Simple Mechanisms for Agents with Non-linear Utilities[*]

Yiding Feng[†] Jason D. Hartline[‡] Yingkai Li[§]

This version: Jan 14, 2024.
newest version

Abstract

We show that many economic conclusions derived from linear utility models (e.g., Bulow and Roberts, 1989) approximately extend to non-linear utility models. Specifically, we quantify the extent to which agents with non-linear utilities resemble agents with linear utilities, and provide a general framework for extending the approximation guarantees of simple mechanisms from agents with linear utilities to agents with non-linear utilities. We apply the framework for the objectives of welfare and revenue on non-linear utility models that include agents with budget constraints and agents with risk aversion, and conclude that simple mechanisms are approximately optimal for these non-linear utility models. Moreover, we show that when the seller only observes the demand functions, not the detailed non-linear utility models, the simple marginal revenue maximization mechanism is robustly optimal.

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

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

[†]Booth School of Business, University of Chicago. Email: yiding.feng@chicagobooth.edu
[‡]Department of Computer Science, Northwestern University. Email: hartline@northwestern.edu
[§]Cowles Foundation for Research in Economics, Yale University. Email: yingkai.li@yale.edu

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 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. Moreover, the multi-agent mechanism design problem can be decomposed as single-agent mechanism design problems through the feasibility constraints in quantile space via the reduced-form approach of Border (1991). Note that for a linear agent, the quantile of an agent is defined to be drawn uniformly between 0 and 1 and a lower quantile implies a higher value for the agent. From Bulow and Roberts (1989), the solution to these single-agent problems for linear agents are (possibly randomized) price postings and the optimal mechanism can be interpreted as marginal revenue maximization.

In environments with a single agent, pricing-based mechanisms can be generalized to non-linear agents by considering *per-unit* prices, i.e., given per-unit price p, an agent can purchase any lottery with winning probability $q \in [0, 1]$ and pay price $p \cdot q$ in expectation. When there are multiple agents, the definition of quantile space is not straightforward in order to implement the pricing-based mechanism in quantile space analogous to linear agents. In particular, the type space of a non-linear agent can be multi-dimensional. For example, for a risk averse agent, her private type could include both her value for the item and

[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".

her risk attitude. In our paper, we provide a natural definition that maps the potentially multi-dimensional type of a non-linear agent to a single-dimensional quantile that captures her willingness to pay based on her demand function given per-unit prices. Such reduction 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 (Theorem 1).

Note that for non-linear agents (e.g., agents with budget constraints), 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. Nonetheless, we introduce a reduction framework to show that these mechanisms are approximately optimal for large families of non-linear agents. For these families we say that the non-linear agents *resemble* linear agents. Specifically, given a pricing-based mechanism that guarantees a β-approximation (i.e., achieves at least $1/\beta$ fraction of the optimal objective) for linear agents and given non-linear agents that are ζ-*resemblant*[3] of linear agents and satisfy the von Neumann-Morgenstern expected utility representation (Morgenstern and von Neumann, 1953), the reduction framework shows that the analogous pricing-based mechanism for non-linear agents guarantees a $\beta\zeta$-approximation to the optimal (Theorem 2). 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), and even for agents that do not satisfy expected utility representation (e.g., agents with endogenous valuation (Gershkov et al., 2021; Akbarpour et al., 2023)). See Section 7 for a detailed discussion.

In the classical linear utility model of mechanism design, there are extensive studies of simple mechanisms with various approximation guarantees. 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 allocating items to agents to maximize the marginal revenue. 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 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 price

[3]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.

discriminate the agents are crucial for revenue maximization. See the survey of Hartline (2012) for detailed discussion of the method of approximation in economics.

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 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 cannot 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.

Among pricing-based mechanisms, the marginal revenue maximization mechanism is particularly noteworthy. This mechanism, which maximizes expected revenue across all pricing-based alternatives, has been shown to be revenue optimal for linear agents (Bulow and Roberts, 1989). When agents have non-linear utilities, while our reduction framework indicates that this mechanism is approximately optimal, it is generally not exactly optimal. However, in many practical applications, the exact Bayesian optimal mechanism may not be implementable due to the principal's limited access to exact utility functions or type distributions of the agents. In often cases, the principal can only estimate the expected demands of the agents using historical purchase data. We show that when the principal only knows the expected demands, the marginal revenue maximization mechanism is robustly optimal, i.e., it 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 3). This robustness result suggests that when 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.

1.1 Discussion of Our Results on ζ-Resemblance

To instantiate our reduction framework to illustrate the approximations of pricing-based mechanisms for non-linear agents, we characterize broad families of non-linear agents that are ζ-resemblant for small constant factors ζ (e.g., 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 ζ-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 ζ-resemblant for modest ζ, non-linearity is a crucial feature that needs specific study for identifying forms of mechanisms lead to good economic outcomes.

The ζ-resemblance property that governs our reduction is based on a single-agent price-posting approximation with ex ante supply constraints (c.f., Bulow and Roberts, 1989). In particular, we say a non-linear agent is ζ-resemblant if given any ex ante allocation constraint $q \in [0, 1]$, the optimal payoff of the principal under the constraint that the ex ante probability an item is sold to the agent is at most q, is at most ζ 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 ζ results in a better approximation of posted pricing in single-agent problems, making non-linear agents more resemblant to linear agents. Note that it is essential to assume ζ-resemblance in our reduction framework of pricing-based mechanisms, as single-agent problems with ex ante supply constraints can be viewed as a specific case of multi-agent problems with symmetric agents.

In our paper, we focus on the objective of revenue maximization and show that for canonical non-linear utility models, parameter ζ is bounded by small constants. We also discuss the extension for welfare maximization and their convex combination in Section 7.1.

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 (CARA), 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 3, for common valuation distributions such as uniform distribution, the revenue loss 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 is her payment exceeds the budget. We show that if the budgets are independent from the values, for regular valuation distributions,[4] a public budget agent is 1-resemblant and a private budget agent is 3-resemblant. This implies that simple mechanisms such as sequential posted pricing are approximately optimal for agents with independent budget distributions, and in particular when the budget is public, the worst case approximation is attained when the budgets never bind, i.e., the agent behaves the same as a linear agent. In contrast, when the values are correlated with budgets, the principal can price discriminate the agents using their budget constraints, thereby achieving almost full surplus in some cases. In those cases, posted pricing mechanisms are not approximately optimal even in single-agent settings and such budgeted agent are not constant resemblant to linear agents.

1.2 Related Work

Frameworks for reducing approximation for non-linear agents to approximation for linear agents has also been studied in Alaei et al. (2013). Their reduction framework converts the marginal revenue mechanism for linear agents to a mechanisms for non-linear agents and general objectives. Unlike our framework which uses single-agent price-posting mechanisms (induced from price-posting payoff curves) as a building-block, Alaei et al. (2013) convert mechanisms for linear agents into mechanisms for non-linear agents with single-agent ex ante optimal mechanisms (induced from optimal payoff curves) as components. From the mechanism designer's perspective, identifying ex ante optimal mechanisms for a single non-linear agents can be much harder than identifying ex ante optimal price-posting mechanisms (e.g., private budget utility, risk averse utility). Furthermore, due to this difference, the implementation of the reduction framework together with its outcome mechanisms in Alaei et al. (2013) is more complex than ours. In general, the framework in Alaei et al. (2013) converts DSIC mechanisms for linear agents into Bayesian incentive compatible mechanisms for non-linear agents.

The optimality of simple mechanisms in complex environments has also been established in Carroll (2015, 2017); Brooks and Du (2021) and 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. In contrast, there is no ambiguity for the principal in our model, and the optimal mechanisms can take complex forms. However, the principal may still wish to adopt simple mechanisms in practice since their performances are close to

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

optimal, and we also get robustness, e.g., to collusion with price posting. See Hartline and Lucier (2015) for more discussions on the benefits of adopting simple mechanisms.

Mechanism design for non-linear agents are widely studied in the literature. In this work, as applications of our general framework, we focus on two specific 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. Baisa (2017) shows that the optimal mechanism for risk averse agents departs from the linear agents, since the optimal mechanism does not allocate to the highest bidder, and can better screen the agents through allocating the item to a group of agents with lotteries. Gershkov et al. (2022) show that if the seller can make positive transfer to the agents, the optimal mechanism features the property that under equilibrium, all agents face no uncertainty in the realized utility.

Laffont and Robert (1996) and Maskin (2000) respectively study the revenue-maximization and welfare-maximization problems for symmetric agents with *public* budgets in single-item environments. 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. Devanur and Weinberg (2017) provide an algorithmic solution to the same problem without distributional assumptions. Richter (2019) shows that 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. For more general settings, no closed-form characterizations are known. However, the optimal mechanism can be solved by a polynomial-time solvable linear program over interim allocation rules (cf., Alaei et al., 2012).

It is well known that simple mechanisms are approximately optimal for linear agents. For single item auction with linear agents, Jin et al. (2019) show that the tight ratio between anonymous pricing and the optimal mechanism is 2.62 under regularity assumption, 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 pointwise 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).

2 Preliminaries

In this paper, we consider general payoff maximization for selling m units of identical items to n non-linear agents where each agent has unit demand for the items. For example, welfare maximization, revenue maximization and their convex combinations are special cases of payoff maximization.

Agent 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. In this paper, we use subscript i to represent agent i in the multi-agent problems, and we will drop the subscripts without ambiguity when we discuss the 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 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 that we are interested in applying our general framework to.

  - **Risk-averse utility:** The agent's type is a pair $t = (v, \varphi)$ where v is the private value and φ is an increasing and concave function that represents the agent's 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

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

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

Mechanisms. In this paper, we consider the sealed-bid mechanisms: in a mechanism $\{(x_i, p_i)\}_{i \in N}$, agents simultaneously submit sealed 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 $p_i(\{b_i\}_{i \in N})$. The outcome of mechanisms is a distribution π over the allocation payment pair (x_i, p_i) for each agent i where the allocation is $x_i \in \{0, 1\}$ and the price is $p_i \in \mathbb{R}_+$.[5] An allocation is feasible if $\sum_i x_i \leq m$ where m is the number of items.[6]

We consider mechanisms that satisfy *Bayesian incentive compatible* (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 rational* (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.

Pricings. We also study mechanisms based on simple per-unit posted pricing.

Definition 2.1. *Posting* per-unit price p *is offering a menu* $\{(z, z \cdot p) : z \in [0, 1]\}$ *to the agent. 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.*[7]

[5] Agents do not observe the realization of π before submitting their bids.

[6] 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 seller. Our results can be extended to multi-unit auctions with more general feasibility constraints. See Appendix A.6 for detailed discussions.

[7] By revelation principle, given any utility function of the agent, any per-unit posted pricing mechanism can also be transformed into a sealed-bid mechanism. In per-unit posted pricing mechanisms, the payment format does not matter for linear agents, but it affects the incentives for non-linear agents. We focus on all-pay format in this paper and show that simple mechanisms are approximately optimal under such payment format. However, our general reduction framework can also be extended to other payment formats such as winner-pays-all.

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

$$d(t, p) \in \max_{z \in [0,1]} z \cdot u(t, 1, z \cdot p) + (1 - z) \cdot u(t, 0, z \cdot p).$$

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 $v \geq p$, and purchase the lottery $d = 0$ otherwise. We assume that utility of the agent is upper semi-continuous in chosen lottery z so that the demand maximizing the utility can always be attained. Moreover, we assume that for any type t, $d(t, 0) = 1$ and there exists $p > 0$ such that $d(t, p) = 0$.

Ordinary Goods. 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 budgeted utility or risk-averse utility.

Assumption 1. *The item is* the ordinary good*, i.e., for any type* t*, the demand* $d(t, p)$ *is upper hemi-continuous, has convex image, and is 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.

Objectives. We assume that the seller is an expected utility maximizer. The seller'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.

3 Examples

In this section, we will focus on revenue maximization for selling m units of identical items to n i.i.d. agents, where $m < n$. We first recap a simple mechanism, sequential posted pricing, for linear agents. Then we illustrate how to extend sequential posted pricing to risk-averse utilities and private-budget utilities, and show that it achieves revenue that is close to the optimal under non-linear utilities through numerical calculations.

Sequential Posted Pricing for Linear Agents. In *sequential posted pricing mechanisms*, there is an order on agents such that the seller approaches them according to this given order. When we restrict to the i.i.d. environments, the order on agents does not matter. Therefore, without loss of generality, we assume that agents are approached by lexicographical order in this section.

Given lexicographical order, there exists a sequence of per-unit prices $\{p_i\}_{i\in[n]}$. For each agent i, the seller offers the agent price p_i if the items have not been sold out to previous $i-1$ agents. Moreover, a linear agent i with value v_i will purchase a unit of item with probability 1 if and only if her value is higher than the price.

3.1 Risk-averse Utility

We consider agents with risk-averse utilities. Recall that a risk-averse agent's type is a pair $t = (v, \varphi)$. We first show that any risk-averse agent, when facing a per-unit price, will behave the same as a linear agent.

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

Proof. Given any per-unit price p, agent's expected utility with demand z is

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

since function φ is concave, and the equality holds when $z \in \{0, 1\}$. Moreover, since φ is monotonically increasing and $z \cdot (v - p)$ is maximized when z is either 0 or 1, the expected utility is maximized at allocation $z \in \{0, 1\}$. □

In the remaining of this section, we focus on the 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.

Numerical Evaluations for Sequential Posted Pricing. We consider a risk averse agent with 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

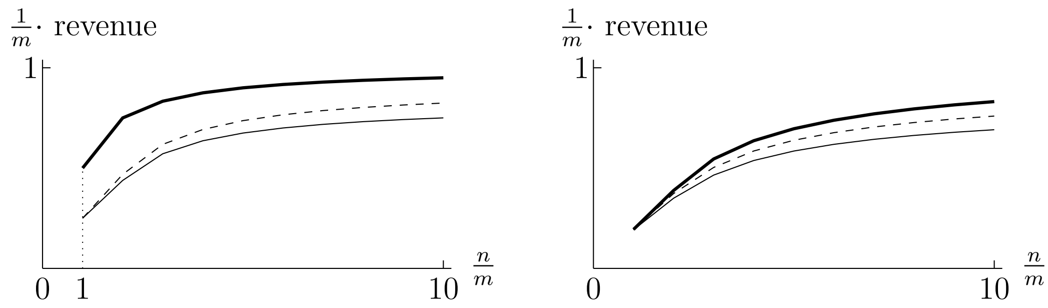

(a) risk-averse agents

(b) private-budget agents

Figure 1: 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 upper bound does not depend on the value of m.

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 price of each agent is set such that the probability each agent accepts the price is exactly $\min\left\{\frac{m}{n}, \frac{1}{2}\right\}$, 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. For each risk-averse agent i with value drawn from uniform distribution $U[0, 1]$, the per-unit price p_i she faces is exactly $1 - \min\left\{\frac{m}{n}, \frac{1}{2}\right\}$.

Note that in general computing the optimal revenue for risk-averse agents is challenging. Therefore, in the comparison of revenues in Figure 1a, instead of plotting the optimal revenue, we will plot an upper bound on the optimal revenue through the technique of ex ante relaxation, which will be elaborated in Section 5. As illustrated in Figure 1a, 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.

3.2 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. In this section, we numerically evaluate the approximation ratio of sequential posted pricing.

Numerical Evaluations for Sequential Posted Pricing. We consider budget constrained agents whose value and budget are drawn independently from uniform distribution in $[0, 1]$. For each agent i with a private budget, when offered a per-unit price p_i, she may purchase a lottery with probability strictly less than 1 in order to satisfy the budget constraint. In this case, we restrict attention to sequential posted pricing mechanisms realize the lottery immediately before offering a price to the next agent.

For agents with private budgets, we will 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 sequential posted pricing are illustrated in Figure 1b. Again, the plot for optimal revenue is replaced by its upper bound using ex ante relaxation. As illustrated in Figure 1b, the gap between the expected revenue from sequential posted pricing and the upper bound of the optimal revenue is small in all cases, especially when m is large. For example, the multiplicative revenue gap when $n = 200$ and $m = 20$ is less than 1.095.

3.3 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 1. In Section 5, we will show that the expected payoffs from a broad family of simple mechanisms are close to the optimal for general non-i.i.d. environments when agents have non-linear utilities. Note that the analyses conducted in those sections are based on worst-case scenarios. The numerical analyses provided in this section serve to underscore 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.

In the construction of sequential posted pricing mechanisms for numerical evaluations, we infer the price posted to each agent by determining the probability each agent will purchase an item given the price. Intuitively, the probability an item is sold to each agent is more important than the value of per-unit prices in the sequential posted pricing mechanisms. In particular, by controlling those probabilities, the items are not sold out with sufficiently large probabilities, and thus the sequential posted pricing mechanism can extract sufficiently high revenue from all agents regardless of the order of their arrival in the mechanism. In Section 4, we will show that this feature of determining the price based on the probability of sale is also a crucial ingredient for other simple mechanisms.

4 Implementation of Pricing-based Mechanisms

In this section we introduce the implementation of *pricing-based mechanisms* for both linear agents and non-linear agents. Our construction relies on the definition of payoff curves, which captures the optimal payoff in the single-agent problems with ex ante supply constraints.

4.1 Single-agent Payoff Curves and Resemblance

In this section, we analyze the single-agent problem with ex ante supply constraints. We define the *optimal payoff curves* and *price-posting payoff curves* as follows, and we refer to them as optimal revenue curves and price-posting revenue curves if the payoff function is the revenue.

Definition 4.1. *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.,*

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

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 define price-posting payoff curves based on mechanisms that post per-unit prices.

Definition 4.2. *The* market clearing price p^q *for the ex ante supply constraint* q *is the highest per-unit price at which there exists a demand function such that the expected demand of the agent is* q*, i.e.,* $\mathbf{E}_{t\sim\bar{F}}[d(t, p^q)] = q$*.*

The existence of the market clearing price is implied by the assumption that the demand correspondence is upper hemi-continuous and has convex image. Moreover, under the ordinary goods assumption, there exists a demand function $d(t, p^q)$ which is weakly increasing in q and satisfies the expected demand constraint for all q. We will focus on such selection rules throughout the rest of the paper.

Definition 4.3. *The* price-posting payoff curve $P(q)$ *is the expected payoff from posting market clearing price* p^q*.*

Price-posting payoff curves are not generally concave, and we can iron it to get the concave hull of the price-posting payoff curves. 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.

Definition 4.4. *The ironed price-posting payoff curve* $\bar{P}$ *is the concave hull of the price-posting payoff curve* P*.*

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 and 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 supply constraints.

ζ-resemblance. We introduce *ζ-resemblance* of an agent to measure how her ironed price-posting payoff curve resembles her optimal payoff curve. This definition is illustrated in Figure 2. Note that based our definition, Lemma 2 implies that a linear agent is 1-resemblant.

Definition 4.5 (ζ-resemblance)**.** *An agent's ironed price-posting payoff curve* $\bar{P}$ *is* ζ-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* ζ-resemblant*.*

Our definition implies that for a ζ-resemblant agent, given any supply constraint q, there exists a distribution over per-unit prices with expected sale probability at most q that

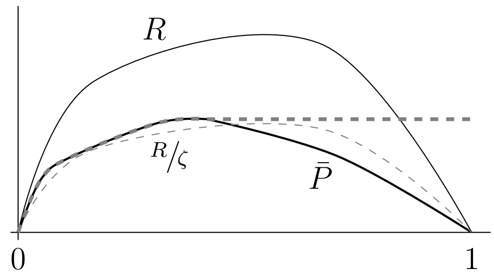

Figure 2: The thin dashed curve is the optimal payoff curve scaled by a factor of ζ, and the thick dashed curve is the monotone extension of the ironed price-posting payoff curve. The definition of ζ resemblance is satisfied if the thick dashed curve is always above the thin dashed line.

achieves at least $\frac{1}{\zeta}$ fraction of the optimal payoff. Note that in this definition, the seller is allowed to not run out the supply q given price-posting mechanisms to approximate the optimal payoff. This is equivalent to an environment of selling items to a continuum of agents with supply constraints, where ζ-resemblance measures the worst-case gap (over all possible supplies) between optimal mechanisms and posted pricing mechanisms for non-linear agents. Since posted pricing mechanisms are optimal mechanisms for an linear agent, smaller ζ implies that non-linear agents resemble linear agents better.

4.2 Pricing-based Mechanisms for Linear Agents

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 even in single-agent environments. For example, to maximize the revenue from a single private-budget agent, the menu size of the optimal mechanism is exponential in the size of the support of the budget distribution (Devanur and Weinberg, 2017).

Among all menu mechanisms, there is a subclass of mechanisms closely related to price posting which allow simple implementations – *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 (approximately optimal) in multi-agent environments (Myerson, 1981; Riley and Zeckhauser, 1983; Yan, 2011; Alaei et al., 2019). Our reduction framework (Theorem 2) in the next section will show that when non-linear agents have small resemblances, by focusing on simple menus mechanisms, i.e., pricing-based mechanisms, the payoff losses due to the non-linearity in utilities are not large.

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

Quantiles for Linear Agents. For agents with linear utilities, 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 uniform randomly if the above interval is not degenerate. 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, the marginal distribution over quantile q is a uniform distribution 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 linear agents, the price-posting payoff curve $P(q)$ is attained by selling to agents with quantiles below q. Therefore, 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) in the value space.[8] 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;
- 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 $q_i \leq Q_i(q_{-i})$.[9]

Essentially, $Q_i(q_{-i})$ is the threshold on allocation in the quantile space. In Section 5.1, we illustrate how to describe the Bayesian optimal mechanism and sequential posted pricing

[8] 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.

[9] 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 $q_i < z$.

mechanisms as pricing-based mechanisms by explicitly defining the corresponding threshold functions $\{Q_i\}$.

Since the distribution over quantiles is a uniform distribution in $[0, 1]$ 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

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

4.3 Pricing-Based Mechanisms for Non-linear Agents

For linear agents, the implementation of pricing-based mechanisms in the quantile space is intuitive since higher values are mapped to lower quantiles according to the cumulative distribution function of the valuation distribution. However, the construction of quantiles for non-linear agents is not obviously since the private types of non-linear agents can be multi-dimensional and there may not exist a natural single-dimensional order on the types that captures the agents' preference for all distributions over outcomes. In this section, we introduce a definition of quantiles for non-linear agents based on their demands given market clearing prices under different supplies.

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 selection 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.[10] Moreover, $H^t(0) = 0$ and $H^t(1) = 1$, and thus H^t can be viewed as a distribution supported on $[0, 1]$ with $H^t(q)$ being the cumulative distribution function for any type t.

Definition 4.6 (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 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.[11] Moreover, given any market clearing price, types with higher demands are mapped to lower

[10]For different selection of demand function $d(t, p^q)$, the defined function $H^t(q)$ may be different for any given t and q. However, our general reduction framework and approximation results in the paper hold for all selection rules.

[11]Note that the definition of quantile for non-linear agents encompasses the definition for linear agents. In particular, for an linear agent with value $v(q)$, the distribution $H^{v(q)}$ is a point-mass distribution at q.

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

Input: Non-linear agents $\{(\mathcal{T}_i, \bar{F}_i, u_i)\}_{i \in N}$; and a pricing-based mechanism $\mathcal{M}$ with threshold function $\{Q\}_{i \in N}$ 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 $q_i \le \hat{q}_i$ 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.[12]

quantiles with higher probabilities, and hence is more likely to win an item in pricing-based mechanisms where items are allocated based on the thresholds in the quantile space.

Another benefit of our definition of quantiles for non-linear agents is that the marginal distribution over quantiles is invariant of the utility models of the agents or the type distributions. Specifically, for any non-linear agent, the cumulative distribution function of quantile q is

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

where the last inequality holds due to the definition of the market clearing price p^z. Therefore, for any non-linear agent, over the randomness of her private type, the marginal distribution over quantile q is also drawn from uniform distribution $U[0, 1]$.

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 4.6, 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*

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

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,

[12] 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.

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 $\mathcal{M}$ defined in the quantile space 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 per-unit price $p^{\hat{q}_i}$ where 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 uniform distribution $U[0,1]$ for all agent i. □

5 Reduction Framework and Approximation Guarantees

We provide 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 Bayesian optimal mechanism for linear agents proposed in Bulow and Roberts (1989) can be extended with approximately optimal performance for non-linear agents.

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 introduce the benchmark of *ex ante relaxation*, 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 $\sum_i q_i \le m$. We denote the set of ex ante feasible quantiles by EAF. The optimal ex ante payoff given a specific collection of 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 2 (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*

- *mechanism $\mathcal{M}$ is a γ-approximation for linear agents, i.e., $\mathcal{M}(\{P_i\}_{i\in N}) \ge {}^1\!/\!_\gamma\cdot\mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$; and*
- *all non-linear agents are ζ-resemblant;*

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

Before the proof of Theorem 2, we first explain the conditions in Theorem 2, and in particular why we refer to condition $\mathcal{M}(\{P_i\}_{i\in N}) \ge {}^1\!/\!_\gamma\cdot\mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$ 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.[13] 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 $\mathcal{M}(\{P_i\}_{i\in N})$ 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 $\text{EAR}(\{\bar{P}_i\}_{i\in N})$. Therefore, the condition that $\mathcal{M}(\{P_i\}_{i\in N}) \geq 1/\gamma \cdot \text{EAR}(\{\bar{P}_i\}_{i\in N})$ can be interpreted as the approximation ratio of pricing-based mechanism $\mathcal{M}$ for the linear agents analog when compared to its ex ante relaxation. This condition is necessary for our reduction framework since linear utility is a special case of non-linear utilities.

Another condition in Theorem 2 is the requirement that all non-linear agents are ζ-resemblant. This condition is also necessary for pricing-based mechanisms to be approximately optimal for non-linear agents. This is because without this assumption, even in single-agent environments with supply constraints, there exists mechanisms with more complicated price discrimination schemes that significantly outperforms simple pricing-based mechanisms.

Proof of Theorem 2. Given the assumption that pricing-based mechanism $\mathcal{M}$ is approximately optimal for linear agents, i.e., $\mathcal{M}(\{P_i\}_{i\in N}) \geq 1/\gamma \cdot \text{EAR}(\{\bar{P}_i\}_{i\in N})$, it is sufficient to show that $\text{EAR}(\{\bar{P}_i\}_{i\in N}) \geq 1/\zeta \cdot \text{EAR}(\{R_i\}_{i\in N})$ when all non-linear agents are ζ-resemblant.

Let $\{q_i^\dagger\}_{i\in N} \in \text{EAF}$ 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 ζ-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,

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

5.1 Application of Reduction Framework

In this section, we will mainly focus on the application of two simple mechanisms, marginal revenue maximization and sequential posted pricing.

[13]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 sequential posted pricing mechanisms and the ex ante relaxation are well defined given the payoff curves, and Theorem 2 holds for payoff curves that does not correspond to any type distributions of the linear agents. Hence, we can refer to the linear agents analog even without the existence of the underlying type distributions.

Sequential Posted Pricing. We first introduce the formal definition of sequential posted pricing mechanisms in quantile space. This definition applies to both linear agents and non-linear agents.

Definition 5.1. *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 quantile corresponding to the per-unit prices to be offered to agents if the items are not sold out to previous agents.*[14]

In sequential posted pricing, agents are approached according to order $\{q_i\}_{i\in N}$ and receive options to purchase the item given personalized prices. In Appendix A.5, we discuss the extension of oblivious posted pricing where the seller only designs the personalized prices and the agents can arrive in arbitrary order.

Remark 5.2. *Every sequential posted pricing mechanism* $\mathcal{M} = (\{\hat{o}_i\}_{i\in N}, \{\hat{q}_i\}_{i\in N})$ *can be described as a pricing-based mechanism with threshold functions* $\{Q_i\}_{i\in N}$*, where* $Q_i(q_{-i}) = \hat{q}_i \cdot \prod_{j:\hat{o}_j<\hat{o}_i} \mathbf{1}[q_j > \hat{q}_j]$.

Note that according to Definition 5.1, the payoff of sequential posted pricing mechanism is uniquely determined by the price-posting payoff curves $\{P_i\}_{i\in N}$ of the agents. Therefore, we denote $\text{SPP}(\{P_i\}_{i\in N})$ as the optimal payoff among the class of sequential posted pricing mechanisms.

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

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* ζ*-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.,* $\text{SPP}(\{P_i\}_{i\in N}) \geq \frac{1}{\zeta}\cdot(1-\frac{1}{\sqrt{2\pi m}})\cdot \text{EAR}(\{R_i\}_{i\in N})$.

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 ζ measures how close the non-linear agents resembles linear agents, and the approximation factor becomes better when non-linear agents become more resemblant to linear agents. Corollary 1 implies that the approximation ratio will be close to 1 in large markets for non-linear agents that closely resembles linear agents. This implies that sequential posted pricing are approximately optimal for those non-linear

[14] In the sequential posted pricing mechanism, each agent may only get a lottery for winning the item. We assume that the lottery is realized immediately after each agent's purchase decision.

agents. The interpretation is that competition and simultaneous implementation are not salient features for payoff maximization even when agents have non-linear utilities.

Marginal Payoff Maximization. Bulow and Roberts (1989) introduce the marginal revenue maximization mechanism and show its revenue-optimality for linear agents. The marginal revenue maximization mechanism can be easily extended to other payoff objectives and we denote its extensions as the marginal payoff maximization mechanisms (MPM). Formally, the marginal payoff maximization mechanism is defined as the one with allocations that maximize the total marginal payoff

$$\sum_{i \in N} x_i \cdot P_i'(q_i)$$

for any profile of quantiles $\{q_i\}_{i \in N}$.[15] Note that the expected payoff of the marginal payoff maximization mechanisms only depend on the payoff curves, not the detailed utility structures of the agent. Therefore, we denote the payoff of MPM for agents with price-posting payoff curves $\{P_i\}_{i \in N}$ as $\text{MPM}(\{P_i\}_{i \in N})$.

Remark 5.3. *The marginal payoff mechanism can be described as a pricing-based mechanism with threshold functions $\{Q_i\}_{i \in N}$, where $Q_i(q_{-i}) = (P_i')^{(-1)}(\max_{j \neq i} P_j'(q_j))$ and $(P_i')^{(-1)}(\cdot)$ is the inverse function of $P_i'(\cdot)$.*

For non-linear agents, although the marginal payoff maximization mechanism is no longer optimal in general, by adopting the same argument as in Bulow and Roberts (1989), it is optimal among the subclass of pricing-based mechanism for any given set of payoff curves. Moreover, our reduction framework shows that it is approximately optimal for non-linear agents that resemble linear agents. Again by Yan (2011), the worst-case approximation factor of marginal payoff maximization to the ex ante payoff is at most $1/(1 - \frac{1}{\sqrt{2\pi m}})$ for selling m units to linear agents. Combining this with Theorem 2, 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 ζ-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., $\text{MPM}(\{P_i\}_{i \in N}) \geq \frac{1}{\zeta} \cdot (1 - \frac{1}{\sqrt{2\pi m}}) \cdot \text{EAR}(\{R_i\}_{i \in N})$.*

Note that although the worst-case approximation guarantee of the marginal payoff maximization mechanism coincides with sequential posted pricing when compared to the ex ante

[15] In the case that $P_i(q_i)$ is not concave in q_i and hence $P_i'(q_i)$ is not monotone in q_i, we can apply the ironing trick in Myerson (1981) by considering allocations that maximize the ironed marginal payoff $\sum_{i \in N} x_i \cdot \bar{P}_i'(q_i)$.

relaxation, the expected payoff from marginal payoff maximization always weakly exceeds the the expected payoff from sequential posted pricing.

Since the demands of risk-averse agents given per-unit prices do not depend on their risk preferences, the implementation of the marginal payoff maximization mechanisms for risk-averse agents coincides with that for linear agents. For example, when all agents have i.i.d. marginal valuation distributions, the marginal payoff maximization mechanism for risk-averse agents can be implemented as the VCG mechanism with an anonymous reserve price for all agents.

In general for non-linear agents such as private-budgeted agents, the implementation of the marginal payoff maximization mechanism (Algorithm 1) differs from that for linear agents. For example, when agents have budget constraints, both the allocation and payment of each agent may be randomized function of the opponents' bids in order for the marginal payoff maximization mechanisms to be budget feasible.

5.2 Resemblance of Revenue Maximization

In the previous section, we have provided a framework showing that pricing-based mechanisms are approximately optimal if the payoff curves of non-linear agents satisfy the resemblance property. This framework only has a bite if we can show that the resemblance property is indeed satisfied for canonical non-linear utilities. In this section, we show that for the objective of revenue maximization, both risk-averse agents and private-budget agents resemble linear agents under mild assumptions of the valuation distributions. In Appendix B.2, we show that those assumptions are necessary for pricing-based mechanisms to be approximately optimal even in single-agent settings without sale constraints. Since we will focus on the single-agent analysis in this section, we will drop the subscript representing the agent in all notations.

5.2.1 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 φ 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 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 such as uniform, exponential, and Gaussian distributions.

Definition 5.4. *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 5.4 implies that the tail probabilities in the valuation distribution is small, i.e., small than the exponential distribution (see, e.g., Allouah and Besbes, 2018).

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 B.1. At a high level, the proof consists of a two-step argument. First, we argue that for every mechanism, its expected revenue is upperbounded 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.

5.2.2 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.[16] However, such extreme price discrimination and full surplus extraction due to correlation structures are rarely observed in many practical applications relevant to agents with budgets. 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 weaker than

[16] We formalize this intuition with Example B.1 in Appendix B.2.

the MHR condition and is satisfied by many common distributions in practice, such as uniform, exponential, and Gaussian distributions.

Definition 5.5. *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.*

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

The proof of Proposition 2 relies on the technique of Lagrangian relaxation on the budget constraint developed in Alaei et al. (2013) and Feng and Hartline (2018), which is deferred to Appendix B.3. Combined with our reduction framework in Theorem 2, the interpretation of Proposition 2 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.

In Appendix B.3, 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.

Private Budget. In this section, we present the resemblance for private-budget agents.

Proposition 3. *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 B.4. At a high level, the proof consists of a two-step argument. First, we argue that every incentive compatible mechanism for a single private-budget agent can be interpreted as a *non-linear pricing scheme* τ, in which the agent chooses a lottery with winning probability z and pays $\tau(z)$. As a sanity check, posting a per-unit price p is a special case of non-linear pricing schemes with $\tau(z) = z \cdot p$. Intuitively, compared with general non-linear pricing schemes, the revenue loss of posting per-unit price p has two components: (i) it does not serve agents with values below p, and (ii) it does not apply price discrimination for agents with values above p. The second step of our argument introduces a decomposition of the ex ante optimal mechanism and shows that the revenue losses from both components are bounded within constant factors of the expected revenue from posted pricing given the independence assumption and the regularity assumption on the marginal value distributions.

We also obtain a constant upper bound on the resemblance property with mild assumption only on the marginal budget distribution. The details of this extension are provided in Appendix B.5.

6 Robust Optimality

In preceding sections, we showed that simple mechanisms such as marginal revenue maximization are approximately optimal for a wide range of non-linear utility functions. This section focus on a robust setting wherein the seller possesses only the knowledge of the buyers' demand functions, lacking detailed information regarding the non-linear utility function or the type distribution. We show that the marginal revenue maximization mechanism is robustly optimal under such environments.

Given any expected demand function $d(p)$ that is a non-increasing function,[17] let $\mathcal{P}(d)$ be the set of feasible agent utility model $(\mathcal{T}, \bar{F}, u)$ such that for any $p \geq 0$,

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

where $d^u(t, p)$ is the demand function induced by utility function u.

For any mechanism $\mathcal{M}$, let $\mathcal{M}(\{(\mathcal{T}_i, \bar{F}_i, u_i)\}_{i \in N})$ be the expected revenue of mechanism $\mathcal{M}$ given a profile of utility models $\{(\mathcal{T}_i, \bar{F}_i, u_i)\}_{i \in N}$. In this section, we assume that the seller can only observe the expected demand function $d_i(p)$ for all agent 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.

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

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

Theorem 3. *For any profile of non-increasing expected demand functions $\{d_i\}_{i \in N}$, if the seller does not observe the detailed profile of utility models, the marginal revenue maximiza-*

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

tion mechanism is the max-min revenue optimal mechanism, i.e.,

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

The interpretation of the robustness result is that if the seller is uncertain about the detailed utility models, the seller should not resort to complex lottery mechanisms to price discriminate the agents.

Note that in the robust environments, the seller does not know the detailed utility model. Therefore, to implement the marginal revenue maximization mechanism, we cannot directly apply the revelation principal to simply ask the agent to report her true type. Instead, we ask the agent to report her demand function given her private type and her utility function. For any agent $i \in N$, let P_i be the price posting revenue curve given demand function d_i. By the construction in Section 4, the expected revenue of marginal revenue maximization mechanism is

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

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

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

Combining the inequalities, Theorem 3 holds.

7 Conclusions and Extensions

This paper provides a general framework for generalizing results from linear agents to non-linear agents. The reduction framework relies on a novel resemblance property which characterizes the gap between optimal mechanism and posted pricing for the single agent problem with supply constraints. As the instantiations of the framework, we analyze the approximation bound for various mechanisms for various non-linear utility model (i.e., risk averse utility, budgeted utility) under the objective of both revenue-maximization. Moreover, the reduction framework allows us to show that the marginal revenue maximization mechanism is robustly optimal when the seller only observes the expected demand function of each agent. Next we discuss several important extensions of our framework.

7.1 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. However, for 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) show 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 exists 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 A.1.

Proposition 4. *An agent with 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 resemble linear agents for both welfare maximization and revenue maximization, then this agent resemble linear agents for any convex combination of the two objectives. This observation relies on the following lemma, with proof provided in Appendix A.2.

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

7.2 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 ζ_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 $\zeta' \geq \zeta \geq 1$, ζ-resemblant implies ζ'-resemblant.*

7.3 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 4 fails to construct an incentive compatible mechanism for agents with 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 5.1), the payoff guarantees for linear agents can be approximately extend to agents with non-expected utilities. The details of the reduction is analogous to Theorem 2 and hence is deferred to Appendix A.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.

7.4 Anonymous Pricing

A desirable property for the multi-agent setting is anonymity. This requires that the per-unit price posted to all agents are the same. However, the approximation guarantees of anonymous pricing for linear agents does not extend to non-linear agents with general payoff functions. 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 A.2 in Appendix A.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 extends 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.

For linear agents, Alaei 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. 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 A.4

Proposition 5. *Fix any set of (non-linear) agents with price-posting revenue curves $\{P_i\}_{i\in N}$ that are ζ-resemblant to their optimal revenue curves $\{R_i\}_{i\in N}$. If the price-posting revenue curves are concave, then anonymous pricing is a ζe-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 instantiation of the reduction framework in Proposition 5, we can show that agents are 1-resemblant and have concave price posting revenue curve when they have public budget and regular valuation distributions, and they are e-resemblant when they are risk averse and their valuation distributions satisfy the monotone hazard rate condition.

7.5 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 implication for the line of work focusing on the design of strategically simple mechanisms (e.g., Chung and Ely, 2007; Li, 2017; 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.

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.

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.

Barlow, R. E. and Marshall, A. W. (1965). Tables of bounds for distributions with monotone hazard rate. *Journal of the American Statistical Association*, 60(311):872–890.

Border, K. C. (1991). Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59(4):1175–1187.

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.

Cheng, Y., Gravin, N., Munagala, K., and Wang, K. (2018). A simple mechanism for a budget-constrained buyer. In *International Conference on Web and Internet Economics*, pages 96–110. Springer.

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.

Gershkov, A., Moldovanu, B., Strack, P., and Zhang, M. (2022). Optimal auctions: Non-expected utility and constant risk aversion. *The review of economic studies*, 89(5):2630–2662.

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.

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

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87.

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.

A Details for Various Extensions

A.1 Welfare Maximization

Proposition 4. *An agent with 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 4 utilizes the price decomposition technique from Abrams (2006) and extends it for welfare analysis.

Fix an arbitrary ex ante constraint q, denote EX as the q ex ante welfare-optimal mechanism, and **Payoff**[EX] as its welfare. We want to decompose EX into two mechanisms $\text{EX}^{\dagger}$ and $\text{EX}^{\ddagger}$ 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 $\text{EX}^{\dagger}$ contains per-unit prices at most the market clearing price, while mechanism $\text{EX}^{\ddagger}$ contains per-unit prices at least the market clearing price. Both mechanisms $\text{EX}^{\dagger}$ and $\text{EX}^{\ddagger}$ satisfy the ex ante constraint q, and the sum of their welfare upper bounds the original ex ante mechanism EX, i.e., $\mathbf{Payoff}[\text{EX}] \leq \mathbf{Payoff}\left[\text{EX}^{\dagger}\right] + \mathbf{Payoff}\left[\text{EX}^{\ddagger}\right]$.

To construct $\text{EX}^{\dagger}$ and $\text{EX}^{\ddagger}$ 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 A.1. *An* allocation-payment function $\tau : [0, 1] \to \mathbb{R}_+$ *is a mapping from the allocation x to the payment p.*

Lemma 5. *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 ∞ 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 5 only characterizes a necessary condition for private-budget utility.[18] 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 3. For agent with budget w, let τ_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 6. *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., $P(q) \geq \textbf{Payoff}[\text{EX}^\dagger]$.*

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

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

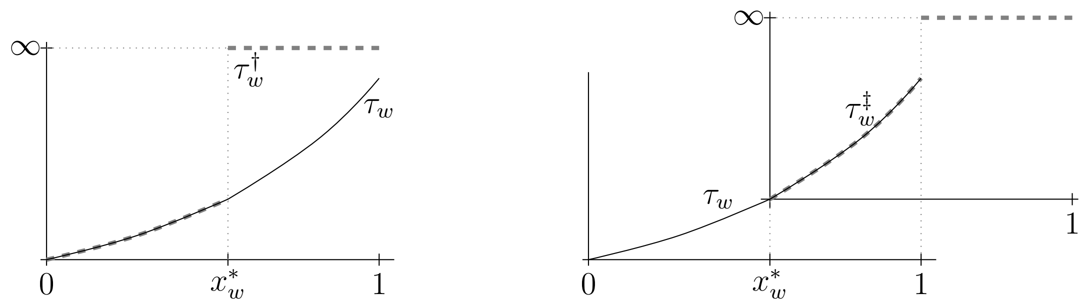

Figure 3: Depicted are allocation-payment function decomposition. The black lines in both figures are the allocation-payment function τ_w in ex ante optimal mechanism EX; the gray dashed lines are the allocation-payment function $\tau_w^\dagger$ and $\tau_w^\ddagger$ in $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$, respectively.

are the same in $\mathrm{EX}^\dagger$ 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.

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., $P(q) \geq \mathbf{Payoff}\left[\mathrm{EX}^\dagger\right]$. □

Lemma 7. *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.,* $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big]$.

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., $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big]$. □

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 4. Combining Lemma 6 and 7, for any quantile q, we have

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

A.2 Convex Combination of Two Payoff Functions

Lemma 3. *If an agent is ζ-resemblant for objective 1 and ζ'-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 $\mathbf{Payoff}_1[\mathrm{EX}]$ be the contribution of objective 1 given mechanism EX and $\mathbf{Payoff}_2[\mathrm{EX}]$ be the contribution of objective 2 given mechanism EX. Let $\mathbf{Payoff}[\mathrm{EX}] = \alpha \cdot \mathbf{Payoff}_1[\mathrm{EX}] + (1-\alpha) \cdot \mathbf{Payoff}_2[\mathrm{EX}]$ 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

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

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

A.3 Reduction Framework for Non-expected Utilities

We first formally introduce the model of non-expected utilities. In Appendix A.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 A.3.2.

Non-expected Utilities. For any agent with type t, her utility for a distribution π 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.[19] We also say such an agent is ζ-resemblant if her ironed price-posting payoff curve $\bar{P}$ is ζ-*resemblant* to her optimal payoff curve R.

A.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 ζ-resemblant to their optimal payoff curves $\{R_i\}_{i \in N}$. If there exists a sequential posted pricing mechanism that is a γ-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* $\mathrm{SPP}(\{P_i\}_{i \in N}) \geq 1/\gamma\zeta \cdot \mathrm{EAR}(\{R_i\}_{i \in N})$.

[19] 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.

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 ζ-resemblant to the optimal payoff curves $\{R_i\}_{i\in N}$, 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. ☐

A.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).[20] 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

[20] 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.

the ex ante relaxation for the objective of both welfare and revenue maximization.

Lemma 8 (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 $x \prec \hat{x}$,*[21] *we have*

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

Proposition 6. *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 $q \leq \hat{q}$ 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 8, 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. □

Proposition 7. *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 6, 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. □

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

A.4 Anonymous Pricing

Example A.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 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 5. *Fix any set of (non-linear) agents with price-posting revenue curves $\{P_i\}_{i\in N}$ that are ζ-resemblant to their optimal revenue curves $\{R_i\}_{i\in N}$. If the price-posting revenue curves are concave, then anonymous pricing is a ζe-approximation to the ex ante relaxation on the optimal revenue curves, i.e.,* $\text{AP}(\{P_i\}_{i\in N}) \geq {}^1\!/\!{\zeta e} \cdot \text{EAR}(\{R_i\}_{i\in N})$.

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 ζ-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

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

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

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

□

A.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 A.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 $\text{OPP}(\{P_i\}_{i\in N}, \{q_i\}_{i\in N})$, and the optimal payoff for the oblivious quantile pricing mechanism is

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

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 ζ-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 γ-approximation to the ex ante relaxation for linear agents analog with price-posting payoff curves $\{P_i\}_{i\in N}$, i.e., $\text{OPP}(\{P_i\}_{i\in N}, \{q_i\}_{i\in N}) \geq {}^{1}/_{\gamma} \cdot \text{EAR}(\{\bar{P}_i\}_{i\in N})$, then this mechanism is also a $\gamma\zeta$-approximation to the ex ante relaxation for non-linear agents, i.e., $\text{OPP}(\{P_i\}_{i\in N}, \{q_i\}_{i\in N}) \geq {}^{1}/_{\gamma\zeta} \cdot \text{EAR}(\{R_i\}_{i\in N})$.*

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).

A.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 $\text{EAF}(\mathcal{X})$. The optimal ex ante payoff given a specific collection of payoff curves $\{R_i\}_{i\in N}$ and feasibility constraint $\mathcal{X}$ is

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

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

$\text{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}, \mathcal{X})$.[22]

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 2) 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.

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 γ-approximation to the ex ante relaxation, i.e., $\text{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}, \mathcal{X}) \geq {}^{1}/_{\gamma} \cdot \text{EAR}(\{\bar{P}_i\}_{i\in N}, \mathcal{X})$ 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 Missing Proofs for Resemblance of Revenue Maximization

B.1 Risk Aversion

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 9. *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 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 φ 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 φ is increasing with $\varphi(0) = 0$, we obtain $\mathbf{E}[vx(v)] \geq \mathbf{E}[p(v)]$ as desired. ☐

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

Lemma 10. *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. □

Lemma 11 (Hartline and Roughgarden, 2009)**.** *For any MHR value distribution F, it is guaranteed that* $\max_{p \in \mathbb{R}_+} p(1 - F(p)) \geq e \cdot \mathbf{E}_{v \in F}[v]$.

As a sanity check, Lemma 11 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 $v(q) \triangleq F^{-1}(q)$ for each quantile $q \in [0, 1]$. Invoking Lemma 9, 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 10 and 11, 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. □

B.2 Necessity of Assumptions

Example B.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 ϵ is incentive compatible and has revenue $O(\ln h)$. However, the revenue of the posted pricing is $O(1)$.

Example B.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.

B.3 Public Budget

Proposition 2. *A public-budget agent is 1-resemblant if her marginal value distribution is regular.*

Proof. 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 λ^* 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 12 (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 12 (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 of values $V(q)$ where $\hat{P}_{\lambda^*}(q) > P_{\lambda^*}(q)$.

For an agent with a general valuation distribution, resemblance follows from the charac-

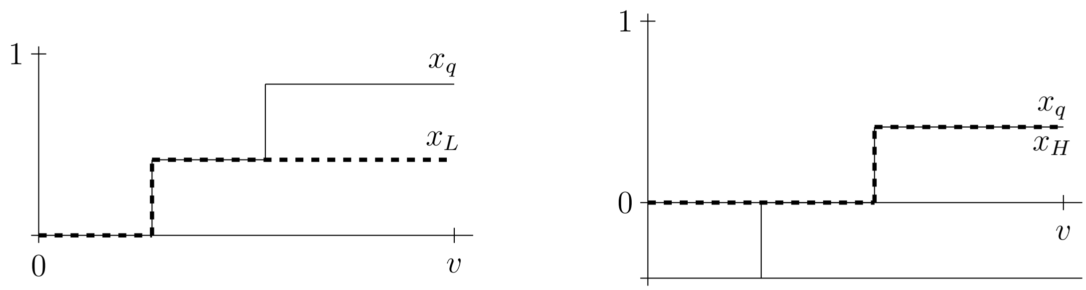

Figure 4: The thin solid line is the allocation rule for the optimal ex ante mechanism. The thick dashed line on the left side is the allocation of the decomposed mechanism with lower price, while the thick dashed line on the right side is the allocation of the decomposed mechanism with higher price.

terization 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 13 (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 8. *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 13, 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 4. Note that both allocation x_L and x_H are (randomized) price-posting allocation rules, and neither allocation violates the allocation constraint q. Thus,

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

B.4 Private Budget

Proposition 3. *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 $\text{EX}^\dagger$ and $\text{EX}^\ddagger$ according to the market clearing price p^q. Intuitively, the per-unit prices in $\text{EX}^\dagger$ 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 is specified in Appendix A.1, and we will bound the revenue from those two mechanisms separately.

Lemma 14. *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., $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^\dagger\big]$.*

Proof. The ex ante allocation of $\mathrm{EX}^\dagger$ is at most the ex ante allocation of 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 15 and 16 are provided at the end of this section.

Lemma 15. *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., $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^\ddagger\big]$.*

Intuitively, Lemma 15 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 16. *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., $2P(q^\dagger) \geq \mathbf{Payoff}\big[\mathrm{EX}^\ddagger\big]$.*

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 5, the revenue from the random pricing

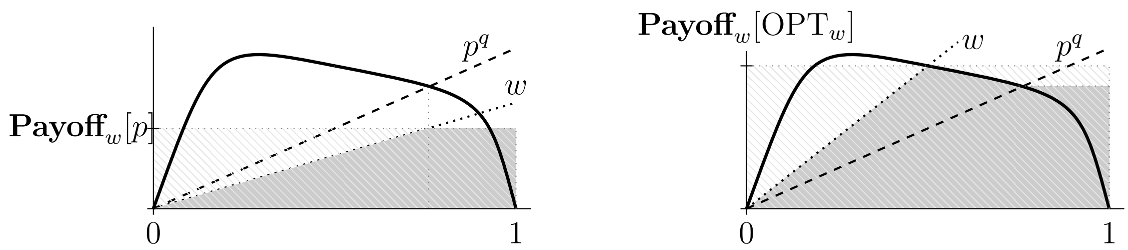

Figure 5: In the geometric proof of Lemma 16, the upper bound on the expected revenue of $\mathrm{EX}^\ddagger$ ($\mathbf{Payoff}_w[p]$ and $\mathbf{Payoff}_w[\mathrm{OPT}_w]$ on the left and right, respectively) is the area of the light gray striped rectangle and the revenue from posting random price $\boldsymbol{p}$ 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., $w \geq p^q$).

$\boldsymbol{p}$ 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.

Proof of Proposition 3. Fix any ex ante constraint q. If the market clearing price $p^q = {}^{P(q)}/_{q}$ is at least the monopoly reserve, Lemma 14 and Lemma 15 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 $q^\dagger = \operatorname{argmax}_{q' \leq q} P(q')$, Lemma 14 and Lemma 16 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 15. 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 16. 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 5, 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 5 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, $\mathbf{Payoff}_w\big[\mathrm{EX}^{\ddagger}\big] \leq \mathbf{Payoff}_w[p^q] = w\,(1 - F(p^q))$ 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 5.

The right side of Figure 5 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 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, $\mathbf{Payoff}_w\big[\mathrm{EX}^{\ddagger}\big] \leq \mathbf{Payoff}_w[\mathrm{OPT}_w]$ as the latter optimizes with relaxed constraints of the former. Laffont and Robert (1996) show that 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 5.

Next, we will show that the revenue of posting the random price $\boldsymbol{p}$ is the grey shaded areas illustrated in Figure 5 (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. $\square$

B.5 Private Budget: Assumption on the Budget Distribution

We consider the assumption that the budget exceeds its expectation with constant probability at least $1/\kappa$. This assumption on budget distribution is also studied in Cheng et al. (2018). Notice that a commonly adopted distribution assumption, monotone hazard rate, is a special case of it with $\kappa = e$ (cf. Barlow and Marshall, 1965).

Proposition 9. *A single agent with private-budget utility has an ironed price-posting revenue curve $\bar{P}$ that is $(1 + 3\kappa - 1/\kappa)$-resemblant to her optimal revenue curve R, if her value and budget are independently distributed, and the probability the budget exceeds its expectation is $1/\kappa$.*

Let w^* denote the expected budget of the agent. For any ex ante constraint q, denote EX as the q ex ante revenue optimal mechanism.

The high level idea here is similar to the analysis for welfare, i.e., the price decomposition technique. Consider the decomposition of EX into three mechanisms $\text{EX}^\dagger$, $\text{EX}^\S$ and $\text{EX}^\ddagger$ such that mechanism $\text{EX}^\dagger$ contains per-unit prices at most the market clearing price, mechanism $\text{EX}^\ddagger$ contains per-unit prices at least the expected budget, while mechanism $\text{EX}^\S$ contains per-unit prices between the market clearing price and the expected budget. All mechanisms satisfy the ex ante constraint q, and the sum of their welfare is upper bounded by the welfare of the original ex ante mechanism EX, i.e., $\textbf{Payoff}[\text{EX}] \leq \textbf{Payoff}\big[\text{EX}^\dagger\big] + \textbf{Payoff}\big[\text{EX}^\S\big] + \textbf{Payoff}\big[\text{EX}^\ddagger\big]$. Note that in the special case where the market clearing price is larger than the expected budget, i.e., $p^q > w^*$, $\text{EX}^\S$ does not exist and mechanism EX is decomposed into $\text{EX}^\dagger$ and $\text{EX}^\ddagger$.

We construct the allocation-payment functions $\tau_w^\dagger$, $\tau_w^\ddagger$ and $\tau_w^\S$ for $\text{EX}^\dagger$, $\text{EX}^\ddagger$, and $\text{EX}^\S$ respectively. For each budget w, let τ_w be the allocation-payment function for types with budget w in mechanism EX, and x_w^* be the utility maximization allocation for the agent with value and budget equal to the market clearing price p^q, i.e., $x_w^* = \text{argmax}\{x : \tau_w'(x) \leq p^q\}$.

Let $x^\sharp_w$ be the utility maximization allocation for the agent with value and budget equal to the expected budget w^*, i.e., $x^\sharp_w = \text{argmax}\{x : \tau'_w(x) \leq w^*\}$. Then the allocation-payment functions $\tau^\dagger_w$, $\tau^\ddagger_w$ and $\tau^\S_w$ are defined respectively as follows,

$$\tau^\dagger_w(x) = \begin{cases} \tau_w(x) & \text{if } x \leq x^*_w, \\ \infty & \text{otherwise;} \end{cases}$$

$$\tau^\S_w(x) = \begin{cases} \tau_w(x^*_w + x) - \tau_w(x^*_w) & \text{if } x \leq x^\sharp_w - x^*_w, \\ \infty & \text{otherwise;} \end{cases}$$

$$\tau^\ddagger_w(x) = \begin{cases} \tau_w(x^\sharp_w + x) - \tau_w(x^\sharp_w) & \text{if } x \leq 1 - x^\sharp_w, \\ \infty & \text{otherwise.} \end{cases}$$

The revenue contribution from $\text{EX}^\dagger$ is bounded in Lemma 14. Next we illustrate how to bound the revenue from $\text{EX}^\ddagger$ and $\text{EX}^\S$ respectively using the revenue from price-posting.

Lemma 17. *For a single agent with private-budget utility, independently distributed value and budget, for any quantile q, there exists $q^\dagger \in [0, q]$ such that $(1 + \kappa - 1/\kappa) \cdot P(q^\dagger) \geq$ $\mathbf{Payoff}\big[\text{EX}^\ddagger\big]$.*

Proof. Let w^* be the expected budget and let $\bar{p} = \max\{w^*, p^q\}$. Let $\bar{q}$ be the quantile corresponding to value $\bar{p}$ and let $q^\dagger = \text{argmax}_{q' \leq q} P(q')$. Thus $P(\bar{q}) \leq P(q^\dagger)$. Moreover, by the construction of the decomposition, the per-unit price in $\text{EX}^\ddagger$ is larger than $\bar{p}$. Similar to the proof of Lemma 15, we only consider the types with value at least $\bar{p}$.

Let $\mathbf{Payoff}_w\big[\tau^\ddagger_w\big]$ be the expected revenue of providing the allocation-payment function $\tau^\ddagger_w$ in $\text{EX}^\ddagger$ to the types with budget w; and let $\mathbf{Payoff}_w[p]$ be the expected revenue of posting price p to the types with budget w. The following three facts allow comparison of $\mathbf{Payoff}\big[\text{EX}^\ddagger\big]$ to $P(q^\dagger)$:

(a) Posting the price $\bar{p}$ makes the budget constraints bind for the types with budget at most w^*, so $\mathbf{Payoff}_w\big[\tau^\ddagger_w\big] \leq \mathbf{Payoff}_w[\bar{p}]$ for all $w \leq w^*$.

(b) $\mathbf{Payoff}_w\big[\tau^\ddagger_w\big] \leq \frac{w}{w^*}\mathbf{Payoff}_{w^*}\big[\tau^\ddagger_w\big]$ for all $w \geq w^*$. This is because if the type (v, w^*) pays her budget w^* (i.e., the budget constraint binds), her payment is a (w/w^*)-approximation to the payment from the type (v, w), since the type (v, w) pays at most w. Moreover, if the type (v, w^*) pays less than her budget w^* (i.e., the unit-demand constraint binds, or the value binds), her allocation is equal to the allocation from the type (v, w) for $w \geq w^*$. Hence, their payments are the same.

(c) Since the revenue of posting price $\bar{p}$ to an agent with budget w^* is at most the revenue to an agent with budget $w > w^*$; with the assumption that budgets exceed the expectation

w^* with probability at least $1/\kappa$, it implies that

$$\mathbf{Payoff}_{w^*}[\bar{p}] \cdot \frac{1}{\kappa} \leq \mathbf{E}\Big[\mathbf{Payoff}_w[\bar{p}] \;\Big|\; w \geq w^*\Big] \cdot \mathbf{Pr}[w \geq w^*] \leq P(\bar{q}).$$

We upper bound the revenue of $\mathrm{EX}^\ddagger$ as follows,

$$\begin{aligned}
\mathbf{Payoff}\big[\mathrm{EX}^\ddagger\big] &= \int_{\underline{w}}^{w^*} \mathbf{Payoff}_w\big[\tau_w^\ddagger\big]\, dG(w) + \int_{w^*}^{\bar{w}} \mathbf{Payoff}_w\big[\tau_w^\ddagger\big]\, dG(w) \\
&\leq \int_{\underline{w}}^{w^*} \mathbf{Payoff}_w[\bar{p}]\, dG(w) + \int_{w^*}^{\bar{w}} \frac{w}{w^*} \mathbf{Payoff}_{w^*}\big[\tau_w^\ddagger\big]\, dG(w) \\
&\leq (1 - \frac{1}{\kappa}) P(\bar{q}) + \frac{\int_{w^*}^{\bar{w}} w dG(w)}{w^*} \mathbf{Payoff}_{w^*}[\bar{p}] \\
&\leq (1 - \frac{1}{\kappa}) P(\bar{q}) + \mathbf{Payoff}_{w^*}[\bar{p}] \leq (1 + \kappa - \frac{1}{\kappa}) P(q^\dagger)
\end{aligned}$$

where the first inequality is due to facts (a) and (b); in the second inequality, the first term is due to $\mathbf{Pr}[w \leq w^*] \leq 1 - 1/\kappa$, the revenue $\mathbf{Payoff}_w[\bar{p}]$ is monotone increasing in w, and by definition $\int_{\underline{w}}^{\bar{w}} \mathbf{Payoff}_w[\bar{p}]\, dG(w) = P(\bar{q})$, and the second term is due to fact (a); and the last inequality is due to $P(\bar{q}) \leq P(q^\dagger)$ and fact (c). □

Lemma 18. *For a single agent with private-budget utility, independently distributed value and budget, when $p^q \leq w^*$, there exists $q^\dagger \leq q$ such that the price-posting revenue from $q^\dagger$ is a $(2\kappa - 1)$-approximation to the revenue from $\mathrm{EX}^\S$, i.e., $(2\kappa - 1) P(q^\dagger) \geq \mathbf{Payoff}\big[\mathrm{EX}^\S\big]$.*

Proof. Let $q^\dagger = \operatorname{argmax}_{q' \leq q} P(q')$. Suppose the support of the budget distribution is from $[\underline{w}, \bar{w}]$. Let $\tilde{p}$ be the price larger than the market clearing price p^q and smaller than the expected budget w^* that maximizes revenue without the budget constraint. Consider the following calculation with justification below.

$$\begin{aligned}
\mathbf{Payoff}\big[\mathrm{EX}^\S\big] &= \int_{\underline{w}}^{w^*} \mathbf{Payoff}_w\big[\tau_w^\S\big]\, dG(w) + \int_{w^*}^{\bar{w}} \mathbf{Payoff}_w\big[\tau_w^\S\big]\, dG(w) \\
&\overset{(a)}{\leq} \int_{\underline{w}}^{w^*} \mathbf{Payoff}_{w^*}\big[\tau_w^\S\big]\, dG(w) + \int_{w^*}^{\bar{w}} \frac{w}{w^*} \mathbf{Payoff}_{w^*}\big[\tau_w^\S\big]\, dG(w) \\
&\overset{(b)}{\leq} \int_{\underline{w}}^{w^*} \mathbf{Payoff}_{w^*}[\tilde{p}]\, dG(w) + \int_{w^*}^{\bar{w}} \frac{w}{w^*} \mathbf{Payoff}_{w^*}[\tilde{p}]\, dG(w) \\
&\overset{(c)}{\leq} (2 - \frac{1}{\kappa})\, \mathbf{Payoff}_{w^*}[\tilde{p}] \\
&\overset{(d)}{\leq} (2\kappa - 1)\, \mathbf{Payoff}[\tilde{p}] \overset{(e)}{\leq} (2\kappa - 1)\, P(q^\dagger).
\end{aligned}$$

Inequality (a) holds because given the allocation payment function $\tau^{\S}_{w}$, the revenue only increases if we increase the budget to w^*, i.e., $\textbf{Payoff}_w\left[\tau^{\S}_{w}\right] \leq \textbf{Payoff}_{w^*}\left[\tau^{\S}_{w}\right]$ for any $w \leq w^*$. Moreover, for any $w > w^*$, given the allocation payment function $\tau^{\S}_{w}$, the revenue is either the same for budget w and w^*, or the budget binds for agent with expected budget w^*. Since the revenue from agent with budget w is at most w, we know that $\textbf{Payoff}_w\left[\tau^{\S}_{w}\right] \leq {}^{w}/_{w^*} \cdot \textbf{Payoff}_{w^*}\left[\tau^{\S}_{w}\right]$. Note that for allocation payment rule $\tau^{\S}_{w}$, per-unit prices are larger than the market clearing price p^q and smaller than the expected budget w^*, and budget does not bind for agents with budget w^*. Therefore, by definition, the optimal per-unit price in this range is $\tilde{p}$, $\textbf{Payoff}_{w^*}\left[\tau^{\S}_{w}\right] \leq \textbf{Payoff}_{w^*}[\tilde{p}]$ and inequality (b) holds. Inequality (c) holds because $\int_{\underline{w}}^{w^*} dG(w) \leq 1 - {}^{1}/_{\kappa}$ by the assumption that the probability the budget exceeds its expectation is at least κ, and $\int_{w^*}^{\bar{w}} \frac{w}{w^*} dG(w) \leq 1$. Inequality (d) holds because $\textbf{Payoff}_{w^*}[\tilde{p}] \leq \kappa \cdot \textbf{Payoff}[\tilde{p}]$ for any randomized prices $\tilde{p}$ according to Cheng et al. (2018). Inequality (e) holds by the definition of the price-posting revenue curve P and quantile $q^\dagger$, the fact that price $\tilde{p}$ is larger than the market clearing price p^q. □

Proof of Proposition 9. Let $q^\dagger = \text{argmax}_{q' \leq q} P(q')$. Combining Lemma 14, 17 and 18, we have

$$\textbf{Payoff}[\text{EX}] \leq \textbf{Payoff}\left[\text{EX}^\dagger\right] + \textbf{Payoff}\left[\text{EX}^\ddagger\right] + \textbf{Payoff}\left[\text{EX}^\S\right] \leq (1 + 3\kappa - {}^{1}/_{\kappa})\, P(q^\dagger). \quad \square$$

C Tie-breaking Rules

In Algorithm 1, it is possible that based on the tie-breaking rule of 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 uniform distribution in $[0, 1]$. Therefore, tie occurs 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 holds 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.