Simple Mechanisms for Agents with Non-linear Utilities*

Yiding Feng†, Jason D. Hartline*, and Yingkai Li*

†Microsoft Research New England.

*Department of Computer Science, Northwestern University.

yidingfeng@microsoft.com, yingkai.li@u.northwestern.edu, hartline@northwestern.edu

Abstract

We show that economic conclusions derived from Bulow and Roberts (1989) for linear utility models 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 we show that the approximation of mechanisms for agents with linear utilities approximately extend for agents with non-linear utilities.

We illustrate the framework for the objectives of revenue and welfare on non-linear models that include agents with budget constraints, agents with risk aversion, and agents with endogenous valuations. and objectives of revenue and welfare. We derive bounds on how much these models resemble the linear utility model and combine these bounds with well-studied approximation results for linear utility models. We conclude that simple mechanisms are approximately optimal for these non-linear agent models.

*This paper combines "Optimal Auctions vs. Anonymous Pricing: Beyond Linear Utility", which is presented at the 20th ACM Conference on Economics and Computation, and "Simple Mechanisms for Non-linear Agents", which is currently available at https://arxiv.org/abs/2003.00545.

1 Introduction

This paper identifies conditions under which the conclusions derived from the simple economics of optimal auctions (e.g., Bulow and Roberts, 1989) approximately extend from linear utility models to general (i.e., non-linear) utility models. For context, optimal mechanisms for agents with non-linear utilities are not simple and therefore difficult to understand precisely. For example, the single-item auction for a single agent with a private budget constraint admits no closed-form characterization (Che and Gale, 2000).[1]

There are extensive studies of simple mechanisms with approximation guarantees in the classical linear utility model of mechanism design. Example 1: Bulow and Roberts (1989) show that the marginal revenue maximization mechanism is revenue optimal. Example 2: Yan (2011) shows that sequential posted pricings, which arrange the agents in an order and offer while-supplies-last posted prices, guarantee an $e/(e-1)$-approximation, i.e., the best order and prices achieves at least 63.2% of the optimal auction revenue. Example 3: Jin et al. (2019) show that when the agents' (non-identical) value distributions satisfy a concavity property, a.k.a., "regular distributions", posting an anonymous price guarantees a 2.62-approximation.

Approximation results allow qualitative conclusions about drivers of good economic outcomes. From Example 1, we see that the drivers of classical microeconomics and auction theory are closely connected. From Example 2, we can conclude that simultaneity and competition are not necessary drivers for revenue maximization. From Example 3, we conclude that, even in asymmetric environments, discrimination across different agents is not a necessary driver of revenue maximization. (In contrast to Example 3, asymmetric bimodal distributions, which do not satisfy the regularity property, generally require discrimination to achieve close to the optimal revenue.) See the survey of Hartline (2012) for detailed discussion of the method of approximation in economics.

We generalize these approximation results from linear agents to non-linear agents.[2] From this generalization, not only do we observe that the main drivers of good mech-

[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".

anisms are similar for non-linear agents, but also that non-linearity itself is not a main concern that necessitates specialized mechanism designs (beyond the approach of our generalization).

Bulow and Roberts (1989), as later interpreted by Alaei et al. (2013), show that to design optimal mechanisms for linear agents, it is without loss to restrict attention to *pricing-based mechanisms*, i.e., mechanisms where the menu offered to each agent is equivalent to a distribution over posted prices. The multi-agent mechanism design problem can be decomposed as single-agent mechanism design problems through the reduced-form approach of Border (1991). 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. Thus, every mechanism for linear agents is equivalent to a pricing-based mechanism.

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. 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 not generally optimal. Nonetheless, we 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. We introduce a reduction framework. 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 transforms the aforementioned pricing-based mechanism for linear agents into an analogous pricing-based mechanism for the non-linear agents. The non-linear agent mechanism guarantees a $\beta\zeta$-approximation bound.

The reduction framework can be combined with approximation results for linear agents to show that simple mechanisms such as marginal revenue maximization, sequential posted pricing, and anonymous pricing are approximately optimal for non-linear agents that resemble linear agents, and the economic lessons (e.g., non-cruciality

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

of simultaneity, competition, discrimination) derived from those mechanisms for linear agents can be lifted to non-linear agents (see Examples 1–3, previously). As an example, agents with independent private budget and regular valuation distribution are 3-resemblant of linear agents, which implies that the approximation of sequential posted pricing for such non-linear agents is $3e/(e-1)$.

The paper characterizes broad families of non-linear agents that are ζ-resemblant for small constant factors ζ (e.g., agents with independent private budget and regular valuation distribution) and families that are not (e.g., agents whose budget and value are correlated). For non-linear agents that are ζ-resemblant, pricing-based mechanisms are approximately optimal wherever they 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.

Our reduction framework can be applied more broadly for non-linear agents beyond the expected utility theory with the restriction to *posted pricing mechanisms*[4] (e.g., sequential posted pricing, anonymous pricing). For instance, non-linear agents with endogenous valuations (Gershkov et al., 2021b) – which do not satisfy expected utility theory – are 1-resemblant under the regularity assumption. Thus, for such agents, sequential posted pricing is approximately optimal and the economic lessons from previous discussions generalize.

1.1 Discussion of Our Results

We define a notion of single-agent approximation by price-posting (see next paragraph) and show that, for non-linear agents that satisfy this definition and the von Neumann-Morgenstern expected utility representation, approximately optimal multi-agent pricing-based mechanisms can be derived from the analogous mechanisms for linear agents. This reduction framework is general – it can be applied to any downward-closed feasibility constraint (e.g., single-item, multi-unit, matroid) and common objectives (e.g., revenue, welfare, or their convex combination) and thus allows many known approximation mechanisms for linear agents to be lifted to non-

[4]Posted pricing mechanisms are pricing-based mechanisms where prices posted to each agent do not depend on actions of other agents.

linear agent environments. The approximation factors we obtain are the product of the single-agent approximation factor of price-posting for non-linear utilities and the approximation factor of the multi-agent mechanisms for linear utilities. Additionally, with the restriction to posted pricing mechanisms, our reduction framework is applicable to non-linear agents without the expected utility presentation.

The single-agent price-posting approximation that governs our reduction is defined as follows. The literature on revenue optimal mechanism design for a single agent under the ex ante constraint defines the *revenue curve* (cf. Bulow and Roberts, 1989) as follows. Fixing any family of mechanisms and a single agent, the revenue curve is a mapping from an ex ante allocation constraint $q \in [0, 1]$ to the revenue of the optimal mechanism in the family that sells the item with the given ex ante probability q. Specifically, the *price-posting revenue curve* is generated by fixing mechanism class to all (single-agent) posted pricing mechanisms, i.e., posting a per-unit price;[5] and the *optimal revenue curve* is generated by allowing all possible mechanisms.[6] In this paper we consider general objectives and general payoff curves that correspond to these objective. For linear agents, the optimal payoff curve is equivalent to the concave hull of the price-posting payoff curve. Motivated by this equivalence, the price-posting approximation for non-linear agents that governs our reduction is the *closeness* between the concave hull of the price-posting payoff curve and the optimal payoff curve. Namely, we say a non-linear agent is ζ-resemblant if price-posting is a ζ-approximation to the (single-agent) optimal mechanism for all ex ante constraints.

It is not hard to invent pathological non-linear agents that do not resemble linear agents. Nonetheless, in our study of three canonical non-linear utility models (i.e., budgeted utility, capacitated utility (i.e., a specific form of risk aversion), and endogenous valuation utility), ζ-resemblance is bounded by small constants for both welfare maximization (Section 5) and revenue maximization (Section 6) problems. See Tables 1 and 2 for summary of the ζ-resemblance results.

• *Budgeted Utility.* We show several of constant-factor resemblance results (i.e., single-agent approximation by price-posting) for public or private budget utility. An agent with independently distributed value and private budget resembles a linear

[5] Given per-unit price p, an agent can purchase any lottery with wining probability $q \in [0, 1]$.

[6] For example, in the revenue maximization problem for a single agent with independent private budget, when the agent's valuation distribution satisfies the decreasing marginal property, the optimal mechanism is not posting a per-unit price, but a menu of lotteries where the lottery with higher winning probability has lower per-unit price (Che and Gale, 2000).

Table 1: Summary of results for ζ-resemblant in revenue maximization problems under the assumption that the valuation distribution F is regular, i.e., $v - \frac{1-F(v)}{f(v)}$ is non-decreasing in v.

	public budget	independent private budget	risk averse, i.e., capacitated utility support $[0, \bar{v}]$, capacity at least C	endogenous valuation
ζ-resemblant	1	3	$2 + \ln \bar{v}/C$	1

Table 2: Summary of results for ζ-resemblant in welfare maximization problems. The public budgeted utility can be thought as a special case of independent private budgeted utility.

	independent private budget	risk averse	endogenous valuation
ζ-resemblant	2	1	1

agent as follows.

For welfare-maximization problems, we identify a constant bound on the closeness between the welfare curves without any assumption on the valuation or budget distributions. For revenue-maximization we show the budgeted agent resembles a linear agent under standard assumptions on the distributions of value and budget. We also construct examples showing the necessity of our assumptions to guarantee the ζ-resemblance for constant ζ.

• *Risk Averse Utility.* It is standard to model risk averse utility as a concave function that maps the agents' wealth to utility. This risk-aversion does not impose challenges in welfare maximization problems since both the optimal mechanism (e.g., VCG mechanism) and the simple price posting mechanisms are deterministic, and agents behave as if they are linear agents. However, for revenue maximization problems, this introduces a non-linearity into the incentive constraints of the agents which in most cases makes mechanism design analytically intractable. In this paper, we restrict attention to the specific form of risk aversion studied in Fu et al. (2013), which is called *capacitated utility*. In particular, the capacitated utility model is parameterized by a *capacity* C, and the utility of an agent in such utility model is the minimum between her wealth and the capacity C. We show that capacitated utility agents are $(2 + \ln {\bar{v}}/{C})$-resemblant where $\bar{v}$ is the upperbound of the support of the valuation distribution. This dependence on $\bar{v}$ and C is tight in worst case.

• *Endogenous Valuation Utility.* In this model, agents can take costly actions to boost their valuations for winning the item in the auction before their interaction with the mechanism. We follow the formalization of the model in Gershkov et al. (2021b), where the authors show that it is equivalent to consider agents with utility linear in payments and convex in the allocation probability. This utility model does not satisfy the expected utility characterization. Gershkov et al. (2021b) show that under regularity conditions, price posting is optimal for the single-agent revenue maximization problem without ex ante constraint. We extend their results to both welfare and revenue maximization problems, and show that price posting is optimal for any ex ante constraint $q \in [0, 1]$, i.e., agents with endogenous valuation are 1-resemblant.

Our resemblance results can be generalized to any convex combination of welfare and revenue as the objective function. For example, if an agent is ζ_1-resemblant for welfare maximization and ζ_2-resemblant for revenue maximization, then this agent is $(\zeta_1 + \zeta_2)$-resemblant for any convex combination of welfare and revenue. This generalization result does not rely on the utility model of the agents or their type distributions.

Our analyses and results of the closeness between the concave hull of the price-posting payoff curve and optimal payoff curve are interesting independently of our reduction framework. The setting of our single-agent analysis with an ex ante constraint is equivalent to the mechanism design problem for a continuum of i.i.d. (non-linear) agents with unit-demand and limited supply. A similar setting has been studied in Richter (2019), who shows that a posted pricing mechanism is optimal in the continuum model for budgeted agents with regular and decreasing density value distributions but, critically, without our unit-demand constraint (which is important for connecting this problem to multi-agent Bayesian mechanism design).

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.

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). This reduction framework converts the marginal revenue mechanism for linear agents to a mechanisms for non-linear agents and general objectives. Their reduction framework is also applicable to other DSIC, IIR, deterministic mechanisms for linear agents. 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.

Mechanism design for non-linear agents is well studied in the literature. In this work, as applications of our general framework, we focus on three specific non-linear models, agents with budget constraints, agents with risk averse attitudes, and agents with endogenous valuation.

Laffont and Robert (1996) and Maskin (2000) study the revenue-maximization and welfare-maximization problems for symmetric agents with *public* budgets in single-item environments. Boulatov and Severinov (2018) generalize their results to agents with i.i.d. values but asymmetric public budgets. 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) consider the single agent problem with private budget and an arbitrary valuation distribution, characterize the optimal mechanism by a linear program, and use an algorithmic approach to construct the solution. Pai and Vohra (2014) generalize the characterization of the optimal mechanism to symmetric agents with uniformly distributed private budgets. 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; Che et al., 2013).

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. (2021a) 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.

The model for agents with endogenous valuation has been studied extensively in Tan (1992); King et al. (1992); Gershkov et al. (2021b); Akbarpour et al. (2021) where agents can make costly investment before the auction. This is a generalization of the model for agents with entry costs (Celik and Yilankaya, 2009). This main focus of the literature is to characterize the optimal mechanisms in restricted settings. For example, Gershkov et al. (2021b) characterize the revenue optimal symmetric mechanism for symmetric buyers.[7] The reduction framework in our paper implies that sequentially offering a price to each agent is a constant approximation for both welfare and revenue maximization when there are multiple asymmetric buyers. Akbarpour et al. (2021) consider approximating the optimal welfare when it is computationally intractable to find the optimal allocation. They show that any algorithm that excludes bossy negative externalities can be converted to a mechanism that guarantees the same approximation ratio to the optimal welfare. They restrict attention to full information equilibrium, while our analysis applies to settings with private valuations.

[7] Gershkov et al. (2021b) also showed that even for symmetric buyers, symmetric mechanism may not be revenue optimal among all possible mechanisms.

It is well known that simple mechanisms generate robust performance guarantees for both welfare maximization (Roughgarden et al., 2017) and revenue maximization (Carroll, 2017; Bei et al., 2019). Moreover, simple mechanisms are approximately optimal under natural assumptions of type distributions. For single item auction and 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 tight approximation ratio is ${}^{e}/_{(e-1)}$ for sequential posted pricing. The approximate optimality of sequential posted pricing can be generalized to multi-item settings when agents have unit-demand valuations (Chawla et al., 2010; Cai et al., 2016). 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). For multiple items, Cheng et al. (2018) shows that selling items separately or as a bundle is approximately optimal for a single agent with additive valuation. Our analyses uses one of their lemmas.

1.3 Organization

In Section 2, we introduce the model and notations used in this paper. In Section 3, we introduce the formal definition of ζ-resemblant and present the reduction framework for posted pricing mechanisms as a warm up. In Section 4, we present our main result – the reduction framework for pricing-based mechanisms. In Sections 5 and 6, we discuss the resemblant of three canonical non-linear utility models for both welfare maximization and revenue maximization. Finally, we finish the paper with the conclusion and extensions in Section 7.

2 Preliminaries

In this paper, we consider general payoff maximization in single-item auction for non-linear agents. 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 $(\mathcal{T}, \Phi, u)$ where $\mathcal{T}, \Phi$, and u are the type space, 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 of each player u is a mapping from her private type and the outcome to her utility for the outcome. There are several specific utility models we are interested in this paper.

- **Linear utility:** For each agent $i \in N$, her private type is her value v_i of the good. Given allocation x and payment p, her utility is $v_i \cdot x - p$. In the following sections, we will drop the subscripts when we discuss the single agent problems.

- **Private-budget utility:** Each agent $i \in N$ has a private value v_i and private budget constraint w_i. We refer to the pair (v_i, w_i) as the private type of the agent. The valuation v_i for each agent i is sampled from the valuation distribution F_i and her budget w_i is sampled from the budget distribution G_i. We assume that F_i and G_i are independent distributions. We also use F_i and G_i to denote the cumulative probability function for the valuation and budget of agent i. Given allocation x and payment p, her utility is $v_i x - p$ if the payment does not exceed her budget, i.e., $p \leq w_i$. Otherwise, her utility is $-\infty$.

  Note that when the support of budget distribution G is a singleton $\{w\}$, it is equivalent to assume that the agent has a (deterministic) public budget w. We name the utility model of such agents as *public-budget utility*.

- **Risk-averse utility:** For each agent $i \in N$, her private type is her value $v_i \in [0, \bar{v}_i]$ of the good. Given allocation x and payment p, the utility function u is a concave function mapping from the wealth $v_i \cdot x - p$ of the agent to her utility. In the later discussion on revenue maximization, we restrict attention to a very specific form of risk aversion studied in Fu et al. (2013), which is both computationally and analytically tractable: utility functions that are linear up to a given capacity C and then flat. Given allocation x and payment p, an agent has utility $\min\{v_i \cdot x - p, C\}$. We refer to this utility function as *capacitated utility*. The capacity C is encoded in the utility function and is not necessarily identical across agents.

- **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. (2021b).[8] 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.

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 of the allocation payment pair (x_i, p_i) for each agent i where the allocation is a probability $x_i \in [0, 1]$ and the price is $p_i \in \mathbb{R}_+$. An allocation is feasible if $\sum_i x_i \leq 1$.[9]

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.

Payoff Curves. The payoff function of the seller is a mapping from the lotteries of each agent, to a real value. We assume that the payoff function satisfies expected utility theory,[10] i.e., the payoff for a distribution over lotteries is the corresponding expected payoff.[11] Moreover, the payoff function is additive separable across different agents. In this paragraph, we define the *payoff curves*, and introduce the *revenue curves* and *welfare curves* as special cases of the payoff curves. Specifically, the

[8]Gershkov et al. (2021b) 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.

[9]Our results can be extended to more general feasibility constraints. See Section 7 for detailed discussion.

[10]In contrast, we do not restrict the agents to satisfy the expected utility theory.

[11]For example, the seller may care about the ex ante welfare of the agents, i.e., the sum of the ex ante utility of the agents when each agent is assigned with a lottery.

revenue contribution from agent i is her expected payment p_i, and the welfare contribution from from agent i is her expected value for realized allocation x_i.[12] In addition, we define the *optimal payoff curves* and *price-posting payoff curves* as follows.

Definition 2.1. *Given ex ante constraint q, the* optimal payoff curve $R(q)$ *is a mapping from quantile q to the optimal ex ante payoff for the single agent problem, i.e., the optimal payoff of the mechanism which in expectation sells the item with probability q.*

Fact 2.1. *The optimal payoff curve is concave.*

Fact 2.1 holds because the space of mechanisms is closed under convex combination. We also study mechanisms based on simple per-unit posted posting.

Definition 2.2. *Posting* per-unit price p *is offering a menu* $\{(x, x \cdot p) : x \in [0, 1]\}$ *to the agent. A budgeted agent with value v and budget w given* per-unit price p *will purchase the lottery* $x = \min\{1, {}^{w}/_{p}\}$ *if $v \geq p$, and purchase the lottery $x = 0$ otherwise.*

Definition 2.3. *The* market clearing price p^q *for the ex ante constraint q is the per-unit price such that the item is sold with probability q.*

Definition 2.4. *Given ex ante constraint q, the* price-posting payoff curve $P(q)$ *is a mapping from quantile q to the optimal price-posting payoff for the single agent problem, i.e., the optimal payoff of the price posting mechanism which sells the item with probability q in expectation over the type distribution and the probabilities of the selected lottery.*

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

Definition 2.5. *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.

[12]Note that there are alternative definitions for welfare of non-linear agents. For example, when agents are risk averse, an alternative definition for welfare contribution from agent i is the sum of her payment p_i and her utility $u_i(x_i, p_i)$. Whether non-linear agents resemble linear agents under this alternative welfare definition is left as an open question.

Lemma 2.2 (Bulow and Roberts, 1989)**.** *The optimal revenue curve R of a linear agent is equal to her ironed price-posting revenue curve $\bar{P}$.*

A similar result holds for the welfare curve. Note that the price-posting welfare curve is always concave for linear agents.

Lemma 2.3. *The optimal welfare curve R of a linear agent is equal to her price-posting welfare curve P, both are concave and $R = P = \bar{P}$.*

In general, for agents with budgets, the optimal payoff (e.g., revenue or welfare) curves and the concave hull of the price-posting payoff curves are not equivalent, and the ex ante optimal mechanism is more complicated and extracts strictly higher payoff than the optimal price posting mechanism and randomizations over price posting mechanisms.

Ex Ante Relaxation. Next we provide the benchmark of our paper, the ex ante relaxation. For auctions with downward-closed feasibility constraints, any sequence of ex ante quantiles $\{q_i\}_{i\in N}$ is ex ante feasible if there exists a randomized, ex post feasible allocation such that the probability agent i receives an item, i.e., marginal allocation probability for agent i, equals q_i. We denote the set of ex ante feasible quantiles by EAF. Note that $\{q_i\}_{i\in N} \in \text{EAF}$ if and only if $\sum_i q_i \leq 1$. The optimal ex ante payoff given a specific collection of payoff curves $\{R_i\}_{i\in N}$ is

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

Pricing-based Mechanisms and Posted Pricing Mechanisms. In Bayesian mechanism design, the taxation principle suggests that it is without loss to focus on menu mechanisms: Fixing any agent, the mechanism offers a menu of outcomes (i.e., her allocation and payment) to the agent, where the menu depends on other agents' bids. Among all such menu mechanisms, there are two subclasses of mechanisms closely related to price posting which allow simple implementations – *pricing-based mechanisms* and *posted pricing mechanisms*. The subclass of *pricing-based mechanisms* consider mechanisms where the menu (offered by the mechanism) is equivalent to posting a per-unit price. Furthermore, a pricing-based mechanism is called a *posted pricing mechanism* if the menu (a.k.a., per-unit price) offered to each agent is invariant of other agents' bids.

3 Reduction Framework for Sequential Posted Pricing

In this section, we introduce the definition of ζ-resemblance to quantify the single-agent approximation by price-posting in non-linear utility models. As a warm up, we introduce a reduction framework which extends approximation results of posted pricing mechanisms for linear agents to non-linear agents that satisfy the definition. In next section, we discuss a more general reduction framework for pricing-based mechanisms.

As we discussed in Section 2, the taxation principle suggests that it is without loss t focus on menu mechanisms in Bayesian mechanism design. For non-linear agents, the menu offered in the Bayesian optimal mechanism are complicated even in single-agent environments. For example, to maximize the revenue from a single agent with private budget, the menu size of the optimal mechanism is exponential to the size of the support of the budget distribution (Devanur and Weinberg, 2017). In contrast, for linear agents, there exist posted pricing mechanisms that is optimal (resp. approximately optimal) in the single-agent (resp. multi-agent) environments (Myerson, 1981; Riley and Zeckhauser, 1983; Yan, 2011; Alaei et al., 2018). Here we introduce a reduction framework that extends the approximation bounds of posted pricing mechanisms for linear agents to non-linear agents.

To simplify the presentation, we focus on the reduction framework on a canonical class of posted pricing mechanisms – *sequential posted pricing mechanisms* (see Definition 3.1 for a formal definition). A generalization of the framework to other posted pricing mechanisms is straightforward and we include more discussions in Section 7.

Note that given the ex ante probability q, the payoff of posting the market clearing price is uniquely determined by the price-posting payoff curve and quantile q. Thus, for simplicity, we define the sequential posted pricing in quantile space.[13]

Definition 3.1. *A* sequential posted pricing mechanism *is parameterized by*

[13] The reason for defining posted pricings in quantile space is that the mapping from quantiles to prices is not generally pinned down by the payoff curve (specifically, for the welfare objective) for non-linear agents. As the actual prices to be posted are not important in our reduction framework, it is convenient to remain in quantile space. Any sequential posted pricing mechanism defined in quantile space can be converted to a sequential posted pricing mechanism in price space (e.g., Chawla et al., 2010). Thus, in this paper, without loss of generality, we will focus on the sequential posted pricing mechanisms in quantile space.

$(\{o_i\}_{i\in N}, \{q_i\}_{i\in N})$ where $\{o_i\}_{i\in N}$ denotes an order of the agents and $\{q_i\}_{i\in N}$ denotes the quantile corresponding to the per-unit prices to be offered to agents if the item is not sold to previous agents.[14]

According to the definition, the payoff of the sequential posted pricing mechanism with parameters $(\{o_i\}_{i\in N}, \{q_i\}_{i\in N})$ is uniquely determined by the price-posting payoff curves $\{P_i\}_{i\in N}$ of the agents. Specifically,

$$\mathrm{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}) = \sum_{i\in N}\left(\prod_{j:o(j)<o(i)}(1-q_j)\right)P_i(q_i)\ .$$

and the optimal payoff among the class of sequential posted pricing mechanisms is

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

As we mentioned above, Yan (2011) shows the following approximation guarantee for sequential posted pricing.

Theorem 3.1 (Yan, 2011)**.** *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 an $e/(e-1)$-approximation to the ex ante relaxation, i.e.,* $\mathrm{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}) \geq (1-1/e)\cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$.

To quantify the extent to which a non-linear agent resemble a linear agent, we start with the following observation. For a linear agent, the ironed price-posting payoff curve equals the optimal payoff curve. However, for a non-linear agent, the Bayesian optimal mechanisms are not posted pricing mechanisms in general. In other words, for a non-linear agent, the ironed price-posting payoff curve is not generally equivalent to the optimal payoff curve. Hence, we introduce ζ-*resemblance* of an agent to measure her ironed price-posting payoff curve resemble her optimal payoff curve.

Definition 3.2 (ζ-resemblance)**.** *An agent's ironed price-posting payoff curve $\bar{P}$ is* ζ-resemblant *to her optimal payoff curve R, if for all $q \in [0,1]$, there exists $q \leq q^\dagger$ such that $\bar{P}(q) \geq 1/\zeta \cdot R(q^\dagger)$. Such an agent is* ζ-resemblant.

[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. The per-unit prices are offered to each agent if and only if the item is not sold to previous agents given the realization.

Smaller ζ-resemblance guarantee implies that such non-linear agents resemble linear agents better, since the approximation guarantee for sequential posted pricing mechanisms for linear agents can be lifted to those non-linear agents with an additional factor ζ (Theorem 3.2). Note that the ζ-resemblant property is equivalent to show the approximation of posted pricing mechanisms for a continuum of i.i.d. (non-linear) agents with unit-demand and limit supply. In Sections 5 and 6, we give small constant bound on this resemblant property under several canonical non-linear utility models for both welfare maximization and revenue maximization.

To extend the approximation of sequential posted pricing mechanisms for linear agents to non-linear agents, we need to reduce a non-linear agent to her linear agent analog as follows.

Definition 3.3. *Fix any set of (non-linear) agents with price-posting payoff curves $\{P_i\}_{i\in N}$. The* linear agents analog *is a set of linear agents whose price-posting payoff curves are $\{P_i\}_{i\in N}$ and the optimal payoff curves are $\{\bar{P}_i\}_{i\in N}$.*

Note that the linear agent analog is well-defined for both welfare maximization and revenue maximization.[15] Based on the definition of ζ-resemblance and the linear agent analog, we present a reduction framework that converts sequential posted pricing mechanisms for linear agents to non-linear agents, and approximately preserves its payoff approximation guarantee.

Theorem 3.2. *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 $(\{o_i\}_{i\in N}, \{q_i\}_{i\in N})$ 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., $\mathrm{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}) \geq {}^1\!/\!_\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$, then this mechanism is also a $\gamma\zeta$-approximation to the ex ante relaxation for non-linear agents, i.e., $\mathrm{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}) \geq {}^1\!/\!_{\gamma\zeta} \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$.*

[15]The price-posting revenue (resp. welfare) curve $P(q)$ of a linear agent uniquely pins down her valuation distribution as $v(q) = \frac{P(q)}{q}$ (resp. $v(q) = P'(q)$). For general payoff function, given the price-posting payoff curves $\{P_i\}_{i\in N}$ of the non-linear agents, there may not exist 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 3.1 holds for payoff curves that does not correspond to any distributions of the agents. Hence, we can refer to the linear agents analog even without the existence of the underlying distributions.

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 $(\{o_i\}_{i\in N}, \{q_i\}_{i\in N})$ only depends on the price posting payoff curves, not on the agents' utility models, we have

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

and Theorem 3.2 holds. $\square$

The reduction framework (Theorem 3.2) seems to be an immediate consequence from the definition of sequential posted pricing and definition of ζ-resemblance. In the later sections, We will discuss its extensions to other (probably more general) classes of mechanisms by adopting the same method. Specifically, in Section 7, we show that how a similar reduction framework hold for other formats of posted pricing mechanisms – oblivious posted pricing where mechanisms cannot control the order of agents, and anonymous pricing where mechanisms need to post an identical price to all agents. In Section 4, we show that when the agents satisfy the expected utility representation, any deterministic, dominant strategy incentive compatible mechanism can be converted to approximately preserve the approximation ratio for non-linear agents.

As an application of the reduction framework in Theorem 3.2, consider (non-linear) agents with private budget utility. Optimal mechanism for agents with private budget utility have been studied in the literature (e.g. Che and Gale, 2000; Devanur and Weinberg, 2017 for single-agent, Pai and Vohra, 2014 for i.i.d. agents and Alaei et al., 2012 for non-i.i.d. agents). The characterization of these optimal mechanisms are complicated even for simple distributions (e.g., value and budget drawn i.i.d. from $[0, 1]$ uniformly). However, with the reduction framework (Theorem 3.2 for posted pricing mechanism and Theorem 4.1 for pricing-based mechanism), due to the resemblance between price-posting payoff curve and optimal payoff curve, we can extend the simple mechanism (i.e., sequential/oblivious posted pricing mechanism and marginal payoff mechanism) from linear agents to private-budgeted agents with good approximation guarantees. See Appendix B for an toy example where we numerically evaluate the resemblance of revenue for private-budgeted agents with uniform values and uniform budgets, and the performance of sequential posted pricing mechanism and for them.

4 Reduction Framework for Pricing-based Mechanisms

Following the discussion in Section 3, in this section we introduce the reduction framework for pricing-based mechanisms. For this reduction framework, we focus on agents satisfying the von Neumann-Morgenstern expected utility representation.

Recall that by taxation principle, it is without loss to consider menu mechanisms. The class of pricing-based mechanisms is ones whose menu offered to each agent is posting a per-unit price. For linear agents, every mechanism (e.g., the Bayesian optimal mechanism) can be implemented as a pricing-based mechanism. Here, our reduction framework extends the approximation bounds of deterministic, dominant strategy incentive compatible (DSIC), interim individual rational (IIR), pricing-based mechanisms for linear agents to non-linear agents whose utility models satisfy the expected utility representation.

Due to the technical reason, we make the following assumption on agents' utility models. Note that this assumption is satisfied for most common utility models, e.g., linear utility, budget utility, risk averse utility.

Assumption 1. *The item is* the ordinary good*, i.e., when offered a per-unit price for the item to the agent, her demand is weakly decreasing in price.*

Based on the definition of ζ-resemblance and linear agent analog, we present the meta-theorem (Theorem 4.1): a reduction framework that converts every deterministic, DSIC, IIR, pricing-based mechanism for linear agents to a DSIC, IIR, pricing-based mechanism for non-linear agents, and approximately preserves its payoff approximation guarantee.

Theorem 4.1 (Reduction Framework)**.** *Fix any set $\mathcal{A}$ of (non-linear) agents with price-posting payoff curves $\{P_i\}_{i\in N}$ and optimal payoff curves $\{R_i\}_{i\in N}$. For any deterministic, DSIC, IIR, pricing-based mechanism $\mathcal{M}_L$ for linear agents, there is a pricing-based mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ that is DSIC, IIR, and satisfies*

i. Identical payoff: mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ has the same payoff as mechanism $\mathcal{M}_L$ for the linear agents analog $\mathcal{A}_L$. Denote the payoff of mechanism $\mathcal{M}$ as $\mathcal{M}(\{P_i\}_{i\in N})$.

ii. Identical feasibility: mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ has the same distribution over outcomes as mechanism $\mathcal{M}_L$ for the linear agents analog $\mathcal{A}_L$.

Denote by γ the approximation of mechanism $\mathcal{M}_L$ for the linear agents analog $\mathcal{A}_L$ to the ex ante relaxation of $\mathcal{A}_L$, i.e., $\mathcal{M}_L(\{P_i\}_{i\in N}) \geq {}^1/_\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$. If each non-linear agent in $\mathcal{A}$ is ζ-resemblant, then mechanism $\mathcal{M}$ for non-linear agents $\mathcal{A}$ is $\gamma\,\zeta$-approximation to the ex ante relaxation of $\mathcal{A}$, i.e., $\mathcal{M}(\{P_i\}_{i\in N}) \geq {}^1/_{\gamma\zeta}\cdot\mathrm{EAR}(\{R_i\}_{i\in N})$.

In Section 4.1, we present the implementation of the reduction framework. In Section 4.2, we show how it achieves the claimed properties in Theorem 4.1. Finally, in Section 4.3, we discuss the consequence of the reduction framework for the marginal payoff mechanism (i.e., the Bayesian optimal mechanism) for linear agents.

4.1 Implementation in Theorem 4.1

Algorithm 1 describes the implementation of Theorem 4.1.[16] This implementation includes two notations $\hat{q}_i^{\mathcal{M}_L}\left(\{q_j\}_{j\in N\setminus\{i\}}\right)$ and $x^{\hat{q}}(t)$ which we define below.

For any deterministic DSIC, IIR mechanism $\mathcal{M}_L$ for linear agents, it can be represented by a mapping from the quantiles of other agents to a threshold quantile for each agent. The agent wins when her quantile is below the threshold and loses when her quantile is above the threshold. We denote the function that maps the profile of other agent quantiles $\{q_j\}_{j\in N\setminus\{i\}}$ to a quantile threshold for agent i as $\hat{q}_i^{\mathcal{M}_L}\left(\{q_j\}_{j\in N\setminus\{i\}}\right)$.

For any non-linear agent model $(\mathcal{T}, F, u)$, the single-agent pricing problem identifies the per-unit (market clearing) price $p^{\hat{q}}$ to offer the agent for any ex ante allocation constraint $\hat{q}$. Denote the allocation probability selected by an agent with type t when offered per-unit price $p^{\hat{q}}$ as $\hat{x}^{\hat{q}}(t)$. For every type t, define function $H_t(q) = \hat{x}^q(t)$. Note

[16] The construction is a simplification of a construction in Alaei et al. (2013).

that under the ordinary good assumption (Assumption 1) $H_t(q)$ is weakly increasing in q for all type t under (Assumption 1), and thus can be viewed as the cumulative density function of a distribution. See Lemma 4.2.

Algorithm 1: Reduction Framework for Pricing-based Mechanism

Input: Non-linear agents $\{(\mathcal{T}_i, F_i, u_i)\}_{i\in N}$; and deterministic, DSIC, IIR mechanism $\mathcal{M}_L$ for linear agents

1 For each agent i with private type t_i, map the type to a random quantile q_i according to the distribution H_{i,t_i} with cdf $H_{i,t_i}(q) = \hat{x}_i{}^q(t_i)$.
`/*` $H_i(q)$ `is well-defined. See Lemma 4.2 */`

2 For each agent i, calculate quantile threshold as $\hat{q}_i = \hat{q}_i^{\mathcal{M}_L}\left(\{q_j\}_{j\in N\setminus\{i\}}\right)$.
`/*` $\hat{q}_i^{\mathcal{M}_L}(\cdot)$ `is well-defined since` $\mathcal{M}_L$ `is deterministic and DSIC. */`

3 For each agent i, set payment $p_i = p^{\hat{q}_i}\, x_i^{\hat{q}_i}(t_i)$, and allocation $x_i = 1$ if $q_i < \hat{q}_i$ and $x_i = 0$ otherwise.

Lemma 4.2. *For an ordinary good (Assumption 1), the allocation probability $x^q(t)$ is weakly increasing in q for all type t.*

Proof. For an ordinary good by definition, the agent's expected allocation probability is weakly decreasing in the price. Thus, the per-unit price in each q ex ante mechanism (with respect to the price-posting payoff curve P) is weakly decreasing in q. Now consider the q ex ante mechanism with respect to the ironed price-posting payoff curve $\bar{P}$ for all quantile q. The per-unit price is monotone (by the previous argument) on quantiles that are not in ironed intervals. Within an ironed interval, the mechanism is a mix over two end-points of non-ironed intervals which linearly interpolates between the end-points and is thus monotone. □

4.2 Proof of Theorem 4.1

We first show the implementation (Algorithm 1) is DSIC, IIR and satisfies both identical payoff and identical feasibility properties.

Lemma 4.3. *Given a deterministic, DSIC, IIR mechanism $\mathcal{M}_L$ for linear agents, the mechanism $\mathcal{M}$ from the implementation (Algorithm 1) is DSIC, IIR, and satisfies identical payoff and identical feasibility properties in Theorem 4.1.*

Proof. Since mechanism $\mathcal{M}_L$ is deterministic and DSIC, Algorithm 1 is well-defined. Since for each agent i, her type t_i is drawn from F_i and q_i is drawn from H_i condition on t_i, the (unconditional) distribution of q_i is uniform on $[0, 1]$. Thus, from each agent i's perspective, the other agents' quantiles are distributed independently and uniformly on $[0, 1]$. This agent faces a distribution over ex ante posted pricing that is identical to the distribution of quantile thresholds in the mechanism $\mathcal{M}_L$. Thus, DSIC and the identical payoff property is satisfied. Since $\mathcal{M}_L$ is IIR, $\mathcal{M}$ is also IIR. Finally, note that the distribution of q_i is uniform on $[0, 1]$, identical feasibility property is satisfied by construction. $\square$

We now show that the implementation extends the approximation guarantee of mechanism $\mathcal{M}_L$ for linear agents. Note that this is immediately implied by the identical payoff property and the following lemma.

Lemma 4.4. *For agents with ironed price-posting payoff curves $\{\bar{P}_i\}_{i\in N}$ and the optimal payoff curves $\{R_i\}_{i\in N}$, if each agent is ζ-resemblant, the ex ante relaxation on the ironed price-posting payoff curve is a ζ-approximation to the ex ante relaxation on the optimal payoff curves, i.e.,* $\mathrm{EAR}(\{\bar{P}_i\}_{i\in N}, \mathcal{X}) \geq {}^1\!/\!_\zeta \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$.

Proof. Let $\{q_i^\dagger\}_{i\in N} \in \mathrm{EAF}(\mathcal{X})$ 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 feasible. Therefore,

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

4.3 Application on Marginal Payoff Mechanism.

In Bulow and Roberts (1989), authors introduce the marginal revenue mechanism and show its revenue-optimality for linear agents. The marginal revenue mechanism can be easily extended to other payoff objectives and we denote its extensions as the marginal payoff mechanisms. The ex ante relaxation gives an upper bound on the Bayesian optimal mechanism. For linear agents, the gap between the ex ante relaxation and the Bayesian optimal mechanisms (i.e., marginal payoff mechanisms) is precisely determined by the optimal payoff curves.

Definition 4.1. *The* ex ante gap *for the optimal payoff curves* $\{R_i\}_{i\in N}$ *is the ratio between the ex ante relaxation* $\text{EAR}(\{R_i\}_{i\in N})$ *and the payoff of the Bayesian optimal mechanism for linear agents* $\text{OPT}(\{R_i\}_{i\in N})$.

In single-item environments, the ex ante gap γ is at most $1/(1 - 1/\sqrt{2\pi})$ (Yan, 2011). By our framework Theorem 4.1 on the marginal payoff mechanisms, we obtain the marginal payoff mechanism for non-linear agents, and its approximation guarantee.

Definition 4.2. *The* marginal payoff mechanism*, denoted by* MPM *(defined in Algorithm 1) corresponds to the linear agent marginal revenue mechanism. Denote the payoff of* MPM *for agents with price-posting payoff curves* $\{P_i\}_{i\in N}$ *as* $\text{MPM}(\{P_i\}_{i\in N})$.

Proposition 4.5. *Given agents with the ironed price-posting payoff curves* $\{\bar{P}_i\}_{i\in N}$ *and the optimal payoff curves* $\{R_i\}_{i\in N}$*, if each agent is* ζ*-resemblant, the worst case ratio between the the marginal payoff mechanism with respect to price-posting payoff curves and the ex ante relaxation on the optimal payoff curves is* $\zeta\gamma$*, i.e.,* $\text{MPM}(\{P_i\}_{i\in N}) \geq 1/\zeta\gamma \cdot \text{EAR}(\{R_i\}_{i\in N})$*, where* γ *is the ex ante gap with curves* $\{\bar{P}_i\}_{i\in N}$.

5 Resemblance of Welfare Maximization

In the previous section, we have provided a framework showing that posted pricing mechanisms are approximately optimal if the payoff curves of the agents satisfy the resemblant property. This framework only has bite if we can show that the resemblance is indeed satisfied in canonical settings for objectives such as welfare or revenue maximization. In this section, we show that the ironed price-posting welfare curves resemble the optimal welfare curves under three canonical non-linear utility models – budgeted utility, risk-averse utility and endogenous valuation utility. Note that the resemblance of welfare curve is a single-agent problem. Thus, we drop subscript of all notations.

5.1 Budgeted Agent

For agents with budget constraints, the ex ante optimal mechanism might be complicated and hard to characterize. However, as we show below, without any assumption on the valuation distribution or the budget distribution except the independence, posting the market clearing price guarantees a 2-approximation in welfare.

Theorem 5.1. *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 Theorem 5.1 generalizes 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 $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$ according to the market clearing price p^q and bound the welfare from those two mechanisms separately. The decomposed mechanism may violate the incentive constraint for budgets, and we refer to this setting as the random-public-budget utility model. Note that the market clearing price is the same in both the private budget model and the random-public-budget utility model. Intuitively, mechanism $\mathrm{EX}^\dagger$ contains per-unit prices at most the market clearing price, while mechanism $\mathrm{EX}^\ddagger$ contains per-unit prices at least the market clearing price. Both mechanisms $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$ satisfy the ex ante constraint q, and the sum of their welfare upper bounds the original ex ante mechanism EX, i.e., $\mathbf{Payoff}[\mathrm{EX}] \leq \mathbf{Payoff}\big[\mathrm{EX}^\dagger\big] + \mathbf{Payoff}\big[\mathrm{EX}^\ddagger\big]$.

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

Definition 5.1. *An* allocation-payment function $\tau : [0, 1] \to \mathbb{R}_+$ *is a mapping from the allocation x to the payment p.*

Lemma 5.2. *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. □

Remark 5.2. *Unlike Myerson's result which give a sufficient and necessary condition for incentive compatible mechanisms for linear agents, Lemma 5.2 only characterizes a necessary condition for private-budget utility.*[17] *This condition is already enough for our arguments in Section 5.1.*

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 1. 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 $\tau^\dagger_w$ and $\tau^\ddagger_w$ for $\text{EX}^\dagger$ and $\text{EX}^\ddagger$ respectively below,

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

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 5.3. *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*

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

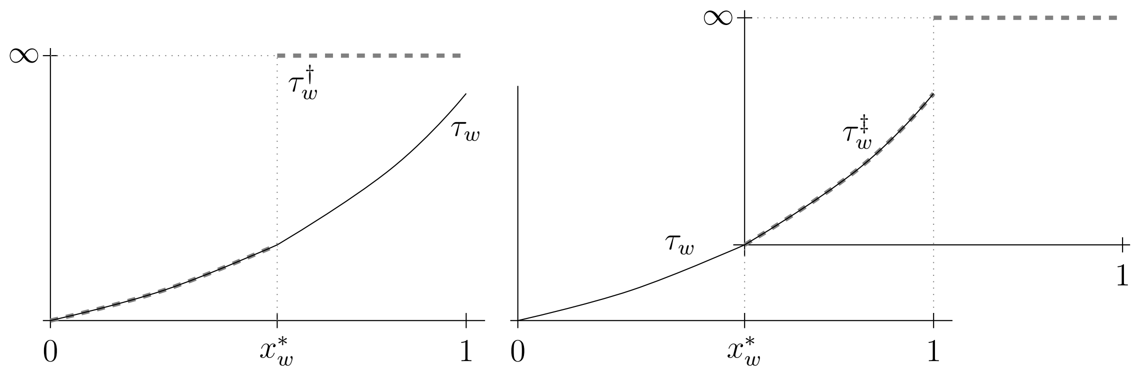

Figure 1: Depicted are allocation-payment function decomposition. The black lines in both figures are the allocation-payment function τ_w in ex ante optimal mechanism EX; the gray dashed lines are the allocation-payment function $\tau_w^\dagger$ and $\tau_w^\ddagger$ in EX† and EX‡, respectively.

market clearing price p^q is at least the welfare from EX†*, i.e.,* $P(q) \geq \textbf{Payoff}\left[\text{EX}^\dagger\right]$.

Proof. Consider agent with type (v, w) and agent with type (v', w), where both value v and v' are higher than the market clearing price p^q. Notice that the allocations for these two types are the same in EX† 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 EX† 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 EX† 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 EX† 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 EX† 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 $v \geq p^q$, $\mathbf{E}_w\big[x^{\dagger}(v,w)\big] \leq \mathbf{E}_w[x^q(v,w)]$ implies that the allocation in market clearing "first order stochastic dominantes" 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., $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^{\dagger}\big]$. $\square$

Lemma 5.4. *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., $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big]$. $\square$

Proof of Theorem 5.1. Combining Lemma 5.3 and 5.4, for any quantile q, we have

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

5.2 Risk Averse Agent

Note that the preference of a risk averse agent coincide with a linear agent when the allocation is deterministic, and the welfare optimal mechanism for the single-agent problem with linear utility is deterministic. Thus it is easy to verify that posting price is optimal for welfare maximization under any ex ante constraint and the price-posting welfare curve is 1-resemblant to the optimal welfare curve. Formally, we have the following theorem, with proof omitted.

Theorem 5.5. *An agent with risk-averse utility has the price-posting welfare curve P that equals (i.e. 1-resemblant) her optimal welfare curve R.*

5.3 Endogenous Valuation

When agents can make investment decisions before the auction, we assume that the investment costs are subtracted from the social welfare, i.e., the welfare contribution from agent i when she chooses investment decision a_i and receives allocation x_i is $v_i(a_i, t_i) \cdot x_i - c_i(a_i)$. Note that for agents with endogenous valuation, to apply Theorem 3.2 it is also 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 welfare maximization.

Lemma 5.6 (Fan and Lorentz, 1954; Gershkov et al., 2021b)**.** *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 $x \prec \hat{x}$,*[18] *we have*

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

Theorem 5.7. *An agent with endogenous valuation has the price-posting welfare curve P that equals (i.e.* 1*-resemblant) her optimal welfare curve R.*

Proof. 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. (2021b), 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 mechanism with allocation x that sells the item with probability $\hat{q}$ satisfies $x \prec \hat{x}$. By Lemma 5.6, the optimal mechanism that is $\hat{q}$ feasible has allocation rule $\hat{x}$, which is posting a deterministic price to the agent. Thus this agent has price-posting welfare curve P that equals (i.e. 1-resemblant) her optimal welfare curve R. $\square$

[18] $x \prec \hat{x}$ 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 $\int_0^1 x(q)\,\mathrm{d}q = \int_0^1 \hat{x}(q)\,\mathrm{d}q$.

6 Resemblance of Revenue Maximization

In this section, we show that the ironed price-posting revenue curves resemble the optimal revenue curves for non-linear agents. We will also drop the subscript representing the agent in all notations.

6.1 Budgeted Agent

In this section we analyze the resemblance of revenue curves for an agent with budget. We show that approximate resemblance is satisfied under weaker assumptions on the valuation distribution or the budget distribution. For simplicity, in this section, we use the notation $\mathbf{Payoff}_w[\cdot]$ to denote the revenue given any mechanism if the budget of the agent is w, and $\mathbf{Payoff}[\cdot]$ to denote the revenue by taking expectation over the budget w.

6.1.1 Public Budget

In this section, we consider the simpler setting where agents have public budgets, i.e., the budget distribution is a point mass. For an agent with a public budget, we show that the ironed price-posting revenue curve is 1-resemblant to her optimal revenue curve if her valuation distribution is regular (Theorem 6.1) and for an agent with general valuation distribution, the ironed price-posting revenue curve is 2-resemblant to her optimal revenue curve (Theorem 6.3).

Theorem 6.1. *An agent with public budget and regular valuation distribution has the ironed price-posting revenue curve $\bar{P}$ that equals (i.e. 1-resemblant) her optimal revenue curve R.*

To prove Theorem 6.1, it is sufficient to show for any quantile $\hat{q} \in [0, 1]$, the $\hat{q}$ ex ante optimal mechanism is a price-posting mechanism, i.e., $R(\hat{q}) = P(\hat{q})$. To show this, we write the ex ante optimal mechanism as an optimization program, and apply Lagrangian relaxation on the budget constraint. This leads to a new optimization program similar to an agent with linear utility but with a Lagrangian objective function. Following the technique that price-posting revenue curve indicates the ex ante optimal mechanism for a linear agent, we consider the *Lagrangian price-posting revenue curve* which characterizes the ex ante optimal mechanism for the Lagrangian objective function. See further discussion about this technique in Alaei

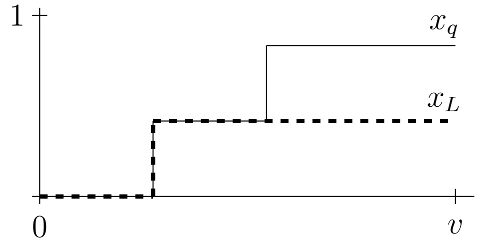

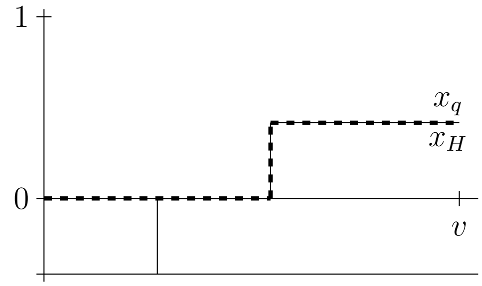

Figure 2: The thin solid line is the allocation rule for the optimal ex ante mechanism. The thick dashed line on the left side is the allocation of the decomposed mechanism with lower price, while the thick dashed line on the right side is the allocation of the decomposed mechanism with higher price.

et al. (2013) and Feng and Hartline (2018). The detailed proof of Theorem 6.1 is deferred to Appendix A.1.

For an agent with a general valuation distribution, resemblance follows from a characterization of the ex ante optimal mechanism from Alaei et al. (2013).

Lemma 6.2 (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.*

Theorem 6.3. *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 6.2, 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 2. Note that both allocation x_L and x_H are (randomized) price-posting allocation rules, and neither allocation violates the allocation constraint q. Thus,

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

6.1.2 Private Budget

In this section, we study the resemblance of the ironed price-posting revenue curve and the optimal revenue curve for agents with private budget. For linear agents, those two curves are equivalent for any valuation distribution. However, for an agent with private budget, the gap between them can be unbounded. Specifically, when the budget distribution is correlated with the valuation distribution, posting prices is not a constant approximation to the optimal revenue for a single agent even with strong regularity assumption on the marginal valuation distribution and budget distribution.

Example 6.1 (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 anonymous pricing is $O(1)$.*

Therefore, in this section, we focus on the case when the budget distribution is independent with the valuation distribution for each agent. Note that even with the independence assumption, without any further assumption on the valuation or the budget distribution, posting prices is not approximately optimal even for a single agent, see the following example as an illustration. Therefore, we consider mild assumption on either the valuation distribution or the budget distribution and show the corresponding resemblant property.

Example 6.2. *Consider 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.,*

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

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

First we show that regularity on the valuation distribution is sufficient to guarantee the resemblance between the ironed price-posting revenue curves and the optimal revenue curve, without further assumption on the budget distribution.

Theorem 6.4. *A single agent with private-budget utility and regular valuation distribution has an ironed price-posting revenue curve* $\bar{P}$ *that is* 3*-resemblant to her optimal revenue curve* R*, if her value and budget are independently distributed.*

Fix an arbitrary ex ante constraint q, denote EX as the q ex ante revenue-optimal mechanism, and **Payoff**[EX] as its revenue. We decompose EX into two mechanisms $\mathrm{EX}^\dagger$ and $\mathrm{EX}^\ddagger$ according to the market clearing price p^q. Intuitively, the per-unit prices in $\mathrm{EX}^\dagger$ for all types are at most the market clearing price and the per-unit prices in $\mathrm{EX}^\ddagger$ for all types are larger than the market clearing price. The details of the decomposition is specified in Section 5.1, and we will bound the revenue from those two mechanisms separately.

Lemma 6.5. *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. $\square$

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.

Lemma 6.6. *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.,* $P(q) \geq \mathbf{Payoff}\big[\mathrm{EX}^\ddagger\big]$.

Proof. 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.

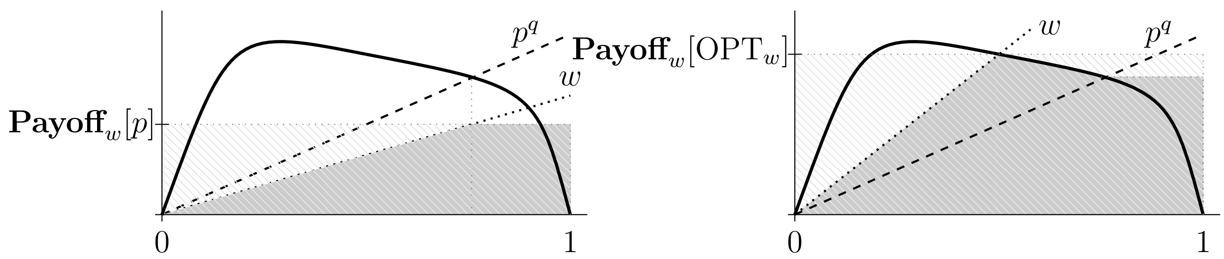

Figure 3: In the geometric proof of Lemma 6.7, the upper bound on the expected revenue of $\mathrm{EX}^{\ddagger}$ ($\mathbf{Payoff}_w[p]$ and $\mathbf{Payoff}_w[\mathrm{OPT}_w]$ on the left and right, respectively) is the area of the light gray striped rectangle and the revenue from posting random price $\boldsymbol{p}$ 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., $w \geq p^q$).

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

Lemma 6.7. *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.,* $2P(q^{\dagger}) \geq \mathbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big]$.

Proof. 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., 2010). For both sides of Figure 3, 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 3 depicts the fixed budgets w that are at most p^q. The area of the light gray striped rectangle upper bounds the revenue of $\mathrm{EX}^{\ddagger}$ as follows. Let $\mathbf{Payoff}_w[p]$ be the expected revenue from posting price p to types with budget w. Under both $\mathrm{EX}^{\ddagger}$ and the market clearing price p^q, types with value below the market clearing price pay zero. For the remaining types, in $\mathrm{EX}^{\ddagger}$ they pay at most their budget and in market clearing they pay exactly their budget. Thus, $\mathbf{Payoff}_w\left[\mathrm{EX}^{\ddagger}\right] \leq \mathbf{Payoff}_w[p^q] = w\,(1 - F(p^q))$ where, recall, $1 - F(p^q)$ is the probability the agent's value is at least the market clearing price p^q. Of course, $w\,(1 - F(p^q))$ is the height and area (its width is 1) of the light gray striped region on the left side of Figure 3.

The right side of Figure 3 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, $\mathbf{Payoff}_w\left[\mathrm{EX}^{\ddagger}\right] \leq \mathbf{Payoff}_w[\mathrm{OPT}_w]$ as the latter optimizes with relaxed constraints of the former. Laffont and Robert (1996) show that 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 3.

Next, we will show that the revenue of posting the random price $\boldsymbol{p}$ is the grey shaded areas illustrated in Figure 3 (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. $\square$

Proof of Theorem 6.4. Fix any ex ante constraint q. If the market clearing price $p^q = {}^{P(q)}/_{q}$ is at least the monopoly reserve, Lemma 6.5 and Lemma 6.6 imply that $\textbf{Payoff}\big[\mathrm{EX}^\dagger\big] \leq P(q)$, and $\textbf{Payoff}\big[\mathrm{EX}^\ddagger\big] \leq P(q)$, thus, $P(q)$ is a 2-approximation to $\textbf{Payoff}\big[\mathrm{EX}^\dagger\big] + \textbf{Payoff}\big[\mathrm{EX}^\dagger\big] = \textbf{Payoff}[\mathrm{EX}]$, 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 6.5 and Lemma 6.7 imply that $\textbf{Payoff}\big[\mathrm{EX}^\dagger\big] \leq P(q) \leq P(q^\dagger)$, and $\textbf{Payoff}\big[\mathrm{EX}^\ddagger\big] \leq 2P(q^\dagger)$, thus, $P(q^\dagger)$ is a 3-approximation to $R(q)$. Thus, the agent is 3-resemblant for ex ante optimization. $\square$

We also 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 common distribution assumption, monotone hazard rate, is a special case of it with $\kappa = e$ (cf. Barlow and Marshall, 1965).

Theorem 6.8. *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 ${}^1/_\kappa$.*

The proof of Theorem 6.8 also uses the similar decomposition technique as in Theorem 5.1 and 6.4, which we deferred to Appendix A.2.

6.2 Risk Averse Agent

In this section, we consider the case when agents are risk averse. Specifically, we consider the risk aversion model in Fu et al. (2013), where each agent's utility function has a capacity constraint. Moreover, following Fu et al. (2013), in this section, we consider the mechanisms that are pointwise individual rational, i.e., losers have no payment, and winners pay at most their reported values. Formally, $x = 0$ implies $p = 0$. In Example 6.5 at the end of this section, we show that price-posting mechanism is not a constant approximation to the optimal mechanism when we allow the winners to be charged more than their reported values, even when the capacity is as large as the support of the value.

We introduce a definition and two lemmas, which are adapted from Fu et al. (2013). Let $(\cdot)^+ \triangleq \max\{\cdot, 0\}$.

Definition 6.3 (Fu et al., 2013)**.** *A mechanism is a two priced mechanism if, when it serves an agent with quantile q and capacity C, the payment is either $V(q)$ or $V(q) - C$. The probability that agent is charged with payment $V(q)$ is denoted by $x^v(q)$, and the probability that agent is charged with payment $V(q) - C$ is denoted by $x^C(q)$.*

Lemma 6.9 (Fu et al., 2013)**.** *The ex ante optimal mechanism for agents with capacitated utility is two priced.*

Lemma 6.10 (Fu et al., 2013)**.** *For any agent with capacity C and price-posting revenue curve P, for two priced mechanism with allocation rule $x(q) = x^v(q) + x^C(q)$, the revenue from that agent is upper bounded as*

$$\mathbf{Payoff}[x] \leq \mathbf{E}\big[(P'(q))^+ \cdot x(q)\big] + \mathbf{E}\big[(P'(q))^+ \cdot x^C(q)\big] + \mathbf{E}\big[(V(q) - C)^+ \cdot x^C(q)\big] \,.$$

Theorem 6.11. *A single agent with capacitated risk averse utility, maximum value $\bar{v}$, and capacity $C \leq \bar{v}$, has a price-posting revenue curve P that is $(2+\ln {}^{\bar{v}}\!/C)$-resemblant to her optimal revenue curve R.*

Proof. For any quantile $\hat{q}$, let x be the optimal allocation that satisfies ex ante allocation constraint $\hat{q}$. By Lemma 6.10,

$$R(\hat{q}) = \mathbf{Payoff}[x] \leq \mathbf{E}\big[(P'(q))^+ \cdot x(q)\big] + \mathbf{E}\big[(P'(q))^+ \cdot x^C(q)\big] + \mathbf{E}\big[(V(q) - C)^+ \cdot x^C(q)\big] \,.$$

Let m^* be the monopoly reserve, and let $q^\dagger = \min\{Q(m^*, P), \hat{q}\}$. By definition, $q^\dagger \leq \hat{q}$. Since the price-posting revenue curve is concave, posting price $V(q^\dagger)$ maximizes expected marginal revenue under ex ante constraint $\hat{q}$. Therefore,

$$\mathbf{E}\big[(P'(q))^+ \cdot x(q)\big] \leq P(q^\dagger)$$

and

$$\mathbf{E}\big[(P'(q))^+ \cdot x^C(q)\big] \le P(q^\dagger).$$

When $q^\dagger = Q(m^*, P)$, for any quantile q, $P(q) \le P(q^\dagger)$. When $q^\dagger = \hat{q} < Q(m^*, P)$, the allocation $x^C(q)$ with ex ante constraint $\hat{q}$ that maximizes $\mathbf{E}\big[(V(q) - C)^+ \cdot x^C(q)\big]$ satisfies that $x^C(q) = 1$ for $q \le q^\dagger$, and $x^C(q) = 0$ for $q > q^\dagger$. Since the price-posting revenue curve is concave, in this case, $P(q) \le P(q^\dagger)$ when $q \le q^\dagger$. Therefore,

$$\begin{aligned}
&\mathbf{E}\big[(V(q) - C)^+ \cdot x^C(q)\big] = \mathbf{E}\left[\left(\frac{P(q)}{q} - C\right)^+ \cdot x^C(q)\right] \\
\le\ & \mathbf{E}\left[\left(\min\left\{\bar{v}, \frac{P(q^\dagger)}{q}\right\} - C\right)^+\right] \\
=\ & \int_{\frac{P(q^\dagger)}{\bar{v}}}^{\min\{1, \frac{P(q^\dagger)}{C}\}} \left(\frac{P(q^\dagger)}{q} - C\right) dq + \int_{\frac{P(q^\dagger)}{\bar{v}}}^{1} (\bar{v} - C) dq \\
\le\ & P(q^\dagger) \ln \frac{\bar{v}}{C}.
\end{aligned}$$

Combining the above inequalities, we have $R(q) \le P(q^\dagger)(2 + \ln \frac{\bar{v}}{C})$. □

In Theorem 6.11, the dependence on $\ln \bar{v}/C$ is necessary even when there is a single agent.

Example 6.4 (necessity of the dependence on $\bar{v}/C$)**.** *Fix a constant $\bar{v}$. Consider a single agent with equal revenue distribution. That is, her value v is drawn from $[1, \bar{v}]$ with a density function $1/v^2$ for $v \in [1, \bar{v})$, and a mass point of probability $1/\bar{v}$ on value $\bar{v}$. The revenue for posting any price is 1. Suppose the agent has capacity constraint $C \ge 1$, Consider the mechanism that always allocates the item to the agent, and charges her 0 if her value v is less than C, and charges her $v - C$ if her value is at least C. The revenue for this mechanism is $\ln \bar{v}/C$.*

Example 6.5 (necessity of the restriction to pointwise individually rational mechanisms)**.** *Fix a constant $\bar{v}$. Consider a single agent with equal revenue distribution as in Example 6.4. The revenue for posting any price is 1. Suppose the agent has capacity constraint $C = \bar{v}$ and consider the mechanism that always allocates the item to the agent, and charges her $v - \bar{v}$ with probability $\frac{1}{2}$, $\bar{v}$ with probability $\frac{1}{2}$. This mechanism is incentive compatible and individually rational. The revenue for this mechanism is* *half of the welfare, which cannot be approximated within a constant fraction by any price-posting mechanism.*

6.3 Endogenous Valuation

For agents with endogenous valuation, we show that posted pricing is optimal for the single agent problem given any ex ante constraint if the type distribution satisfies the regularity condition.

Theorem 6.12. *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. (2021b), the function $L(x, q)$ is supermodular in (x, q) and convex in x if the type distribution is regular. Similar to Theorem 5.7, for any quantile $\hat{q}$, the optimal mechanism for maximizing the expected virtual value that sells the item with probability at most $\hat{q}$ is posted pricing. Since the expected revenue equals the expected virtual value, this agent has price-posting revenue curve $\bar{P}$ that equals (i.e. 1-resemblant) her optimal revenue curve R. □

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 resemblant property which characterizes the gap between the concave hull of the price-posting payoff curve and the ex ante payoff curve for the single agent problem. As the instantiations of the framework, we analyze the approximation bound for various mechanisms for various non-linear utility model (i.e., budgeted utility, risk averse utility, endogenous valuation utility) under the objective of both revenue-maximization and welfare-maximization. Next we discuss several important extensions of our framework.

7.1 Convex Combination of Welfare and Revenue Maximization

One common objective of the designer considered in the literature is to maximize the convex combination of welfare and revenue of the mechanism. Formally, given any $\alpha \in (0, 1)$, the objective of the designer is to maximize $\alpha \cdot \text{Wel} + (1-\alpha) \cdot \text{Rev}$. We can extend our results in Section 5 and 6 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. The argument holds by applying the following lemma since both Wel and Rev are non-negative.

Lemma 7.1. *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 $\mathbf{Payoff}_1[\text{EX}]$ be the contribution of objective 1 given mechanism EX and $\mathbf{Payoff}_2[\text{EX}]$ be the contribution of objective 2 given mechanism EX. Let $\mathbf{Payoff}[\text{EX}] = \alpha \cdot \mathbf{Payoff}_1[\text{EX}] + (1-\alpha) \cdot \mathbf{Payoff}_2[\text{EX}]$ be the convex combination of the contributions given $\alpha \in (0, 1)$. Let $q_1 = \text{argmax}_{q' \leq q} \bar{P}_1(q')$ and $q_2 = \text{argmax}_{q' \leq q} \bar{P}_2(q')$, 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

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

Thus this agent is $(\zeta+\zeta')$-resemblant for the 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, then oblivious posted pricing for these agents is a $2\max_i\{\zeta_i\}$-approximation to the optimal ex ante relaxation.

Lemma 7.2. *For any $\zeta' \geq \zeta \geq 1$, ζ-resemblant implies ζ'-resemblant.*

7.3 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 7.1. *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 3.2, we have the following reduction framework for oblivious posted pricing for non-linear agents. The proof is identical to Theorem 3.2, hence omitted here.

Theorem 7.3. *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.,* $\mathrm{OPP}(\{P_i\}_{i\in N}, \{q_i\}_{i\in N}) \geq {}^1\!/\!_\gamma \cdot \mathrm{EAR}(\{\bar{P}_i\}_{i\in N})$*, then this mechanism is also a* $\gamma\zeta$*-approximation to the ex ante relaxation for non-linear agents, i.e.,* $\mathrm{OPP}(\{P_i\}_{i\in N}, \{q_i\}_{i\in N}) \geq {}^1\!/\!_{\gamma\zeta} \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$*.*

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

7.4 Anonymous Pricing

A desirable property for the multi-agent setting is anonymity. This requires that the price posted to all agents are the same. Note that for welfare maximization, although anonymous pricing guarantees 2-approximation for linear agents (Lucier, 2017), it may lead to huge welfare loss for non-linear agents. This is illustrated in the following example.

Example 7.2. *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.*

Thus we will focus on revenue maximization for anonymous pricing. Alaei et al. (2018) showed that for linear agents, 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 payoff curves are $\{P_i\}_{i\in N}$.

Theorem 7.4. *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 the price-posting payoff curves are concave, then anonymous pricing is a* ζe*-approximation to the ex ante relaxation on the optimal payoff curves, i.e.,* $\mathrm{AP}(\{P_i\}_{i\in N}) \geq {}^1\!/\!_{\zeta e} \cdot \mathrm{EAR}(\{R_i\}_{i\in N})$*.*

Proof. Let $\{q_i\}_{i\in N}$ be the optimal ex ante relaxation for ex ante revenue curves $\{R_i\}_{i\in N}$, and let $q_i^\dagger$ be the quantile assumed to exist by 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 payoff curves are concave, we have $\{P_i\}_{i\in N} = \{\bar{P}_i\}_{i\in N}$, and

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

By Alaei et al. (2018), $e \cdot \mathrm{AP}(\{P_i\}_{i\in N}) \geq \mathrm{EAR}(\{P_i\}_{i\in N})$ if the price-posting payoff curves $\{P_i\}_{i\in N}$ are concave. Combining the inequalities, we have

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

As instantiation of the reduction framework in Theorem 7.4, 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 $(2+\ln {}^{\bar{v}}/_{C})$-resemblant and have concave price posting revenue curve when they have capacitated risk averse utility with maximum value $\bar{v}$, capacity $C \leq \bar{v}$, and regular valuation distributions.

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

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

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

$\mathcal{X}$ is denoted by $\text{SPP}(\{P_i\}_{i\in N}, \{o_i\}_{i\in N}, \{q_i\}_{i\in N}, \mathcal{X})$.[19]

It is easy to verify that the reduction framework for (sequential) posted pricing mechanisms (Theorem 3.2) and the reduction framework for pricing-based mechanism (Theorem 4.1) 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.

Theorem 7.5 (Agrawal et al., 2010; Yan, 2011)**.** *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, and* $\gamma = {}^1\!/\!(1 - {}^1\!/\!\sqrt{2\pi k})$ *for k-unit auctions.*

References

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

Agrawal, S., Ding, Y., Saberi, A., and Ye, Y. (2010). Correlation robust stochastic optimization. In *Proc. 21st ACM Symp. on Discrete Algorithms*, pages 1087–1096. Society for Industrial and Applied Mathematics.

Akbarpour, M., Kominers, S. D., Li, S., and Milgrom, P. R. (2021). Investment incentives in near-optimal mechanisms. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 26–26.

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.

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

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

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

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.

Bei, X., Gravin, N., Lu, P., and Tang, Z. G. (2019). Correlation-robust analysis of single item auction. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 193–208. SIAM.

Border, K. C. (1991). Implementation of reduced form auctions: A geometric approach. *Econometrica: Journal of the Econometric Society*, pages 1175–1187.

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

Boulatov, A. and Severinov, S. (2018). Optimal mechanism with budget constrained buyers. Technical report, Mimeo, October.

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

Cai, Y., Devanur, N. R., and Weinberg, S. M. (2016). A duality-based unified approach to bayesian mechanism design. *ACM SIGecom Exchanges*, 15(1):71–77.

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

Celik, G. and Yilankaya, O. (2009). Optimal auctions with simultaneous and costly participation. *The BE Journal of Theoretical Economics*, 9(1).

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.

Che, Y.-K., Kim, J., and Mierendorff, K. (2013). Generalized reduced-form auctions: A network-flow approach. *Econometrica*, 81(6):2487–2520.

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. (2010). Revenue maximization with a single sample. In *ECOM10*.

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., Hartline, J., and Hoy, D. (2013). Prior-independent auctions for risk-averse agents. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 471–488.

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. (2021a). Optimal auctions: Non-expected utility and constant risk aversion. *Accepted at Review of Economic Studies*.

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

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

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

Jin, Y., Lu, P., Qi, Q., Tang, Z. G., and Xiao, T. (2019). Tight approximation ratio of anonymous pricing. In *Proceedings of the 51th Annual ACM Symposium on Theory of Computing, to appear*.

King, I., Welling, L., and McAfee, R. P. (1992). Investment decisions under first and second price auctions. *Economics Letters*, 39(3):289–293.

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.

Roughgarden, T., Syrgkanis, V., and Tardos, E. (2017). The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101.

Tan, G. (1992). Entry and R&D in procurement contracting. *Journal of Economic Theory*, 58(1):41–60.

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

Online Appendix

A Missing Proofs for Resemblance of Revenue Maximization

A.1 Public Budget

Theorem 6.1. *An agent with public budget and regular valuation distribution has the ironed price-posting revenue curve $\bar{P}$ that equals to (i.e. 1-resemblant) her optimal revenue curve R.*

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

$$\begin{aligned} \max_{(x,p)} \quad & \mathbf{E}_v[p(v)] \\ s.t. \quad & (x,p) \text{ are IC, IR}, \\ & \mathbf{E}_v[x(v)] = \hat{q}, \\ & p(\bar{v}) \leq w. \end{aligned} \tag{1}$$

where $\bar{v}$ is the highest possible value of the agent. Consider the Lagrangian relaxation of the budget constraint in (1),

$$\begin{aligned} \min_{\lambda \geq 0} \max_{(x,p)} \quad & \mathbf{E}_v[p(v)] + \lambda w - \lambda p(\bar{v}) \\ s.t. \quad & (x,p) \text{ are IC, IR}, \\ & \mathbf{E}_v[x(v)] = \hat{q}. \end{aligned} \tag{2}$$

Let λ^* be the optimal solution in program (2). If we fix $\lambda = \lambda^*$ in program (2), 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 A.1 (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 A.1 (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 $\hat{P}_{\lambda^*}(q) > P_{\lambda^*}(q)$.

A.2 Private Budget

Theorem 6.8. *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.

Our analysis here is similar to the analysis for welfare, i.e., the price decomposition technique. Consider the decomposition of EX into three mechanisms $\mathrm{EX}^{\dagger}$, $\mathrm{EX}^{\S}$ and $\mathrm{EX}^{\ddagger}$ such that mechanism $\mathrm{EX}^{\dagger}$ contains per-unit prices at most the market clearing price, mechanism $\mathrm{EX}^{\ddagger}$ contains per-unit prices at least the expected budget, while mechanism $\mathrm{EX}^{\S}$ 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., $\textbf{Payoff}[\mathrm{EX}] \leq \textbf{Payoff}\big[\mathrm{EX}^{\dagger}\big] + \textbf{Payoff}\big[\mathrm{EX}^{\S}\big] + \textbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big]$. Note that in the special case where the market clearing price is larger than the expected budget, i.e., $p^q > w^*$, $\mathrm{EX}^{\S}$ does not exist and mechanism EX is decomposed into $\mathrm{EX}^{\dagger}$ and $\mathrm{EX}^{\ddagger}$.

We construct the allocation-payment functions $\tau_w^{\dagger}$, $\tau_w^{\ddagger}$ and $\tau_w^{\S}$ for $\mathrm{EX}^{\dagger}$, $\mathrm{EX}^{\ddagger}$, and $\mathrm{EX}^{\S}$ 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., $x_w^* = \operatorname{argmax}\{x : \tau_w'(x) \leq p^q\}$. Let $x_w^{\sharp}$ be the utility maximization allocation for the agent with value and budget equal to the expected budget w^*, i.e., $x_w^{\sharp} = \operatorname{argmax}\{x : \tau_w'(x) \leq w^*\}$. Then the allocation-payment functions $\tau_w^{\dagger}$, $\tau_w^{\ddagger}$ and $\tau_w^{\S}$ are defined respectively as follows,

$$\tau_w^{\dagger}(x) = \begin{cases} \tau_w(x) & \text{if } x \leq x_w^*, \\ \infty & \text{otherwise;} \end{cases} \qquad \tau_w^{\S}(x) = \begin{cases} \tau_w(x_w^* + x) - \tau_w(x_w^*) & \text{if } x \leq x_w^{\sharp} - x_w^*, \\ \infty & \text{otherwise;} \end{cases}$$

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

The revenue contribution from $\mathrm{EX}^{\dagger}$ is bounded in Lemma 6.5. Next we illustrate how to bound the revenue from $\mathrm{EX}^{\ddagger}$ and $\mathrm{EX}^{\S}$ respectively using the revenue from price-posting.

Lemma A.2. *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 $(1 + \kappa - {}^{1}\!/\!_{\kappa}) \cdot P(q^{\dagger}) \geq \textbf{Payoff}\big[\mathrm{EX}^{\ddagger}\big]$.*

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 $q^{\dagger} = \operatorname{argmax}_{q' \leq q} P(q')$. Thus $P(\bar{q}) \leq P(q^{\dagger})$. Moreover, by the construction of the decomposition, the per-unit price in $\mathrm{EX}^{\ddagger}$ is larger than $\bar{p}$. Similar to the proof of Lemma 6.6, we only consider the types with value at least $\bar{p}$.

Let $\mathbf{Payoff}_w\left[\tau_w^{\ddagger}\right]$ be the expected revenue of providing the allocation-payment function $\tau_w^{\ddagger}$ in $\mathrm{EX}^{\ddagger}$ 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 $\mathbf{Payoff}\left[\mathrm{EX}^{\ddagger}\right]$ to $P(q^{\dagger})$:

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

(b) $\mathbf{Payoff}_w\left[\tau_w^{\ddagger}\right] \leq \frac{w}{w^*}\mathbf{Payoff}_{w^*}\left[\tau_w^{\ddagger}\right]$ 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,

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

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). $\square$

Lemma A.3. *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 $\tau_w^\S$, the revenue only increases if we increase the budget to w^*, i.e., $\mathbf{Payoff}_w\big[\tau_w^\S\big] \leq \mathbf{Payoff}_{w^*}\big[\tau_w^\S\big]$ for any $w \leq w^*$. Moreover, for any $w > w^*$, given the allocation payment function $\tau_w^\S$, 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 $\mathbf{Payoff}_w\big[\tau_w^\S\big] \leq {}^w\!/\!{}_{w^*} \cdot \mathbf{Payoff}_{w^*}\big[\tau_w^\S\big]$. Note that for allocation payment rule $\tau_w^\S$, 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 $\tilde{p}$, $\mathbf{Payoff}_{w^*}\big[\tau_w^\S\big] \leq \mathbf{Payoff}_{w^*}[\tilde{p}]$ 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 $\mathbf{Payoff}_{w^*}[\tilde{p}] \leq \kappa \cdot \mathbf{Payoff}[\tilde{p}]$ 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. $\square$

Proof of Theorem 6.8. Let $q^\dagger = \operatorname{argmax}_{q' \leq q} P(q')$. Combining Lemma 6.5, A.2 and A.3, we have

$$\mathbf{Payoff}[\mathrm{EX}] \leq \mathbf{Payoff}\big[\mathrm{EX}^\dagger\big] + \mathbf{Payoff}\big[\mathrm{EX}^\ddagger\big] + \mathbf{Payoff}\big[\mathrm{EX}^\S\big] \leq (1 + 3\kappa - {}^1\!/\!_\kappa)\, P(q^\dagger).$$

$\square$

B Numerical Result for Uniformly Distributed Private-budgeted Agents

In this section, we discuss the numerical results of the approximation ratios of revenue-maximization for i.i.d. private-budgeted agents with value and budget drawn uniformly from $[0, 1]$ independently. This example and the optimal mechanisms have been studied in Che and Gale (2000) for a single agent and Pai and Vohra (2014) for multiple agents. For both scenarios, the optimal mechanisms are complicated. However, Figure 4a suggests that for a single agent, posting a single price is a good approximation to the optimal mechanism for all ex ante probability constraint; Figure 4b suggests that for multi-agents, simple pricing based mechanisms (i.e. oblivious posted pricing and marginal payoff maximization) achieve good approximation to the optimal mechanism. Next, we explain how the numerical results are computed.

First we focus on the single agent problem, i.e., the calculation of the price-posting revenue curve and ex ante revenue curve illustrated in Figure 4a. For the price-posting revenue curve, we directly compute the probability the item is sold and the corresponding revenue for any price p. Thus, we can have the closed-form characterization for the mapping from the ex ante allocation constraint to the optimal price-posting revenue. For the ex ante revenue curve, by approximating the continuous uniform distribution with a discretized uniform distribution, we can write this optimization problem as a finite dimensional linear program, which allows us to numerically evaluate the optimal ex ante revenue given any ex ante allocation constraint q. By evaluating the curve on quantiles $q \in \{0, {}^1\!/\!_{50}, \ldots, 1\}$ with grid size ${}^1\!/\!_{50}$, we have the numerical figure for the ex ante revenue curve.

For the multi-agent problem, since both oblivious posted pricing and marginal payoff mechanism are pricing based mechanism, the revenues of both mechanisms for private-budgeted agents are equivalent to the revenues of both mechanisms for linear agents with the same price-posting revenue curve. By the above paragraph, we have the closed-form for the price-posting revenue curve, which pins down the value distribution of such linear agents. First note that since agents are i.i.d., the revenue from oblivious posted pricing (OPP) is the same as sequential posted pricing (SPP). We compute the revenue for both OPP and SPP using an dynamic programming (i.e. backward induction). For i.i.d. regular linear agents, the revenue of the marginal payoff mechanism is the same as the revenue of the second price auction with monopoly reserve, which can be solved analytically. Finally, we can numerical calculate the optimal ex ante relaxation using the ex ante revenue curve for a single agent, and evaluate the approximation ratio for both mechanisms when number of agents ranges from 1 to 15.

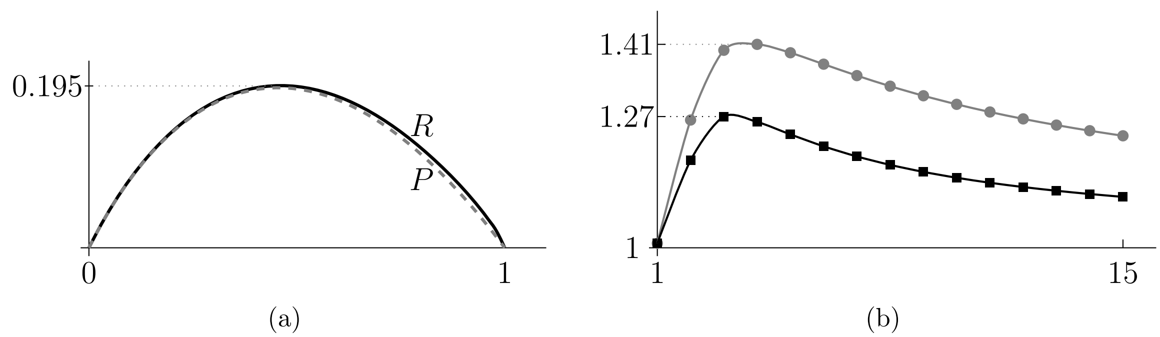

Figure 4: Figure 4a illustrates the comparison between the price-posting revenue curve (dashed line) and the ex ante revenue curve (solid line) for selling a single item to a private-budgeted agent with value and budget both drawn uniformly from $[0, 1]$. The x-axis is the ex ante probability and the y-axis is the expected revenue. The price-posting revenue curve for this uniform budgeted agent is 1.02-resemblant to her ex ante revenue curve.

Figure 4b illustrates the comparison between approximation ratio of optimal oblivious posted pricing (grey line) and marginal payoff mechanism (black line) to the ex ante relaxation for selling a single item to i.i.d. private-budgeted agents with value and budget both drawn uniformly from $[0, 1]$. The x-axis is the number of agents and the y-axis is the approximation ratio. When there are 15 agents, the approximation ratio for oblivious posted pricing is 1.23 and the approximation ratio for marginal payoff mechanism is 1.11.