated. A mechanism which charges the agent $v - 2\epsilon$ with probability $1-\frac{\epsilon}{h}$, or $w$ with probability $\frac{\epsilon}{h}$ for sufficient small positive $\epsilon$ is incentive compatible and has revenue $O(\ln h)$. However, the revenue of the posted pricing is $O(1)$.*

**Example A.2.** *Consider a budgeted agent where the budget distribution is the discrete equal revenue distribution, i.e., $g(i) = {}^{1}/_{\varpi\cdot i^2}$, where $\varpi = {}^{\pi^2}/_{6}$. Let the quantile function of the valuation distribution be $q(i) = {}^{1}/_{\ln i}$. The optimal price posting revenue is a constant. Next consider the pricing function $\tau(x) = \frac{1}{1-x}$. From this pricing function, the value $v_i$ corresponding to payment $i$*

*is $v_i = i^2$. Note that the revenue from this payment function is infinity, i.e.,*

$$\begin{aligned}\textbf{Payoff}[\tau] &\geq \lim_{m\to\infty}\sum_{i=1}^{m}(i\cdot q(v_i)\cdot g(i))\\ &= \frac{1}{2\varpi}\lim_{m\to\infty}\sum_{i=1}^{m}\frac{1}{i\cdot\ln i}\\ &= \frac{1}{2\varpi}\lim_{m\to\infty}\ln\ln m\to\infty.\end{aligned}$$

*Therefore, the gap between price posting and the optimal mechanism is infinite.*

## A.3 Public Budget

We first consider a simpler setting in which agents have deterministic public budgets. This situation can be seen as a special case of private-budget agents, where we consider a point-mass budget distribution. As a sanity check, the independence assumption holds for public-budget agents.

**Proposition 5.** *A public-budget agent is $1$-resemblant if her marginal value distribution is regular.*

The proof of Proposition 5 relies on the technique of Lagrangian relaxation on the budget constraint developed in Alaei et al. (2013) and Feng and Hartline (2018). Combined with our reduction framework in Theorem 3, the interpretation of Proposition 5 is that for public-budget agents with regular marginal value distributions, the worst-case approximation ratio of pricing-based mechanisms to the ex ante relaxation occurs when the public budgets do not bind. Moreover, in Proposition 6, we also show that the regularity assumption is not essential for approximation. That is, without the assumption of regular valuation distribution, a public-budget agent is still $2$-resemblant.

*Proof of Proposition 5.* For an agent with public budget $w$, the $\hat{q}$ ex ante optimal mechanism is the solution of the following program,

$$\begin{aligned}\max_{(x,p)}\quad & \mathbf{E}_v[p(v)]\\ \text{s.t.}\quad & (x,p)\text{ are IC, IR},\\ & \mathbf{E}_v[x(v)]=\hat{q},\\ & p(\bar{v})\leq w.\end{aligned}\tag{2}$$

where $\bar{v}$ is the highest possible value of the agent. Consider the Lagrangian relaxation of the budget constraint in (2),

$$\begin{aligned} \min_{\lambda \geq 0} \max_{(x,p)} \quad & \mathbf{E}_v[p(v)] + \lambda w - \lambda p(\bar{v}) \\ \text{s.t.} \quad & (x,p) \text{ are IC, IR,} \\ & \mathbf{E}_v[x(v)] = \hat{q}. \end{aligned} \tag{3}$$

Let $\lambda^*$ be the optimal solution in program (3). If we fix $\lambda = \lambda^*$ in program (3), its inner maximization program can be thought as a $\hat{q}$ ex ante optimal mechanism design for a linear agent with Lagrangian objective function $\mathbf{E}_v[p(v)] - \lambda^* p(\bar{v})$. Thus, we define the Lagrangian price-posting revenue curve $P_{\lambda^*}(\cdot)$ where $P_{\lambda^*}(q)$ is the maximum value of the Lagrangian objective $\mathbf{E}_v[p(v)] - \lambda^* p(\bar{v})$ in price-posting mechanism with per-unit price $V(q)$. For any $q \in (0, 1]$, by the definition, $P_{\lambda^*}(q) = qV(q) - \lambda^* V(q)$. For $q = 0$, notice that the agent with $\bar{v}$ is indifferent between purchasing or not purchasing. Thus, by the definition, $P_{\lambda^*}(q) = 0$ if $q = 0$.

Now, we consider the concave hull of the Lagrangian price-posting revenue curve $P_{\lambda^*}(\cdot)$ which we denote as $\hat{P}_{\lambda^*}(\cdot)$. Let $q^\dagger$ be the smallest solution of equation $P_{\lambda^*}(q) = qP'_{\lambda^*}(q)$. Since $P_{\lambda^*}(0) \leq 0$, $P_{\lambda^*}(1) = 0$ and $P_{\lambda^*}(\cdot)$ is continuous, $q^\dagger$ always exists. Then, for any $q \leq q^\dagger$, $\hat{P}_{\lambda^*}(q) = q\, P'_{\lambda^*}(q^\dagger)$. For any $q \geq q^\dagger$, we show $\hat{P}_{\lambda^*}(q) = P_{\lambda^*}(q)$ by the following arguments. First notice that $P_{\lambda^*}(q^\dagger) \geq 0$, and hence $q^\dagger \geq \lambda^*$. Consider $P''_{\lambda^*}(q) = V''(q)(q - \lambda^*) + 2V'(q)$. Clearly, $V'(q) \leq 0$. If $V''(q) \leq 0$, then $P''_{\lambda^*}(q) \leq 0$. If $V''(q) > 0$, then $P''_{\lambda^*}(q) = V''(q)(q - \lambda^*) + 2V'(q) \leq qV''(q) + 2V'(q) \leq 0$, where $qV''(q) + 2V'(q)$ is non-positive due to the regularity of the valuation distribution.

To summarize, $\hat{P}_{\lambda^*}(\cdot)$, the concave hull of the Lagrangian price-posting revenue curve satisfies

$$\hat{P}_{\lambda^*}(q) = \begin{cases} q\, P'_{\lambda^*}(q^\dagger) & \text{if } q \in [0, q^\dagger] \\ P_{\lambda^*}(q) & \text{if } q \in [q^\dagger, 1] \end{cases}$$

Therefore, use the similar ironing technique based on the revenue curves for linear agents with irregular valuation distribution (e.g. Myerson, 1981; Bulow and Roberts, 1989; Alaei et al., 2013), Lemma 8 (stated below) suggests that the $\hat{q}$ ex ante optimal mechanism irons quantiles between $[0, q^\dagger]$ under $\hat{q}$ ex ante constraint, which is still a posted-pricing mechanism. □

**Lemma 8** (Alaei et al., 2013)**.** *For incentive compatible and individual rational mechanism* $(x(\cdot), p(\cdot))$ *and an agent with any Lagrangian price-posting revenue curve* $P_{\lambda^*}(q)$*, the expected Lagrangian objective of the agent is upper-bounded by her expected marginal Lagrangian objective of the same allocation rule, i.e.,*

$$\mathbf{E}_v[p(v)] + \lambda^* p(\bar{v}) \leq \mathbf{E}_q\left[\hat{P}'_{\lambda^*}(q) \cdot x(V(q))\right].$$

*Furthermore, this inequality holds with equality if the allocation rule* $x(\cdot)$ *is constant all intervals*

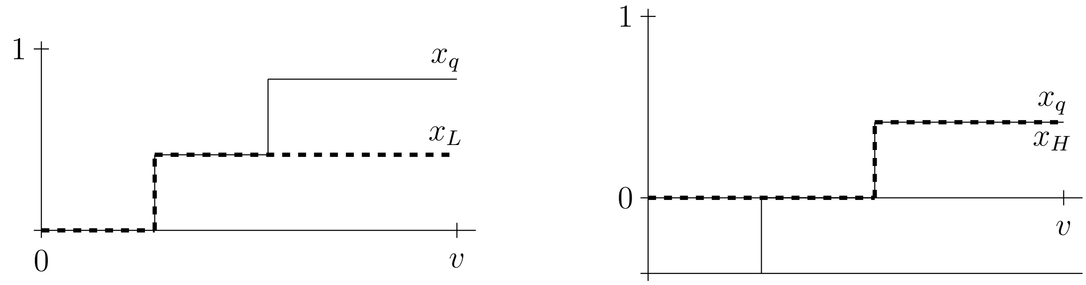


Figure 3: The thin solid line is the allocation rule for the optimal ex ante mechanism. The thick dashed line on the left side is the allocation of the decomposed mechanism with lower price, while the thick dashed line on the right side is the allocation of the decomposed mechanism with higher price.

*of values $V(q)$ where $\hat{P}_{\lambda^*}(q) > P_{\lambda^*}(q)$.*

For an agent with a general valuation distribution, resemblance follows from the characterization that the ex ante optimal mechanism has menu size at most 2 (Alaei et al., 2013). Thus offering the optimal single menu pricing mechanism, i.e., posting the market clearing price, is a 2-approximation.

**Lemma 9** (Alaei et al., 2013)**.** *For a single agent with public budget, the $q \in [0, 1]$ ex ante optimal mechanism has a menu with size at most two.*

**Proposition 6.** *An agent with public budget has the ironed price-posting revenue curve $\bar{P}$ that is $2$-resemblant to her optimal revenue curve $R$.*

*Proof.* By Lemma 9, the allocation rule $x_q$ of the ex ante revenue maximization mechanism for the single agent with public budget has a menu of size at most two. We decompose its allocation into $x_L$ and $x_H$ as illustrated in Figure 3. Note that both allocation $x_L$ and $x_H$ are (randomized) price-posting allocation rules, and neither allocation violates the allocation constraint $q$. Thus,

$$R(q) = \textbf{Payoff}[x_q] = \textbf{Payoff}[x_L] + \textbf{Payoff}[x_H] \leq 2 \max_{q^\dagger \leq q} \bar{P}(q^\dagger). \qquad \square$$

## A.4 Private Budget

**Proposition 2.** *A private-budget agent is 3-resemblant if her value and budget are independently distributed, and her marginal value distribution is regular.*

Fix an arbitrary ex ante constraint $q$, denote EX as the $q$ ex ante revenue-optimal mechanism, and **Payoff**[EX] as its revenue. We decompose EX into two mechanisms $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$ according to the market clearing price $p^q$. Intuitively, the per-unit prices in $\mathrm{EX}^\dagger$ for all types are at most the market clearing price and the per-unit prices in $\mathrm{EX}^\ddagger$ for all types are larger than the market clearing price. The details of the decomposition are specified in Appendix B.1, and we will bound the revenue from those two mechanisms separately.

**Lemma 10.** *For a single agent with random-public-budget utility, independently distributed value and budget, and any ex ante constraint $q$; the revenue of $\mathrm{EX}^\dagger$ is at most the revenue from posting the market clearing price, i.e., $P(q) \geq$ **Payoff**$\left[\mathrm{EX}^\dagger\right]$.*

*Proof.* The ex ante allocation of $\mathrm{EX}^\dagger$ is at most the ex ante allocation of $\mathrm{EX}$, i.e., $q$. Combining with the fact that the per-unit prices in $\mathrm{EX}^\dagger$ for all types are weakly lower than the market clearing price, its revenue is at most the revenue of posting the market clearing price. □

For the revenue bound of $\mathrm{EX}^\ddagger$, we consider two different cases: (1) the market clearing price is larger than the monopoly reserve; and (2) the market clearing price is smaller than the monopoly reserve. The proofs of Lemma 11 and 12 are provided at the end of this section.

**Lemma 11.** *For a single private-budget agent with independently distributed value and budget and regular value distribution, if the market clearing price $p^q = {}^{P(q)}/_q$ is larger than the monopoly reserve, i.e., $p^q = P(q)/q \geq m^*$, the revenue of posting the market clearing price is at least the revenue of $\mathrm{EX}^\ddagger$, i.e., $P(q) \geq$ **Payoff**$\left[\mathrm{EX}^\ddagger\right]$.*

Intuitively, Lemma 11 is analyzed separately for each realized budget under two cases. If the budget is lower than the market clearing price, then the budget binds and the revenues are the same in both cases. If the budget is higher than the market clearing price, then the budget does not bind, and the revenue comparison is analogous to the setting without budget constraints. Since the revenue curve is concave and the market clearing price is larger than the monopoly reserve, higher price in $\mathrm{EX}^\ddagger$ leads to lower expected revenue.

**Lemma 12.** *For a single private-budget agent with independently distributed value and budget and regular value distribution, if the market clearing price $p^q = {}^{P(q)}/_q$ is smaller than the monopoly reserve, there exists $q^\dagger \leq q$ such that the market clearing revenue from $q^\dagger$ is a 2-approximation to the revenue from $\mathrm{EX}^\ddagger$, i.e., $2P(q^\dagger) \geq$ **Payoff**$\left[\mathrm{EX}^\ddagger\right]$.*

In fact, we show there is a distribution over interim feasible random pricing that attains at least half of the optimal ex ante revenue. Of course, the optimal deterministic price that is at least $p^q$ is only better than the random price, and hence the lemma is shown. Specifically, consider posting a random price $\boldsymbol{p} = \max\{p^q, \boldsymbol{p}_0\}$ with $\boldsymbol{p}_0$ drawn identically to the agents value distribution. As shown in Figure 4, the revenue from the random pricing $\boldsymbol{p}$ corresponds to the dark grey area, while the optimal ex ante revenue corresponds to the light grey area. By the concavity of the revenue curve, the geometry of the figure directly implies the 2-approximation on expected revenue.

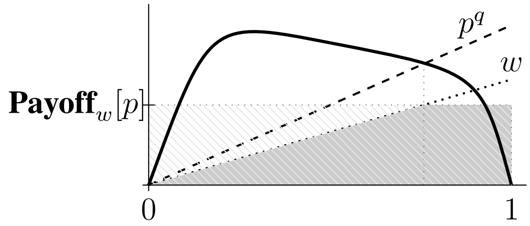

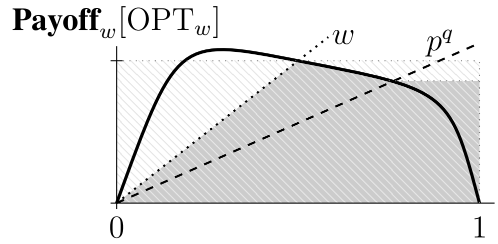


Figure 4: In the geometric proof of Lemma 12, the upper bound on the expected revenue of $\mathrm{EX}^{\ddagger}$ ($\mathbf{Payoff}_w[p]$ and $\mathbf{Payoff}_w[\mathrm{OPT}_w]$ on the left and right, respectively) is the area of the light gray striped rectangle and the revenue from posting random price $\boldsymbol{p}$ is the area of the dark gray region. By geometry, the latter is at least half of the former. The black curve is the price-posting revenue curve with no budget constraint $P^L$. The figure on the left depicts the small-budget case (i.e., $w < p^q$), and the figure on the right depicts the large-budget case (i.e., $w \geq p^q$).

*Proof of Proposition 2.* Fix any ex ante constraint $q$. If the market clearing price $p^q = {}^{P(q)}/_q$ is at least the monopoly reserve, Lemma 10 and Lemma 11 imply that $\mathbf{Payoff}\big[\mathrm{EX}^{\dagger}\big] \leq P(q)$, and $\mathbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big] \leq P(q)$, thus, $P(q)$ is a 2-approximation to $\mathbf{Payoff}\big[\mathrm{EX}^{\dagger}\big] + \mathbf{Payoff}\big[\mathrm{EX}^{\dagger}\big] = \mathbf{Payoff}[\mathrm{EX}]$, i.e., $R(q)$. If the market clearing price $p^q$ is smaller than the monopoly reserve, let $q^{\dagger} = \mathrm{argmax}_{q' \leq q} P(q')$, Lemma 10 and Lemma 12 imply that $\mathbf{Payoff}\big[\mathrm{EX}^{\dagger}\big] \leq P(q) \leq P(q^{\dagger})$, and $\mathbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big] \leq 2P(q^{\dagger})$, thus, $P(q^{\dagger})$ is a 3-approximation to $R(q)$. Thus, the agent is 3-resemblant for ex ante optimization. □

*Proof of Lemma 11.* In both $\mathrm{EX}^{\ddagger}$ and the mechanism that posts the market clearing price, the types with value lower than the market clearing price $p^q$ will purchase nothing, so we only consider the types with value at least $p^q$ in this proof. Each budget level is considered separately.

For types with budget $w \leq p^q$, by posting the market clearing price $p^q$, those types always pay their budgets $w$, which is at least the revenue from those types in $\mathrm{EX}^{\ddagger}$.

For types with budget $w > p^q$, by posting the market clearing price $p^q$, those types always pay $p^q$. Since the budget constraints do not bind for these types, it is helpful to consider the price-posting revenue curve without budget, which we denote by $P^L$. The regularity of the valuation distribution guarantees that $P^L$ is concave. The concavity of $P^L$ implies that higher prices above $m^*$ extracts lower revenue than $p^q$. Since the per-unit prices in $\mathrm{EX}^{\ddagger}$ for all types are at least $p^q$, the concavity of $P^L$ guarantees that the expected revenue of posting $p^q$ for types with budget larger than $p$ is at least the expected revenue for those types in $\mathrm{EX}^{\ddagger}$. Combining these bounds above, we have $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big]$. □

*Proof of Lemma 12.* Note that any price that is at least $p^q$ is feasible for the ex ante constraint $q$. We consider posting a random price $\boldsymbol{p} = \max\{p^q, \boldsymbol{p}_0\}$ with $\boldsymbol{p}_0$ drawn identically to the agents value distribution. Fixing the budget of the agent $w$, consider the following geometric argument (cf. Dhangwatnotai et al., 2015). For both sides of Figure 4, the area of the light gray stripped rectangle upper bounds the revenue of $\mathrm{EX}^{\ddagger}$ and the area of the dark gray region is the expected revenue from posting random price $\boldsymbol{p}$. Consequently, concavity of the price-posting revenue curve with no budget constraint $P^L$ (by regularity of the value distribution) implies that a triangle with half the area of the light gray rectangle is contained within the dark gray region and, thus, the random price is a 2-approximation. As the random price does not depend on the budget $w$, the same bound holds when $w$ is random. Of course, the optimal deterministic price that is at least $p^q$ is only better than the random price and the lemma is shown. The remainder of this proof verifies that the geometry of the regions described above is correct.

The left side of Figure 4 depicts the fixed budgets $w$ that are at most $p^q$. The area of the light gray striped rectangle upper bounds the revenue of $\mathrm{EX}^{\ddagger}$ as follows. Let $\mathbf{Payoff}_w[p]$ be the expected revenue from posting price $p$ to types with budget $w$. Under both $\mathrm{EX}^{\ddagger}$ and the market clearing price $p^q$, types with value below the market clearing price pay zero. For the remaining types, in $\mathrm{EX}^{\ddagger}$ they pay at most their budget and in market clearing they pay exactly their budget. Thus, $\mathbf{Payoff}_w\left[\mathrm{EX}^{\ddagger}\right] \leq \mathbf{Payoff}_w[p^q] = w\,(1 - F(p^q))$ where, recall, $1 - F(p^q)$ is the probability the agent's value is at least the market clearing price $p^q$. Of course, $w\,(1 - F(p^q))$ is the height and area (its width is 1) of the light gray striped region on the left side of Figure 4.

The right side of Figure 4 depicts the fixed budgets $w$ that are at least $p^q$. The area of the light gray striped rectangle upper bounds the revenue of $\mathrm{EX}^{\ddagger}$ as follows. Let $\mathrm{OPT}_w$ be the optimal mechanism to types with budget $w$ without ex ante constraint and $\mathbf{Payoff}_w[\mathrm{OPT}_w]$ be its expected revenue from these types. Clearly, $\mathbf{Payoff}_w\left[\mathrm{EX}^{\ddagger}\right] \leq \mathbf{Payoff}_w[\mathrm{OPT}_w]$ as the latter optimizes with relaxed constraints of the former. Laffont and Robert (1996) show that $\mathrm{OPT}_w$ posts the minimum between budget $w$ and the monopoly reserve $m^*$ when the agent has public budget and regular valuation. As the budget does not bind for this price, its revenue is given by the price-posting revenue curve with no budget constraint, i.e., $\mathbf{Payoff}_w[\mathrm{OPT}_w] = P^L(1 - F(\min\{w, m^*\}))$. Of course, this revenue is the height and area (its width is 1) of the light gray striped region on the right side of Figure 4.

Next, we will show that the revenue of posting the random price $\boldsymbol{p}$ is the grey shaded areas illustrated in Figure 4 (in both cases). A random price from the value distribution, i.e., $\boldsymbol{p}_0$, corresponds to a uniform random quantile constraint, i.e., drawing uniformly from the horizontal axis. Since we truncate the lower end of the price distribution at the market clearing price $p^q$, the revenue from quantiles greater than $q$ equals the revenue from the market clearing price. For any fixed $w$, when $\boldsymbol{p} \in [p^q, w]$, the budget does not bind and the revenue of posting price $\boldsymbol{p}$ is $P^L(\boldsymbol{q})$ where $P^L$ is the price-posting revenue curve without budget; and when $\boldsymbol{p} > w$, the revenue of posting price $\boldsymbol{p}$ is $w\boldsymbol{q}$. Thus, the revenue from a random price is given by the integral of the area under the curve defined by $\boldsymbol{q}w$ when $\boldsymbol{p} \geq w$, by $P^L(\boldsymbol{q})$ when $\boldsymbol{p} \in [w, p^q]$ and this interval exists, and by $\min(w, p^q)$

when $\boldsymbol{p} = p^q$, i.e., when $\boldsymbol{p}_0 \leq p^q$. This area is the dark gray region. $\square$

# References


Abrams, Z. (2006). Revenue maximization when bidders have budgets. In *Proceedings of the 17th annual ACM-SIAM Symposium on Discrete Algorithm*, pages 1074–1082.

Akbarpour, M., Kominers, S. D., Li, K. M., Li, S., and Milgrom, P. (2023). Algorithmic mechanism design with investment. *to appear at Econometrica*.

Alaei, S., Fu, H., Haghpanah, N., and Hartline, J. (2013). The simple economics of approximately optimal auctions. In *Proc. 54th IEEE Symp. on Foundations of Computer Science*, pages 628–637. IEEE.

Alaei, S., Fu, H., Haghpanah, N., Hartline, J., and Malekian, A. (2012). Bayesian optimal auctions via multi-to single-agent reduction. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, page 17.

Alaei, S., Hartline, J., Niazadeh, R., Pountourakis, E., and Yuan, Y. (2019). Optimal auctions vs. anonymous pricing. *Games and Economic Behavior*, 118:494–510.

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.

Baisa, B. (2017). Auction design without quasilinear preferences. *Theoretical Economics*, 12(1):53–78.

Bajari, P., Nekipelov, D., Ryan, S. P., and Yang, M. (2015). Machine learning methods for demand estimation. *American Economic Review*, 105(5):481–485.

Balkanski, E. and Hartline, J. D. (2016). Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web*, pages 189–203.

Börgers, T. and Li, J. (2019). Strategically simple mechanisms. *Econometrica*, 87(6):2003–2035.

Boulatov, A. and Severinov, S. (2021). Optimal and efficient mechanisms with asymmetrically budget constrained buyers. *Games and Economic Behavior*, 127:155–178.

Brooks, B. and Du, S. (2021). Optimal auction design with common values: An informationally robust approach. *Econometrica*, 89(3):1313–1360.

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

Carroll, G. (2015). Robustness and linear contracts. *American Economic Review*, 105(2):536–63.

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

Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. (2010). Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM.

Chawla, S., Malec, D. L., and Malekian, A. (2011). Bayesian mechanism design for budget-constrained agents. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 253–262. ACM.

Che, Y.-K. and Gale, I. (2000). The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic theory*, 92(2):198–233.

Che, Y.-K. and Gale, I. (2006). Revenue comparisons for auctions when bidders have arbitrary types. *Theoretical Economics*, 1(1):95–118.

Chen, Y.-C., He, W., Li, J., and Sun, Y. (2019). Equivalence of stochastic and deterministic mechanisms. *Econometrica*, 87(4):1367–1390.

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

Devanur, N. R. and Weinberg, S. M. (2017). The optimal mechanism for selling to a budget constrained buyer: The general case. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 39–40. ACM.

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

Fan, K. and Lorentz, G. G. (1954). An integral inequality. *The American Mathematical Monthly*, 61(9):626–631.

Feldman, M., Svensson, O., and Zenklusen, R. (2016). Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. SIAM.

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.

Fu, H., Liaw, C., Lu, P., and Tang, Z. G. (2018). The value of information concealment. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2533–2544. SIAM.

Gershkov, A., Goeree, J. K., Kushnir, A., Moldovanu, B., and Shi, X. (2013). On the equivalence of bayesian and dominant strategy implementation. *Econometrica*, 81(1):197–220.

Gershkov, A., Moldovanu, B., Strack, P., and Zhang, M. (2021). A theory of auctions with endogenous valuations. *Journal of Political Economy*, 129(4):1011–1051.

Guo, Y. and Shmaya, E. (2023). Regret-minimizing project choice. *Econometrica*, 91(5):1567–1593.

Hartline, J., Johnsen, A., and Li, Y. (2020). Benchmark design and prior-independent optimization. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 294–305. IEEE.

Hartline, J. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proc. 10th ACM Conf. on Electronic Commerce*.

Hartline, J. D. (2012). Approximation in mechanism design. *American Economic Review*, 102(3):330–36.

Hartline, J. D. and Lucier, B. (2015). Non-optimal mechanism design. *American Economic Review*, 105(10):3102–24.

Holt Jr, C. A. (1980). Competitive bidding for contracts under alternative auction procedures. *Journal of Political Economy*, 88(3):433–445.

Jin, Y., Jiang, S., Lu, P., and Zhang, H. (2021). Tight revenue gaps among multi-unit mechanisms. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 654–673.

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*, pages 674–685.

Kazumura, T., Mishra, D., and Serizawa, S. (2020). Mechanism design without quasilinearity. *Theoretical Economics*, 15(2):511–544.

Kushnir, A. and Liu, S. (2019). On the equivalence of bayesian and dominant strategy implementation for environments with nonlinear utilities. *Economic Theory*, 67:617–644.

Laffont, J.-J. and Robert, J. (1996). Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186.

Lucier, B. (2017). An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47.

Maskin, E. (2000). Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, 44(4–6):667–681.

Maskin, E. and Riley, J. (1984). Optimal auctions with risk averse buyers. *Econometrica*, pages 1473–1518.

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

Morgenstern, O. and von Neumann, J. (1953). *Theory of games and economic behavior*. Princeton university press.

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

Pai, M. M. and Vohra, R. (2014). Optimal auctions with financially constrained buyers. *Journal of Economic Theory*, 150:383–425.

Richter, M. (2019). Mechanism design with budget constraints and a population of agents. *Games and Economic Behavior*, 115:30–47.

Riley, J. and Zeckhauser, R. (1983). Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, pages 267–289.

Yan, Q. (2011). Mechanism design via correlation gap. In *Proc. 22nd ACM Symp. on Discrete Algorithms*, pages 710–719. SIAM.

ONLINE APPENDIX FOR ONLINE PUBLICATION

# B Details for Various Extensions

## B.1 Welfare Maximization

**Proposition 3.** *An agent with a private budget has the price-posting welfare curve $P$ that is 2-resemblant to her optimal welfare curve $R$ if the budget is drawn independently from the valuation.*

The proof of Proposition 3 utilizes the price decomposition technique from Abrams (2006) and extends it for welfare analysis.

Fix an arbitrary ex ante constraint $q$, denote $\mathrm{EX}$ as the $q$ ex ante welfare-optimal mechanism, and **Payoff**$[\mathrm{EX}]$ as its welfare. We want to decompose $\mathrm{EX}$ into two mechanisms $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$ according to the market clearing price $p^q$ and bound the welfare from those two mechanisms separately. The decomposed mechanism may violate the incentive constraint for budgets, and we refer to this setting as the random-public-budget utility model. Note that the market clearing price is the same in both the private budget model and the random-public-budget utility model. Intuitively, mechanism $\mathrm{EX}^\dagger$ contains per-unit prices at most the market clearing price, while mechanism $\mathrm{EX}^\ddagger$ contains per-unit prices at least the market clearing price. Both mechanisms $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$ satisfy the ex ante constraint $q$, and the sum of their welfare upper bounds the original ex ante mechanism $\mathrm{EX}$, i.e., **Payoff**$[\mathrm{EX}] \leq$ **Payoff**$\left[\mathrm{EX}^\dagger\right] +$ **Payoff**$\left[\mathrm{EX}^\ddagger\right]$.

To construct $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$ that satisfy the properties above, we first introduce a characterization of all incentive compatible mechanisms for a single agent with private-budget utility, and her behavior in the mechanisms.

**Definition B.1.** *An* allocation-payment function $\tau : [0, 1] \to \mathbb{R}_+$ *is a mapping from the allocation $x$ to the payment $p$.*

**Lemma 13.** *For a single agent with private-budget utility, in any incentive compatible mechanism, for all types with any fixed budget, the mechanism provides a convex and non-decreasing allocation-payment function, and subject to this allocation-payment function, each type will purchase as much as she wants until the budget constraint binds, or the unit-demand constraint binds, or the value binds (i.e., her marginal utility becomes zero).*

*Proof.* Myerson (1981) show that any mechanisms $(x, p)$ for a single linear agent is incentive compatible (the agent does not prefer to misreport her value) if and only if a) $x(v)$ is non-decreasing; b) $p(v) = vx(v) - \int_0^v x(t)dt$. Thus, given any non-decreasing allocation $x$, the payment $p$ is uniquely pined down by the incentive constraints.

Comparing with the linear utility, the incentive compatibility in the private-budget utility guarantees that the agent does not prefer to misreport either her value or budget. If we relax the incentive constraints such that she is only allowed to misreport her value, Myerson result already shows that for any fixed budget level $w$, the allocation $x(v, w)$ is non-decreasing in $v$ and the payment $p(v, w) = vx(v, w) - \int_0^v x(t, w)dt$ is uniquely pined down. We define the allocation-payment function $\tau_w(\hat{x}) = \max\{p(v, w) + v \cdot (\hat{x} - x(v, w)) : x(v, w) \leq \hat{x}\}$ if $\hat{x} \leq x(\bar{v}, w)$; and $\infty$ otherwise. Given the characterization of allocation and payment above, this allocation-payment function is well-defined, non-decreasing and convex. □

This characterization is shown by relaxing the agent's incentives for misreporting the budget. Unlike Myerson's result which give a sufficient and necessary condition for incentive compatible mechanisms for linear agents, Lemma 13 only characterizes a necessary condition for private-budget utility.[22] This condition is already enough for our arguments.

Now we give the construction of $\text{EX}^\dagger$ and $\text{EX}^\ddagger$ by constructing their allocation-payment functions. The decomposition is illustrated in Figure 5. For agent with budget $w$, let $\tau_w$ be the allocation-payment function in mechanism EX, and $x_w^*$ be the utility maximization allocation for a linear agent with value equal to the market clearing price $p^q$, i.e., $x_w^* = \text{argmax}\{x : \tau'_w(x) \leq p^q\}$. For agents with budget $w$, we define the allocation-payment functions $\tau_w^\dagger$ and $\tau_w^\ddagger$ for $\text{EX}^\dagger$ and $\text{EX}^\ddagger$ respectively below,

$$\tau_w^\dagger(x) = \begin{cases} \tau_w(x) & \text{if } x \leq x_w^*, \\ \infty & \text{otherwise;} \end{cases} \quad \tau_w^\ddagger(x) = \begin{cases} \tau_w(x_w^* + x) - \tau_w(x_w^*) & \text{if } x \leq 1 - x_w^*, \\ \infty & \text{otherwise.} \end{cases}$$

By construction, for each type of the agent, the allocation from EX is upper bounded by the sum of the allocation from $\text{EX}^\dagger$ and $\text{EX}^\ddagger$, which implies that the welfare from EX is upper bounded by the sum of the welfare from $\text{EX}^\dagger$ and $\text{EX}^\ddagger$, and the requirements for the decomposition are satisfied.

As sketched above, we separately bound the welfare in $\text{EX}^\dagger$ and $\text{EX}^\ddagger$ by the welfare from posting the market clearing price.

**Lemma 14.** *For a single agent with random-public-budget utility, independently distributed value and budget, and any ex ante constraint $q$, the welfare from posting the market clearing price $p^q$ is at least the welfare from $\text{EX}^\dagger$, i.e., $P(q) \geq$* **Payoff**$\left[\text{EX}^\dagger\right]$.

*Proof.* Consider agent with type $(v, w)$ and agent with type $(v', w)$, where both value $v$ and $v'$ are higher than the market clearing price $p^q$. Notice that the allocations for these two types are the same in $\text{EX}^\dagger$ and in market clearing, since the per-unit price in both mechanisms is at most $p^q$ which makes the mechanisms unable to distinguish these two types.

[22]This characterization is only necessary because it relaxes the incentive constraints for misreporting the private budget.

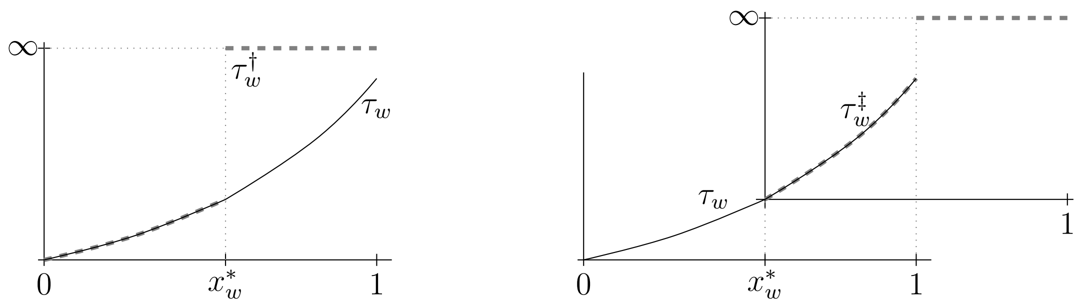


Figure 5: Depicted are allocation-payment function decomposition. The black lines in both figures are the allocation-payment function $\tau_w$ in ex ante optimal mechanism EX; the gray dashed lines are the allocation-payment function $\tau_w^\dagger$ and $\tau_w^\ddagger$ in $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$, respectively.

Let $x^\dagger$ be the allocation rule in $\mathrm{EX}^\dagger$ and let $x^q$ be the allocation rule in posting the market clearing price $p^q$. For any value $v \geq p^q$, the expected allocation for types with value $v$ is lower in $\mathrm{EX}^\dagger$ than in market clearing, i.e., $\mathbf{E}_w\left[x^\dagger(v,w)\right] \leq \mathbf{E}_w[x^q(v,w)]$. Otherwise suppose the types with value $v^*$ has strictly higher allocation in $\mathrm{EX}^\dagger$ for some value $v^* \geq p^q$, i.e, $\mathbf{E}_w\left[x^\dagger(v^*,w)\right] > \mathbf{E}_w[x^q(v^*,w)]$. By the fact stated in previous paragraph, we have that for any budget $w$ and any value $v, v^* \geq p^q$, $x^q(v,w) = x^q(v^*,w)$, $x^\dagger(v,w) = x^\dagger(v^*,w)$, and the expected allocation in $\mathrm{EX}^\dagger$ is

$$\begin{aligned}\mathbf{E}_{v,w}\left[x^\dagger(v,w)\right] &\geq \Pr[v \geq p^q] \cdot \mathbf{E}_{v,w}\left[x^\dagger(v,w) \mid v \geq p^q\right]\\ &= \Pr[v \geq p^q] \cdot \mathbf{E}_w\left[x^\dagger(v^*,w)\right]\\ &> \Pr[v \geq p^q] \cdot \mathbf{E}_w[x^q(v^*,w)]\\ &= \Pr[v \geq p^q] \cdot \mathbf{E}_{v,w}[x^q(v,w) \mid v \geq p^q] = q,\end{aligned}$$

where the qualities hold due to the independence between the value and the budget. Note that this implies that $\mathrm{EX}^\dagger$ violates the ex ante constraint $q$, a contradiction. Further, for any type with value $v \geq p^q$, $\mathbf{E}_w\left[x^\dagger(v,w)\right] \leq \mathbf{E}_w[x^q(v,w)]$ implies that the allocation in market clearing "first order stochastic dominates" the allocation in $\mathrm{EX}^\dagger$, i.e., for any threshold $v^\dagger$, the expected allocation from all types with value $v \geq v^\dagger$ in market clearing is at least the expected allocation from those types in $\mathrm{EX}^\dagger$. Taking expectation over the valuation and the budget, the expected welfare from market clearing is at least the welfare from $\mathrm{EX}^\dagger$, i.e., $P(q) \geq \textbf{Payoff}\left[\mathrm{EX}^\dagger\right]$. $\square$

**Lemma 15.** *For a single agent with random-public-budget utility, independently distributed value and budget, and any ex ante constraint $q$; the welfare from market clearing is at least the welfare from $\mathrm{EX}^\ddagger$, i.e., $P(q) \geq \textbf{Payoff}\left[\mathrm{EX}^\ddagger\right]$.*

*Proof.* In both $\mathrm{EX}^\ddagger$ and market clearing, types with value lower than $p^q$ will purchase nothing, so we only consider the types with value at least $p^q$ in this proof. Consider any type $(v, w)$ where $v \geq p^q$, its allocation in market clearing is at least its allocation in $\mathrm{EX}^\ddagger$, because the per-unit price in $\mathrm{EX}^\ddagger$ is higher. Thus, the welfare from market clearing is at least the welfare from $\mathrm{EX}^\ddagger$, i.e., $P(q) \geq \textbf{Payoff}\left[\mathrm{EX}^\ddagger\right]$. $\square$

Intuitively, in $\mathrm{EX}^\dagger$, we are able to show that for any value higher than $p^q$, the expected allocation of the agent in $\mathrm{EX}^\dagger$ is weakly lower compared to market clearing in order for the ex ante feasibility constraint to be satisfied. Thus allocation given the market clearing price "first order stochastic dominates" the allocation in $\mathrm{EX}^\dagger$, and hence expected welfare in $\mathrm{EX}^\dagger$ is lower compared to posting market clearing price $p^q$.

In $\mathrm{EX}^\ddagger$, types with value lower than $p^q$ will purchase nothing, and types with value higher than $p^q$ will less compared to market clearing price as the per-unit price in $\mathrm{EX}^\ddagger$ is higher. Thus the expected welfare in $\mathrm{EX}^\ddagger$ is also lower compared to posting market clearing price $p^q$.

*Proof of Proposition 3.* Combining Lemma 14 and 15, for any quantile $q$, we have

$$R(q) = \textbf{Payoff}[\mathrm{EX}] \leq \textbf{Payoff}\left[\mathrm{EX}^\dagger\right] + \textbf{Payoff}\left[\mathrm{EX}^\ddagger\right] \leq 2P(q) \leq \max_{q' \leq q} 2\bar{P}(q'). \qquad \square$$

## B.2 Convex Combination of Two Payoff Functions

**Lemma 3.** *If an agent is $\zeta$-resemblant for objective 1 and $\zeta'$-resemblant for objective 2 with non-negative values, then this agent is $(\zeta + \zeta')$-resemblant for any convex combination of the two objectives.*

*Proof.* For any quantile $q$, let EX be the $q$ ex ante optimal mechanism for the convex combination of the objectives. Let $\textbf{Payoff}_1[\mathrm{EX}]$ be the contribution of objective 1 given mechanism EX and $\textbf{Payoff}_2[\mathrm{EX}]$ be the contribution of objective 2 given mechanism EX. Let $\textbf{Payoff}[\mathrm{EX}] = \alpha \cdot \textbf{Payoff}_1[\mathrm{EX}] + (1 - \alpha) \cdot \textbf{Payoff}_2[\mathrm{EX}]$ be the convex combination of the contributions given $\alpha \in (0, 1)$. Let $q_1 = \mathrm{argmax}_{q' \leq q} \bar{P}_1(q')$ and $q_2 = \mathrm{argmax}_{q' \leq q} \bar{P}_2(q')$, where $\bar{P}_1$ and $\bar{P}_2$ are the concave hull of price posting payoff curves for objectives 1 and 2 respectively. Let $\bar{P}$ be the concave hull of price posting payoff curves for the convex combination of objectives 1 and 2. Then, we have

$$\begin{aligned}
&\textbf{Payoff}[\mathrm{EX}] = \alpha \cdot \textbf{Payoff}_1[\mathrm{EX}] + (1 - \alpha) \cdot \textbf{Payoff}_2[\mathrm{EX}] \\
&\leq \alpha\zeta \cdot \bar{P}_1(q_1) + (1 - \alpha)\zeta' \cdot \bar{P}_2(q_2) \leq \zeta \cdot \bar{P}(q_1) + \zeta' \cdot \bar{P}(q_2) \leq (\zeta + \zeta') \max_{q' \leq q} \bar{P}(q').
\end{aligned}$$

Thus this agent is $(\zeta + \zeta')$-resemblant for the convex combination of the two objectives. $\square$

## B.3 Reduction Framework for Non-expected Utilities

We first formally introduce the model of non-expected utilities. In Appendix B.3.1, we present the reduction framework that converts sequential posted pricing mechanisms for linear agents to non-linear agents without expected utility representations, and approximately preserves its payoff approximation guarantee. We illustrate the application of our framework by showing that agents with endogenous valuations are 1-resemblant for both welfare and revenue maximization in Appendix B.3.2.

**Non-expected Utilities.** For any agent with type $t$, her utility for a distribution $\pi$ over allocations and payments is $u(t, \pi)$. Given a per-unit price $p$, letting $\pi_p(z)$ be the distribution over allocations and payments given demand $z$, the optimal demand of the agent is

$$d(t,p) \in \max_{z\in[0,1]} \; u(t, \pi_p(z)).$$

We assume that the maximum demand is always attainable. Moreover, we assume the demand correspondence $d(t,p)$ is upper hemi-continuous and has convex image for any type $t$. Therefore, the price-posting payoff curves $P(q)$, ironed price-posting payoff curve $\bar{P}$, and optimal payoff curves $R(q)$ in Section 2 can be analogously defined for agents with non-expected utilities.[23] We also say such an agent is $\zeta$-resemblant if her ironed price-posting payoff curve $\bar{P}$ is $\zeta$-*resemblant* to her optimal payoff curve $R$.

### B.3.1 Reduction Framework for Sequential Posted Pricing

We present a reduction framework that converts sequential posted pricing mechanisms for linear agents to non-linear agents without expected utility representations, and approximately preserves its payoff approximation guarantee.

**Theorem 4.** *Fix any set of non-linear agents with price-posting payoff curves $\{P_i\}_{i\in N}$ that are $\zeta$-resemblant to their optimal payoff curves $\{R_i\}_{i\in N}$. If there exists a sequential posted pricing mechanism that is a $\gamma$-approximation to the ex ante relaxation for linear agents analog with price-posting payoff curves $\{P_i\}_{i\in N}$, i.e.,* $\mathrm{SPP}(\{P_i\}_{i\in N}) \geq {}^1\!/\!_\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$*, then this mechanism is also a $\gamma\zeta$-approximation to the ex ante relaxation for non-linear agents, which implies that* $\mathrm{SPP}(\{P_i\}_{i\in N}) \geq {}^1\!/\!_{\gamma\zeta} \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$.

*Proof.* Let $\{q_i^\dagger\}_{i\in N}$ be the profile of optimal ex ante quantiles for optimal payoff curves $\{R_i\}_{i\in N}$. Since the ironed price-posting payoff curves $\{\bar{P}_i\}_{i\in N}$ are $\zeta$-resemblant to the optimal payoff curves

[23] For agents with non-expected utilities, the ironed price-posting payoff curve can still be attained by randomizing over at most two per-unit prices, and letting the agent know about the realizations of the per-unit prices before making her demand choices.

$\{R_i\}_{i \in N}$, there exists a sequence of quantiles $\{q_i^\ddagger\}_{i \in N}$ such that for any agent $i$, $q_i^\ddagger \leq q_i^\dagger$ and $\bar{P}(q_i^\ddagger) \geq 1/\zeta \cdot R(q_i^\dagger)$. Note that since $\sum_i q_i^\ddagger \leq \sum_i q_i^\dagger \leq 1$, $\{q_i^\ddagger\}_{i \in N}$ is also feasible for ex ante relaxation. Therefore,

$$\mathrm{EAR}(\{R_i\}_{i \in N}) = \sum_{i \in N} R_i(q_i^\dagger) \leq \zeta \cdot \sum_{i \in N} \bar{P}_i(q_i^\ddagger) \leq \zeta \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i \in N}).$$

Since the expected payoff of the sequential posted pricing mechanism only depends on the price posting payoff curves, not on the agents' utility models, we have

$$\mathrm{SPP}(\{P_i\}_{i \in N}) \geq 1/\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i \in N}) \geq 1/\gamma\zeta \cdot \mathrm{EAR}(\{R_i\}_{i \in N}),$$

and Theorem 4 holds. □

#### B.3.2 Endogenous Valuation

Each agent $i \in N$ can make costly investments before the auction by taking action $a_i \in \mathbb{R}$. For agent $i$ with private type $t_i$, the cost for action $a_i$ is $c_i(a_i)$ and the value for the item is $v_i(a_i, t_i) = a_i + t_i$. Given allocation $x$ and payment $p$, agent $i$ taking action $a_i$ has utility $x \cdot v_i(a_i, t_i) - p - c_i(a_i)$. This is the model presented in Gershkov et al. (2021).[24] Note that in this endogenous utility model, the agent can be equivalently modeled as one with convex preference over allocations, which does not satisfy the expected utility characterization.

For agents with endogenous valuation, in order to apply our reduction framework, it is important to specify the timeline for agents to exert costly efforts as it affects the equilibrium payoff of any given mechanism. In this paper, we assume that the agent can delay the investment decision until she sends a message to the seller. In the case of sequential posted pricing mechanisms, for each agent $i$, the agent makes the investment decisions after she sees the realized price offered by the seller. Note that the price is infinite if the item is sold to previous agents and agent $i$ will not make any investment given this price.

Under this timeline of the model, we can show that agents with endogenous valuation are 1-resemblant for both welfare maximization and revenue maximization under regularity conditions. This implies that the worst-case approximation guarantee of sequential posted pricing for agents with endogeous valuations is the same as linear agents when compared to the ex ante relaxation for the objective of both welfare and revenue maximization.

[24] Gershkov et al. (2021) characterized the single-agent revenue optimal mechanism for slightly more general classes of valuation functions. To simplify the presentation, in this paper, we only illustrate the proof for this special form of valuation function, and the same technique can be easily extended to broader settings.

**Lemma 16** (Fan and Lorentz, 1954; Gershkov et al., 2021). *For any function $L : \mathbb{R}^2 \to \mathbb{R}$ such that $L(x, q)$ is supermodular in $(x, q)$ and convex in $x$, for any pair of allocations $x \prec \hat{x}$,*[25] *we have*

$$\int_0^1 L(x(q), q)\,\mathrm{d}q \le \int_0^1 L(\hat{x}(q), q)\,\mathrm{d}q.$$

**Proposition 7.** *An agent with endogenous valuation has the price-posting welfare curve $P$ that equals (i.e. 1-resemblant) her optimal welfare curve $R$.*

*Proof.* For an agent with endogenous valuation, her type $t$ is mapped to quantile $q$ according to type distribution $\bar{F}$ in the same way as a linear agent, i.e., the quantile is in essence one minus the inverse of the cumulative distribution function. Let $L(x, q)$ be the welfare of the agent with type corresponding to quantile $q$ when she makes optimal investment decision given allocation $x$. By Gershkov et al. (2021), the function $L(x, q)$ is supermodular in $(x, q)$ and convex in $x$. For any quantile constraint $\hat{q}$, let $\hat{x}$ be the allocation such that $\hat{x}(q) = 1$ for any $q \le \hat{q}$ and $\hat{x}(q) = 0$ otherwise. Any incentive compatible mechanism with allocation $x$ that sells the item with probability $\hat{q}$ satisfies $x \prec \hat{x}$. By Lemma 16, the optimal mechanism that is $\hat{q}$ feasible has allocation rule $\hat{x}$, which is posting a deterministic price to the agent. Thus this agent has price-posting welfare curve $P$ that equals (i.e. 1-resemblant) her optimal welfare curve $R$. $\square$

**Proposition 8.** *An agent with endogenous valuation and regular type distribution has the ironed price-posting revenue curve $\bar{P}$ that equals (i.e. 1-resemblant) her optimal revenue curve $R$.*

*Proof.* Let $L(x, q)$ be the virtual value of the agent given allocation $x$ and type with quantile $q$. By Gershkov et al. (2021), the function $L(x, q)$ is supermodular in $(x, q)$ and convex in $x$ if the type distribution is regular. Similar to Proposition 7, for any quantile $\hat{q}$, the optimal mechanism for maximizing the expected virtual value that sells the item with probability at most $\hat{q}$ is posted pricing. Since the expected revenue equals the expected virtual value, this agent has price-posting revenue curve $\bar{P}$ that equals (i.e. 1-resemblant) her optimal revenue curve $R$. $\square$

## B.4 Anonymous Pricing

**Example B.2.** *Consider the single-item setting with two budget agents. Let $v$ be a sufficiently large number. Agent 1 has value $v$ and no budget constraint while agent 2 has value $v^2$ and budget 1. The welfare optimal mechanism allocates the item to agent 2, with welfare $v^2$. However, if the anonymous price is at most $v$, then agent 1 will buy the whole item and if the anonymous price is*

[25] $x \prec \hat{x}$ means that for any $\hat{q} \in [0, 1]$, $\int_0^{\hat{q}} x(q)\,\mathrm{d}q \le \int_0^{\hat{q}} \hat{x}(q)\,\mathrm{d}q$ and $\int_0^1 x(q)\,\mathrm{d}q = \int_0^1 \hat{x}(q)\,\mathrm{d}q$.

*larger than $v$, the item is sold with probability at most $\frac{1}{v}$. Thus anonymous pricing can guarantee welfare at most $v$, with approximation factor at least $v$, which is unbounded.*

**Proposition 4.** *Fix any set of (non-linear) agents with price-posting revenue curves $\{P_i\}_{i\in N}$ that are $\zeta$-resemblant to their optimal revenue curves $\{R_i\}_{i\in N}$. If the price-posting revenue curves are concave, then anonymous pricing is a $\zeta e$-approximation to the ex ante relaxation on the optimal revenue curves, i.e.,* $\mathrm{AP}(\{P_i\}_{i\in N}) \geq {}^1/_{\zeta e} \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$.

*Proof.* Let $\{q_i\}_{i\in N}$ be the optimal ex ante relaxation for ex ante revenue curves $\{R_i\}_{i\in N}$, and let $q_i^\dagger$ be the quantile assumed to exist by $\zeta$-resemblance such that $q_i^\dagger \leq q_i$ and $\bar{P}_i(q_i^\dagger) \geq \frac{1}{\zeta} R_i(q_i)$ for each $i$. Since the price-posting revenue curves are concave, we have $\{P_i\}_{i\in N} = \{\bar{P}_i\}_{i\in N}$, and

$$\mathrm{EAR}(\{P_i\}_{i\in N}) = \mathrm{EAR}(\{\bar{P}_i\}_{i\in N}) \geq \sum\nolimits_i \bar{P}_i(q_i^\dagger) \geq \frac{1}{\zeta} \sum\nolimits_i R(q_i) = \frac{1}{\zeta} \mathrm{EAR}(\{R_i\}_{i\in N}).$$

By Alaei et al. (2019), $e \cdot \mathrm{AP}(\{P_i\}_{i\in N}) \geq \mathrm{EAR}(\{P_i\}_{i\in N})$ if the price-posting revenue curves $\{P_i\}_{i\in N}$ are concave. Combining the inequalities, we have

$$\zeta e \cdot \mathrm{AP}(\{P_i\}_{i\in N}) \geq \mathrm{EAR}(\{R_i\}_{i\in N}). \qquad \square$$

## B.5 Oblivious Posted Pricing

For oblivious posted pricing mechanisms (e.g. Chawla et al., 2010), we show how to apply resemblant property between the ironed price-posting payoff curve and optimal payoff curve to obtain approximation results for agents with general utility. Similar to sequential posted pricing, we will define the oblivious posted price in quantile space.

**Definition B.3.** *An* oblivious posted pricing mechanism *is $(\{q_i\}_{i\in N})$ where the adversary chooses an ordering $\{o_i\}_{i\in N}$ of the agents, and $\{q_i\}_{i\in N}$ denotes the quantile corresponding to the per-unit prices to be offered to agents at the time they are considered according to the order $\{o_i\}_{i\in N}$ if the item is not sold to previous agents. Note that quantiles $\{q_i\}_{i\in N}$ can be dynamic and depends on both the order and realization of the past agents.*

Given the definition of the oblivious quantile pricing mechanism, we denote the payoff of the oblivious quantile pricing mechanism $(\{q_i\}_{i\in N})$ for agents with a collection of price-posting payoff curves $\{P_i\}_{i\in N}$ by $\mathrm{OPP}(\{P_i\}_{i\in N}, \{q_i\}_{i\in N})$, and the optimal payoff for the oblivious quantile pricing mechanism is

$$\mathrm{OPP}(\{P_i\}_{i\in N}) = \max_{\{q_i\}_{i\in N}} \mathrm{OPP}(\{P_i\}_{i\in N}, \{q_i\}_{i\in N}).$$

Similar to Theorem 4, we have the following reduction framework for oblivious posted pricing for non-linear agents. The proof is identical to Theorem 4, hence omitted here.

**Theorem 5.** *Fix any set of (non-linear) agents with price-posting payoff curves $\{P_i\}_{i \in N}$ that are $\zeta$-resemblant to their optimal payoff curves $\{R_i\}_{i \in N}$. If there exists an oblivious posted pricing mechanism ($\{q_i\}_{i \in N}$) that is a $\gamma$-approximation to the ex ante relaxation for linear agents analog with price-posting payoff curves $\{P_i\}_{i \in N}$, i.e.,* $\mathrm{OPP}(\{P_i\}_{i \in N}, \{q_i\}_{i \in N}) \geq {}^1/_\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i \in N})$*, then this mechanism is also a $\gamma\zeta$-approximation to the ex ante relaxation for non-linear agents, i.e.,* $\mathrm{OPP}(\{P_i\}_{i \in N}, \{q_i\}_{i \in N}) \geq {}^1/_{\gamma\zeta} \cdot \mathrm{EAR}(\{R_i\}_{i \in N})$.

For the single item setting, there exists an oblivious posted pricing mechanism that is a 2-approximation to the ex ante relaxation for linear agents (Feldman et al., 2016). In addition, if the price-posting payoff curves are the same for all gents, the approximation ratio is improved to ${}^1/_{(1 - {}^1/\sqrt{2\pi})}$ (Yan, 2011).

## B.6 General Feasibility Constraint

Our results can be generalized to multi-unit auctions with downward closed feasibility constraints. Let $\mathcal{X}$ be the set of feasible allocation profiles. The set $\mathcal{X}$ is downward closed if for any $\{x_i\}_{i \in N} \in \mathcal{X}$, we have $\{x'_i\}_{i \in N} \in \mathcal{X}$ if $x'_i \leq x_i$ for any $i \in N$. We denote the set of ex ante feasible quantiles with respect to feasibility constraint $\mathcal{X}$ by EAF($\mathcal{X}$). The optimal ex ante payoff given a specific collection of payoff curves $\{R_i\}_{i \in N}$ and feasibility constraint $\mathcal{X}$ is

$$\mathrm{EAR}(\{R_i\}_{i \in N}, \mathcal{X}) = \max_{\{q_i\}_{i \in N} \subseteq \mathrm{EAF}(\mathcal{X})} \sum\nolimits_{i \in N} R_i(q_i).$$

Given feasibility constraint $\mathcal{X}$, the sequential posted pricing mechanism $(\{o_i\}_{i \in N}, \{q_i\}_{i \in N})$ offers each agent $i$ the price corresponding to quantile $q_i$ according to order $\{o_i\}_{i \in N}$ if it is feasible to serve agent $i$ given the allocation of previous agents. The payoff achieved by the sequential posted pricing mechanism $(\{o_i\}_{i \in N}, \{q_i\}_{i \in N})$ for agents with a specific collection of price-posting payoff curves $\{P_i\}_{i \in N}$ given feasibility constraint $\mathcal{X}$ is denoted by $\mathrm{SPP}(\{P_i\}_{i \in N}, \{o_i\}_{i \in N}, \{q_i\}_{i \in N}, \mathcal{X})$.[26]

It is easy to verify that the reduction framework for (sequential) posted pricing mechanisms (Theorem 4) and the reduction framework for pricing-based mechanism (Theorem 3) directly apply when there is a downward closed feasibility constraint $\mathcal{X}$. In addition, it is shown in the literature that for general class of feasibility constraints, posted pricing mechanisms are approximately optimal for linear agents. We formally state the reduction result for sequential posted pricing in the following theorem.

[26]Here we only formally discuss the extension for sequential posted pricing mechanisms. The generalizations for other posted pricing mechanisms hold similarly.

**Theorem 6.** *Given feasibility constraint $\mathcal{X}$, for linear agents with the price-posting payoff curves $\{P_i\}_{i\in N}$, there exists a sequential posted pricing mechanism $(\{o_i\}_{i\in N}, \{q_i\}_{i\in N})$ that is a $\gamma$-approximation to the ex ante relaxation, i.e.,* $\mathrm{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}, \mathcal{X}) \geq {}^{1}\!/_{\gamma} \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N}, \mathcal{X})$ *where $\gamma = {}^{e}\!/_{(e-1)}$ if $\mathcal{X}$ is a matroid (Yan, 2011), $\gamma = {}^{1}\!/_{(1-{}^{1}\!/_{\sqrt{2\pi}})}$ if $\mathcal{X}$ is a knapsack (Balkanski and Hartline, 2016), and $\gamma = {}^{1}\!/_{(1-{}^{1}\!/_{\sqrt{2\pi k}})}$ for $k$-unit auctions (Yan, 2011).*

## B.7 Examples for Approximation Guarantees

In this section, we provide the details for numerical evaluations. Recall that we will focus on selling $m$ units of identical items to $n > m$ i.i.d. agents with the objective of revenue maximization.

### B.7.1 Risk-averse Utility

Recall that we focus on a simple example where each agent has constant absolute risk aversion (CARA) utility function, i.e.,

$$\varphi(z) = \frac{1}{a}\left(1 - \exp\left(-az\right)\right)$$

for risk parameter $a > 0$, where $\exp(\cdot)$ is the exponential function. We assume that the risk averse agent has private value drawn from uniform distribution in $[0, 1]$.

By Claim 1, a risk-averse agent behave the same as a linear agent when facing a per-unit price and hence the risk attitude of the agent does not matter under posted pricing. Therefore, we can apply the same sequential posted pricing mechanism for risk-averse agents as the one for linear agents, and achieve the same expected revenue as if all agents have linear utilities. In particular, we will numerically evaluate the revenue of sequential posted pricing mechanism where the per-unit price $p_i$ of each agent is set such that the probability the price is accepted equals $\min\left\{\frac{m}{n}, \frac{1}{2}\right\}$, where $\frac{1}{2}$ is the probability the item is sold in the single-agent problem when the seller posts a monopoly price of $\frac{1}{2}$ to maximize the expected revenue for uniform distributions. Thus, the per-unit price $p_i$ equals $1 - \min\left\{\frac{m}{n}, \frac{1}{2}\right\}$ for each agent $i$. The revenue guarantees of this particular sequential posted pricing are illustrated in Figure 2a.

### B.7.2 Private-budget Utility

Recall that we consider budget constrained agents whose value and budget are drawn independently from uniform distribution in $[0, 1]$. We numerically evaluate the revenue of sequential posted pricing mechanism where the per-unit price of each agent is set such that the probability each agent wins a unit of item is $\min\{\frac{m}{n}, \frac{1+\sqrt{3}}{6}\}$, where $\frac{1+\sqrt{3}}{6}$ is the probability an item is sold in the single-agent problem when the seller posts a per-unit price to maximize the expected revenue. Therefore, the per-unit price for each budgeted agent $i$ in this example is

$$p_i = \frac{1}{2}\left(3 - \sqrt{9 - 8(1 - \min\{\frac{m}{n}, \frac{1+\sqrt{3}}{6}\})}\right).$$

The revenue guarantees of this particular sequential posted pricing are illustrated in Figure 2b.

# C Tie-breaking Rules

In Algorithm 1, it is possible that based on the tie-breaking rule of the pricing-based mechanism, agent $i$ may not receive the item if her quantile equals the threshold quantile. In this case, the probability of agent $i$ with type $t_i$ for winning an item when others have quantiles $q_{-i}$ can be strictly smaller than $d(t_i, p^{\hat{q}_i})$. Denote this probability as $z_i(t_i, q_{-i})$. In this case, the payment of agent $i$ with type $t_i$ in mechanism $\mathcal{M}$ is $p^{\hat{q}_i} \cdot z_i(t_i, q_{-i})$ regardless of the allocation.

Note that based on our construction, the marginal quantile for each agent is drawn from the uniform distribution in $[0, 1]$. Therefore, ties occur with measure zero and hence the expected payoff of mechanism $\mathcal{M}$ for non-linear agents is not affected.

Now it is sufficient to show that both DSIC and IIR properties hold in mechanism $\mathcal{M}$ for non-linear agents hold even for the measure zero event that

$$z_i(t_i, q_{-i}) < d(t_i, p^{\hat{q}_i}).$$

Note that for each agent $i$ with type $t_i$, her quantile is drawn from cumulative distribution function $H_i^{t_i}$. Therefore, such event happens only if distribution $H_i^{t_i}$ has a point mass at threshold quantile $\hat{q}_i$. Moreover,

$$z_i(t_i, q_{-i}) = \lim_{q \to \hat{q}_i^-} d(t_i, p^q).$$

Since the demand correspondence is upper hemi-continuous (Assumption 1), lottery $z_i(t_i, q_{-i})$ is also a best response of agent $i$ with type $t_i$ given market clearing price $p^{\hat{q}_i}$. Therefore, the constructed mechanism is also DSIC and IIR for the measure zero event.