PROCUREMENT WITHOUT PRIORS: A SIMPLE MECHANISM AND ITS NOTABLE PERFORMANCE

By

Dirk Bergemann, Tibor Heumann and Stephen Morris

December 2025

COWLES FOUNDATION DISCUSSION PAPER NO. 2479

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Procurement without Priors: A Simple Mechanism and its Notable Performance *

Dirk Bergemann† Tibor Heumann‡ Stephen Morris§

December 8, 2025

Abstract

How should a buyer design procurement mechanisms when suppliers' costs are unknown, and the buyer does not have a prior belief? We demonstrate that notably simple mechanisms – those that share a *constant fraction* of the buyer utility with the seller – allow the buyer to realize a guaranteed positive fraction of the efficient social surplus across all possible costs. Moreover, a judicious choice of the share based on the known demand maximizes the surplus ratio guarantee that can be attained across all possible (arbitrarily complex and nonlinear) mechanisms and cost functions. Results apply to related nonlinear pricing and optimal regulation problems.

JEL CLASSIFICATION: D44, D47, D83, D84.

KEYWORDS: Regulation, Procurement, Nonlinear Pricing, Second Degree Price Discrimination, Surplus Guarantee, Profit Guarantees, Competitive Ratio.

*We acknowledge financial support from NSF grants SES-2001208 and SES-2049744, and ANID Fondecyt Regular 1241302. An Extended Abstract of an early version of this paper appeared under the title "Cost-Based Nonlinear Pricing" in the Conference Proceedings of ACM-EC '23. We thank seminar audiences at Cambridge, Chicago, Essex, Gerzensee, LSE, Oxford, Stanford, Texas, UBC, UCL, and UCLA for helpful discussions. We have benefitted from comments by Jerry Anunrojwong, Costas Arkolakis, Omar Besbes, Ben Brooks, Yang Cai, Eduardo Davila, Daniel Garrett, Jason Hartline, Yingkai Li, Zvika Neeman, Xavier Vives and Jinzhao Wu. We thank Hongcheng Li and David Wambach for excellent research assistance.

†Yale University, dirk.bergemann@yale.edu

‡Pontificia Universidad Católica de Chile, tibor.heumann@uc.cl

§Massachusetts Institute of Technology, semorris@mit.edu

1 Introduction

1.1 Motivation

Procurement policies face a fundamental challenge in both public and private sector operations. The buyer must offer a purchase mechanism without knowing the supplier's true costs. A government agency awarding infrastructure contracts, a firm purchasing components, or a regulator setting utility prices—all must commit to payment rules with limited information about production technologies, input costs, or operational efficiencies. Bayesian mechanism design, pioneered by Myerson (1981), offers a solution: if the buyer has a Bayesian prior over suppliers' costs, optimal mechanisms can be characterized precisely. But this prescription demands substantial informational sophistication. In practice, buyers often lack reliable probability distributions over cost structures, particularly when: (i) procuring novel products or services, (ii) operating across heterogeneous contexts (local, regional, federal levels), and (iii) facing rapidly evolving technologies or input markets. This paper asks: What can a buyer guarantee themselves when they know their own demand but possess minimal information about suppliers' costs?

We propose a remarkably simple class of mechanisms—*constant utility share mechanisms*—that offer powerful performance guarantees. Under such mechanisms, the buyer commits to paying the supplier a fixed share of the buyer's gross utility for any quantity supplied. Despite knowing nothing about the supplier's cost function, the buyer can guarantee themselves a definite positive fraction of what could be achieved as the social welfare under complete information. Thus, the buyer can guarantee himself, with a simple mechanism and unknown supply, a constant fraction of the efficient social surplus that could be attained with an optimal mechanism and complete information about supply. Furthermore, a judicious choice of the share (calibrated to the demand) can guarantee the maximum share of the efficient social surplus across all, possibly arbitrarily complex, mechanisms.

The maximum share of the efficient social surplus under complete information that the buyer can guarantee themselves is referred to as the *competitive ratio* in theoretical computer science and operations research, see El-Yaniv and Borodin (1998). Thus, we show that a simple mechanism attains the competitive ratio for our problem. In other words, we establish that constant share mechanisms maintain their performance guarantees across all possible supplier cost functions; furthermore, they provide robust procurement efficiency in the sense that they attain the competitive ratio. We view this as a supporting argument in favor of the simple mechanisms we study.

We first present our approach in the case where the supplier has a constant marginal cost of production, and the buyer has a power utility function. The demand generated by any specific power utility function generates a constant price elasticity of demand. We introduce constant share mechanisms and identify a specific constant share mechanism that generates the maximum surplus ratio for the buyer (Proposition 1). We then show that the maximum surplus ratio attained by a constant share mechanism is also the maximum surplus ratio attained by *any* mechanism and thus attains the competitive ratio. We thus establish that constant share mechanisms solve the max min problem, where the buyer chooses a mechanism to maximize the surplus ratio, subject to nature choosing the distribution of costs to minimize the surplus ratio (Proposition 2). We can provide a closed form for the optimal constant share mechanism and thus the competitive ratio as a function of the demand elasticity: the competitive ratio decreases from 1 to $1/e$ as the magnitude of the demand elasticity varies from 1 to ∞. We then show that there exists a sequence of probability distributions over constant marginal costs such that a constant share mechanism is the limit of the solutions to the classical optimal Bayesian mechanism for that sequence of beliefs. This allows us to establish a saddle point property of the optimal constant share mechanism: the optimal constant share mechanism also solves the min max problem where nature first chooses a distribution over (constant) marginal costs to minimize the surplus ratio, subject to the buyer choosing an optimal mechanism for that distribution (Proposition 3).

All these results extend from the constant marginal cost case to our general setting with weakly increasing and weakly convex cost functions. Theorem 1 establishes the competitive ratio as a simple function of the buyer's demand elasticity. As the magnitude of the demand elasticity becomes larger, the share of the surplus conceded to the seller increases and the competitive ratio decreases. We also show how our results can be extended to allow variable elasticity of demand and non-convex costs, although now the interpretation of results is more subtle, as the lower bound is not the solution of the max min problem anymore (Theorem 2).

We present our results for procurement mechanisms, but our approach extends to other screening problems. We describe the related environments of regulation (as in Baron and Myerson (1982)) and nonlinear pricing (as in Mussa and Rosen (1978) and Maskin and Riley (1984)). In the regulatory problem, the regulator assigns a positive weight to the profit of the seller, but possibly smaller than the weight given to the buyer surplus. Theorem 3 confirms that the constant share mechanism continues to perform well. To the extent that the regulator integrates the profit of the seller into the weighted social surplus, the objective of the regulator becomes closer to the benchmark, the social welfare with equal weights given to the buyer and seller. As a consequence, the competitive ratio improves and the optimal share of surplus conceded to the seller increases as well.

In the nonlinear pricing environment, we focus on power cost functions. This allows us to express the cost environment of the firm in terms of a single cost elasticity parameter. Proposition 4 shows that a constant mark-up mechanism tailored to the cost elasticity attains the competitive ratio. Thus, a *cost-based* nonlinear pricing approach solves the nonlinear pricing problem in the same way that a *demand-based* pricing solves the procurement problem.

The competitive ratio, similar to other objectives such as regret minimization or maximin utility, has emerged in the analysis of robust decision-making in the absence of a Bayesian prior distribution. We view the competitive ratio as relevant and appropriate in the current context for two reasons: First, it is a way to obtain a quantitative insight into the robust policies. We compare these alternative criteria in Section 6 and identify the advantages of the competitive ratio in the present setting. Similar to earlier influential contributions in the theory of regulation, see Rogerson (2003) and Chu and Sappington (2007), we view the competitive ratio as a practical quantity that provides guidance in the absence of a Bayesian prior distribution. Second, in procurement settings, a procuring agency has to articulate a policy that applies to all procuring events, whether of small or large scale, and thus requires a policy that is invariant to the scale of the procurement. The scale independence of the competitive ratio – a well-known feature – is of particular relevance for the design of procurement and regulatory policy. Procurement policies, by public or private procurement agencies, have to define policies that attain good performance across a wide scale of procurement events. By comparing the attained buyer surplus relative to the efficient social surplus across all possible scales, the resulting policy is performing well across all scales. A high-performing policy has to realize a substantial share of the gains from trade at every level of trade. Thus, the mechanism cannot exclude trades at low quantities, nor does it attempt to aggressively discount at large quantities. While we focus much of our discussion on procurement and regulation, the emphasis on scale invariance is equally relevant for optimal selling mechanisms. By measuring the success of the selling policy against the social surplus, the selling policy is measured against the possible social surplus, thus the size of the total market. For a private seller, say a start-up or a seller of a new product, the selling policy just focuses on the attainable social surplus rather than the expected profit on the initial expectation that may be biased downwards relative to the eventual size of the market.

1.2 Related Literature

Our work connects to several strands of literature. First, we contribute to the literature on simple versus optimal mechanisms, exemplified by Rogerson (2003) and Chu and Sappington (2007) who analyzed linear contracts in procurement settings. While previous work has focused on comparing simple mechanisms to optimal ones under specific distributional assumptions, we characterize conditions under which simple mechanisms perform well across all possible cost structures.

The regulation setting of Baron and Myerson (1982) with constant marginal cost is closely related to our setting and we obtain competitive ratios for the model of Baron and Myerson (1982). Rogerson (2003) offers an analysis of simple contracts, allowing only for a binary menu of either a fixed price contract or a cost-reimbursement contract (FCBR menu) within the framework of Laffont and Tirole (1986) which contains elements of adverse selection and moral hazard. The seller has a cost level that is private information (adverse selection) and reduces the unit cost by an additional effort (moral hazard). Rogerson (2003) shows that the optimal simple menu can attain $3/4$ of the unrestricted optimal menu, which consists of a menu of a continuum of linear contracts in a parametrized version of quadratic cost of effort and uniformly distributed private information (Proposition 2 and Section 3 of Rogerson (2003)). Chu and Sappington (2007) provides a more complete analysis of the performance of simple contracts by allowing a larger class of distributions and a larger class of contracts. They obtain a competitive ratio of the simple relative to the optimal regulation solution of $2/e$.

In the theory of regulation, a concern for robustness is inherent to the problem. Garrett (2014) considers the model of Laffont and Tirole (1986) and shows that the simple menu proposed by Rogerson (2003) is the optimal robust menu when the regulator is uncertain about the cost reduction technology of the firm. Garrett (2021) pursues this line of research further and asks what an analyst who is uncertain of the cost reduction technology can assert about the payoff implications for the regulator and the firm when the regulator is with knowledge about the technology offers an optimal incentive compatible regulation.

Guo and Shmaya (2025) consider the regret-minimizing regulation policy in the setting of Baron and Myerson (1982) when the regulator is uncertain about the demand and supply for the product or services of a monopoly. They consider the regret-minimizing policy rather than the competitive ratio.

We derive performance guarantees and robust pricing policies that secure these guarantees for large classes of second-degree price discrimination problems as introduced by Mussa and Rosen

(1978) and Maskin and Riley (1984). We do this without imposing any restriction on the distribution of the values, such as regularity or monotonicity assumptions, or finite support conditions. We only require that the social surplus of the allocation problem has a finite expectation.

The optimal monopoly pricing problem for a single object is a special limiting case of our model when the marginal cost of increasing supply becomes infinitely large. The analysis of the competitive ratio in the single-unit case is also a special case of a result of Neeman (2003). He investigates the performance of English auctions with and without reserve prices. The case of the optimal monopoly pricing is a special case of the auction environment with a single bidder. He establishes a tight bound for the single-bidder case that is given by a "truncated Pareto distribution" (Theorem 5). The bound that he derives is a function of a parameter that is given by the ratio of the bidder's expected value and the bidder's maximal value. Without the introduction of this ratio as a parameter, the bound is zero: as this ratio converges to zero, so does the bound. Similarly, Eren and Maglaras (2010) and Hartline and Roughgarden (2014) establish that for distributions with support $[1, h]$, the competitive ratio of the optimal pricing problem is $1 + \ln h$. Thus, as h grows, the approximation becomes arbitrarily weak (see Theorem 2.1 of Hartline and Roughgarden (2014)). By contrast, our results obtain a constant approximation for every convex cost function. Thus, the introduction of a convex cost function (or a concave utility function) leads to a noticeable strengthening of the approximation quality.

Carrasco, Luz, Monteiro, and Moreira (2019) consider a different robust version of the nonlinear allocation problem. The seller faces a privately informed buyer and only knows the first moment of the distribution and its finite support (taken to be the unit interval). As in Bergemann and Schlag (2008) they solve the problem by characterizing equilibria of an auxiliary zero-sum game played by the seller and an adversarial nature. Their main result is that in the optimal robust mechanism the ex-post payoff of the seller has to be linear in the buyer's realized type. This property, which can be traced to the moment constraint on the mean, does not hold for our solution. Carroll (2017) considers a robust version of the multi-item pricing problem. The buyer has an additive value for finitely many items and has private information about the value of each item. There, the seller knows the marginal distribution for every item but is uncertain about the joint value distribution. Carroll (2017) solves a minmax problem and shows that separate item-by-item pricing is the robustly optimal pricing policy. Deb and Roesler (2024) and Che and Zhong (2025) consider a related problem under informational robustness. There, the joint distribution of values is given by a commonly known distribution, but nature or the buyer can choose the optimal information structure. In the corresponding solution of the mechanism design problem, they show that the optimal mechanism is always one of pure bundling. Mishra, Patil, and Pavan (2025) considers a robust regulation problem where the buyer is uncertain both about their own utility function and the distribution over (constant) marginal cost of the seller. The buyer is then following a two-step procedure. They first identify the set of mechanisms that provides the highest welfare guarantee against the set of all possible value and cost functions, and then in a second step choose from the short list of mechanisms the one that maximizes the welfare against the conjectured value and cost. The setting of uncertainty, as well as the analysis, is motivated by ambiguity and misspecification concerns on both sides of the market, and thus quite distinct from our perspective.

2 Model

2.1 Payoffs

A (large) buyer procures quantity $q \in \mathbb{R}_+$ of a product from a seller. The buyer's utility gross of transfers is increasing and concave, denoted by:

$$u(q) = \frac{q^\sigma}{\sigma}, \tag{1}$$

with $\sigma \in (0,1)$. The variable $q \in \mathbb{R}_+$ can also be interpreted as the quality of a product. The seller's cost of providing quantity q is given by a cost function:

$$c : \mathbb{R}_+ \to \mathbb{R}_+.$$

The cost function is assumed to be (weakly) increasing and (weakly) convex. Buyer and seller have net utility and profit functions that are quasi-linear in transfers t. We maintain these assumptions on the utility and cost functions throughout the paper, except in Section 4.2 where we examine the robustness of our results.

The buyer's demand $D(p)$ for the good at uniform price p is denoted by:

$$D(p) \triangleq \arg\max_q \{u(q) - pq\}.$$

From the first-order condition, we obtain that this is given by:

$$D(p) = (u')^{-1}(p),$$

The power utility function means the buyer demand has a constant (negative) elasticity of demand:

$$\varepsilon_D(p) = -\frac{\frac{D(p)}{dp}}{\frac{D(p)}{p}} = \frac{1}{1-\sigma} \in (1, \infty).$$

Analogously, the seller's supply function is

$$S(p) = (c')^{-1}(p).$$

That is, at per-unit price p the seller's supplied quantity is $S(p)$. The assumption that the cost is convex corresponds to an increasing supply.

2.2 Private Information and Mechanism

The seller has private information about their cost function. The buyer does not know the seller's cost function. We denote by $\mathcal{C}$ the set of feasible cost functions that the buyer considers possible. For example, the class $\mathcal{C}$ of feasible cost functions could consist of all linear or all convex cost functions. For now, we only require that the cost is not identical to 0 so that the efficient social surplus is finite.

The buyer chooses a mechanism that elicits the private information of the seller and determines the level of production q. By the taxation principle (see Proposition 1, Guesnerie and Laffont (1984)), it is without loss to consider the indirect mechanism described by a nonlinear payment rule:

$$t : \mathbb{R}_+ \to \mathbb{R}_+.$$

We assume that $t(0) = 0$, which guarantees that the indirect mechanism satisfies the participation constraint. We write T for the set of mechanisms.

Now the seller chooses a quantity q depending on her cost function c and the mechanism t:

$$q(c, t) \triangleq \arg\max_{q \in \mathbb{R}_+} \{t(q) - c(q)\}.$$

If multiple maximizers exist, we assume that the seller chooses the smallest one. Hence, the choice is adversarial from the buyer's perspective. The precise way in which the seller breaks indifferences makes no substantive difference for the analysis.

2.3 The Buyer Surplus Ratio Guarantee

The buyer does not know the true cost function, only knowing that the true cost c is within the class $\mathcal{C}$ of cost functions. In particular, the buyer does not have any prior belief over cost functions.

In the absence of a prior distribution over feasible cost functions $\mathcal{C}$, the buyer cannot compute the *expected* buyer surplus of any given mechanism. In the absence of a prior distribution, we will compare the realized buyer surplus with the social surplus under complete information. The buyer, therefore, seeks to identify the mechanism t that guarantees them the largest possible share of the efficient social surplus as buyer surplus. The efficient social surplus is given by the maximum attainable utility net of cost under complete information:

$$W(c) \triangleq \max_{q \in \mathbb{R}_+} \{u(q) - c(q)\}.$$

The efficient social surplus depends only on the realized cost function c but not the mechanism t. The efficient quantity supplied is denoted by:

$$q^*(c) \triangleq \arg\max_{q \in \mathbb{R}_+} \{u(q) - c(q)\},$$

which will be uniquely defined.

The buyer surplus U is determined by the supply $q(c,t)$ and the payment $t(q)$ for the supplied quantity q:

$$U(c,t) \triangleq u(q(c,t)) - t(q(c,t)). \tag{2}$$

The seller surplus, or profit, Π is given by

$$\Pi(c,t) \triangleq t(q(c,t)) - c(q(c,t)). \tag{3}$$

We evaluate the performance of a mechanism t using the buyer surplus ratio guarantee:

$$\min_{c \in \mathcal{C}} \frac{U(c,t)}{W(c)}.$$

If there was no private information regarding the cost c, then the buyer would be able to extract the full efficient social surplus $W(c)$ and the buyer surplus ratio would attain its maximum possible value of 1. The buyer surplus ratio guarantee is the proportion of the efficient social surplus that the mechanism is guaranteed to provide.

The competitive ratio is the maximum possible buyer surplus ratio guarantee:

$$\max_{t \in T} \min_{c \in \mathcal{C}} \frac{U(c,t)}{W(c)}. \tag{4}$$

A mechanism t is optimal if it attains the competitive ratio. A strictly positive competitive ratio offers the buyer a guarantee that is proportional to the efficient social surplus, irrespective of the scale of the problem. In Section 6, we provide a detailed discussion relating the competitive ratio to other decision criteria used in the presence of non-Bayesian uncertainty.

The max min competitive ratio has an associated min max that is weakly larger:

$$\max_{t\in T}\min_{c\in\mathcal{C}}\frac{U\left(c,t\right)}{W\left(c\right)}\leq\min_{c\in\mathcal{C}}\max_{t\in T}\frac{U\left(c,t\right)}{W\left(c\right)}. \tag{5}$$

With deterministic price schedules and costs, there is typically a gap between max min and min max, but when there is no gap, often requiring stochastic strategies, then we refer to the resulting solution as a *saddle point*. After we provide Theorem 1 we discuss the performance of stochastic mechanisms.

3 The Linear Cost Environment

We first illustrate our results with the class of linear cost functions. The seller produces quantity $q \in \mathbb{R}_+$ with a linear cost function

$$c\left(q\right)=c\cdot q,\quad c\in\mathbb{R}_{+}.$$

We abuse notation by writing c for the constant marginal cost in this Section. The constant marginal cost $c \in \mathbb{R}_+$ is known to the seller, but unknown to the buyer. In this environment, the relevant properties of the buyer and seller are summarized in a single parameter each, the utility exponent σ of the utility function (see (1)) and the marginal cost c, respectively.

We proceed as follows in this Section. We first restrict attention to *constant share mechanisms*, where the transfer is a constant fraction of the buyer's utility. We show that this class of mechanisms gives rise to a positive ratio guarantee. We then show that the ratio guarantee from the constant share mechanism is the maximum that can be attained by *any* mechanism, and thus attains the competitive ratio. Finally, we exhibit a specific class of distributions over the cost parameter c, and show that the constant share mechanism can be rationalized as the Bayes optimal mechanism against this distribution. As a by-product, we prove the existence of a saddle point. This provides an indirect proof that a constant share mechanism attains the competitive ratio. Hence, we provide two distinct arguments for the optimality of constant share mechanisms. The direct argument is closer to the arguments we use in the general case, so it is a better illustration of the general arguments we later provide, even if, for this case, we have a shorter (indirect) proof using a specific Bayesian distribution over cost.

The main takeaway from this section is that constant share mechanisms perform well. We also provide strong economic intuitions for why they perform well.

3.1 Constant Share Mechanisms

We first ask whether a constant surplus share mechanism allows the buyer to realize a significant buyer surplus in the absence of information about the true cost of the seller. A constant share mechanism takes the form

$$t(q) = z \cdot u(q), \; z \in (0,1), \tag{6}$$

so the buyer compensates the seller with a constant fraction $z \in (0,1)$ of his utility. Under this constant share mechanism, the buyer retains $(1-z)\,u(q)$ as their buyer surplus.

The seller does not know that the mechanism is a constant share mechanism (i.e., the seller knows t but does not care about z and u). But the given the constant share mechanism, we know that seller surplus will equal :

$$\max_{q \in \mathbb{R}_+} \left\{ \frac{z}{\sigma} q^{\sigma} - cq \right\}, \tag{7}$$

Thus the quantity supplied as a function of the seller's realized cost c and the share z is

$$q(c,z) = \left(\frac{z}{c}\right)^{\frac{1}{1-\sigma}}. \tag{8}$$

The optimal quantity supplied $q(c,z)$ is decreasing in the marginal cost c and increasing in the seller share z. The resulting buyer surplus is

$$U(c,z) = \frac{1-z}{\sigma} \left(\frac{z}{c}\right)^{\frac{\sigma}{1-\sigma}}. \tag{9}$$

The constant-share mechanism has several noteworthy features. It guarantees that the seller with any finite marginal cost makes a sale, as $q(c,z) > 0$ for all $c \in \mathbb{R}_+$. Thus the mechanism (6) does not exclude any seller. Moreover, the supply of q is increasing in the share z of the surplus conceded as indicated by (8). The seller surplus is given by:

$$\Pi(c,z) = \frac{z(1-\sigma)}{\sigma} \left(\frac{z}{c}\right)^{\frac{\sigma}{1-\sigma}}.$$

The expression for buyer surplus (9) indicates that there is a clear trade-off in the choice of the optimal share z. A larger share z results in larger supply of q but the buyer surplus is a smaller share $1-z$ of the gross buyer utility. The share z that maximizes the buyer surplus ratio for any realized cost c is given by:

$$z^* = \sigma, \tag{10}$$

that is, the optimal share z^* matches the exponent σ of the power utility function.

Thus, if we *were to restrict* attention to constant share mechanisms, the optimal share for the buyer is given by (10). In other words, for any realization of c, the largest (absolute) buyer surplus is attained by $z^* = \sigma$. But of course, for any realized cost c, there exists a non-constant mechanism that could guarantee a higher buyer surplus. If the buyer knew the realization of c, he could offer a compensation equal to the cost of providing the socially efficient quantity, as long as this quantity is supplied, and zero otherwise. This is the outcome that would prevail in *first-degree price discrimination.*

But in the absence of the information about the realized cost c or any distributional information about the cost, it is difficult to say how well the buyer does by adopting the best constant sharing rule $z^* = \sigma$ relative to any other possible non-constant share mechanism.

3.2 The Ratio Guarantee with Constant Share Mechanisms

We assume the buyer compares the buyer surplus with the social surplus under complete information. Social surplus is maximized by quantity

$$q^*(c) = \left(\frac{1}{c}\right)^{\frac{1}{1-\sigma}},$$

and the efficient social surplus is given by:

$$W(c) = \frac{1-\sigma}{\sigma}\left(\frac{1}{c}\right)^{\frac{\sigma}{1-\sigma}}. \tag{11}$$

Now the buyer surplus ratio of (9) and (11), respectively, is:

$$\frac{U(c,z)}{W(c)} = \frac{\frac{1-z}{\sigma}\left(\frac{z}{c}\right)^{\frac{\sigma}{1-\sigma}}}{\frac{1-\sigma}{\sigma}\left(\frac{1}{c}\right)^{\frac{\sigma}{1-\sigma}}} = \frac{1-z}{(1-\sigma)\, z^{\frac{\sigma}{\sigma-1}}}. \tag{12}$$

This ratio does not depend on the cost realization c. Therefore, we can find a sharing rule z^* which maximizes the ratio and only depends on the exponent σ of the utility function. We summarize these findings in the proposition below.

Proposition 1 (Constant Share Mechanisms)
For all constant share mechanisms, $t(q) = z \cdot u(q)$, we have that

$$\min_{c\in\mathbb{R}_+} \frac{U(c,z)}{W(c)} \leq \sigma^{\frac{\sigma}{1-\sigma}}.$$

The inequality is an equality if and only if $z^* = \sigma$, in which case the bound is attained for all $c \in \mathbb{R}_+$.

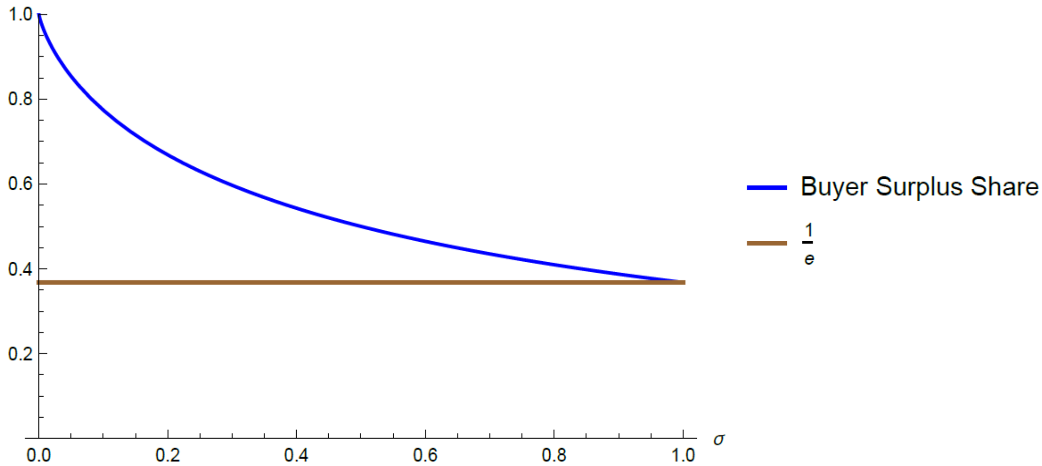

Figure 1: Buyer Surplus Guarantee as Share of Social Surplus

We thus find that the ratio guarantee for the optimal constant share mechanism. The buyer surplus share is bounded below by $1/e \approx 0.37$ as $\sigma^{\frac{\sigma}{1-\sigma}}$ is decreasing in σ and converges to $1/e$ as $\sigma \to 1$. In this limit, the demand becomes perfectly elastic. This is illustrated in Figure 1. While in this setting the ratio guarantee is at least $1/e$, we will see that the ratio guarantee converges to 0 as the σ converges to 1 when the buyer considers all convex cost functions.

In fact, the constant share mechanism also provides a remarkable performance in terms of efficiency. The *joint surplus* that buyer and seller can achieve with this constant share mechanism is bounded below by $2/e \approx 0.74$, thus leaving buyer and seller with a substantial part of the social surplus despite the inefficiency due to private information. In Figure 2 we display the joint behavior of buyer and seller surplus.

3.3 The Competitive Ratio and General Mechanisms

In this previous section, we restricted attention to constant share mechanisms. Within this class, we identified a particular sharing rule $z^* = \sigma$ that maximizes the absolute buyer surplus and thus the buyer surplus ratio relative to the complete information social surplus. We now consider all possible indirect mechanisms. We show that even in this unrestricted class of mechanisms, the constant share mechanism $z^* = \sigma$ maximizes the ratio guarantee and thus attains the competitive ratio.

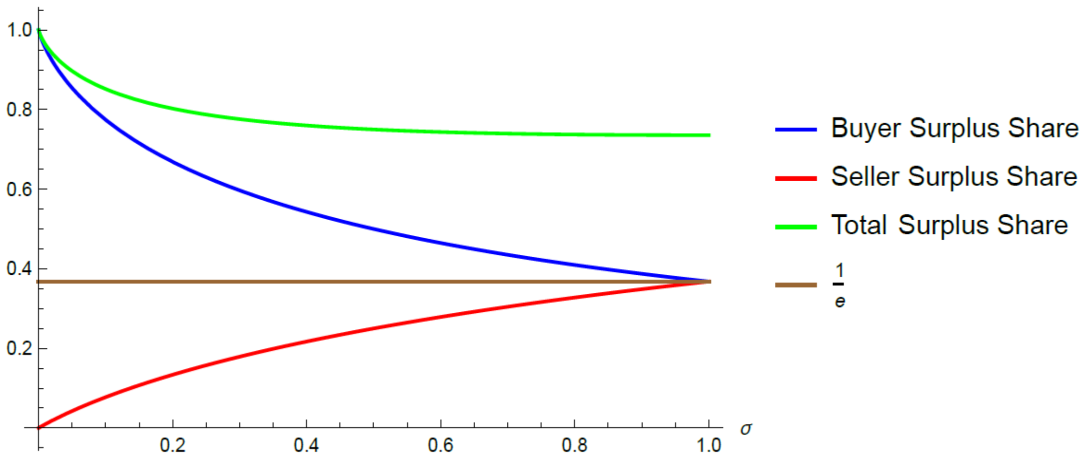

Figure 2: Buyer and Seller Surplus Guarantee as Share of Social Surplus

Proposition 2 (Competitive Ratio)

For all mechanisms t, we have that:

$$\min_{c\in\mathbb{R}_+}\frac{U(c,t)}{W(c)}\leq\sigma^{\frac{\sigma}{1-\sigma}}. \tag{13}$$

The inequality is tight if and only if $t\left(q\right)=z^{*}u\left(q\right)$.

We provide a sketch of this result (a rigorous proof is given in Proposition 3 following a different route). That is, we prove that, for any $t(q)$ such that $t(q)\neq\sigma u\left(q\right)$ for some q, satisfies:

$$\min_{c\in\mathbb{R}_+}\frac{U(c,t)}{W(c)}<\sigma^{\frac{\sigma}{1-\sigma}}$$

To simplify the exposition, we consider only the case that the transfer $t(q)$ is concave, and the share of surplus conceded to the seller

$$z(q)\triangleq\frac{t(q)}{u(q)}$$

is strictly decreasing at some $q=\widehat{q}$. (The argument we provide next mirrors the central argument used to prove Theorem 1, which we later provide. There we also consider the cases omitted in current argument and allow for general nonlinear cost functions).

To make the notation more compact, we define:

$$\widehat{z}\triangleq z(\widehat{q})\text{ and }\widehat{c}\triangleq t'(\widehat{q}).$$

Hence, by construction, when the cost is $c(q) = \widehat{c}q$, we have that:

$$\widehat{q} = \arg\max\{t(q) - c(q)\}$$

We obtain this because by assumption $t(q)$ is concave so the objective function is concave, and by construction $\widehat{q}$ satisfies the first-order condition (that is, $c'(\widehat{q}) = t'(\widehat{q})$). That is, $\widehat{z}$ is the share of surplus shared by the buyer if the quantity sold is $\widehat{q}$ and $\widehat{c}$ is the marginal cost that rationalizes $\widehat{q}$ as the optimal supply of quality.

We now note that $z(q)$ is strictly decreasing at $\widehat{q}$, thus $z'(\widehat{q}) < 0$, and so by construction:

$$t'(\widehat{q}) = \widehat{z}u'(\widehat{q}) + z'(\widehat{q})\, u(\widehat{q}) < \widehat{z}u'(\widehat{q}).$$

This, in turn, implies that:

$$q(\widehat{c}, t) < q(\widehat{c}, \widehat{z}). \tag{14}$$

In other words, under transfer t and cost $c(q) = \widehat{c}q$, the supply is smaller under transfer $t(q)$ than under the constant share mechanism $\widehat{t}(q) = \widehat{z}u\,(q)$. We thus have that:

$$\frac{U(\widehat{c}, t)}{W(\widehat{c})} < \frac{U(\widehat{c}, \widehat{z})}{W(\widehat{c})} \leq \frac{U(\widehat{c}, z^*)}{W(\widehat{c})} = \sigma^{\frac{\sigma}{1-\sigma}}.$$

The strict inequality follows from (14) while the weak inequality follows from the definition of z^* (note that the denominator does not depend on $\widehat{z}$). Finally, the equality follows from Proposition 1.

This proves the result for the specific case we set out to prove. We assumed that $z(q)$ is decreasing at some q; when this is not satisfied one can repeat a similar argument by considering the limits of $q \to \{0, \infty\}$. Addressing the case in which $t(q)$ is not concave is technically more cumbersome, but the arguments based on concavification remain very similar. Overall, the intuition is that whenever the buyer departs from a constant share mechanism, we can find a cost realization $\widehat{c}$ in which the seller supplies a quantity smaller than $q(c, \widehat{z})$.

3.4 The Constant Share Mechanism is a Bayes Optimal Mechanism

We have shown that the constant sharing mechanism attains the highest guarantee among all feasible mechanisms. We now ask whether the constant share mechanism may, in fact, constitute an optimal solution for a common prior over the cost function. We now propose a specific distribution over linear cost functions. We find that the Bayes optimal rule for this distribution is again the simple sharing rule $z^* = \sigma$. With the existence of the Bayes solution, we can then establish an equivalence between the maximin and the minimax problem as described earlier in (5).

We consider a power distribution over the marginal cost c:

$$F(c) = \left(\frac{c}{\overline{c}}\right)^{\alpha},$$

with finite support $c \in [0, \overline{c}]$ and a positive exponent $\alpha > 0$. We first note that the expected efficient social welfare is finite if and only if $\alpha > \sigma/(1-\sigma)$. That is:

$$\int W(c)dF(c) < \infty,$$

if and only if $\alpha > \sigma/(1-\sigma)$. We denote the Bayes-optimal mechanism under the power distribution with exponent α by:

$$t_\alpha \triangleq \max_t \int U(c,t)dF(c).$$

In the procurement problem, the virtual utility determines the optimal allocation pointwise:

$$u'(q) - \frac{F(c)}{f(c)} - c = 0.$$

It is then possible to compute the optimal mechanism $t_\alpha(q)$ as the indirect mechanism from incentive and participation constraints. In particular, we find that:

$$t_\alpha(q) = \frac{\alpha}{1+\alpha}u(q) - \frac{1-\sigma}{\sigma}\left(\frac{\alpha}{1+\alpha}\right)^{\frac{1}{1-\sigma}}\left(\frac{1}{\overline{c}}\right)^{\frac{\sigma}{1-\sigma}}.$$

Hence, the Bayes optimal transfer is a constant share mechanism minus a constant that is constructed to leave a seller with marginal cost $\overline{c}$ with zero transfer (and sellers with higher cost would be excluded, but they have zero probability under the considered distribution). As we take the limit $\overline{c} \to \infty$, the constant converges to zero and the mechanism converges pointwise to a constant share mechanism.

Proposition 3 (Bayes Optimal Mechanism)
The ratio of buyer surplus to social welfare attained by the Bayes optimal mechanism satisfies:

$$\lim_{\alpha \downarrow \frac{\sigma}{1-\sigma}} \frac{\int U(c,t_\alpha)dF(c)}{\int W(c)dF(c)} = \sigma^{\frac{\sigma}{1-\sigma}}. \tag{15}$$

Furthermore, the optimal mechanism satisfies:

$$\lim_{\substack{\alpha \downarrow \frac{\sigma}{1-\sigma} \\ \overline{c} \to \infty}} t_\alpha(q) = z^* u(q).$$

We thus obtain that the optimal constant share mechanism is indeed the Bayes optimal mechanism for some distribution over marginal cost (by taking the appropriate limits). Furthermore, the ratio of the efficient social surplus attained by the Bayes optimal rule is the same as that obtained by the optimal constant share mechanism. Now, we know that in general, the solution of the above max min problem is (weakly) below the solution of the corresponding min max problem:

$$\max_{t} \min_{F \in \Delta \mathbb{R}_+} \frac{U(F,t)}{W(F)} \leq \min_{F \in \Delta \mathbb{R}_+} \max_{t} \frac{U(F,t)}{W(F)}.$$

Hence, (13) follows as an immediate implication from (15). One can then conclude the proof of Proposition 2 by noting that (13) is tight when the mechanism is a constant share mechanism with sharing constant z^* (Proposition 1).

4 The Nonlinear Cost Environment

We now analyze the performance of constant share mechanisms in an environment with general nonlinear cost functions. We first consider the set of all (weakly) convex cost functions and show that the constant share mechanism guarantees the maximum possible share of the efficient social surplus– that is, a constant share mechanism attains the competitive ratio. Hence, we provide the natural extension of Proposition 2. In this general environment, a saddle point does not exist in deterministic strategies; that is, Proposition 3 does not extend to this setting. Yet, if we allow the buyer to randomize over mechanisms, we can find a saddle point; we relegate the details of this generalization of Proposition 3 to the appendix (as Proposition 5), and discuss some details of the saddle point at the end of Section 4.1.

In Section 4.2, we provide the most general version of our model: we relax the assumption of iso-elastic demand and convex cost. In this more general setting, we show that constant share mechanisms still guarantee a positive ratio guarantee. However, we do not show that a constant share mechanism can attain the competitive ratio. Hence, while in this more general setting our results are significantly weaker, we can still establish that constant share mechanisms have desirable properties.

4.1 The Environment with Convex Cost

We denote by $\mathcal{C}_{cx}$ the set of all convex functions:

$$\mathcal{C}_{cx} \triangleq \{c : \ c(q) \text{ is increasing and convex in } q\}.$$

Unless we explicitly state otherwise, all notions of monotonicity and convexity are weak rather than strict. Throughout this subsection, we assume that the buyer considers all convex cost functions to be plausible, that is, $\mathcal{C} = \mathcal{C}_{cx}$.

Given the large class of possible cost functions that the buyer faces, namely C_{cx}, it will be critical to determine which cost functions may minimize the ratio guarantee. Given a mechanism t and a cost function, the quantity supplied by the seller $q(c,t)$ is determined by the optimality condition that marginal revenue equals marginal cost:

$$t'\left(q(c,t)\right) = c'\left(q(c,t)\right).$$

The denominator of the ratio guarantee is the social welfare $W\left(c\right)$ and hence to compute this ratio, the whole cost function is relevant, and not just the marginal cost at the margin (i.e., $c'(q(c,t))$). For any given mechanism t we find the lowest convex cost $\widehat{c}$ that is consistent with keeping the quantity supplied unchanged. That is, we find:

$$\widehat{c} \triangleq \arg\max_{\widetilde{c} \in \mathcal{C}_{cx}} W(\widetilde{c}) \qquad \text{subject to: } q\left(c,t\right) = q\left(\widetilde{c},t\right). \tag{16}$$

To make the notation more compact, we define the equilibrium choices under c and t:

$$\widehat{q} \triangleq q\left(c,t\right) \text{ and } \widehat{\gamma} \triangleq c'(\widehat{q}).$$

That is, $\widehat{q}$ is the targeted quantity supplied under $\widehat{c}$ and $\widehat{\gamma}$ is the marginal cost at the quantity supplied. Given the quantity supplied, the mechanism, and the cost, we can graphically find the buyer surplus, profits, and deadweight loss, as illustrated in Figure 3. In the figure, we plot the marginal utility, marginal transfer, and marginal cost, so the buyer surplus, profits, and deadweight loss correspond to the appropriate areas between the curves.

The solution to (16) is:

$$\widehat{c}(q) \triangleq \begin{cases} 0, & \text{if } q < \widehat{q}; \\ \widehat{\gamma}(q - \widehat{q}), & \text{if } q \geq \widehat{q}; \end{cases} \tag{17}$$

which is the lowest convex cost that is consistent with the quantity supplied being $\widehat{q} = q(c,t)$. That is, the marginal cost of providing $q \leq \widehat{q}$ is zero (and hence so is the total cost), and afterwards it increases at constant rate $\widehat{\gamma}$. Here the marginal cost at quantities $q \geq q(c,t)$ are disciplined by the fact that the cost is convex, so the marginal cost cannot be smaller than $\widehat{\gamma}$. In Figure 4 we the critical cost function and the additional profits and deadweight loss relative to c. Note that by construction the supply is the same as with cost c, so buyer surplus stays the same. Hence, the set

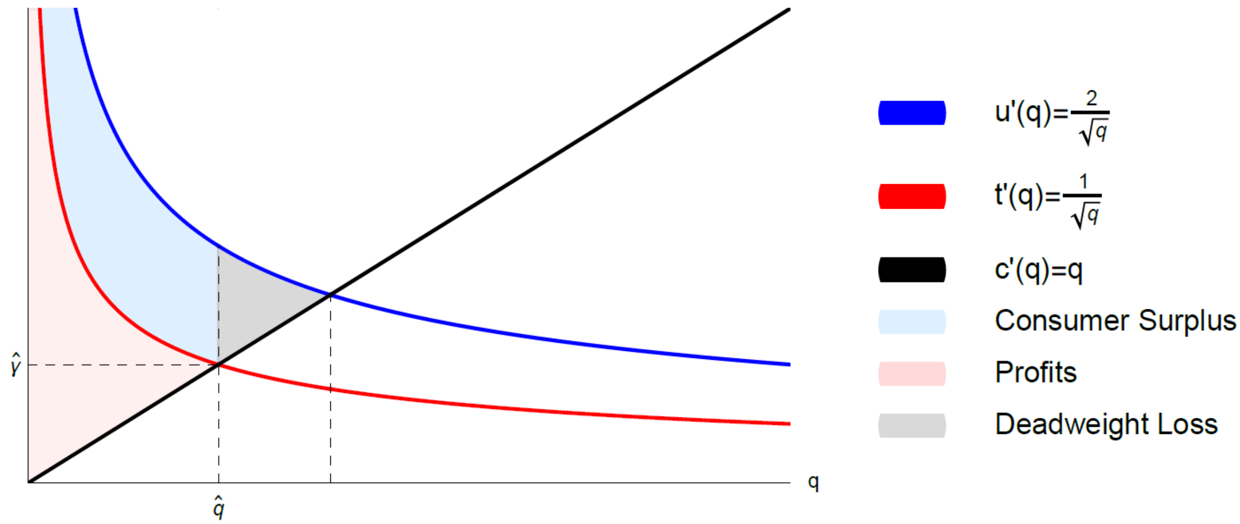

Figure 3: Quadratic Cost Function and Piecewise Linear Cost Function

of critical cost functions that minimize the competitive ratio will consist of piece-wise linear cost functions with a single kink at some $\widehat{q}$ with zero cost of providing $q \leq \widehat{q}$ and marginal cost $t'(\widehat{q})$ of providing $q \geq \widehat{q}$.

With this conjecture regarding the critical cost functions, for now, argued informally but eventually proven formally, we can now construct a bound based on a constant share mechanism. For any constant share mechanism

$$t(q) = z \cdot u(q),$$

the ratio guarantee can be computed with the above piecewise linear cost function, similar to the case of the linear cost function in Section 3.2. It is given by:

$$\frac{U(c,t)}{W(c)} = \frac{1-z}{(1-\sigma)z^{\frac{\sigma}{\sigma-1}} + \sigma z}.$$

Relative to the ratio earlier (12), it contains an additional term in the denominator, which reflects the zero cost of the initial units, which increases the efficient social surplus.

This leads us now to consider the following mechanism:

$$t(q) = z^*(\sigma)u(q), \tag{18}$$

where $z^*(\sigma)$ solves the problem of maximizing the competitive ratio under the specific piecewise linear cost function given by (17):

$$z^*(\sigma) \triangleq \arg\max_{z\in[0,1]} \frac{1-z}{(1-\sigma)z^{\frac{\sigma}{\sigma-1}} + \sigma z}. \tag{19}$$

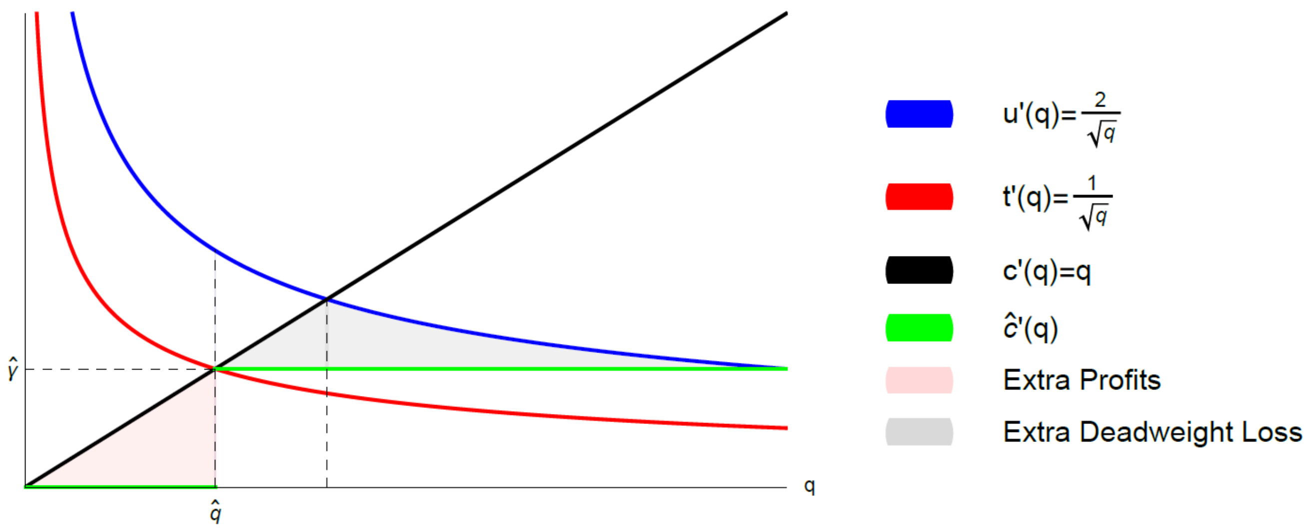

Figure 4: Buyers' Surplus and Social Surplus

The above ratio is strictly quasi-concave, so $z^*(\sigma)$ is uniquely defined. Hence, the candidate mechanism is a constant surplus sharing rule. To make the notation more compact, it is useful to provide notation for the value attained by (19):

$$B(\sigma) = \frac{1 - z^*(\sigma)}{(1-\sigma)(z^*(\sigma))^{\frac{\sigma}{\sigma-1}} + \sigma z^*(\sigma)}. \tag{20}$$

We now show that $B(\sigma)$ is indeed an upper bound on the competitive ratio and that this bound is tight.

Theorem 1 (Buyer Surplus Guarantee)
For every mechanism t,

$$\inf_{c \in \mathcal{C}_{cx}} \frac{U(c,t)}{W(c)} \leq B(\sigma). \tag{21}$$

Furthermore, the inequality is attained as an equality if and only if

$$t(q) = z^*(\sigma)u(q).$$

The result provides a mechanism that can guarantee a fraction of the efficient social surplus regardless of the cost function. Furthermore, the mechanism is optimal in the sense that it maximizes the share of the efficient total surplus across all possible mechanisms. The transfer rule is simple in the sense that it consists of sharing a constant fraction of the utility $u\,(q)$.

To prove Theorem 1, we first prove that, if $t(q) \neq z^*(\sigma)u(q)$ then we can find $(\widehat{\gamma}, \widehat{q})$ such that the ratio evaluated at $\widehat{c}$ (see (17)) is strictly smaller than $B(\sigma)$. We then show that, if $t = z^* u(q)$ then the ratio is always (weakly) larger than $B(\sigma)$.

We are then left with providing intuition for the fraction of utility that the buyer shares with the seller, that is, $z^*(\sigma)$. We plot the behavior of the surplus sharing rule as a function of the exponent σ below in Figure 5. As an approximation, we can see that $z^*(\sigma) \approx \sigma$, and in fact,

$$|z^*(\sigma) - \sigma| \leq 0.12.$$

Hence, the optimal rule prescribes that the buyer should share with the seller a fraction that is almost the same as the exponent of the utility function. For intuition, suppose first that the exponent $\sigma \approx 0$. Then, an ε supply of the product can guarantee a utility of (almost) 1 to the buyer, which is essentially the maximum utility that the buyer can attain. Hence, the buyer can offer a small transfer as long as there is some positive supply, which is attained by $t(q) \approx \sigma u(q)$. Note that $u'(0) = \infty$, so the buyer will always obtain some ε amount of the product. By contrast, suppose that the exponent is $\sigma \approx 1$ and thus the utility is near linear $u(q) = q$. Suppose the buyer offers $t(q) = mq$, for some $m < 1$. Then, the buyer surplus would be zero if the cost is $c(q) = (1 + m)q/2$, but the social surplus would still be positive (in fact, infinite if $u(q)$ is exactly linear). Hence, we can see that to guarantee a positive ratio as $\sigma \to 1$, we must have that $t(q)$ converges to $u(q)$.

We plot the competitive ratio guarantee in Figure 6. We can see that $B(\sigma)$ is decreasing in σ. Furthermore, a linearly decreasing function in the exponent σ seems like a good approximation:

$$B(\sigma) \approx 1 - \sigma.$$

More precisely, one can check numerically that:

$$|B(\sigma) - (1 - \sigma)| \leq 0.16.$$

Hence, as an order of magnitude it is useful to think that the competitive ratio guarantee is almost linear in σ. In contrast to Proposition 2, we now have that the competitive ratio converges to 0 as σ approaches 1. We provide an intuition after we provide Theorem 2 because it will be easier to explain this difference in a more general environment.

Theorem 1 provides the competitive ratio, which is attained by a constant share mechanism. This simple policy generated a positive ratio guarantee even when the buyer was restricted to

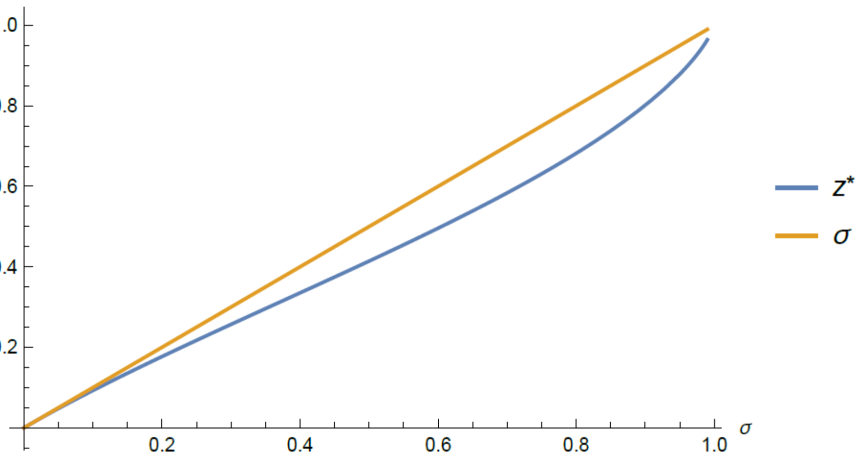

Figure 5: Surplus share $z^*(\sigma)$ as function of the elasticity σ.

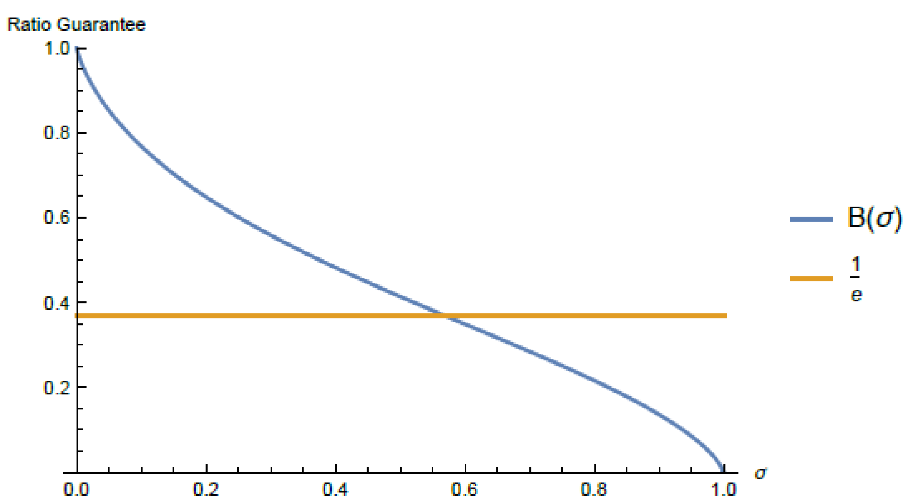

Figure 6: Surplus Share Guarantee as a Function of Elasticity

use a deterministic mechanism. In the Appendix (Section 9) we provide an alternative version of this result in which we allow the buyer to randomize over all possible mechanisms. We find in Proposition 5 that the mechanism that attains the competitive ratio is a randomization over constant share mechanisms. Furthermore, we characterize the saddle point, analogously to Proposition 3. Somewhat surprisingly, the competitive ratio when the buyer is allowed to randomize does not differ qualitatively from the one characterized in Theorem 1.

4.2 General Cost and Utility Functions

Theorem 1 used the assumption that the utility of the quantity is given by a power function and the cost function is weakly convex (in addition to being increasing). The constant elasticity of demand and the weak convexity have a natural economic interpretation if the units of q are meaningful, e.g., they represent quantity. But if there are no natural units, there may not be a natural economic interpretation, e.g., if q represents quality. We now relax the assumptions on utility and cost and replace them with an assumption that is unit-free. The basic idea of this section is to obtain a result in this more general environment by reducing it to arguments we have established in Theorem 1.

In this section, we let u be any increasing twice differentiable function, and let $\mathcal{C}$ consist of all almost everywhere differentiable increasing functions.

For any given utility function $u(q)$ of the buyer, we suggest a change of variable that leads to a new utility function that can be stated as a power utility. The change of variable induced by the utility function then generates cost functions under the new variable for every cost function $c \in \mathcal{C}$ in the feasible set $\mathcal{C}$. We shall derive conditions, essentially conditions on demand and supply elasticity, under which the cost functions under the new variable preserve monotonicity and weak convexity. This allows us to apply the arguments provided in Theorem 1 much beyond the environment that we consider there.

In more detail, we consider an increasing utility function $u(q)$ and define a new variable based on it:

$$x \triangleq (\sigma u(q))^{\frac{1}{\sigma}} \tag{22}$$

for some $\sigma \in (0,1)$. From the definition, we obtain that q as a function of the newly defined variable x:

$$q = u^{-1}\left(\frac{x^{\sigma}}{\sigma}\right).$$

Hence, the utility function in terms of this newly defined variable can be written as a power utility function:

$$\widehat{u}\left(x\right) \triangleq \frac{x^{\sigma}}{\sigma}. \tag{23}$$

We can then describe any utility function u as a power function under the right change of variables.

The bound in Theorem 1 depended on the exponent σ of the utility function. However, in the change of variables suggested by (23) we have not yet imposed any discipline on the possible values of the exponent σ. However, we cannot apply Theorem 1 unless (after the change of variables) every cost function $c \in \mathcal{C}$ is convex. We thus have to find which values of σ in (23) are consistent with a convex cost function. More precisely, we consider the cost function:

$$\widehat{c}\left(x\right) \triangleq c\left(u^{-1}\left(\frac{x^{\sigma}}{\sigma}\right)\right), \tag{24}$$

which is the cost function written in terms of the variable x instead of q. The question then arises as to whether under the transformation suggested by x, any newly defined cost function $\widehat{c}\left(x\right)$ based on $c\left(q\right)$ is convex in x. We provide a condition based on the curvature of the underlying functions $u\left(q\right)$ and $c\left(q\right)$ under which $\widehat{c}\left(x\right)$ preserves monotonicity and convexity.

For a given cost and utility function, $c\left(q\right)$ and $u\left(q\right)$ we consider a joint index of the curvature of the cost and utility function defined by:

$$\delta(q,c) \triangleq \frac{u(q)}{qu'(q)}\left(\frac{c''\left(q\right)q}{c'\left(q\right)} - \frac{qu''(q)}{u'(q)}\right). \tag{25}$$

A desirable property of this measure of curvature is that it is a unit-free measure. That is, for any increasing function, say $q = h(x)$, if we measure quality by units of x rather than q the measure δ does not change. More precisely, we define:

$$\widetilde{u}(x) \triangleq u(h(x)) \text{ and } \widetilde{c}(x) \triangleq c(h(x)),$$

and $\widetilde{\delta}$ appropriately defined in terms of x (using $\widetilde{u}$). It is simple to verify that:

$$\widetilde{\delta}(x,\widetilde{c}) = \delta(q,c).$$

Hence, non-linear transformations of the units in which q is measured do not change the measure of curvature (it is easy to verify this is not true for other operations, like derivatives or elasticities). We can then write δ more conveniently as follows:

$$\delta(q,c) = \frac{\overline{u}\frac{d^2 c(u^{-1}(\overline{u}))}{d\overline{u}^2}}{\frac{dc(u^{-1}(\overline{u}))}{d\overline{u}}} \tag{26}$$

Hence, δ provides the curvature of the cost– measure by the elasticity of marginal cost– when measured in utils. In fact, this is also a measure of the curvature of the social welfare when measured in utils because, by construction, the utility function is linear when measuring quality in utils. We denote by $\underline{\delta}$ the lowest value attained by $\delta(q,c)$:

$$\underline{\delta} \triangleq \inf_{(q,c)\in\mathbb{R}_+\times\mathcal{C}} \delta(q,c) \text{ with } (q,c) \text{ such that } q \leq \overline{q}(c).$$

In other words, we find the infimum over quantity and costs, considering only quantities below the efficient level. We assume that:

$$\underline{\delta} > 0. \tag{27}$$

This implies that the law of diminishing returns applies to our setup when cost is measured in utils (that is, (26) is positive). We can now discipline the choice of the exponent σ in the change of variable given by (23).

Lemma 1 (Convexity of the Transformed Cost)
The cost function $\widehat{c}(x)$ (see(24)) is convex for every $c \in \mathcal{C}$ and at every $x \in \mathbb{R}_+$ if and only if

$$\sigma \geq \frac{1}{1+\underline{\delta}}. \tag{28}$$

We can now apply the results in Theorem 1 by performing the change of variables in (22) subject to σ satisfying (28), and in fact, we will get the sharpest bound when (28) is satisfied with equality. We thus define:

$$\widehat{\sigma} \triangleq \frac{1}{1+\underline{\delta}} \tag{29}$$

Thus, a strongly positive index allows for a lower exponent $\widehat{\sigma}$ and in turn for an improved bound $B(\widehat{\sigma})$.

Theorem 2 (Competitive Ratio with General Cost and Utility Functions)
The mechanism $t(q) = z^*(\widehat{\sigma})u(q)$ attains at least ratio:

$$\min_{c\in\mathcal{C}} \frac{U(c,t)}{W(c)} \geq B(\widehat{\sigma})$$

This theorem generalizes Theorem 1 by showing that a constant surplus sharing rule provides a positive ratio guarantee in general utility and cost functions environments. In particular, our simple mechanism can guarantee a share $B(\widehat{\sigma})$ of the efficient social surplus (with $\widehat{\sigma}$ appropriately defined in terms of $\underline{\delta}$). We note that (27) guarantees that $\widehat{\sigma} \in (0,1)$, which in turn guarantees that the ratio guarantee is indeed strictly positive. When (27) is not satisfied, $\widehat{\sigma}$ attains values greater than 1. In this case, our results in Section 4.1 do not apply, and consequently, we are not able to guarantee of positive ratio.

Intuitively, the buyer can guarantee a positive ratio only if social surplus is sufficiently concave in q (measured by the concavity of (24)). In contrast to Theorem 1, we now only show that $t = z^*(\widehat{\sigma})u(q)$ attains $B(\widehat{\sigma})$ but we do not show that this is optimal (that is, that any other mechanism attains less). The reason is that we do not necessarily have that for every convex function $\widehat{c} \in \mathcal{C}_{cx}$, there exists $c \in \mathcal{C}$ such that (24) is satisfied. In other words, the set of feasible cost functions $\mathcal{C}$ might be too small to guarantee that any other mechanism performs worse than $t = z^*(\widehat{\sigma})u(q)$. It is easy to verify that if the set:

$$\widehat{C} \triangleq \{\widehat{c} \in C_{cx} : \text{ there exists } c \in \mathcal{C} \text{ such that } \widehat{c} \text{ if given by (24).}\}$$

satisfies that $\widehat{C} = C_{cx}$ (note that by definition $\widehat{C} \subseteq C_{cx}$) then $t = z^*(\widehat{\sigma})u(q)$ indeed attains the competitive ratio.

5 Related Applications: Nonlinear Pricing and Regulation

We now study two problems closely related to the procurement problem. First, we examine the optimal mechanism for a regulator who maximizes a linear combination of buyer surplus and profits. Hence, unlike the procurement setting, the regulator also places positive weight on profits. Second, we examine the pricing rule for a seller facing unknown demand, the nonlinear pricing problem as defined by Mussa and Rosen (1978). In both cases, we find that a simple rule can guarantee a positive fraction of the efficient social surplus; furthermore, the simple mechanism is optimal in the sense of the competitive ratio.

5.1 Regulation

We now examine situations in which there is a regulator who maximizes a linear combination of buyer surplus and profits. The buyer surplus U and seller surplus Π are determined by the allocation $q(c, t)$ and the mechanism $t(q(c, t))$ for the output, as given by (2) and (3). We assume the regulator places a weight of 1 on buyer surplus and a weight α on profits:

$$\alpha \in [0, 1].$$

In the absence of a common prior over the class $\mathcal{C}$ of permissible cost functions, the regulator cannot compute the *expected* regulator surplus of any given mechanism across all possible cost functions. Therefore, we consider the ratio between the regulator surplus under incomplete information and the efficient social surplus under complete information. Formally, the regulator is choosing a mechanism t while nature is choosing a cost function c:

$$\sup_t \inf_c \frac{U(c,t) + \alpha \Pi(c,t)}{W(c)}.$$

The social welfare $W(c)$ depends on the realized cost function c but not on the mechanism. The regulator seeks to identify the mechanism that attains the largest ratio of weighted regulator surplus against the efficient social surplus across all possible cost functions $c \in \mathcal{C}$.

To provide our next result, we go back to our baseline model in Section 4.1. That is, we assume that $u(q) = q^\sigma/\sigma$ and $\mathcal{C}$ is the class of all increasing and convex functions. As before, we start with a constant share mechanism:

$$t(q) = z_\alpha^* u(q),$$

where now z_α^* solves:

$$z_\alpha^* \triangleq \arg\max_{z \in [0,1]} \frac{1 - (1-\alpha)z}{(1-\sigma)z^{\frac{\sigma}{\sigma-1}} + \sigma z}. \tag{30}$$

The objective function is strictly quasi-concave, so z_α^* is uniquely defined. Hence, the price schedule is a constant buyer surplus share mechanism. As before, to make the notation more compact, it is useful to provide notation for the value attained by (30):

$$B_\alpha(\sigma) \triangleq \frac{1 - (1-\alpha)z_\alpha^*}{(1-\sigma)(z_\alpha^*)^{\frac{\sigma}{\sigma-1}} + \sigma z_\alpha^*}$$

We now show that $B_\alpha(\sigma)$ is the upper bound on the ratio and that this bound is tight, thus attaining the competitive ratio.

Theorem 3 (Competitive Ratio of Regulation)
For every mechanism $t(q)$,

$$\min_{c(q)} \frac{U(c,t) + \alpha \Pi(c,t)}{W(c)} \leq B_\alpha(\sigma)$$

Furthermore, the inequality is tight if and only if $t(q) = z_\alpha^* u(q)$.

Theorem 3 establishes that Theorem 1 extends to situations in which the buyer is a regulator who seeks to maximize a linear combination of buyer surplus and profits. The current model of regulation follows closely the work of Baron and Myerson (1982) who ask how to regulate a monopolist with private information about their cost. There, the regulator, as the public agency, seeks to maximize the buyer surplus or a weighted sum of buyer surplus and monopoly profit. As the weight $\alpha \in [0, 1]$ that the regulator assigns to the profit of the firm increases, the objective of the regulator becomes closer to the benchmark of the social welfare. A natural comparative static result then emerges.

Corollary 1 (Comparative Statics of Competitive Ratio)
As the weight on profit increases, the competitive ratio $B_\alpha(\sigma)$ increases and converges to 1 as α converges to 1 for all $\sigma \in (1, \infty)$.

5.2 Nonlinear Pricing

We now consider the nonlinear pricing problem. Here, the (representative) buyer has private information about their willingness to pay and the seller has to offer a tariff in the presence of uncertainty about the willingness to pay. Thus, we invert the roles of buyer and seller. Namely there is a seller with a known cost function:

$$c(q) = \frac{\sigma - 1}{\sigma} \frac{q^\sigma}{\sigma},$$

with $\sigma > 1$. The buyer has a utility function:

$$u(q) = vq,$$

for some value (willingess-to-pay) $v \in \mathbb{R}_+$. The value v is known to the buyer but unknown to the seller. The seller offers a transfer $t(q)$ to the buyer, who then chooses a quantity that maximizes their net utility:

$$q(v, t) = \max\{u(q) - t(q)\}.$$

The profits are given by:

$$\Pi(v, t) = t(q(v, t)) - c(q(v, t)).$$

Following the same notation as before, we denote by $W(v)$ the efficient social surplus when the buyer's value is v. We thus obtain the same model as Mussa and Rosen (1978), with the additional constraint that the cost function is a power function. However, instead of characterizing the Bayes optimal mechanism, we provide a mechanism that attains a positive ratio guarantee. Furthermore, it attains the competitive ratio.

Proposition 4 (Competitive Ratio for Nonlinear Pricing)
For every mechanism $t(q)$*,*

$$\min_{v\in\mathbb{R}_+} \frac{\Pi(v,t)}{W(v)} \leq \sigma^{\frac{\sigma}{1-\sigma}}.$$

Furthermore, the inequality is tight if and only if $t(q) = \sigma c(q)$*.*

We thus obtain the same result as in Proposition 2. While the expression for the competitive ratio in terms of σ remains the same, we now have that $\sigma > 1$, so we now obtain a lower competitive ratio relative to the procurement problem. But similar to the procurement problem with linear cost, the competitive ratio in the nonlinear pricing environment satisfies the saddle-point property and can be interpreted as the solution to a Bayesian optimal nonlinear pricing problem. Thus, an analogue to Proposition 3 exists and can be found in an earlier working paper version (Bergemann, Heumann, and Morris (2023)).

As prescribed by Proposition 2, we now have that the optimal pricing is:

$$t(q) = \sigma c(q).$$

We now have that $\sigma > 1$, so as expected, the transfer is always greater than the cost. Finally, note that by construction:

$$\frac{t(q) - c(q)}{t(q)} = \frac{\sigma - 1}{\sigma}.$$

Hence, the optimal pricing rule is a constant markup rule where the principal always demands a fixed markup over cost. The common use of a constant markup rule was also already observed by Hall and Hitch (1939) who document from interviews, the use of "full-cost" or "cost-plus" pricing. Scherer and Ross (1990) notes that "Hall and Hitch and later analysts found several reasons why businessmen use cost-based rules of thumb in their pricing decisions. . . .". One is that "it is a way of coping with (essentially by ignoring) uncertainties in the estimation of demand function shapes and elasticities".

6 Alternative Criteria

We proposed a class of simple mechanisms– constant utility share mechanisms– for environments where the buyer has minimal information about the seller. The conventional approach of Bayesian utility maximization was therefore not available. We suggested that the competitive ratio may be a suitable criterion to evaluate the performance of a simple rule. In this final section, we briefly discuss

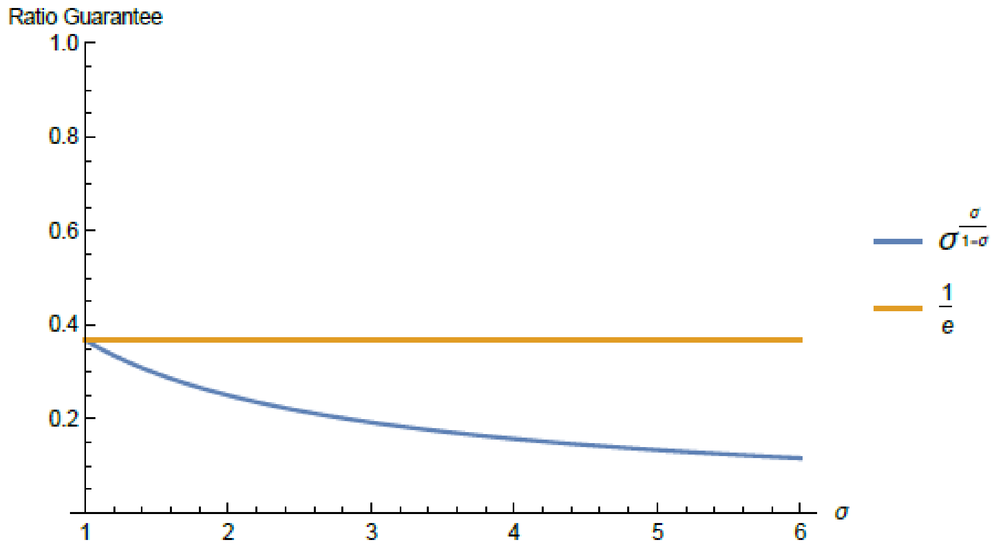

different but conceptually related criteria used in the absence of a Bayesian prior distribution. We conclude by analyzing a variation of the competitive ratio itself, where we replace the benchmark of the social surplus by the Bayes optimal utility. We show that these variations lead to qualitatively as well as quantitatively very similar results in support of constant share mechanisms.

6.1 Alternative Robustness Criteria

The *maximin ratio* is a commonly used criterion in the absence of a Bayesian prior distribution. There are related criteria such as the *maximin utility* or *maximin regret.* Anunrojwong, Balseiro, and Besbes (2025) provide a unifying framework that they refer to as λ regret and show that all of the above robust decision criteria can be obtained as special cases by choosing the weight λ in an additive objective:

$$\max_{t \in \mathcal{T}} \min_{c \in \mathcal{C}} \left\{ U\left(c, t\right) - \lambda W\left(c\right) \right\},$$

with $\lambda \in [0, 1]$. The maximin regret is obtained with $\lambda = 1$, the maximin utility with $\lambda = 0$, and the maximin ratio solution is obtained by the value of $\lambda \in (0, 1)$ at which the above max min problem attains the value of zero. We now briefly discuss the advantages of the competitive ratio in the procurement setting.

For concreteness we consider our baseline payoff environment introduced in Section 2 and change the criteria used to evaluate different mechanisms. That is, we consider a buyer who is procuring q

units of a product from a seller. The buyer's utility is:

$$u(q) = \frac{q^\sigma}{\sigma},$$

and the seller has a cost function c that is increasing in the quantity q. Let $\mathcal{C}$ be a set of feasible cost functions– possibly different than the ones considered thus far– and consider the following maximin utility problem:

$$\max_{t \in \mathcal{T}} \min_{c \in \mathcal{C}} U(c,t), \tag{31}$$

If $\mathcal{C}$ is the set of all convex cost functions, $\mathcal{C} = \mathcal{C}_{cx}$, then we clearly will have that:

$$\inf_{c(q) \in \mathcal{C}_{cx}} U(c,t) = 0$$

for all mechanisms t. After all, the adversarial nature can choose a cost function sufficiently large so that the socially efficient allocation is arbitrarily close to zero. In this case, even the best mechanism cannot generate any positive net buyer utility. Moreover, all payment functions that are below the gross buyer surplus, i.e., $t\left(q\right) \leq U\left(q\right)$ are "optimal". Thus, the prediction arising from the maximin utility approach is very weak. Suppose now that the set of feasible cost functions $\mathcal{C}$ is compact and that they generate a minimum positive social welfare $\underline{W}$:

$$\underline{W} = \min_{c(q) \in \mathcal{C}} W(c).$$

That is, $\underline{W}$ is the minimum social surplus across all cost functions. A mechanism that solves (31) is:

$$t(q) = u(q) - \underline{W}. \tag{32}$$

In other words, the buyer gives the efficient social surplus to the seller, except for the fixed component $\underline{W}$.

To check that (32) indeed solves (31) it suffices to check that the buyer cannot guarantee himself more than $\underline{W}$ and that the above mechanism indeed guarantees $\underline{W}$. Since there is a non-trivial fixed price, it is necessary to satisfy the participation constraint. However, by construction, the seller will solve:

$$\max_q \{t(q) - c(q)\} = W(c) - \underline{W} \geq 0,$$

where the inequality follows from the definition of $\underline{W}$.

The discussion illustrates that there are two important drawbacks to using maximin mechanisms that are not present when using the competitive ratio.

First, from a normative perspective, the optimal maximin rule is determined by the lowest type (in this case, the highest possible cost in $\mathcal{C}$). Hence, it is very sensitive to the assumptions about the family of cost structures $\mathcal{C}$ that the buyer must consider possible. If the buyer makes a mistake in specifying $\mathcal{C}$ and does not consider a cost function that allows generating ε less surplus, then the optimal rule will guarantee zero buyer surplus. Hence, the buyer is left with a complicated problem of finding a class of cost functions that is sufficiently large to prevent any misspecification, but also restrictive enough that allows to guarantee of some buyer surplus. Ball and Kattwinkel (2025) study the discontinuous impact of the domain of uncertainty in maximin problems, and show that this is a more general phenomenon when using a maximin approach.

Second, and from a positive perspective, there is a plethora of optimal maximin mechanisms. Hence, it is difficult to provide sharp predictions about the optimal policy, and simple mechanisms do not naturally arise as a unique optimal solution. An alternative to selecting an optimal maximin rule is by considering the one that performs the best against some parametrized Bayesian class of cost functions. In this case, the rule will perform well against the worst-case scenario, and can perform the best possible against some parametrized class of cost functions. This is the approach undertaken by Mishra, Patil, and Pavan (2025).

While we have argued that the competitive ratio serves as a more instructive benchmark than maximin utility, we emphasize that the discussion here is phrased for the specific problem that we analyze. For example, in other economic environments one does not need to bound or restrict the space of uncertainty to obtain a well-defined maximin problem. This is usually the case when the space of uncertainty does not have such a direct impact on the scale of the objective function. For example, when the uncertainty is about higher order beliefs (see Brooks and Du (2021), Brooks and Du (2025)) or about the correlation in the values across many goods (see Carroll (2017)), then the realization of the non-Bayesian uncertainty does not change the efficient social surplus. Hence, in these examples, there is no clear "lowest type" in the problem. Hence, there is no unique type that determines the optimal maximin rule. In this type of problem, maximin objectives also sometimes lead to sharp predictions about the optimal rule, and frequently to simple mechanisms. Hence, what is crucial about the competitive ratio that makes it useful in this context is that it is a scale-free benchmark, which is useful in the context of our problem.

While we have focused our discussion on optimal maximin mechanisms, the same kind of concerns arise when using mechanisms that minimize regret (see, for example, Guo and Shmaya (2025)). The main difference is that mechanisms that minimize regret are tailored to perform well against the highest type (lowest cost), but the main ideas of the discussion go through unchanged. The key benefit of using the competitive ratio is that it is a scale-free benchmark, so it can be applied without any restrictions on the set of problems considered. Hence, the rule we find is naturally simple, as there are no parameters about the unknown supply that can be used to tailor the optimal mechanism.

6.2 Alternative Versions of the Competitive Ratio

So far, we have evaluated the performance of constant share mechanisms by computing the buyer's surplus relative to the efficient social surplus, and we have argued that this is a good benchmark in the absence of Bayesian uncertainty. We now discuss an alternative notion of competitive ratio with an alternative motivation.

Suppose now the buyer had Bayesian uncertainty about the seller's cost, but the buyer was constrained to using a constant share mechanism. We now compare the performance of constant share mechanisms relative to the Bayes-optimal mechanism. We interpret this constraint as capturing the fact that computing the optimal transfer when the Bayesian prior is over all possible convex cost functions may be an intractable problem.

To formalize this problem, for any $G \in \Delta(\mathcal{C}_{cx})$ we define:

$$U^*(G) \triangleq \sup_t \int U(c,t)dG(c).$$

That is, $U^*(G)$ is the buyer surplus generated by the Bayes-optimal mechanism when the distribution over cost is G. The performance of the constant share mechanism relative to the Bayes-optimal mechanism is defined as follows:

$$\min_{G \in \Delta(\mathcal{C}_{cx})} \max_{z \in [0,1]} \frac{\int U(c,z)dG(c)}{U^*(G)}. \tag{33}$$

Thus, we compute the performance of the simple rule relative to the Bayes-optimal rule. We emphasize that in this case, we allow for the constant share mechanism to be adapted to the distribution over cost. This is consistent with the current motivation for choosing a constant share mechanism, which is the impossibility of implementing a Bayesian-optimal mechanism due to its complexity (instead of the absence of a Bayesian prior).

In the Appendix (Section 9) we characterize the value of the competitive ratio (Proposition 6). Somewhat surprisingly, we obtain the same value as the competitive ratio when we consider our original benchmark– that is, when the social surplus is in the denominator– but we allow for randomized policies (which is also the case when we can construct a saddle point for the original problem, see Proposition 5). While the problems are distinct, the value of the competitive ratio is qualitatively similar to the one obtained in Theorem 1.

Some of the literature has also considered the buyer surplus generated by the optimal rule–as in(33)– instead of the efficient social welfare. For example, Rogerson (2003) and Chu and Sappington (2007) compare the performance of simple mechanisms relative to the Bayes optimal mechanisms in specific parametric environments. This approach is also pursued in Cai, Devanur, and Weinberg (2021) with notable applications to multi-unit optimal pricing and multi-unit auctions. These are problems where the solution to the optimal Bayesian problem is either unknown or computationally complex, and hence the question arises whether a simple solution adapted to the distribution of uncertainty can attain a good approximation in the sense of the ratio guarantee.

‘

7 Conclusion

This paper establishes that simple mechanisms can resolve complex screening problems. When buyers lack distributional knowledge about supplier costs, constant surplus sharing based solely on demand elasticity achieves optimal worst-case ratio guarantees. The elegance of this solution—sharing a fraction of buyer surplus proportional to the demand elasticity—provides both theoretical insight and practical guidance.

The optimal sharing constants and competitive ratios depend on a single parameter: the elasticity of demand (or supply in selling problems). This parsimony is striking. A procurement officer need not estimate cost distributions, but merely understand their organization's demand responsiveness to price. Similarly, regulators need only observe the demand elasticity.

Our results extend substantially beyond the baseline procurement setting. The same principles generate optimal mechanisms for sellers facing unknown demand and regulators balancing buyer and producer welfare. In each setting, simple rules depending only on known elasticities achieve competitive ratios that cannot be improved by any mechanism, regardless of complexity.

While our sharpest results assume constant elasticity and convex costs, Theorem 2 demonstrates that a positive ratio guarantee obtains under much weaker conditions—essentially, that social surplus exhibits diminishing returns when measured in utility units. This "unit-free" condition accommodates variable elasticity and non-constant returns to scale.

The theoretical foundations we provide help explain the prevalence of cost-plus contracts and markup pricing in practice. These seemingly naive mechanisms emerge as optimal responses to distributional uncertainty. By characterizing performance guarantees as functions of elasticities, we offer concrete guidance for mechanism design: more elastic environments require larger concessions but yield weaker guarantees.

8 Appendix: Omitted Proofs

The Appendix collects the proofs. Before we begin with the proof of Theorem 1, we introduce some notation and provide an auxiliary lemma. We denote by

$$\text{cav}[t]$$

the concavification of the function $t : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, that is, the smallest concave function that is everywhere larger than t. For any $\widehat{q}$ such that $\text{cav}[t](\widehat{q}) > t(\widehat{q})$ we denote by q_1, q_2 the support of the concavification at $\widehat{q} \in [q_1, q_2]$, that is, $\text{cav}[t](q)$ is linear in $[q_1, q_2]$, $\text{cav}[t](q_1) = t(q_1)$, and $\text{cav}[t](q_2) = t(q_2)$. Finally, we define:

$$z(q) \triangleq \frac{t(q)}{u(q)} \text{ and } Z(q) \triangleq \frac{\text{cav}[t](q)}{u(q)}.$$

We provide some properties of these two functions, $z(q)$ and $Z(q)$, which will be later used in the main part of the proof.

Lemma 2 (Concavification of t)

1. *If $\widehat{q}$ is such that $\text{cav}[t]'(\widehat{q}) = c'(\widehat{q})$ then*

$$\widehat{q} \in \arg\max_{q \in \mathbb{R}} \ \{\text{cav}[t](q) - c(q)\}. \tag{34}$$

2. *If $\widehat{q}$ is such that $\text{cav}[t](\widehat{q}) = t(\widehat{q})$ and (34) is satisfied, then*

$$\widehat{q} \in \arg\max_{q \in \mathbb{R}} \ \{t(q) - c(q)\}. \tag{35}$$

3. *If $\widehat{q}$ is such that $\text{cav}[t](\widehat{q}) > t(\widehat{q})$ and $Z'(\widehat{q}) < 0$ then, $Z'(q_1) < 0$ or $Z'(q_2) < 0$.*

4. *If $\widehat{q}$ is such that $\text{cav}[t](\widehat{q}) > t(\widehat{q})$ and $Z'(\widehat{q}) \geq 0$ then,*

$$\widehat{q} Z'(\widehat{q}) \geq \min \{q_1 Z'(q_1), q_2 Z'(q_2)\}. \tag{36}$$

Proof. We establish the above properties of t and $\text{cav}[t]$ serially.

(1.) By construction $\text{cav}[t](q)$ is concave, while by assumption c is convex. Hence, the objective function is concave, so the first-order condition is sufficient for optimality.

(2.) This follows from the fact that by construction $\text{cav}[t](q) \geq t(q)$ pointwise, so if $\widehat{q}$ solves (34) and this attains the same value when replacing with t, then it must also solve (35).

(3.) We note that:

$$Z'(q) = \frac{u'(q)}{(u(q))^2}\left(\frac{q}{\sigma}\text{cav}[t]'(q) - \text{cav}[t](q)\right).$$

We now note that, if $\text{cav}[t](q) > t(q)$, then the term inside the parenthesis is linear in q for all $q \in [q_1, q_2]$. Hence, if the term inside the parentheses is negative at $\widehat{q}$, then it must also be negative at q_1 or q_2.

(4.) We prove that, if $\text{cav}[t](\widehat{q}) > t(\widehat{q})$ and $Z'(\widehat{q}) \geq 0$, then $qZ'(q)$ is concave in $q \in [q_1, q_2]$. If $\text{cav}[t](\widehat{q}) > t(\widehat{q})$, then $\text{cav}[t](q)$ is linear for all $q \in [q_1, q_2]$, so for some $\alpha, \beta \in \mathbb{R}$:

$$\text{cav}[t](q) = \alpha + \beta q.$$

We note that $\text{cav}[t]$ is increasing, concave and $\text{cav}[t](0) = 0$. Hence, we must have that $\alpha, \beta \geq 0$. We now note that:

$$\frac{d^2}{dq^2}(qZ'(q)) = -\sigma^2 q^{-\sigma-2}\left(\alpha\sigma(\sigma+1) + \beta q(1-\sigma)^2\right) \leq 0.$$

We thus obtain that $qZ'(q)$ is concave in $q \in [q_1, q_2]$, which implies that (36) is satisfied. ■

Proof of Theorem 1. First, we prove that the inequality (21) is satisfied for every mechanism $t(q)$ and that the inequality (21) is strict whenever $t(q) \neq z^*(\sigma)u(q)$. Second, we establish that the inequality (21) is attained as an equality with the transfer rule $t(q) = z^* u(q)$. Thus, we establish that the bound (20) is an upper bound for the competitive ratio, and then that the upper bound (20) can be attained by a constant share mechanism (18). To make the notation more compact, we omit the argument of $z^*(\sigma)$, and write simply z^*.

(Part I: Upper Bound) It is without loss of generality to assume that $t(q)$ is monotone increasing (as the seller will never choose a dominated quality). We assume without loss of generality that $t(q)$ is differentiable; we can approximate any monotonic function by a differentiable function pointwise arbitrarily well, and the ratio will converge appropriately. We fix some $\widetilde{q} \in \mathbb{R}_+$ that we will specify later and define:

$$\widetilde{z} \triangleq \frac{t(\widetilde{q})}{u(\widetilde{q})} \text{ and } \widetilde{\gamma} \triangleq t'(\widetilde{q}).$$

Hence, the tilde indicates that we are evaluating at $\widetilde{q}$. Following the definition of $z(q)$ we can write $\widetilde{\gamma}$ as follows:

$$\widetilde{\gamma} = \widetilde{z}u'(\tilde{q}) + z'(\tilde{q})u(\tilde{q}). \tag{37}$$

We show that the inequality (21) is satisfied when the cost function is a piecewise linear cost function of the form:

$$c(q) = \begin{cases} 0, & \text{if } q < \widetilde{q}; \\ \widetilde{\gamma}(q - \widetilde{q}), & \text{if } q \geq \widetilde{q}. \end{cases} \tag{38}$$

where (as before) $\widetilde{q}$ is specified later. By construction $\widetilde{q}$ satisfies the first-order condition:

$$t'(\widetilde{q}) - c'(\widetilde{q}) = 0.$$

And, following Lemma 2, if $t(\widetilde{q}) = \text{cav}[t](\widetilde{q})$, we then have that:

$$\widetilde{q} \in \arg\max_{q \in [0,1]} t(q) - c(q);$$

in other words, in this case, we thus have that $\widetilde{q} = q(c, t)$.

We now compute the efficient quality and the respective social surplus when the cost is (38). The first-order condition is given by:

$$u'(\overline{q}(c)) - c'(\overline{q}(c)) = 0.$$

We thus have that:

$$\overline{q}(c) = \widetilde{\gamma}^{\frac{1}{\sigma-1}}.$$

Hence, the efficient social surplus is given by:

$$W(c) = u(\overline{q}(c)) - c(\overline{q}(c)) = \frac{1-\sigma}{\sigma} \left(\frac{1}{\widetilde{\gamma}} \right)^{\frac{\sigma}{1-\sigma}} + \widetilde{\gamma}\widetilde{q}. \tag{39}$$

We analyze three cases. We first analyze the case in which $Z(q)$ is strictly decreasing in some part of the domain. We then analyze the case in which $Z(q)$ is weakly increasing where we distinguish the cases in which:

$$\lim_{q \to \infty} Z(q) \neq z^* \text{ or } \lim_{q \to 0} Z(q) \neq z^*.$$

We show that in each of the cases the bound is satisfied.

(Case 1) Suppose first that there exists $\widetilde{q}$ such that $Z'(\widetilde{q}) < 0$. Following Lemma 2.3, without loss of generality we can take $\widetilde{q}$ such that $t(\widetilde{q}) = \text{cav}[t](\widetilde{q})$, and so (also following Lemma 2) we have that $\widetilde{q} = q(c, t)$. Using that $Z'(\widetilde{q}) = z'(\widetilde{q}) < 0$, (37) implies that:

$$\widetilde{\gamma} < \widetilde{z}(\widetilde{q})^{\sigma-1} \leq (\widetilde{q})^{\sigma-1}. \tag{40}$$

We note that:

$$\frac{1-\sigma}{\sigma}\left(\frac{1}{\gamma}\right)^{\frac{\sigma}{1-\sigma}}+\gamma\widetilde{q}$$

is strictly quasi-convex in γ (fixing $\widetilde{q}$) with a unique minimum at $\gamma=\widetilde{q}^{\sigma-1}$. We thus have that:

$$W(c)=\frac{1-\sigma}{\sigma}\left(\frac{1}{\widetilde{\widetilde{\gamma}}}\right)^{\frac{\sigma}{1-\sigma}}+\widetilde{q}\widetilde{\gamma}>(1-\sigma)\frac{\widetilde{q}^{\sigma}}{\sigma}\widetilde{z}^{\frac{\sigma}{\sigma-1}}+\widetilde{z}\sigma\frac{\widetilde{q}^{\sigma}}{\sigma}=u(q(c,t))\left(\widetilde{z}\sigma+(1-\sigma)\widetilde{z}^{\frac{\sigma}{\sigma-1}}\right). \quad (41)$$

The inequality comes from the quasi-concavity of (39) and (40); the second equality comes from the fact that $\widetilde{q}=q(c,t)$. We thus have that

$$\frac{U(c,t)}{W(c)}=\frac{u(q(c,t))-t(q(c,t))}{u(\overline{q}(c))-c(\overline{q}(c))}<\frac{(1-\widetilde{z})}{(1-\sigma)\widetilde{z}^{\frac{\sigma}{\sigma-1}}+\widetilde{z}\sigma}\leq\frac{1-z^*}{(1-\sigma)(z^*)^{\frac{\sigma}{\sigma-1}}+\sigma z^*}. \quad (42)$$

where the first inequality follows from (41) while the second inequality follows from the definition of z^*.

(Case 2) We now consider the case that $t(q)\neq z^*u(q)$, $Z'(q)\geq 0$ for all q and

$$\lim_{q\to\infty} Z(q)\neq z^*.$$

Since $Z(q)\in[0,1]$ we must have that $Z'(q)$ must converge to 0 fast enough as $q\to\infty$, in fact,

$$\liminf_{q\to\infty}\ Z'(q)q=0.$$

Following Lemma 2.4, we also have that:

$$\liminf_{q\to\infty}\ z'(q)q=0.$$

We can thus find some $\widetilde{q}$ large enough such that:

$$\frac{(1-\widetilde{z})}{(1-\sigma)\left(\tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}\right)^{\frac{\sigma}{\sigma-1}}+\left(\tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}\right)\sigma}<\frac{1-z^*}{(1-\sigma)(z^*)^{\frac{\sigma}{\sigma-1}}+\sigma z^*}.$$

We now fix this $\widetilde{q}$ and prove the result.

Replacing (37) into (39), we get:

$$\begin{aligned}u(\overline{q}(c))-c(\overline{q}(c))=&\frac{1-\sigma}{\sigma}\left(\frac{1}{\tilde{z}u'(\widetilde{q})+z'(\widetilde{q})u(\widetilde{q})}\right)^{\frac{\sigma}{1-\sigma}}+q(c,t)\left(\tilde{z}u'(\widetilde{q})+z'(\tilde{q})u(\widetilde{q})\right)\\=&\frac{1-\sigma}{\sigma}\widetilde{q}^{\sigma}\left(\frac{1}{\tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}}\right)^{\frac{\sigma}{1-\sigma}}+\widetilde{q}^{\sigma}\left(\tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}\right).\end{aligned} \quad (43)$$

We thus have that:

$$\frac{u(q(c,t))-t(q(c,t))}{u(\overline{q}(c))-c(\overline{q}(c))}=\frac{(1-\widetilde{z})}{(1-\sigma)\left(\ \tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}\right)^{\frac{\sigma}{\sigma-1}}+\left(\tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}\right)\sigma}<\frac{1-z^*}{(1-\sigma)(z^*)^{\frac{\sigma}{\sigma-1}}+\sigma z^*}. \quad (44)$$

We thus prove the result.

(Case 3) We now consider the case that $t(q)\neq z^*u(q)$, $Z'(q)\geq 0$ for all q and

$$\lim_{q\to 0} Z(q)\neq z^*.$$

We note that we must have that:

$$\liminf_{q\to 0}\ Z'(q)q=0.$$

Following Lemma 2.4, we also have that:

$$\liminf_{q\to 0}\ z'(q)q=0.$$

We thus find some $\widetilde{q}$ small enough such that:

$$\frac{(1-\widetilde{z})}{(1-\sigma)\left(\ \tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}\right)^{\frac{\sigma}{\sigma-1}}+\left(\tilde{z}+z'(\widetilde{q})\frac{\widetilde{q}}{\sigma}\right)\sigma}<\frac{1-z^*}{(1-\sigma)(z^*)^{\frac{\sigma}{\sigma-1}}+\sigma z^*}.$$

We can reach (43) in the same way as before, which in turn allows us to also reach (44) the same as before.

(Part II: Attainment of Upper Bound through $t(q)=z^*u(q)$) We now prove that, if $t(q)=z^*u(q)$ then

$$\frac{U(c,t)}{W(c)}\geq B(\sigma),$$

for all cost functions, with equality if and only if the marginal cost satisfies:

$$c'(q)\qquad\begin{cases}=0, & q\in[0,\widetilde{q}];\\ =z^*u'(\widetilde{q}), & q\in[\widetilde{q},(z^*)^{\frac{1}{\sigma-1}}\widetilde{q}];\\ \geq z^*u'(\widetilde{q}), & q\in[(z^*)^{\frac{1}{\sigma-1}}\widetilde{q},\infty).\end{cases}$$

for some $\widetilde{q}$. For any mechanism $t(q)=zu(q)$ for some generic $z\in[0,1]$ we get the following. The seller will choose quantity $q(c,t)$ satisfying:

$$z(q(c,t))^{\sigma-1}-c'(q(c,t))=0.$$

The buyer surplus is is a fraction $(1 - z^*)$ of the buyer's gross utility, that is,

$$U(c,t) = (1-z)\frac{(q(c,t))^\sigma}{\sigma}. \tag{45}$$

The seller surplus is at most:

$$\Pi(c,t) \le t(q(c,t)) = z\frac{(q(c,t))^\sigma}{\sigma}. \tag{46}$$

The inequality is satisfied with equality if and only if $c(q) = 0$ for all $q \in [0, q(c,t)]$. The deadweight loss due to inefficiencies can also be bounded:

$$DWL = \int_{q(c,t)}^{\overline{q}(c)} (u'(q) - c'(q))\, dq \le u(\overline{q}(c)) - u(q(c,t)) - c'(q(c,t))(\overline{q}(c) - q(c,t)), \tag{47}$$

where we recall that $\overline{q}(c)$ is the efficient quantity and c is convex so $c'(q) \ge c'(q(c,t))$ for all $q \in [q(c,t), \overline{q}(c)]$. The efficient quantity $\overline{q}(c)$ can be bounded from above as follows:

$$\overline{q}(c) = (c'(\overline{q}(c)))^{\frac{1}{\sigma-1}} \le (c'(q(c,t)))^{\frac{1}{\sigma-1}} = z^{\frac{1}{\sigma-1}} q(c,t).$$

We thus have that:

$$DWL \le u(z^{\frac{1}{\sigma-1}} q(c,t)) - u(q(c,t)) - c'(q(c,t))(\overline{q}(c) - q(c,t)) \le u(q(c,t))\left((1-\sigma)z^{\frac{\sigma}{\sigma-1}} + \sigma z - 1\right).$$

Furthermore, this inequality is satisfied with equality if and only if $c'(q)$ is constant for all $q \in [q(c,t), \overline{q}(c)]$. The ratio (4):

$$\frac{U}{W} = \frac{U}{U + \Pi + DWL},$$

after inserting the above terms (45)-(47) satisfies

$$\frac{U}{U + \Pi + DWL} \ge \frac{1-z}{(1-\sigma)z^{\frac{\sigma}{\sigma-1}} + \sigma z}. \tag{48}$$

Thus, for any $t(q) = zu(q)$ and any convex cost function, we obtain the inequality (48). Now, replacing z with z^* we get a lower bound of which we know by the first part of the proof that it is also an upper bound. Finally, we note that (48) is satisfied with equality if and only if $c(q) = 0$ for all $q \in [0, q(c,t)]$ and $c'(q)$ is constant for all $q \in [q(c,t), \overline{q}(c)]$. This concludes the proof. ■

Proof of Lemma 1 and Theorem 2. The theorem is proved by performing a change of variables to recover our original power utility model. For the given utility function $u(q)$, we define:

$$x = (\widehat{\sigma} u(q))^{\frac{1}{\widehat{\sigma}}},$$

with $\widehat{\sigma}$ defined as in (29) (the relevant case is when $\underline{\delta} > 0$). To make the notation more compact, throughout the proof we write σ instead of $\widehat{\sigma}$. We then have that the utility of the buyer and the cost of the seller in terms of the variable x is:

$$\widehat{u}(x) = \frac{x^\sigma}{\sigma} \text{ and } \widehat{c}(x) = c\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right).$$

We then prove Theorem 2 by applying Theorem 1. The proof reduces to showing that $\widehat{c}$ is increasing and convex.

We first verify that $\widehat{c}$ is always increasing:

$$\widehat{c}'(x) = c'\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right) u^{-1\prime}\left(\frac{x^\sigma}{\sigma}\right) x^{\sigma-1} \geq 0.$$

For the second derivative, we get:

$$\begin{aligned} \widehat{c}''(x) &= \frac{\sigma}{x^{2-\sigma}} c'\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right) u^{-1\prime}\left(\frac{x^\sigma}{\sigma}\right) \times \\ & \left(\frac{c''\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right) u^{-1}\left(\frac{x^\sigma}{\sigma}\right)}{c'\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right)} \frac{u^{-1\prime}\left(\frac{x^\sigma}{\sigma}\right)\frac{x^\sigma}{\sigma}}{u^{-1}\left(\frac{x^\sigma}{\sigma}\right)} + \frac{x^\sigma}{\sigma}\frac{u^{-1\prime\prime}\left(\frac{x^\sigma}{\sigma}\right)}{u^{-1\prime}\left(\frac{x^\sigma}{\sigma}\right)} - \frac{(1-\sigma)}{\sigma}\right). \end{aligned}$$

We now note that:

$$\frac{c''\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right) u^{-1}\left(\frac{x^\sigma}{\sigma}\right)}{c'\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right)} \frac{u^{-1\prime}\left(\frac{x^\sigma}{\sigma}\right)\frac{x^\sigma}{\sigma}}{u^{-1}\left(\frac{x^\sigma}{\sigma}\right)} = \frac{qc''(q)}{c'(q)}\frac{u(q)}{qu(q)}$$

and

$$\frac{u^{-1\prime}\left(\frac{x^\sigma}{\sigma}\right)}{u^{-1\prime\prime}\left(\frac{x^\sigma}{\sigma}\right)}\frac{\sigma}{x^\sigma} = \frac{u^{-1\prime}(u(q))}{u^{-1\prime\prime}(u(q))}\frac{1}{u(q)} = -\frac{u'(q)}{qu''(q)} \cdot \frac{qu'(q)}{u(q)}.$$

We can thus write the second derivative as follows:

$$\widehat{c}''(x) = \frac{\sigma}{x^{2-\sigma}} c'\left(u^{-1}\left(\frac{x^\sigma}{\sigma}\right)\right) u^{-1\prime}\left(\frac{x^\sigma}{\sigma}\right)\left(\delta(q,c) - \frac{(1-\sigma)}{\sigma}\right),$$

where we use $\delta(c,q)$ as defined in (25). Since $(1-\sigma)/\sigma = \underline{\delta}$, we have that $\widehat{c}''(x) \geq 0$. Applying Theorem 1, we get the result. ■

Proof of Theorem 3. The proof of Theorem 1 goes through without any significant difference. We now have that:

$$u(q(c,t)) - t(q(c,t)) + \alpha\left(t(q(c,t)) - c(q(c,t))\right) = u(\widetilde{q})(1 - \widetilde{z} + \alpha\widetilde{z}),$$

where the cost continues to be as in (38). We now explain how Parts I and II of the proof need to be modified to account for the weight on profits.

For Part I of the proof, we just have to add αz to the numerator in (42), and analogously to Cases 2 and 3 analyzed therein. Note that here we provide a cost and show that the competitive ratio always satisfies the corresponding bound when we add the profits (weighted by α) to the numerator.

For Part II of the proof, we first note that (4) is at least equal to α. To prove this, we note that when $t(q) = u(q)$, the seller chooses the efficient quality and extracts the efficient social surplus, so we get that (4) is given by:

$$\frac{u(q(c,t)) - t(q(c,t)) + \alpha(t(q(c,t)) - c(q(c,t)))}{u(\overline{q}(c)) - c(\overline{q}(c))} = \alpha.$$

We also note that, for any given mechanism t, the efficient surplus is the sum of buyer surplus, profits and deadweight loss:

$$u(\overline{q}(c)) - c(\overline{q}(c)) = U + \Pi + DWL.$$

Since, the ratio is greater than α, we have that:

$$\frac{U + \alpha\Pi}{U + \Pi + DWL}$$

is decreasing in Π. We thus have that (46) continues to be the relevant bound. The rest of the proof remains the same (by adding αz to the numerator). ■

Proof of Proposition 4. The proof follows the same way as the analysis in Section 3 with the change of variables $c = 1/v$. Since in Section 3 the constant share mechanism is found to be the Bayesian-optimal mechanism when the distribution of cost is a power function, for the optimal selling mechanism, the constant markup mechanism will be the Bayesian-optimal mechanism when the distribution of values $v = 1/c$ follows a Pareto distribution (the reciprocal of a random variable that follows a power distribution follows a Pareto distribution). The rest of the analysis follows the same way. ■

9 Supplemental Appendix: Additional Material

Our main result in the nonlinear cost environment, Theorem 1, provides the competitive ratio for deterministic mechanisms, which is attained by a constant share mechanism. We now relax the restriction to deterministic policies and, in consequence, obtain an improved competitive ratio. Yet, every policy in the support of the stochastic mechanism remains a constant share rule. And surprisingly, the relaxation to stochastic policies does not yield a noticeable improvement in the competitive ratio.

In a second relaxation, we weaken the benchmark–the denominator–in the competitive ratio. So far, the benchmark in the competitive ratio was the efficient social surplus. We weaken the benchmark and compare the performance of the optimal constant share mechanism with the revenue of the Bayes optimal (unconstrained) mechanism. We find that the constant surplus sharing in fact attains the same bound as against the social surplus. This section thus provides additional support to the notion that simple rules perform well in the procurement environment.

Stochastic Policies We now allow for randomizations over transfer policies, thus allowing the buyer to choose a distribution F over transfer policies, thus: $F \in \Delta\{t \,|\, t : \mathbb{R} \to \mathbb{R}\}$. We thus seek to identify a competitive ratio:

$$\underline{B}(\sigma) \triangleq \max_{F \in \Delta\{t:t:\mathbb{R}\to\mathbb{R}\}} \min_{c \in C_{cx}} \frac{\int U(c,t)dF(t)}{W(c)}.$$

As we allow for stochastic policies, we can also ask whether the solution from the related min max problem:

$$\overline{B}(\sigma) \triangleq \min_{G \in \Delta C_{cx}} \max_{\{t:\mathbb{R}\to\mathbb{R}\}} \frac{\int U(c,t)dG(c)}{\int W(c)dG(c)}$$

coincides with the max min value, that is whether $\underline{B}(\sigma) = \overline{B}(\sigma)$. A *saddle point* is a pair of distributions (F^*, G^*) such that:

$$\min_{G \in \Delta C_{cx}} \frac{\int\int U(c,t)dG(c)dF^*(t)}{\int W(c)dG(c)} = \max_{F \in \Delta\{t:t:\mathbb{R}\to\mathbb{R}\}} \frac{\int\int U(c,t)dG^*(c)dF(t)}{\int W(c)dG^*(c)}. \tag{49}$$

That is, these distributions attain the values of the saddle point. We answer both of these questions in the affirmative in the next result. To make the notation more compact, we define:

$$\widehat{B}(\sigma) \triangleq \frac{1}{\sigma^{\frac{-\sigma}{1-\sigma}} - \sigma - \log(1-\sigma)},$$

which will be the value of the saddle point.

Proposition 5 (Saddle Point Property with Stochastic Policies)
A saddle point (49) exists and the competitive ratio is:

$$\underline{B}(\sigma) = \overline{B}(\sigma) = \widehat{B}\,(\sigma)\,.$$

The saddle point is attained by distributions (F^*, G^*) that only place positive weights on constant share mechanisms.

Proof. To find the saddle point, we consider a parametrized class of cost functions of the following form:

$$c_{\widetilde{\gamma}}(q) = \begin{cases} 0, & \text{if } q \in [0, q_0]; \\ u(q) - u(q_0) - \kappa \log(\frac{q}{q_0}), & \text{if } q \in [q_0, q_1]; \\ u(q_1) - u(q_0) - \kappa \log(\frac{q_1}{q_0}) + \widetilde{\gamma}(q - q_1), & \text{if } q \geq q_1, \end{cases}$$

The cost functions are parametrized by $\widetilde{\gamma}$ and the rest of the parameters are determined as follows:

$$q_0 \triangleq (1-\sigma)^{\frac{1}{\sigma}} \left(\frac{\sigma}{\widetilde{\gamma}}\right)^{\frac{1}{1-\sigma}} \quad ; \quad q_1 \triangleq \left(\frac{\sigma}{\widetilde{\gamma}}\right)^{\frac{1}{1-\sigma}} \quad ; \quad \kappa \triangleq (1-\sigma) \left(\frac{\sigma}{\widetilde{\gamma}}\right)^{\frac{\sigma}{1-\sigma}}.$$

We illustrate this class of cost functions in Figure 7 by showing the marginal cost $c'_{\widetilde{\gamma}}(q)$. We can see that the marginal cost is 0 for all $q \in [0, q_0]$, then it is increasing for all $q \in [q_0, q_1]$ (hence $c_{\widetilde{\gamma}}$ is strictly convex in this range), and finally the marginal cost is equal to $\widetilde{\gamma}$ for all $q \in [q_1, \infty)$. We now show that this class of cost functions and the constant share mechanisms form a saddle point (with the appropriate mixing).We prove the result in two steps, first prove the maximin problem attains $\widehat{B}$ and then prove that the minmax problem also attains $\widehat{B}$.

(Step 1: $\widehat{B}(\sigma) = \underline{B}(\sigma)$) Suppose the buyer uses a payment rule:

$$t(q) = zu(q),$$

but randomizes z according to distribution

$$F(z) = \begin{cases} -\widehat{B}(\sigma) \log(1-z), & \text{if } z < \sigma; \\ 1, & \text{if } z \geq \sigma. \end{cases} \tag{50}$$

Hence, F is absolutely continuous in $[0, \sigma)$ and has an atom at $z = \sigma$.

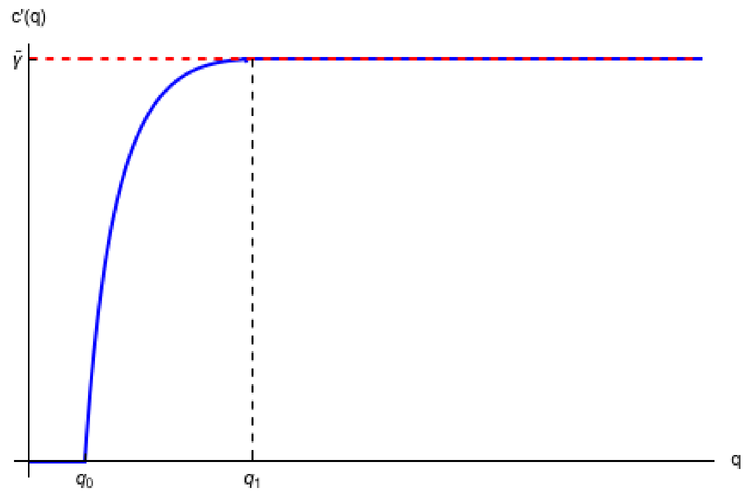

Figure 7: Illustration of Marginal Cost Function $c'_{\gamma}(q)$

We first note that, for any linear cost $c(q) = \gamma q$, the efficient social surplus and buyer surplus are:

$$W(c) = \frac{1-\sigma}{\sigma}\left(\frac{1}{\gamma}\right)^{\frac{\sigma}{1-\sigma}} \quad \text{and} \quad U(q,z) = \frac{(1-z)}{\sigma}\left(\frac{z}{\gamma}\right)^{\frac{\sigma}{1-\sigma}}.$$

We thus have that for any linear cost function $c(q) = \gamma q$ the expected competitive ratio is:

$$\frac{\int_0^{\sigma} U(c,z)dF(z)}{W(c)} = \widehat{B}(\sigma).$$

We remark that to compute the integral, we need to account for the fact that F has an atom at $z = \sigma$. We now show that the linear cost function minimizes the competitive ratio:

$$\widehat{B}(\sigma) = \min_{c} \frac{\int_0^{\sigma} U(c,z)dF(z)}{W(c)},$$

hence the proposed randomization over constant share mechanisms indeed maximizes the competitive ratio.

We denote by $\mathcal{C}_b$ the set of convex cost functions that have constant marginal cost for high enough q:

$$\mathcal{C}_b \triangleq \{c : c \text{ is increasing, convex, and } c''(q) = 0 \text{ for all } q \text{ such that } c'(q) \geq \sigma u'(q)\}.$$

In other words, $\mathcal{C}_b$ consists of increasing and convex functions with constant marginal cost whenever the marginal cost exceeds $\sigma u'(q)$. Given that the distribution over z is given by (50) it is without loss of generality to consider cost functions in $\mathcal{C}_b$. More precisely, for all $c \notin C_b$, there exists $\widehat{c} \in C_b$ such that:

$$U(c, z) = U(\widehat{c}, z) \text{ and } W(c, z) \leq W(\widehat{c}, z),$$

for all $z \in [0, \sigma]$. The inequality comes from the fact that lowering the marginal cost of units that are never sold to the buyer always increases the efficient social welfare without changing the buyer surplus.

We now fix $c \in \mathcal{C}_b$ and prove that:

$$\frac{\int_0^\sigma U(c, z) dF(z)}{W(c)} = \widehat{B}(\sigma).$$

For this, we consider a class of cost functions $c_x(q)$ parametrized by $x \in [0, 1]$ with the following properties: (i) $c_0'(q) = c'(\overline{q}(c))$ for all $q \in \mathbb{R}$, (ii) $c_1(q) = c(q)$ for all $q \in \mathbb{R}$, and (iii) $c_x'(q)$ strictly decreasing in q and x, and (iv) continuous in $x \in [0, 1]$. The first property states that $c_0(q)$ is a linear cost function with marginal cost equal to the marginal cost of c at the efficient q. The second property states that at $x = 1$ the parametrized function coincides with c. The third property states how c_x changes with x. Since $c \in \mathcal{C}_b$, we have that for all $q \geq q(c, \sigma)$, $c_x'(q)$ is constant in x. To make the notation more compact we define:

$$r(x) \triangleq \frac{\int U(c_x, z) dF(z)}{W(c_x)}.$$

We note that:

$$\frac{\partial r(x)}{\partial x} = \frac{1}{W(c_x)} \left(\frac{\partial \int U(c_x, z) dF(z)}{\partial x} - r(x) \frac{\partial W(c_x)}{\partial x} \right).$$

We now prove this derivative is equal to 0. Finally, to make the notation more compact, we denote $q(x, z) \triangleq q(c_x, z)$.

The total surplus is given by:

$$W(c_x) = \int_0^\infty \max\{u'(q) - c_x'(q), 0\} dq.$$

We compute the derivative of total surplus with respect to x:

$$\frac{\partial W(c_x)}{\partial x} = -\int_{q(x,0)}^{q(x,\sigma)} \frac{\partial c_x'(q)}{\partial x} dq.$$

We now note that:

$$\frac{\partial q(x, z)}{\partial z} = -u'(q(x, z)) \frac{1}{z u''(q(x, z)) - c_x''(q(x, z))}.$$

Thus,

$$\frac{\partial W(c_x)}{\partial x} = \int_0^\sigma \frac{\partial c_x'(q)}{\partial x} \frac{u'(q(x,z))}{zu''(q(x,z)) - c_x''(q(x,z))} dz.$$

We now compute the derivative of the buyer surplus with respect to x.

The expected buyer surplus:

$$\int_0^1 U(x,z)F(z)dz = \widehat{B}(\sigma) \int_0^\sigma u(q(x,z))dz + (1 + \log(1-\sigma)\widehat{B}(\sigma))(1-\sigma)u(q(x,\sigma)).$$

We compute the derivative with respect to x:

$$\frac{\partial}{\partial x} \int U(x,z)f(z)dz = \widehat{B}(\sigma) \int_0^\sigma u'(q(x,z)) \frac{\partial q(x,z)}{\partial x} dz.$$

Note that $q(x,\sigma)$ does not change with x so the atom $z = \sigma$ does not appear when computing the derivative (recall that for all $q \geq q(c,\sigma)$, $c_x'(q)$ is constant in x). We now note that:

$$\frac{\partial q(x,z)}{\partial x} = \frac{\partial c_x'(q(x,z))}{\partial x} \frac{1}{zu''(q(x,z)) - c_x''(q(x,z))}.$$

We thus get that:

$$\frac{\partial}{\partial x} \int U(x,z)f(z)dz = \widehat{B}(\sigma) \int_0^\sigma \frac{\partial c_x'(q(x,z))}{\partial x} \frac{u'(q(x,z))}{zu''(q(x,z)) - c_x''(q(x,z))} dz.$$

Hence, we have that:

$$\frac{\partial}{\partial x} \int U(x,z)f(z)dz = \widehat{B}(\sigma) \frac{\partial W(c_x)}{\partial x},$$

which we can use to compute the derivative of the competitive ratio.

We have that:

$$\frac{\partial r(x)}{\partial x} = \frac{\partial W(c_x)}{\partial x} \frac{1}{W(c)} (\widehat{B}(\sigma) - r(x)).$$

Hence $r(x)$ is determined by a linear differential equation. Since $r(x) = \widehat{B}(\sigma)$ is a solution to this differential equation; this is indeed the unique solution. Hence, we obtain the result.

(Step 2: $\widehat{B}(\sigma) = \overline{B}(\sigma)$) We consider a distribution over cost functions $c_{\widetilde{\gamma}}$ according to a cumulative distribution

$$F(\widetilde{\gamma}) = \widetilde{\gamma}^\theta,$$

for some $\theta > \sigma/(1-\sigma)$.

Following standard techniques the buyer surplus is given by:

$$U(c,t) = \int_0^1 \max_q \left\{ \left(u(q) - c_{\widetilde{\gamma}}(q) - \frac{F(\widetilde{\gamma})}{f(\widetilde{\gamma})} \frac{\partial c_{\widetilde{\gamma}}(q)}{\partial \widetilde{\gamma}} \right) \right\} dF(\widetilde{\gamma}).$$

To make the notation more compact, we define:

$$V_{\widetilde{\gamma}}(q) \triangleq \left(u(q) - c_{\widetilde{\gamma}}(q) - \frac{F(\widetilde{\gamma})}{f(\widetilde{\gamma})} \frac{\partial c_{\widetilde{\gamma}}(q)}{\partial \widetilde{\gamma}} \right),$$

which is the virtual surplus. Computing explicitly, we get:

$$V_{\widetilde{\gamma}}(q) = \begin{cases} 0, & \text{if } q \in [0, q_0]; \\ \frac{\left(\frac{\sigma}{\widetilde{\gamma}}\right)^{\frac{\sigma}{1-\sigma}} \left(\theta(1-\sigma)+(\theta-\frac{\sigma}{1-\sigma})\left(\sigma(1-\sigma)\log(q)+\sigma\log\left(\widetilde{\gamma}^{\frac{1-\sigma}{\sigma}}\right)-\log(1-\sigma)\right)\right)}{\theta\sigma}, & \text{if } q \in [q_0, q_1]; \\ (\sigma\frac{1+\theta}{\theta} - 1) \left(\frac{\sigma}{\widetilde{\gamma}}\right)^{\frac{\sigma}{1-\sigma}} \left(1 + \frac{\log(1-\sigma)}{\sigma}\right) - \frac{(\theta+1)}{\theta}\widetilde{\gamma}q + \frac{q^\sigma}{\sigma}, & \text{if } q \geq q_1. \end{cases}$$

We note that the optimum is at

$$\left(\frac{\theta}{\widetilde{\gamma}(1+\theta)}\right)^{\frac{1}{1-\sigma}} = \arg\max_{q\in\mathbb{R}} V_{\widetilde{\gamma}}(q).$$

We thus have that:

$$\max_{q\in\mathbb{R}} V_{\widetilde{\gamma}}(q) = (\sigma\frac{1+\theta}{\theta} - 1)\left(\frac{\sigma}{\widetilde{\gamma}}\right)^{\frac{\sigma}{1-\sigma}} \left(1 + \frac{\log(1-\sigma)}{\sigma}\right) + \frac{(1-\sigma)}{\sigma}\left(\frac{\theta}{\widetilde{\gamma}(1+\theta)}\right)^{\frac{\sigma}{1-\sigma}},$$

and

$$\max_{t}\ U(c,t) = \left((\sigma\frac{1+\theta}{\theta} - 1)\,(\sigma)^{\frac{\sigma}{1-\sigma}} \left(1 + \frac{\log(1-\sigma)}{\sigma}\right) + \frac{(1-\sigma)}{\sigma}\left(\frac{\theta}{(1+\theta)}\right)^{\frac{\sigma}{1-\sigma}} \right) \left(\frac{\theta}{\theta - \frac{\sigma}{1-\sigma}}\right).$$

Computing now the total surplus, we get:

$$W(c,t) = \int_0^1 \max_q \{u(q) - c_{\widetilde{\gamma}}(q)\}\, dF(\widetilde{\gamma}).$$

The efficient social surplus is:

$$\overline{q}(\widetilde{\gamma}) = \left(\frac{1}{\widetilde{\gamma}}\right)^{\frac{1}{1-\sigma}}.$$

Thus, the total surplus is given by:

$$W = \left((\sigma - 1)\,(\sigma)^{\frac{\sigma}{1-\sigma}} \left(1 + \frac{\log(1-\sigma)}{\sigma}\right) + \frac{(1-\sigma)}{\sigma} \right) \left(\frac{\theta}{\theta - \frac{\sigma}{1-\sigma}}\right).$$

We thus have that the ratio is given by:

$$\max_t \frac{U(c,t)}{W(c)} = \frac{\left((\sigma\frac{1+\theta}{\theta} - 1)\,(\sigma)^{\frac{\sigma}{1-\sigma}} \left(1 + \frac{\log(1-\sigma)}{\sigma}\right) + \frac{(1-\sigma)}{\sigma}\left(\frac{\theta}{(1+\theta)}\right)^{\frac{\sigma}{1-\sigma}} \right)}{\left((\sigma - 1)\,(\sigma)^{\frac{\sigma}{1-\sigma}} \left(1 + \frac{\log(1-\sigma)}{\sigma}\right) + \frac{(1-\sigma)}{\sigma} \right)}.$$

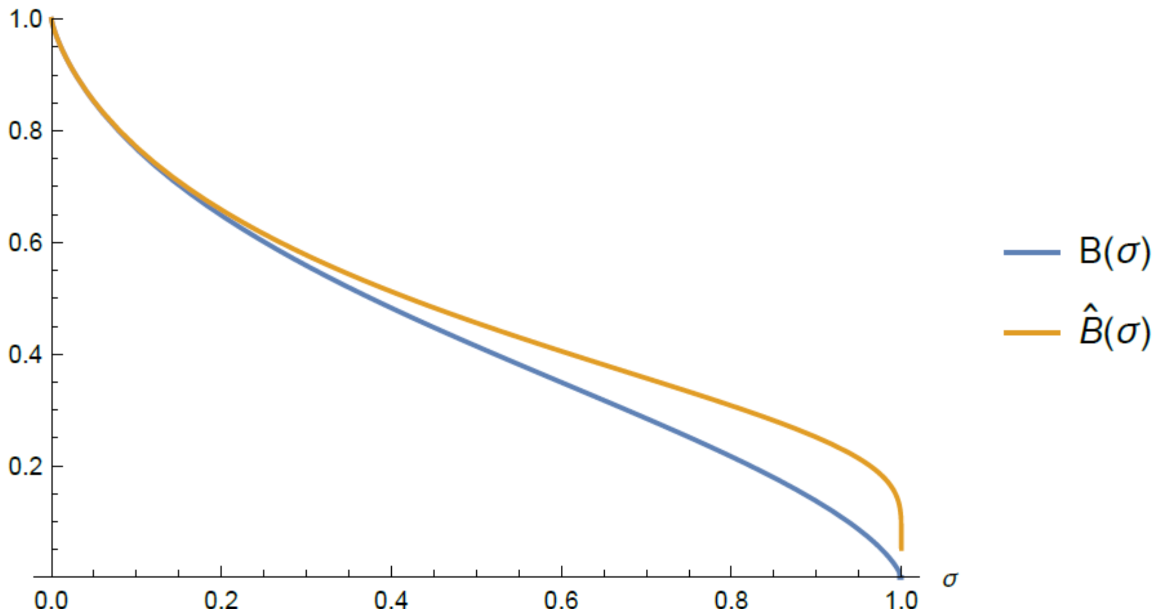

Figure 8: Competitive Ratio Guarantee with Deterministic (blue) and Stochastic (yellow) Transfer Policies

We now note that:

$$\lim_{\theta \downarrow \frac{1-\sigma}{\sigma}} \frac{W}{U} = \frac{1}{\left(-\sigma - \log(1-\sigma) + \sigma^{\frac{-\sigma}{1-\sigma}}\right)} = \widehat{B}(\sigma).$$

We thus prove the result. ■

With stochastic policies, the class of critical cost functions becomes larger. It now consists of cost functions that have three terms, linear terms, log terms, and utility terms. We illustrate the value attained by the saddle point in Figure 8. We can see that the value of the saddle point is quite similar to the bound we found with deterministic policies $B(\sigma)$: they are both equal to 1 when σ is small and they both converge to 0 when σ converges to 1. The only part where both values are significantly different is near 1 because in the limit $\sigma \to 1$, we have that $\widehat{B}'(\sigma) \to -\infty$, while $B'(\sigma)$ remains bounded.

Simple vs Bayes Optimal Mechanisms Suppose now the buyer had Bayesian uncertainty about the seller's cost, but the buyer was constrained to using a constant share mechanism. So far, we have compared the performance of the simple rule relative to the efficient social welfare. We now compare the performance of the constant share mechanisms relative to the Bayes-optimal rule. We interpret this constraint as capturing the fact that computing the optimal transfer when the

Bayesian prior is over all possible convex cost functions may be an intractable problem.

We recall the basic definitions introduced in Section 6.2. For any $G \in \Delta(\mathcal{C}_{cx})$ we define:

$$U^*(G) \triangleq \sup_t \int U(c,t)dG(c).$$

That is, $U^*(G)$ is the buyer surplus generated by the Bayes-optimal mechanism when the distribution over cost is G. The performance of the simple rule relative to the Bayes-optimal rule is defined as follows:

$$\min_{G \in \Delta(\mathcal{C}_{cx})} \max_{z \in [0,1]} \frac{\int U(c,z)dG(c)}{U^*(G)}. \tag{51}$$

Thus, we compute the performance of the simple rule relative to the Bayes-optimal rule. We emphasize that in this case, we allow for the constant share mechanism to be adapted to the distribution over cost.

Proposition 6 (Using the Bayesian-Optimal Mechanism as Benchmark)
The minmax problem (51) *attains a value equal to the saddle point:*

$$\widehat{B}(\sigma) = \min_{G \in \Delta(\mathcal{C}_{cx})} \max_{z \in [0,1]} \frac{\int U(c,z)dG(c)}{U^*(G)}.$$

Proof. To prove the result, we first note that:

$$\widetilde{B}(\sigma) \geq \min_{G \in \Delta(\mathcal{C}_{cx})} \max_{z \in [0,1]} \frac{\int U(c,z)dG(c)}{\int W(c)dG(c)} = \underline{B}(\sigma). \tag{52}$$

The inequality follows from the fact that we are replacing the term in the denominator with a strictly larger term; the equality follows from the fact that:

$$\min_{G \in \Delta(\mathcal{C}_{cx})} \max_{z \in [0,1]} \frac{\int U(c,z)dG(c)}{\int W(c)dG(c)} = \min_{G \in \Delta(\mathcal{C}_{cx})} \max_{F \in \Delta[0,1]} \frac{\int U(c,z)dG(c)dF(z)}{\int W(c)dG(c)}.$$

That is, by allowing randomizations over constant share mechanisms we do not change the value of the minmax problem. We now note that:

$$\min_{G \in \Delta(\mathcal{C}_{cx})} \max_{F \in \Delta[0,1]} \frac{\int U(c,z)dG(c)dF(z)}{\int W(c)dG(c)} \leq \min_{G \in \Delta(\mathcal{C}_{cx})} \max_{F \in \Delta\{t:t:\mathbb{R}\to\mathbb{R}\}} \frac{\int U(c,t)dG(c)dF(t)}{\int W(c)dG(c)} = \underline{B}(\sigma).$$

The inequality follows from the fact that by minimizing over all transfers rather than constant share mechanisms we get a weakly larger minimization; the equality follows from the fact that this is the saddle point characterized in Proposition 5. Following Proposition 5, by randomizing over constant share mechanisms the buyer can guarantee himself $\underline{B}(\sigma)$, so

$$\min_{G\in\Delta(\mathcal{C}_{cx})}\max_{F\in\Delta[0,1]}\frac{\int U(c,z)dG(c)dF(z)}{\int W(c)dG(c)}\geq\underline{B}(\sigma).$$

We thus get (52). We now note that:

$$\widetilde{B}(\sigma)\leq\min_{c\in\mathcal{C}_{cx}}\max_{z\in[0,1]}\frac{U(c,z)}{W(c)}\leq\overline{B}(\sigma). \tag{53}$$

The first inequality follows from the fact that, in contrast to (51), we allow the minimization to be only over deterministic cost functions (rather than distributions over cost). In the denominator we replaced $U^*(F)$ with $W(c)$ but this is without loss of generality because for any degenerate distribution F we have that

$$U^*(F)=W(c).$$

That is, if there is no uncertainty about cost, the optimal mechanism extracts the efficient social surplus. The second inequality in (53) follows from the fact that for any $\widetilde{\gamma}\in\mathbb{R}$ and $z\in[0,\sigma]$:

$$\frac{U(c_{\widetilde{\gamma}},z)}{W(c_{\widetilde{\gamma}})}=\overline{B}(\sigma).$$

Hence, by minimizing over all costs we get a weakly smaller number. Following (52)-(53) and Proposition 5, we obtain the result. ■

We thus obtain the same competitive ratio as in Proposition 5. We consider the buyer surplus generated by the optimal rule (instead of the efficient social welfare) in (51) following some of the literature. For example, Rogerson (2003) and Chu and Sappington (2007) compare the performance of simple mechanisms relative to the Bayes optimal mechanisms in specific parametric environments. This approach is also pursued in Cai, Devanur, and Weinberg (2021) with notable applications to multi-unit optimal pricing and multi-unit auctions. These are problems where the solution to the optimal Bayesian problem is either unknown or computationally complex, and hence the question arises whether a simple solution adapted to the distribution of uncertainty can attain a good approximation in the sense of the competitive ratio.

References

Anunrojwong, J., S. Balseiro, and O. Besbes (2025): "Robust Auction Design with Support Information," *Management Science*, forthcoming.

Ball, I., and D. Kattwinkel (2025): "Robust Robustness," Discussion paper.

Baron, D., and R. Myerson (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 911–930.

Bergemann, D., T. Heumann, and S. Morris (2023): "Cost-Based Non-Linear Pricing," Discussion Paper 2368, Cowles Foundation for Research in Economics.

Bergemann, D., and K. Schlag (2008): "Pricing Without Priors," *Journal of the European Economic Association Papers and Proceedings*, 6, 560–569.

Brooks, B., and S. Du (2021): "Optimal Auction Design with Common Values: An Informationally Robust Approach," *Econometrica*, 89, 1313–1360.

——— (2025): "Informational Robustness in Mechanism Design: A Case for (and Against) the Worst Case," Discussion paper, University of Chicago and University of Claifornia at San Diego.

Cai, Y., N. R. Devanur, and S. M. Weinberg (2021): "A Duality-based Unified Approach to Bayesian Mechanism Design," *SIAM Journal on Computing*, 50, STOC16–160–STOC16–200.

Carrasco, V., V. F. Luz, P. Monteiro, and H. Moreira (2019): "Robust Mechanisms: The Curvature Case," *Economic Theory*, 68, 203–222.

Carroll, G. (2017): "Robustness and Separation in Multidimensional Screening," *Econometrica*, 85, 453–488.

Che, Y.-K., and W. Zhong (2025): "Robustly Optimal Mechanisms for Selling Multiple Goods," *Review of Economic Studies*, 92, 2923âĂŞ2951.

Chu, L., and D. Sappington (2007): "Simple Cost-Sharing Contracts," *American Economic Review*, 97, 419–428.

Deb, R., and A. Roesler (2024): "Multi-dimensional Screening: Buyer-Optimal Learning and Informational Robustness," *Review of Economic Studies*, 91, 2744–2770.

El-Yaniv, R., and A. Borodin (1998): *On-Line Computation and Competitive Analysis*. Cambridge University Press, Cambridge.

Eren, S., and C. Maglaras (2010): "Monopoly Pricing with Limited Demand Information," *Journal of Revenue and Pricing Management*, 9, 23–48.

Garrett, D. (2014): "Robustness of Simple Menus of Contracts in Cost-Based Procurements," *Games and Economic Behavior*, 87, 631–641.

——— (2021): "Payoff Implications of Incentive Contracting," *Theoretical Economics*, 16, 1281–1312.

Guesnerie, R., and J. Laffont (1984): "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm," *Journal of Public Economics*, 25, 329–369.

Guo, Y., and E. Shmaya (2025): "Robust monopoly regulation," *American Economic Review*, 115, 599–634.

Hall, R. L., and C. J. Hitch (1939): "Price theory and business behaviour," *Oxford Economic Papers*, 2, 12–45.

Hartline, J., and T. Roughgarden (2014): "Optimal Platform Design," Discussion paper, Northwestern University.

Laffont, J., and J. Tirole (1986): "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94, 614–641.

Maskin, E., and J. Riley (1984): "Monopoly with Incomplete Information," *RAND Journal of Economics*, 15, 171–196.

Mishra, D., S. Patil, and A. Pavan (2025): "Robust Procurement Design," Discussion paper.

Mussa, M., and S. Rosen (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–317.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Neeman, Z. (2003): "The Effectiveness of English Auctions," *Games and Economic Behavior*, 43, 214–238.

Rogerson, W. (2003): "Simple Menus of Contracts in Cost-based Procurement and Regulation," *American Economic Review*, 93, 919–926.

Scherer, F., and D. Ross (1990): *Industrial market structure and economic performance.* Houghton Mifflin.