THE UNIVERSITY OF CHICAGO

ESSAYS ON CONSTRAINED MECHANISM DESIGN AND REGULATION

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
CHRISTOPH SCHLOM

CHICAGO, ILLINOIS
AUGUST 2024

To my parents, Wendy and Darrell

TABLE OF CONTENTS

LIST OF FIGURES v

ACKNOWLEDGMENTS vi

1 PRICE DISTRIBUTION REGULATION 1
 1.1 References 40

2 SHARPENING WINKLER'S EXTREME POINT THEOREM: ECONOMIC APPLICATIONS 66
 2.1 References 77

3 CONCAVIFICATION BOUNDS AND MECHANISM SIMPLICITY 78
 3.1 References 87

LIST OF FIGURES

Figure 1.1 The allocative effects of the Mussa-Rosen distortion and various price-based regulations. 5

Figure 1.2 Proof idea for Theorem 1.2. 19

Figure 1.3 The regulator's optimal quality schedule for various values of α. 22

Figure 1.4 Comparison of the regulator's optimal quality schedules under different equity weights, with α fixed. 27

Figure 3.1 The optimal mechanism and the optimal posted-price mechanism under an arbitrarily chosen 2-point distribution. 84

Figure 3.2 The optimal mechanism and the optimal posted-price mechanisms under the worst-case distribution. 86

ACKNOWLEDGMENTS

I am deeply grateful to many academic mentors throughout my career: Phil Reny, Emir Kamenica, Alex Frankel, Doron Ravid, Joel Sobel, Elliot Lipnowski, Casey Mulligan, Ben Brooks, Joe Root, Eric Budish, Modibo Camara, Lars Stole and Juan-Pablo Xandri. I have also learned a great deal from fellow students: Zizhe Xia, Takuma Habu, Andy Choi, Elena Istomina, Karen Wu, John Mori, Kai Hao Yang, Sulagna Dasgupta, Sota Ichiba and Hazen Eckert.

On a personal level, I would also like to thank my family: Wendy, Darrell and Rob; my wonderful partner, Hannah Martin; and my close friends during my time in Chicago: Lisa Hillas and Kenan Diab.

CHAPTER 1

PRICE DISTRIBUTION REGULATION

Price Distribution Regulation

Christoph Schlom*

University of Chicago, Booth School of Business

July 25, 2024

I propose Price Distribution Regulation (PDR), an optimal price-based regulatory framework for differentiated product (Mussa-Rosen) monopolists. In PDR, a regulator sets a target probability distribution for transacted prices, which the monopolist's mechanism must meet. Since PDR only depends on price data, it is useful when product quality is difficult to directly regulate. I show that, while PDR is sufficient to fully restore allocative efficiency, a regulator with type-weighted utilitarian preferences will optimally distort qualities. Specifically, (1) a higher regulatory preference for consumer surplus to government revenue will lead the regulator to distort qualities upwards, whereas (2) a higher regulatory preference for low-type to high-type consumer surplus will reverse (1), harming *all* consumer types and increasing government revenue. Additionally, I show that PDR provides a novel incentive for monopolists to use mechanisms featuring price randomization, and characterize the regulations under which such randomization will occur.

*I and this paper have benefited greatly from the wisdom and patience of my advisor, Philip Reny. I would also particularly like to thank my committee: Alex Frankel, Emir Kamenica and Doron Ravid; as well as Lars Stole and Ben Brooks, who have all offered invaluable guidance. I would like to thank for helpful conversations: Modibo Camara, Andy Choi, Sulagna Dasgupta, Hazen Eckert, Takuma Habu, Marina Halac, Lisa Hillas, Elena Istomina, Elliot Lipnowski, Andrew McClellan, Casey Mulligan, Daniel Rappoport, Joseph Root, Chad Syverson, Zizhe Xia, and seminar audiences at UChicago, Booth and UC Davis.

1 Introduction

This paper studies the optimal price-based regulation of a differentiated product (Mussa-Rosen) monopolist (Mussa and Rosen, 1978). Such a monopolist screens consumers on their taste for quality, by offering a menu of quality/price pairs. Absent regulation, the quality obtained by each consumer taste type is inefficiently low, especially for low types. Appropriate regulation may help correct this inefficiency.

It is common for firms to practice such quality-based screening, and may become more so as new technologies lend themselves more readily to customization. Examples occur in air travel (first class/coach), cable television (different bundles of channels), broadband internet access (fast and slow connection speeds) and streaming services (ad quantity and/or content selection on Hulu or Spotify). Two differentiated product settings of particular economic importance are health insurance, where plans differ substantially according to breadth of coverage,[1] and low-income housing, where apartment qualities differ.

Economic analyses of the optimal regulation of such differentiated product monopolists have followed two main approaches. The first, taken by Goldman, Leland, and Sibley (1984) and Wilson (1993) posits that, through meticulous regulations, a regulator would be able to dictate the monopolist's entire quality/price menu. Such regulations are evidently sufficient to reverse the Mussa-Rosen allocative quality distortion (since the regulator can directly set each consumer type's menu item to the efficient quality level); the analysis in these papers typically focuses on a Ramsey problem analysis of how to maximize social surplus subject to a minimum profit guarantee for the monopolist.[2]

An important issue with this approach is that, in many settings, product quality is difficult or undesirable to explicitly control. For example, it would likely be difficult to write or enforce regulations regarding the quality of food or service on airplanes, the irritatingness of ads on streaming services, or the durability of various products. Further, one would expect

[1] See, e.g., Tebaldi, Torgovitsky, and Yang (2023).

[2] This analysis can also be thought of as describing optimal government provision of a differentiated product; indeed, these papers treat government provision and regulation interchangeably.

quality regulations to be especially plagued by Hayekian knowledge problems, if firms better understand the costs and consumer WTP associated with various aspects of quality than do regulators. For instance, it would seem unwise for a regulator to dictate the selection of TV shows available on Hulu, or the set of providers a particular health insurance plan must include in its coverage network.

Given this, a second approach to differentiated product regulation has been to explicitly specify particular *price-based* regulations,[3] and to analyze their effect on monopolist conduct. The two key regulations of this type which have been previously studied are a price cap (Besanko, Donnenfeld, and White, 1987), which specifies that no products (of any quality level) may transact at a price exceeding the cap; and an ad valorem subsidy (Krishna, 1990), which specifies a rebate to the monopolist, paid each time a transaction occurs, of a (constant) percentage of the transacted price.

In this paper, I propose a generalization of the second approach, which I call price distribution regulation. In price distribution regulation, the regulator mandates the distribution of transacted prices (as well as the number of transactions) the monopolist's mechanism must generate, and also sets a lump-sum monetary transfer to or from the monopolist. In practice, price distribution regulation might be implemented by recording transacted prices over a window of time, and offering the monopolist a reward or punishment proportional to the distance between that empirical distribution and the target distribution.

Price distribution regulation is the most powerful possible price-based regulation, and so provides an upper-bound against which to evaluate simpler price-based regulations, such as the cap or the ad-valorem subsidy. In a standard Mussa-Rosen numerical example, shown in Figure 1, price distribution regulation fully reverses the allocative inefficiency of the Mussa-Rosen distortion, and so restores 100% of the "lost" total surplus, whereas the cap and the ad-valorem subsidy restore 0.4% and 40% of this lost surplus, respectively.

I analyze price distribution regulation in two main thrusts. In the first thrust, I assume

[3] I.e., regulations that can be assessed using only price data.

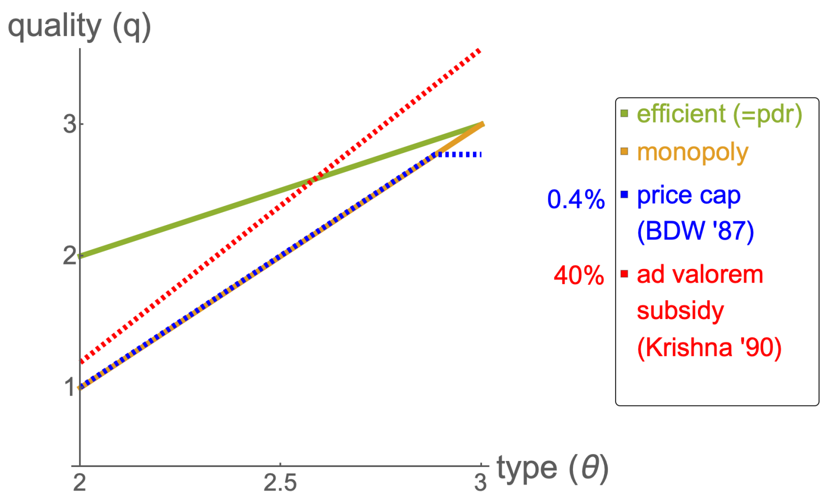

Figure 1: The allocative effects of the Mussa-Rosen distortion and various price-based regulations. Green: The efficient qualities for each consumer type, also obtained by price distribution regulation. Orange: Mussa-Rosen (unregulated monopoly) qualities. Blue: Price cap (Besanko et al., 1987). Red: Ad valorem subsidy (Krishna, 1990). Example parameters: consumer taste type $\theta \sim U(2,3)$; monopolist cost of quality $c(q) = q^2/2$.

that the monopolist will use a deterministic mechanism in response to any price distribution regulation. This assumption is standard in the Mussa-Rosen literature, but is typically without loss for the monopolist; in this setting, it is sometimes with loss. Under this no randomization assumption, I show that price distribution regulation is powerful enough that the regulator can force the monopolist to use any IC and IR (for consumers) price/quality schedule, with the important caveat that that schedule must give the lowest type consumer zero surplus. This caveat comes from the possibility of the monopolist undetectably degrading product quality.

In particular, this means that the regulator can always force the efficient quality schedule, as long as it is paired with the incentive-compatible price schedule which gives the lowest type zero surplus. Thus, a regulator interested only in total surplus can achieve her first-best. However, the "zero surplus at the bottom" friction rears its head when the regulator is interested in more fine-grained welfare measures, as is the case if she weights consumer and government surplus differently, or puts different welfare weights on different consumer types.

In this case, I show that a regulator who puts more weight on consumer than government surplus optimally forces the monopolist to use an inefficiently *high* quality schedule, reversing the usual Mussa-Rosen distortion. Perhaps counterintuitively, I additionally show that a "more equitable" regulator (one who puts more weight on low consumer types and less on high types) optimally forces the monopolist to use a lower quality schedule, increasing government revenue, but hurting *all* consumer types.

Additionally, in Appendix B, I consider an extension to a Baron-Myerson style regulatory environment (Baron and Myerson, 1982) in which the regulator does not have perfect information about the monopolist's costs, and can use an arbitrary price-based regulatory mechanism. I show that the regulator's problem decomposes into an "inner problem" of choosing the optimal cost-constrained price distribution regulation for each monopolist type, and an "outer" Baron-Myerson style problem of screening the monopolist.

In the second thrust, I relax the assumption that the monopolist use a deterministic mechanism. In this case, the monopolist may find it profitable to randomize over *prices* to satisfy the terms of the price distribution regulation. I show that this ability to "shroud" his mechanism with randomness allows him to effectively choose between a broader set of deterministic mechanisms: namely, those that induce a price distribution which is a mean-preserving contraction of the target price distribution. I then introduce a new result in mathematical optimization to characterize the regulated price distributions under which the monopolist would not want to randomize prices. When the regulator maximizes total surplus, her first-best price distribution satisfies this no-randomization condition whenever a certain condition related only to the distribution of consumer types is satisfied. This condition resembles but is distinct from Myerson regularity; like Myerson regularity, it is satisfied by many "common" distributions, including the uniform distribution.

The aforementioned mathematical optimization result concerns the minimization of a convex function over a majorization set (the set of distributions which are mean-preserving contractions or spreads of a given distribution). In this way, it is a convex optimization counterpart to recent fruitful work in economic theory on the extreme points of such majorization sets by Kleiner, Moldovanu, and Strack (2021), which are useful for linear optimization. The idea of my result is that, to rule out the profitability of any mean-preserving contraction of the given distribution, it suffices to rule out the profitability of a much smaller set: the "local" mean-preserving contractions. A local mean-preserving contraction of a discrete distribution is one which moves two adjacent points in the distribution slightly closer together; for a continuous distribution, it is defined in a limiting sense via a sequence of discretizations.

Taken together, my results should be thought of as providing a positive and a negative result for practical regulation. The positive result is that regulations which pertain *only* to price are sufficient[4] to reverse the *allocative* Mussa-Rosen inefficiency (i.e., price distribution regulation attains the first-best when total surplus is the regulatory objective). Under this total surplus objective, my paper thus "closes the gap" between the abstract regulatory approach taken by Goldman, Leland, and Sibley (1984) and Wilson (1993) and the ad-hoc price based approach taken by Besanko, Donnenfeld, and White (1987) and Krishna (1990). The negative result describes two frictions inherent to any price-based regulation: (1) a *quality degradation* friction, that results in the lowest consumer type necessarily receiving 0 surplus, and (2) a *mechanism design* friction, under which the monopolist can bypass some regulations by incorporating randomness directly into his mechanism.

1.1 Related literature

The starting point for this paper is the classic literature on optimal screening by a monopolist, by Mussa and Rosen (1978) and Maskin and Riley (1984). Those papers identify a key inefficiency: that the monopolist optimally lowers the quality of goods obtained by low-type consumers. This paper attempts to use regulation to correct that inefficiency.

A follow-up literature has taken on the question of how to regulate a screening monopolist. Besanko, Donnenfeld, and White (1987) consider two natural policies: a price cap, and a

[4] In the deterministic case.

quality floor. Both of those regulations monitor only a single product attribute (quality or price). This is also the case in my setting (and is the source of the regulatory friction that the low-type consumer must get 0 surplus); my price distribution regulation generalizes their price cap. Similarly, Krishna (1990) studies the impact of specific (per-item) and ad-valorem (per-dollar) taxes on these monopolists; price distribution regulation (plus a lump-sum transfer) also generalizes these policies.

Other papers in this literature consider regulations that operate on quality and price simultaneously. This is appropriate in contexts where quality is easily assessed, such as electricity markets, where we can interpret a household's service quality as a function of the kilowatt hours they consume. One natural question (see Spence (1977), Roberts (1979), Goldman, Leland, and Sibley (1984) and Wilson (1993)) is to find the social welfare-maximizing screening schedule, subject to a minimum profit constraint for the monopolist (to offset fixed costs). This exercise can be interpreted as describing the conduct of a regulator who can make arbitrary rules regarding quality and price simultaneously. Armstrong, Cowan, and Vickers (1995) assess more practical regulations in this framework, including a cap on the mean price (mirroring the policy of average revenue regulation used to regulate utilities in the UK).[5][6] More recently, Akbarpour, Dworczak, and Kominers (2023) and Nikzad (2022) study an extension of such a model (which they interpret as a government provision problem); they are particularly interested in questions regarding optimal government rationing, which occurs due to a distinct mechanism from monopolist randomization in my model.

Finally, a cost-based regulatory scheme for screening monopolists was proposed by Kim and Jung (1995) and iterated on by Lee (1997). In line with cost-based regulations in traditional settings,[7] this scheme has the advantage of not requiring any regulatory knowledge

[5]Since this mean is taken per unit quantity (e.g., per kilowatt hour), it differs economically from mean price regulation in my setting.

[6]More recently, Wong (2012) studies the conduct of a screening monopolist, subject to exogenous restrictions, for instance on the number of service classes. Though one could interpret these restrictions as originating from a regulator, that paper's results focus on an implementability question, rather than optimal regulation.

[7]The "Incremental Surplus Subsidy" scheme of Sappington and Sibley (1988) is the counterpart of this scheme for non-screening monopolists.

of the monopolist's cost structure (though realized costs must be observable), but the disadvantage of introducing a cost-bloating incentive.

More broadly, this paper relates to the literature on contracting and regulation with endogenous and unobservable quality. A number of papers, including Frankena (1975), Kihlstrom and Levhari (1977) and Mulligan and Tsui (2016) analyze quality degradation under price controls in competitive markets. Hart, Shleifer, and Vishny (1997) analyze the optimal assignment of control rights between a government and a firm in an incomplete contracts model, where product quality is observable but not contractible.

In unpublished work, I analyze monopoly regulation with unregulable product quality, in a "homogenous" market, where all consumers value quality equally. I show that, whereas quality degrades under price controls, it does not under quantity targets.[8] Since the price distribution regulation I analyze in this paper is powerful enough to specify quantity targets, this same effect rears its head, and explains why price distribution regulation often *raises* prices, rather than lowering them:[9] the consumer-base expanding work traditionally done by price caps is now wholly accomplished via quantity targets.

My analysis of a regulator with imperfect information about the monopolist's cost is a straightforward extension of Baron and Myerson (1982) to the setting of regulating a screening monopolist; and my main result in this analysis suggests a fundamental parallel between the two settings.

Finally, my analysis of potential monopolist randomization under price distribution regulation connects to a broader theoretical literature on the structure of optimal selling mechanisms, and particularly the benefits of randomization. The initial result in this literature, Riley and Zeckhauser (1983) is a negative one: an unrestricted monopolist selling an indivisible good never benefits from randomization. This conclusion is reversed when the monopolist "irons," as he might in the presence of explicit (Loertscher and Muir, 2022) or de-facto (My-

[8] In a classic paper, Weitzman (1974) also compares price and quantity controls, but from the perspective of imperfect regulatory information, rather than latent quality.

[9] In Appendix C, I show that mean price regulation always increases the mean price above the unregulated monopoly level.

erson, 1981; Bulow and Roberts, 1989) quantity constraints. In unpublished work, I show that, much like a quantity constraint, a regulatory price constraint can introduce an incentive to randomize; Loertscher and Muir (2023) take up similar themes.

Under the linear utility specification of the Mussa-Rosen model that I use in this paper, a similar negative result holds – that an unrestricted monopolist would never use a randomized mechanism – though Strausz (2006) shows that randomization may be optimal under more general utility specifications. Through this lens, the need to comply with price-distribution regulation is a novel impetus for monopolist randomization; such randomization is important to understand as an "unintended consequence" of regulation.

The mathematical analysis of this potential randomization builds on the theory of mean-preserving spreads introduced to economics by Rothschild and Stiglitz (1971). In particular, my notion of a local mean-preserving contraction is closely related to the "atomic" discrete mean-preserving contractions proposed in that paper, and clarified by Leshno, Levy, and Spector (1997). I then use these local mean-preserving contractions to prove a convex optimization result for moment sets; this complements the work of Kleiner, Moldovanu, and Strack (2021), which is critical for linear optimization over these sets.

This convex optimization result gives me a condition, similar to but distinct from Myerson regularity, under which the monopolist does not benefit from randomization under the efficient price distribution regulation. Loosely speaking, under this condition, the monopolist's preferences are sufficiently aligned with the regulator's that he wishes to assign higher types higher qualities while respecting the price distribution. This facilitates a comparison between my distributional constraint and other papers in which a distributional rule is imposed by a principal or naturally arises. In Jackson and Sonnenschein (2007), one could say that this alignment is always satisfied, whereas in Frankel (2014) and Lin and Liu (2022) this alignment is satisfied given natural monotonicity conditions between the two agents' preferences.

2 Model

I will first describe the underlying Mussa-Rosen preferences, and then the regulatory and mechanism design environment.

Preferences. A monopolist (he) sells an item of variable quality to a large number of consumers (they). Each consumer is described by a privately known taste-for-quality type, $\theta \in \Theta := [\underline{\theta}, \overline{\theta}]$, with $\underline{\theta} \geq 0$, drawn iid from a distribution with CDF F, which is known by the monopolist. I assume that F is full-support over the interval Θ, and admits a PDF, f, with $f(\overline{\theta}) > 0$. If a consumer receives an item of quality q and pays price t, their payoff is given by:

$$u = \theta q - t.$$

All qualities are non-negative. If a consumer does not buy an item, they receive 0.

The monopolist finds it costly to provide quality, and has a cost of quality function, $c : \mathbb{R}^+ \to \mathbb{R}^+$. If he sells an item of quality q to the consumer at price t, he obtains profit:

$$\pi = t - c(q).$$

I assume that $c(0) = 0$, and that c is strictly increasing and strictly convex on $(0, \infty)$. A quality of 0 represents non-service, and a discontinuity in c at 0 represents a fixed cost of service.

Regulatory environment. First, the regulator (she), who knows F and c, announces a mandated price distribution, G, supported on $\mathbb{R}^+$, as well as a lump-sum transfer, τ, from the monopolist to the government. (Under the regulatory objectives I will consider, the regulator will use this transfer to fully expropriate the monopolist's profits.) Then, the monopolist chooses whether to produce or not. If he does not produce, he receives an outside option of value 0.[10] If he does produce, he must choose a selling mechanism satisfying the properties

[10] For the regulatory objectives I will consider, the analysis would not change if the monopolist had a non-zero outside option.

I describe below.

The idea is that this mechanism must satisfy three economically substantive restrictions: it must be individually rational for the consumers, induce the distribution G, and not involve randomization (I relax the restriction on randomization in Section 5). The other properties are without loss by a revelation principle, as I show in Appendix A.

Properties of the selling mechanism

The monopolist must choose a direct, deterministic selling mechanism (q, t), with $q : \Theta \to \mathbb{R}^+$ and $t : \Theta \to \mathbb{R}^+$, and where t is Lebesgue measurable. Further, (q, t) must satisfy the incentive compatibility (IC) and individual rationality (IR) properties for the consumer:

$$\theta q(\theta) - t(\theta) \geq \theta q(\hat{\theta}) - t(\hat{\theta}) \text{ for all } \theta, \hat{\theta} \in \Theta \tag{IC}$$

$$\theta q(\theta) - t(\theta) \geq 0 \text{ for all } \theta \in \Theta. \tag{IR}$$

Additionally, (q, t) must satisfy the regulatory constraint, which states that the induced price distribution must coincide with G:

$$G(\hat{t}) = \mu_F(\theta : t(\theta) \leq \hat{t}) \text{ for all } \hat{t} \in \mathbb{R}, \tag{G Reg}$$

where μ_F is the measure associated with F. Notice that the regulator is able to observe prices of 0, even when they arise from non-service ($q = 0$). The interpretation is that the regulator can count transactions, and mandate a target quantity of transactions in addition to the distribution over transacted prices.[11]

Thus, facing mandated price distribution G, the monopolist solves:

$$\max_{E,q,t} E\Big(\int [t(\theta) - c(q(\theta))] dF(\theta) - \tau \Big) \tag{1}$$
$$\text{s.t. (IC), (IR), (G Reg),}$$

[11]The ability to regulate transaction quantity is in-keeping with the spirit of giving the regulator access to all possible price data.

where $E \in \{0, 1\}$ represents the monopolist's decision to produce. For the moment, I will leave the regulator's preferences unspecified.

2.1 Discussion: interpretation of Θ

The space of consumer types, Θ, has two natural interpretations. First, it might represent the range of tastes of all potential consumers. But, in practice, the regulator might only learn about the distribution of consumers' tastes by observing the behavior of the unregulated monopolist. In this case, she would know little about the taste distribution among unserved consumers. In particular, any monopolist's pricing strategy is always consistent with there being no consumers in the population with type below the lowest type served by the monopolist. (In an empirical paper, Luo, Perrigne, and Vuong (2018) estimate a Mussa-Rosen model in the cell phone service industry, and encounter exactly this limitation; see Section IV, C. of their paper.)

Practically, then, perhaps we should think that Θ represents the range of tastes of consumers served by the unregulated monopolist. In this case, the optimal quantity target, for all regulator preferences I will consider, takes the form of a full-service requirement: that the regulated monopolist continue to serve all of Θ.

3 Preliminary analysis: enactability

Let us simplify the monopolist's problem, (1). First, we will appeal to the standard result that a schedule (q, t) is implementable for the monopolist if and only if it is non-decreasing and satisfies an envelope formula. Specifically, in this section, the envelope formula for q in terms of t will be useful, rather than the (perhaps more familiar) envelope formula for t in terms of q.

Lemma 1 (t-implementability)**.** *For a schedule (q, t), the following are equivalent:*

1. *The schedule satisfies (IC) and (IR).*

2. *The schedule satisfies t-implementability:*

   (a) t *is non-decreasing*

   (b) *The envelope formula,*

$$q(\theta) = \frac{t(\theta)}{\theta} + \int_{\underline{\theta}}^{\theta} \frac{t(x)}{x^2} dx + B \tag{2}$$

holds, for some $B \geq 0$.

Next, for any price distribution, G, the following lemma shows that there is a unique non-decreasing price schedule that induces it. For intuition, notice that if G is full-support on some interval, then there is a unique strictly increasing t that induces it. To see this, notice that, for any θ, the regulatory constraint gives:

$$G(t(\theta)) = P_{\hat{\theta} \sim F}(t(\hat{\theta}) \leq t(\theta)) = F(\theta),$$

and so we must have:

$$t(\theta) = G^{-1}(F(\theta)).$$

When G need not be full-support and t can be weakly increasing, a similar result still holds, but we must be a bit more careful. All proofs, including the proof of this lemma, are found in Appendix E.

Lemma 2 (Invertibility)**.** *Suppose that price schedule* t *is non-decreasing and satisfies (G Reg) for the price distribution* G*. Then,*

$$t(\theta) = G^{-1}(F(\theta)),$$

except possibly on a set of Lebesgue measure 0, where G^{-1} *is the generalized inverse CDF of* G*.*[12]

[12]Throughout this paper, if H is a CDF then H^{-1} denotes the generalized inverse CDF of H.

Now, using Lemma 1 and Lemma 2, we can rewrite the monopolist's problem under mandated price distribution G (if he produces) as:

$$\max_{q,t} \int [t(\theta) - c(q(\theta))] dF(\theta) \tag{3}$$

$$\text{s.t. } t(\theta) = G^{-1}(F(\theta))$$

$$q(\theta) = \frac{t(\theta)}{\theta} + \int_{\underline{\theta}}^{\theta} \frac{t(x)}{x^2} dx + B, \text{ for some } B \geq 0.$$

The solution to this problem is trivial: the monopolist sets $B = 0$, since any higher B strictly increases his costs. (Notice that this choice of B results in a non-negative quality schedule, since G is a distribution over non-negative prices.) Given this, the constraints entirely pin down (q, t).

Let us now discuss what schedules the regulator can induce the monopolist to use. First, it will be useful to have a term to describe an IC and IR schedule (which the monopolist must necessarily use).

Definition 1. *Say that a schedule, (q, t), is a* ***market outcome*** *if it satisfies (IC) and (IR).*

Next, we will call a market outcome "enactable" if the regulator can induce it.

Definition 2. *Say that a market outcome, (q, t), is* ***enactable*** *if there exists a price distribution G, and a transfer, τ such that $(E = 1, q, t)$ solves (1) under G, τ.*

The following theorem characterizes the enactable market outcomes. It says that a market outcome is enactable if and only if it delivers 0 surplus to the lowest consumer type.

Theorem 1 (Enactability)**.** *A market outcome (q, t) is enactable if and only if $\underline{\theta} q(\underline{\theta}) - t(\underline{\theta}) = 0$.*

Theorem 1 follows easily from the previous discussion. The idea is that, by Lemma 2, the regulator can induce any non-decreasing t by mandating the price distribution induced by that t; this t must be accompanied by the implementable q with $B = 0$. Importantly, the regulator cannot enact a schedule which gives the low type consumer positive surplus, because the monopolist could "invisibly" lower q by a constant, decreasing his costs without violating the price distribution regulation.

An important consequence of Theorem 1 is that price distribution regulation is powerful enough to completely undo the allocative inefficiency of the Mussa-Rosen distortion. To show this, consider the efficient quality schedule,

$$q^{\text{eff}}(\theta) := \arg\max_{q} \theta q - c(q). \tag{4}$$

Notice that a regulator who wishes to maximize total surplus (defined as the sum of consumer surplus, monopolist profits and government revenues) achieves her first-best whenever the monopolist's quality schedule is q^{eff}, since this maximizes allocative efficiency.

Theorem 1 then implies that a regulator who wishes to maximize total surplus can achieve her first-best, by enacting the efficient quality schedule (paired with the 0-surplus-at-the-bottom IC price schedule).

Proposition 1. *The market outcome* $(q^{\text{eff}}, t^{\text{eff}})$*, where* $t^{\text{eff}}(\theta) = \theta q^{\text{eff}}(\theta) - \int_{\underline{\theta}}^{\theta} q^{\text{eff}}(x)dx$*, is enactable. Therefore, a total surplus maximizing regulator can achieve her first-best.*

(Observe that $(q^{\text{eff}}, t^{\text{eff}})$ is indeed a market outcome, because q^{eff} is non-decreasing by the concavity of c, and t^{eff} satisfies the envelope formula for t in terms of q.)

To conclude our discussion of enactability, let me point out that Theorem 1 is really a result about one-sided regulability. That is, in an alternative model where the regulator could set a quality distribution, but not regulate price, Theorem 1 (and therefore all of my analysis in Section 4) would still hold. The key argument with unregulable price parallels the argument with unregulable quality: now the regulator cannot deliver positive surplus to the low type, because the monopolist could invisibly raise the price (rather than lower the quality) for all types.

4 Consumer surplus and distributional concerns

So far, we have seen that a regulator interested only in allocative efficiency can achieve her first-best using price distribution regulation. However, this focus on allocative efficiency sidesteps distributional considerations, including the funds raised (or expended) by the government, and the surplus accruing to each consumer type. In this section, I will consider a regulator who puts differential weight on each of these stakeholders. Now, the "0 surplus at the bottom" enactability friction meaningfully rears its head, in some cases constraining the regulator beyond the demands of consumer IC and IR. We will examine comparative statics describing how the regulator's optimal regulation changes with the weight she puts on each stakeholder.

4.1 Consumer surplus

First, let us consider a regulator who distinguishes between consumer surplus (CS) and government revenue (GR), and maximizes the quantity $\alpha\text{CS} + \text{GR}$, for some non-negative α.[13] Such a regulator will choose τ to fully extract the monopolist's profits. (That is, for any G, she will optimally set τ equal to the monopolist's profits under the market outcome which G enacts.) Using Theorem 1, the regulator's problem can be written as:

$$\max_{q \text{ incr.}} \int \alpha[\theta q(\theta) - t(\theta)] + [t(\theta) - c(q(\theta))]dF(\theta), \tag{5}$$

where $t(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(x)dx$, and we are using the standard implementability characterization in terms of q. Notice that the second term in the integral corresponds to τ.

To aid in the analysis of this problem, let us introduce the standard virtual value function, and state the key fact about it.

[13]Notice that a regulator who instead puts weight α on consumer surplus, 1 on government revenue and $\alpha_m < 1$ on monopolist revenue will act as though she has the above preferences, since she will also use the lump-sum transfer to fully extract the monopolist's profits.

Definition 3. *The virtual value function, $\phi : \Theta \to \mathbb{R}$, is given by:*

$$\phi(\theta) := \theta - \frac{1 - F(\theta)}{f(\theta).}$$

Lemma 3 (Myerson's identity)**.** *Let q and t be functions on Θ, with $t(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(x)dx$. Then,*

$$\int t(\theta)dF(\theta) = \int \phi(\theta)q(\theta)dF(\theta).$$

Using Myerson's identity, we can rewrite (5) as:

$$\max_{q \text{ incr.}} \int [q(\theta)\phi_\alpha(\theta) - c(q(\theta))]dF(\theta), \tag{6}$$

where

$$\phi_\alpha(\theta) := \theta - (1 - \alpha)\frac{1 - F(\theta)}{f(\theta)}$$

To build intuition for the regulator's problem, notice that, while the regulator cannot deliver positive surplus to the lowest type, she can raise higher types' utilities by raising the quality for lower types. (This can be seen from the standard IC identity $u'(\theta) = q(\theta)$, where $u(\theta)$ gives type θ's utility under schedule q.) This delivers two natural conjectures. First, we might expect that a regulator who puts higher weight on consumer surplus will enact a higher quality schedule. Second, since each consumer type only benefits from quality increases to strictly lower types, we might expect that (modulo monotonicity issues) the optimal quality schedule for any regulator will remain undistorted at the top. Theorems 2 and 1 establish these results.

Proposition 2. *Let $\alpha \leq \tilde{\alpha}$, and suppose that q^α is a solution to* (6) *with consumer surplus weight α. Then, there exists a quality schedule, $q^{\tilde{\alpha}}$, that solves* (6)*, with consumer surplus weight $\tilde{\alpha}$, and such that $q^{\tilde{\alpha}}(\theta) \geq q^\alpha(\theta)$, for all $\theta \in \Theta$.*

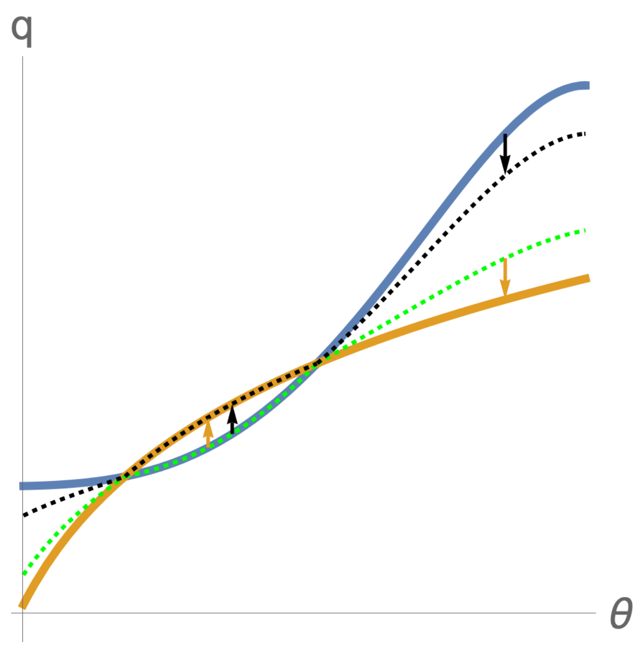

Figure 2: Proof idea for Theorem 2. Blue: $q^{\tilde{\alpha}}$. Orange: q^{α}. Green: q_{γ}. Black: $q^{\tilde{\alpha}} + h$. If moving from green to orange increases surplus, then so does moving from blue to black.

Figure 2 illustrates the idea behind the proof of Proposition 2. If $q^{\tilde{\alpha}}$ (blue) ever falls below q^{α} (orange), we can define the function q_{γ} (green), as:

$$q_{\gamma}(\theta) := (1 - \gamma) \min\{q^{\alpha}(\theta), q^{\tilde{\alpha}}(\theta)\} + \gamma q^{\tilde{\alpha}}(\theta),$$

where $\gamma \in [0, 1]$ is such that $PS(q_{\gamma}) = PS(q^{\alpha})$, where $PS(q) := \int [t(\theta) - c(q(\theta))] dF(\theta)$ is the government's revenue (equal to the monopolist's profit without the transfer). Such a γ exists by an intermediate value theorem argument. Clearly, q_{γ} is monotone; therefore, it was feasible for the α regulator, but was not chosen. It must therefore have lower CS than q^{α} (since both have the same PS).

Now, define the function $h(\theta) := q^{\alpha}(\theta) - q_{\gamma}(\theta)$. We will argue that $q^{\tilde{\alpha}} + h$ (black) has both higher CS and higher PS than $q^{\tilde{\alpha}}$. Intuitively, we wanted to add h at q_{γ} (yielding q^{α}), and so also want to add h at $q^{\tilde{\alpha}}$. CS is straightforward: CS is linear in q, so it increases by the same amount when h is added to *any* schedule. On the other hand, PS is composed of a linear revenue component in q, and a convex cost component. Of course, the revenue changes by the same amount as it changed from q_γ to q^α, but by convexity, the costs decrease *more* (they change by the same amount on the "lens" where $h > 0$, and decrease by more outside that region). Thus, $q^{\tilde{\alpha}} + h$ dominates $q^{\tilde{\alpha}}$; we can also show that it is monotone and hence feasible. This contradicts the optimality of $q^{\tilde{\alpha}}$ for the $\tilde{\alpha}$ regulator.

We can obtain the standard Mussa-Rosen downward quality distortion result from Proposition 2 by taking $\alpha = 0$ and $\tilde{\alpha} = 1$. In this case, q^α gives the monopoly schedule, while $q^{\tilde{\alpha}}$ gives the efficient schedule. The following corollary to Proposition 2 shows that a regulator more interested in consumer surplus than government revenue optimally enacts an *upward* quality distortion, reversing the Mussa-Rosen distortion.

Corollary 1. *For $\alpha \geq 1$ there exists a quality schedule, q^α, that solves* (6)*, with consumer surplus weight α, and such that:*

$$q^\alpha(\theta) \geq q^{\mathit{eff}}(\theta),$$

for all $\theta \in \Theta$.

Corollary 1 follows by taking $\alpha = 1$ and $\tilde{\alpha} \geq 1$ in Proposition 2.

The regulatory distortion in Corollary 1 is a consequence of the "no surplus at the bottom" friction. The regulator is only able to improve consumer surplus by distorting qualities to increase consumers' information rents; she is not able to, for instance, induce the monopolist to sell the efficient qualities at a discount (relative to the lowest type's WTP). This contrasts with the analysis in e.g., Goldman, Leland, and Sibley (1984) and Wilson (1993). As I argue in detail in Section 4.1.1, those analyses should be thought of as describing *government provision*, rather than *regulation* of a differentiated product; a key contribution of this paper is to demarcate a distinction between those two concepts.

To prove the "no distortion at the top" result, we will introduce a regularity condition, analogous to Myerson regularity. First, the solution to the relaxation of (6) in which monotonicity is dropped is given by:

$$q^*(\theta) = \arg\max_q \phi_\alpha(\theta)q - c(q). \tag{7}$$

This is non-decreasing as long as ϕ_α is monotone, which the following condition imposes.

Definition 4 (CS regularity)**.** *Say that* (α, F) *is CS regular if:*

$$\theta - (1-\alpha)\frac{1-F(\theta)}{f(\theta)}$$

is non-decreasing in θ*.*

CS regularity is easy to satisfy for α near 1, since θ is increasing. For α large, it is difficult to satisfy. For instance, the uniform distribution does not satisfy CS regularity for $\alpha > 2$.

If CS regularity holds, the optimal solution to (6) is given by (7). Comparing (7) with (4), since $\phi_\alpha(\overline{\theta}) = \overline{\theta}$, we find that q^* satisfies the no distortion at the top condition, $q^*(\overline{\theta}) = q^{\text{eff}}(\overline{\theta})$. I record this as the following observation.

Observation 1. *Suppose* (α, F) *satisfies CS regularity. Then, the schedule*

$$q^*(\theta) = \arg\max_q \phi_\alpha(\theta)q - c(q).$$

solves (6)*, and satisfies* $q^*(\overline{\theta}) = q^{\mathit{eff}}(\overline{\theta})$*.*

Figure 3, which plots the regulator's optimal quality schedule for $\alpha = 0$ (bottom) through $\alpha = 4$ (top), illustrates the results in this section. As Proposition 2 shows, the optimal quality schedule is higher for higher α, and, as Corollary 1 shows, features upward distortion relative to the efficient schedule (in green) for $\alpha > 1$. Finally, for the blue schedules ($\alpha \leq 2$), CS regularity is satisfied, and the schedules feature no distortion at the top (Observation 1), while for the orange schedules, CS regularity is not satisfied, and there is distortion at the top.

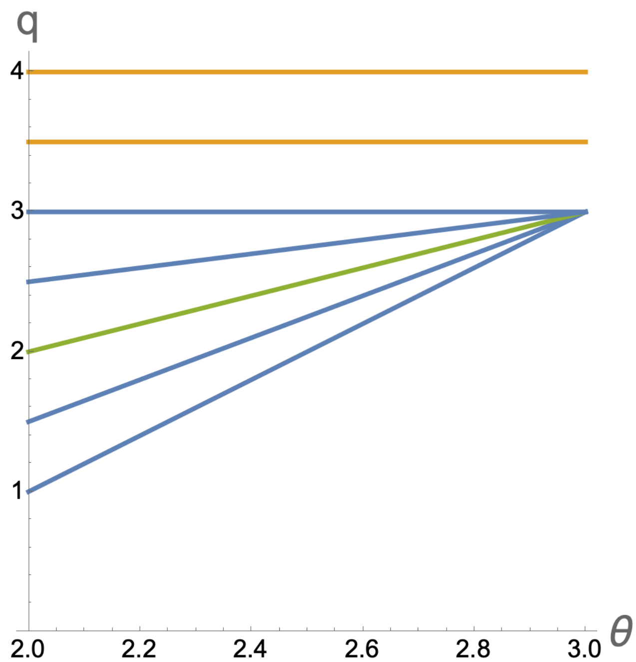

Figure 3: Bottom to top: The regulator's optimal quality schedule for $\alpha = 0, 0.5, 1, 1.5, 2, 3, 4$. Green: Efficient quality schedule ($\alpha = 1$). Blue: CS regularity holds. Orange: CS regularity fails. Throughout, $F \sim U(2, 3)$, $c(q) = q^2/2$.

4.1.1 Comparison to government provision

Now, let us consider a model, studied by Goldman, Leland, and Sibley (1984) and Wilson (1993), in which a government planner provides the differentiated product directly. The government chooses the market outcome, (q, t), pays the production costs, and collects the revenues. We will additionally stipulate that the government's mechanism must generate some minimum profit, π_0, perhaps to cover its operating costs.[14]

If, as in Section 4.1, the government maximizes $\alpha\text{CS} + \text{GR}$, and if $\alpha \leq 1$, then, given any q, government will optimally *choose* t to give the lowest consumer type 0 surplus. This is because, once q has been chosen, IC pins down t up to a constant; and, since $\alpha \leq 1$, the government would like to keep this constant as low as possible. In this case, our models

[14] This prevents the government from setting an infinitely *negative* price for all goods when $\alpha > 1$. Another approach (see Goldman, Leland, and Sibley (1984)) would would be to treat $\alpha > 1$ as an inadmissable government preference. I disagree with this approach: in some settings, say low income housing, it seems that a government would in fact like to transfer some money at a 1:1 rate to the population in question. A third approach would be to restrict prices to be positive; this approach would lead to a third distinct analysis.

align, and PDR attains the government provision upper bound.[15]

On the other hand, if $\alpha > 1$, the government's optimal mechanism sets $q = q^{\text{eff}}$, and $t(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(x)dx - B$, with B chosen so that the mechanism yields exactly π_0 profit. (This must be optimal, since it maximizes allocative efficiency while holding the government to its lowest allowed profit level.) This contrasts with the upward distortion described in Corollary 1, which arises from the "0 surplus at the bottom" friction. In that analysis, the regulator cannot transfer surplus to consumers "directly," by decreasing prices while holding qualities fixed, and instead must do so indirectly, by increasing their information rents.

Practically, the message of my paper in relation to this literature is the following. When the regulator does not have an especially high interest in consumer welfare in the regulated market ($\alpha \leq 1$), price-based regulation is preferable to government provision. This is because the two are theoretically equivalent, but there are likely to be unaccounted for bureaucratic and know-how costs associated with government provision. On the other hand, when the regulator has an especially high interest in consumer welfare ($\alpha > 1$), there may be more justification for direct government provision. This seems to be the case in public housing: the government is able to use its simultaneous control over prices and qualities to ensure that the cheapest units are priced substantially below their recipients' WTP.

4.2 Distributional concerns

A regulator may wish to give different ethical consideration to the different consumer types. For instance, if higher type consumers are on average wealthier, a utilitarian regulator might discount their surplus, compared to that of lower types. I augment the regulator's preferences from Section 4.1 to study such a regulator's optimal regulation.

My main result in this section is again a comparative static: this time, holding the regulator's aggregate concern for consumer surplus constant, and varying her within-consumer preferences. I will say that one regulator is more equitable than another if she puts higher

[15] Despite this alignment, the comparative statics I derive in Sections 4.1 and 4.2 have, to my knowledge, not been previously observed in this government provision literature.

welfare weight on lower type consumers and lower weight on higher type consumers. The result then is that a more equitable regulator enacts a schedule that lowers quality for all consumer types. Such a schedule is worse for *all* consumers, but raises more government revenue.

The intuition again comes from the "0 surplus at the bottom" friction. Since it is difficult for the regulator to enact a schedule which significantly benefits low type consumers, the more equitable regulator substitutes towards raising more government revenue. Surprisingly, even for low consumer types, this effect outweighs the regulator's move towards a regulation that comparatively benefits these low types more.

Formally, let us suppose the regulator puts Pareto weight 1 on the government, and $\alpha\rho(\theta)$ on consumer type θ, where $\rho : \Theta \to \mathbb{R}^+$ is normalized so that $\int \rho(\theta)dF(\theta) = 1$. The residual weight, α, represents the regulator's payoff if all consumer types receive 1 dollar of surplus. We will compare regulator equitability by comparing the ρ functions in the following way, and holding α fixed.

Definition 5. *Given weighting functions $\rho, \tilde{\rho} : \Theta \to \mathbb{R}^+$, say that $\tilde{\rho}$ is **more equitable** than ρ if, for all θ in Θ, we have:*

$$\int_{\underline{\theta}}^{\theta} \rho(x)dF(x) \leq \int_{\underline{\theta}}^{\theta} \tilde{\rho}(x)dF(x).$$

This notion of equitability resembles first-order stochastic dominance: it demands that the aggregate welfare weight up to any type for the more equitable regulator exceed that for the less equitable regulator. It is a relatively weak condition: for instance, it is implied by a decreasing ratio of welfare weights between the two regulators.

As in Section 4.1, using Theorem 1, the regulator solves:

$$\max_{q \text{ incr.}} \int \alpha\rho(\theta)[\theta q(\theta) - t(\theta)] + [t(\theta) - c(q(\theta))]dF(\theta) \tag{8}$$

where $t(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(x)dx$. To analyze (8), the following adaptation of Myerson's identity will be useful.

Lemma 4 (Weighted Myerson's identity)**.** *Let q and t be functions on Θ, with $t(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(x)dx$. Then,*

$$\int \rho(\theta)[\theta q(\theta) - t(\theta)]dF(\theta) = \int q(\theta)\frac{1 - F_\rho(\theta)}{f(\theta)}dF(\theta),$$

where

$$F_\rho(\theta) := \int_{\underline{\theta}}^{\theta} \rho(\theta)dF(\theta).$$

Then, using Lemma 4 and the weighted Myerson's identity, we can rewrite (8) as:

$$\max_{q \text{ incr.}} \int [q(\theta)\phi_{\alpha,\rho}(\theta) - c(q(\theta))]dF(\theta), \tag{9}$$

where

$$\phi_{\alpha,\rho}(\theta) := \theta - \frac{1 - F(\theta)}{f(\theta)} + \alpha\frac{1 - F_\rho(\theta)}{f(\theta)}.$$

We are now ready to state the main result of this section.

Proposition 3. *Fix $\alpha \geq 0$, and let $\rho, \tilde{\rho} : \Theta \to \mathbb{R}^+$ be weighting functions with $\tilde{\rho}$ more equitable than ρ. Let q^ρ be an optimal solution to* (9) *with weighting function ρ. Then, there exists an optimal solution, $q^{\tilde{\rho}}$, to* (9) *with weighting function $\tilde{\rho}$, such that $q^{\tilde{\rho}}(\theta) \leq q^\rho(\theta)$ for all θ in Θ.*

Proposition 3 says that the more equitable regulator enacts a market outcome which gives lower qualities to all consumer types. From the IC identity, $u'(\theta) = q(\theta)$, (and the fact that $u(\underline{\theta}) = 0$) we can immediately see that this market outcome also yields lower utility to all consumer types:

Corollary 2. *Let $\alpha, \rho, \tilde{\rho}, q^\rho, q^{\tilde{\rho}}$ be as defined in Proposition 3, and let u^ρ and $u^{\tilde{\rho}}$ be the utility schedules corresponding to q^ρ and $q^{\tilde{\rho}}$. Then, $u^{\tilde{\rho}}(\theta) \leq u^\rho(\theta)$, for all $\theta \in \Theta$.*

The proof of Proposition 3 follows a similar variational approach to the proof of Proposition 2. Notice again that it does not require any regularity assumptions.

To analyze the problem in more detail, let us introduce a regularity condition, analogous to the one in Section 4.1. As in that section, in the relaxed problem where monotonicity is dropped, the optimal schedule is given by:

$$q^{**}(\theta) = \arg\max_q \phi_{\alpha,\rho}(\theta)q - c(q).$$

That schedule is monotone if the following condition is satisfied.

Definition 6 (Equity regularity)**.** *Say that (α, ρ, F) is equity regular if:*

$$\theta - \frac{1-F(\theta)}{f(\theta)} + \alpha\frac{1-F_\rho(\theta)}{f(\theta)}$$

is non-decreasing in θ.[16]

Notice that equity regularity is equivalent to CS regularity when ρ is the constant function $\rho(\theta) = 1$.

If this regularity condition is satisfied, we can say more about the form of the regulator's optimal schedule:

Observation 2. *Fix $\alpha \geq 0$, and let $\rho, \tilde{\rho} : \Theta \to \mathbb{R}^+$ be weighting functions. Suppose that (α, ρ, F) and $(\alpha, \tilde{\rho}, F)$ satisfy equity regularity. Then,*

$$q^\rho(\theta) := \arg\max_q \phi_{\alpha,\rho}(\theta)q - c(q) \quad \textit{and} \quad q^{\tilde{\rho}}(\theta) := \arg\max_q \phi_{\alpha,\tilde{\rho}}(\theta)q - c(q)$$

are solutions to (9) *with weighting functions ρ and $\tilde{\rho}$ respectively, and satisfy:*

$$q^\rho(\overline{\theta}) = q^{\tilde{\rho}}(\overline{\theta}) \quad \textit{and} \quad q^\rho(\underline{\theta}) = q^{\tilde{\rho}}(\underline{\theta}).$$

[16]Equity regularity may be more difficult to satisfy than CS regularity. For instance, if $F \sim U(0,1)$, recall that (α, F) is CS regular iff $\alpha \leq 2$. In contrast, if $\rho(\theta) = \frac{5}{4} - \frac{\theta}{2}$, we have that (α, ρ, F) is equity regular iff $\alpha \leq \frac{8}{5}$.

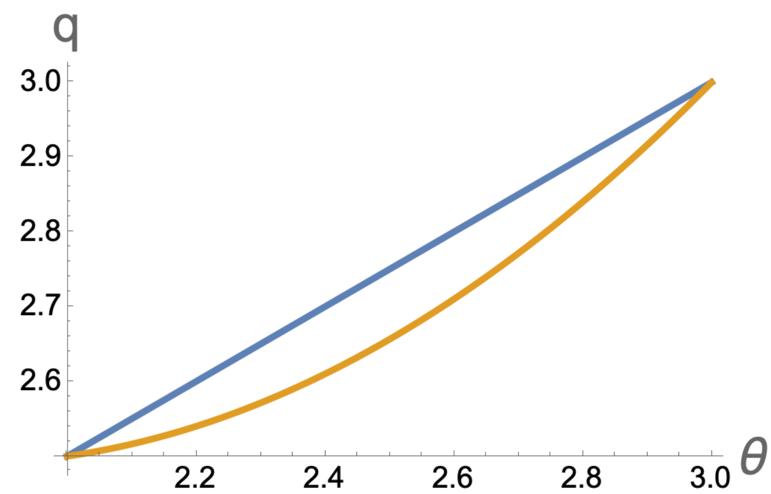

Figure 4: Comparison of the regulator's optimal quality schedules under different equity weights, with α fixed. Blue: Less equitable, $\rho(\theta) = 1$. Orange: More equitable, $\tilde{\rho}(\theta) = -\frac{1}{2}(\theta - \frac{5}{2}) + 1$. Throughout, $\alpha = 3/2, F \sim U(2,3)$, $c(q) = q^2/2$.

Observation 2 says that, under regularity, the regulator's optimal schedules for any two weighting functions (holding α fixed) agree at the top and the bottom. In particular, this implies that there is no distortion at the top (taking $\rho = 1$ and applying Observation 1). Observation 2 follows immediately from the fact that $F_\rho(\underline{\theta}) = 0, F_\rho(\overline{\theta}) = 1$, for any weighting function ρ.

Figure 4 illustrates the results in this section. The more equitable regulator's schedule ($\tilde{q}$, orange) is pointwise lower than the less equitable regulator's (q, blue), as indicated by Proposition 3. Further, since regularity is satisfied, $\tilde{q}$ and q agree at the top and bottom, as indicated by Observation 2.

5 Monopolist randomization: finite types

Let us now consider the possibility of the monopolist using a randomized selling mechanism in response to price distribution regulation. At present, I am only able to prove my results in a model with finite types. In this section, I will present and analyze such a model. Then, in Section 6, I will state the conjectured analog of my main technical result for the continuum types model we have analyzed so far, and show its implications for monopolist randomization in that model.

5.1 Model

I will begin by presenting the finite types model. The model primitives are the same, except that each consumer's type can take on values from the finite set $\Theta := \{\theta_1, ..., \theta_n\}$, with each $\theta_i \leq \theta_{i+1}$, each of which occurs with probability $1/n$.

Then, instead of the monopolist being restricted to use a deterministic mechanism following an announced price distribution regulation, G, he may use any mechanism. Formally, a mechanism is an arbitrary extensive form game between the monopolist and a consumer, the outcomes of which are distributions over the quality and the payment, as well as a public commitment to a strategy in that game for the monopolist. The solution concept is the monopolist's choice among optimal (potentially mixed) strategies for the consumer.

The mechanism is constrained to satisfy the consumers' (interim) IR condition, that it give each consumer type at least 0 payoff on expectation. Further, it must satisfy the regulatory constraint. A mechanism, $\mathcal{M}$, for the monopolist induces some ex-ante distribution over consumer payments (taking into account both the randomly drawn consumer type and any randomness in the mechanism), $H_{\mathcal{M}}$. The regulatory constraint says that $H_{\mathcal{M}}$ must match the mandated distribution.

Using a revelation principle, we can think of the monopolist without loss as using a direct randomized mechanism, (Q, T) with $Q, T \in \Delta(\mathbb{R}^+)^n$, which satisfies IC, IR and the regulatory constraint.[17] (The tuples Q and T represent each reported type's assignment – for instance, reported type θ_i receives the quality distribution Q_i and pays the transfer distribution T_i. As usual, I will also use Q_i and T_i to refer to the CDFs of these distributions.)

The following definition will be useful.

[17] I will not prove this revelation principle in this draft, but the proof is similar to the deterministic revelation principle in Appendix A.

Definition 7. *Given a distribution $H \in \Delta(\mathbb{R}^+)$, define the expectation operator as $\overline{H} := \mathrm{E}_{x\sim H}[x]$. Similarly, given a sequence of distributions, $H \in \Delta(\mathbb{R}^+)^n$, define the expectation operator as $\overline{H} := (\mathrm{E}_{x\sim H_1}[x], \dots, \mathrm{E}_{x\sim H_n}[x])$.*

Since the cost function, c, is strictly convex, any stochastic mechanism, (Q, T), where Q_i is strictly stochastic for some i, can be strictly improved by changing it to $(\overline{Q}, T)$. (This does not affect consumer incentives or the induced price distribution.) Thus, without loss, the monopolist chooses a direct mechanism (q, T), with $q \in \mathbb{R}^{+n}$ and $T \in \Delta(\mathbb{R}^+)^n$. I will call such a mechanism **semi-stochastic**. The monopolist's problem, given regulation G is, if he enters:

$$\max_{q,T} \sum_i \overline{T}_i - c(q_i) \tag{10}$$

$$\text{s.t. } \theta_i q_i - \overline{T}_i \geq \theta_i q_j - \overline{T}_j, \text{ for all } \theta_i, \theta_j \in \Theta \tag{IC}$$

$$\theta_i q_i - \overline{T}_i \geq 0 \text{ for all } \theta_i \in \Theta \tag{IR}$$

$$G(t) = \frac{1}{n} \sum_i T_i(t) \text{ for all } t \in \mathbb{R}. \tag{G Sto}$$

I will say that a deterministic mechanism, (q, t), solves (10) if $(q, t^\bullet)$ solves (10), where $t^\bullet \in \Delta(\mathbb{R}^+)^n$ denotes the sequence of distributions such that $t_i^\bullet$ puts probability 1 on t_i.

5.2 Ex-interim vs ex-post IR

I have modeled the consumers' IR constraint as being assessed ex-interim – that is, after the consumer learns their type, but before they learn the realization of randomness from the price mechanism. This can be thought of as describing a contract which the monopolist asks consumers to sign, in which they agree to pay the random price, whatever its realization.

One might instead be interested in a more restrictive form of monopolist randomization, described by an ex-post IR constraint. In this case, at the interim stage, the consumer chooses a menu item (possibly specifying some price randomization), but may walk away after the random price has realized. (Critically, the monopolist can prevent them from returning to re-sample the randomized price.) Clearly, the ex-post IR constraint imposes more restrictions on the monopolist, and thus allows the regulator at least as much flexibility in what she can enact.

Importantly, in the limit of finely spaced consumer types (and in the continuous types model) the ex-post IR constraint also allows the regulator *no more* flexibility than the ex-interim IR constraint. More specifically, Corollary 1 holds under the ex-post IR constraint, and Theorem 4 holds under the ex-post IR constraint, in the limit of finely space consumer types. The idea behind these results is that, for the regulator to be able to enact a mechanism, she must guard against "local" deviations, which are still available to the monopolist under the ex-post IR constraint, because most consumer types have strictly positive information rents. Please check back soon for a formal presentation of these results.

5.3 The deterministic case

Let us first state the problem of a monopolist who is constrained to use a deterministic mechanism. This is the analog of (1) in the discrete model (if the monopolist produces).

$$\max_{q,t} \sum_i t_i - c(q_i) \tag{11}$$

$$\text{s.t. } \theta_i q_i - t_i \geq \theta_i q_j - t_j, \text{ for all } \theta_i, \theta_j \in \Theta \tag{IC}$$

$$\theta_i q_i - t_i \geq 0 \text{ for all } \theta_i \in \Theta \tag{IR}$$

$$G(\hat{t}) = \frac{1}{n}\#\{i : t_i \leq \hat{t}\} \text{ for all } \hat{t} \in \mathbb{R}, \tag{G Reg}$$

where # denotes set cardinality.

The definition and characterization of (deterministic) enactability also extend to this discrete model.

Definition 8. *Say that a deterministic direct mechanism, (q, t) is **enactable** if there exists* *a price distribution, G, and a transfer, τ, such that (q, t) solves (11) and the monopolist produces.*

Theorem 2. *A deterministic direct mechanism, (q, t), is enactable if and only if:*

1. *q and t are non-decreasing.*

2. *The mechanism satisfies the discrete envelope formulae,*

$$t_i = \theta_i q_i - \sum_{j=1}^{i-1} q_j(\theta_{j+1} - \theta_j)$$

$$q_i = \frac{t_i}{\theta_i} + \sum_{j=1}^{i-1} \left(\frac{1}{\theta_j} - \frac{1}{\theta_{j+1}}\right) t_j,$$

for all $\theta_i \in \Theta$.

Theorem 2 says that a (discrete) deterministic direct mechansism is enactable if and only if (1) it satisfies IC and IR and (2) it satisfies 0 surplus at the bottom *and downward incentive constraints bind*. The necessity of downward incentive constraints binding is again due to a quality degradation argument: if type θ_i strictly preferred to report truthfully over reporting θ_{i-1}, the monopolist could profitably lower $q_i, ..., q_n$ by a small constant, $\varepsilon > 0$.

5.4 Example

Let me illustrate via an example how the monopolist's ability to randomize reduces the regulator's power. Suppose the model primitives are such that an unconstrained monopolist would optimally use the schedule (q^*, t^*), where $t^* = (2, 3, 4, 5)$, which induces the price distribution $(\frac{1}{4} \circ 2, \frac{1}{4} \circ 3, \frac{1}{4} \circ 4, \frac{1}{4} \circ 5)$. Suppose further that the regulator mandates the price distribution $G = (\frac{1}{2} \circ 2, \frac{1}{2} \circ 5)$. Notice that if the monopolist could not randomize, he would have to use the price schedule $t(\theta_1) = t(\theta_2) = 2$, $t(\theta_3) = t(\theta_4) = 5$, and the corresponding two-level quality schedule.

However, the monopolist can use price randomization to effectively implement his ideal mechanism, while complying with the regulation. To do this, he can offer a menu consisting of four options, given by:

$$\left\{\{q_1^*, 2\}, \{q_2^*, (\frac{2}{3} \circ 2, \frac{1}{3} \circ 5)\}, \{q_3^*, (\frac{1}{3} \circ 2, \frac{2}{3} \circ 5)\}, \{q_4^*, 5\}\right\}.$$

The idea is that the monopolist effectively implements (q^*, t^*), but asks for payment in probabilistic units of 2 and 5. It is easy to check that the resulting ex-ante distribution of payments matches G.

It is no coincidence that the monopolist was able to effectively implement a deterministic mechanism, (q^*, t^*), that induced a price distribution which was a mean-preserving contraction (MPC) of G. As we will see, this MPC condition characterizes the monopolist's options when facing a regulated price distribution.

5.5 Analysis

Now, let us analyze the model with monopolist randomization. We will build to an enactability result. First, I will introduce notation for the price distribution induced by a price schedule.

Definition 9. *Given a deterministic price schedule, $t \in \mathbb{R}^{+n}$, denote the induced price distribution by $G_t(\hat{t}) := \frac{1}{n}\#\{i : t_i \leq \hat{t}\}$, for all $\hat{t} \in \mathbb{R}$.*

Given a stochastic price schedule, $T \in \Delta(\mathbb{R}^+)^n$, denote the induced price distribution by $G_T(\hat{t}) := \frac{1}{n}\sum_i T_i(\hat{t})$, for all $\hat{t} \in \mathbb{R}$.

Next, I will say that a deterministic price schedule is compatible with a regulation if the monopolist can effectively use it while abiding by the regulation. In our example, $(2, 3, 4, 5)$ was compatible with $(\frac{1}{2} \circ 2, \frac{1}{2} \circ 5)$.

Definition 10. *Say that a deterministic price schedule, $t \in \mathbb{R}^{+n}$ is **compatible** with a* *distribution G if there exists a stochastic price schedule, $T \in \Delta(\mathbb{R}^+)^n$ such that (1) $\overline{T} = t$ and (2) $G_T = G$.*

Next, I will present the standard definition of a mean-preserving contraction, due to Rothschild and Stiglitz (1971).

Definition 11. *Say of two distributions, $H_1, H_2 \in \Delta(\mathbb{R})$ that H_1 is a **mean-preserving contraction** of H_2 if there exist random variables X_1, X_2, Z,[18] such that:*

1. *X_1 has distribution H_1 and X_2 has distribution H_2.*
2. *$\mathrm{E}[Z|X_1] = 0$.*
3. *$X_2 = X_1 + Z$.*

I will denote "H_1 is a mean-preserving contraction of H_2" by $H_1 \preceq_{MPS} H_2$.

As suggested by our example, the next result shows that there is a tight connection between compatible schedules and MPCs of the mandated price distribution.

Lemma 5. *A deterministic price schedule, $t \in \mathbb{R}^{+n}$, is compatible with a distribution G if and only if $G_t \preceq_{MPS} G$.*

Lemma 5 implies that the monopolist's value from (10) is the same as from the following, where q, t are functions from Θ to $\mathbb{R}^+$:

$$\max_{q,t} \sum_i t_i - c(q_i) \tag{12}$$

$$\text{s.t. } \theta_i q_i - t_i \geq \theta_i q_j - t_j, \text{ for all } \theta_i, \theta_j \in \Theta \tag{IC}$$

$$\theta_i q_i - t_i \geq 0 \text{ for all } \theta_i \in \Theta \tag{IR}$$

$$G_t \preceq_{MPS} G. \tag{G MPC}$$

[18] These are measurable maps from the standard probability space to $\mathbb{R}$, where by the standard probability space I mean $(\Omega, \mathcal{F}, P)$, where Ω is the unit interval, $\mathcal{F}$ is the Borel σ-algebra, and P is the Lebesgue measure.

(To see this, notice that if (q, T) solves (10), then $(q, \overline{T})$ is feasible for (12), while if (q, t) solves (12), then there is a (q, T) with identical value that is feasible for (10).)

Next, I will define the notion of enactability when the monopolist is able to randomize.

Definition 12. *Say that a semi-stochastic direct mechanism,* (q, T)*, is* ***robustly enactable*** *if there exists a distribution* G*, such that, under* G*, the mechanism* (q, T) *solves* (10)*, and the monopolist produces.*

Say that a deterministic direct mechanism, (q, t)*, is* ***robustly enactable*** *if there exists a distribution* G*, such that, under* G*, the mechanism* (q, t) *solves* (10)*, and the monopolist produces.*

The next lemma shows that, due to the nested structure of MPCs, it is without loss for the regulator to focus on robustly enacting deterministic mechanisms.

Lemma 6. *If the semi-stochastic direct mechanism* (q, T) *is robustly enactable, then so is the deterministic direct mechanism* $(q, \overline{T})$*.*

To state my main result, we will need two more pieces of notation. First, let us define the majorization relation, due to Hardy, Littlewood, and Pólya (1929).

Definition 13. *Say of two non-decreasing vectors,* $t, \tilde{t} \in \mathbb{R}^n$ *that* $\tilde{t}$ ***majorizes*** t *if for all* $k \in \{1, \ldots, n\}$*, we have:*

$$\sum_{i=1}^{k} \tilde{t}_i \geq \sum_{i=1}^{k} t_i,$$

with equality for $k = n$*.*

It is a standard result, due to Rothschild and Stiglitz (1971), that (a non-decreasing) $\tilde{t}$ majorizes (a non-decreasing) t if and only if $G_{\tilde{t}} \preceq_{MPS} G_t$. I therefore also use the notation $\tilde{t} \preceq_{MPS} t$ to denote "$\tilde{t}$ majorizes t."

Second, let us define the monopolist's cost function from using the (unique) enactable mechanism with price schedule t. Define $\Gamma : \mathbb{R}^n \to \mathbb{R}$ by:

$$\Gamma(t) := \frac{1}{n} \sum_{i=1}^{n} c \left(\frac{t_i}{\theta_i} + \sum_{j=1}^{i-1} \left(\frac{1}{\theta_j} - \frac{1}{\theta_{j+1}} \right) t_j \right).$$

(The argument of c is simply the envelope formula for q_i in terms of t.) Notice that Γ is a convex, continuously differentiable function. This convexity is the reason that Γ is a useful way to record costs (rather than, say, as a function of q).

Now, I will state two supporting results: a lemma and a technical theorem. The lemma combines what we know about enactability with our MPC characterization of compatibility in Lemma 5.

Lemma 7. *Suppose (q, t) is enactable. Then, it is robustly enactable if and only if t solves:*

$$\min_{\hat{t} \preceq_{MPS} t} \Gamma(\hat{t}).$$

Lemma 7 tells us that, to check whether an enactable mechanism (q, t) is robustly enactable, we must check whether *any* majorizing $\hat{t}$ lowers the monopolist's costs. Since the majorization set is large, this seems infeasible. Theorem 3 significantly reduces our search space. It says that we need only rule out what I call "local contractions" of t. A local contraction of a non-decreasing vector, t, slightly increases t_i and slightly decreases t_{i+1}, for some $i < n$. The proof of Theorem 3 is somewhat involved. I build to and present it in Appendix D.

Theorem 3. *Suppose $\Gamma : \mathbb{R}^n \to \mathbb{R}$ is a convex, continuously differentiable function, and suppose that $t \in \mathbb{R}^n$ is non-decreasing and satisfies the no local contractions condition,*

$$\Gamma_i(t) - \Gamma_{i+1}(t) \geq 0,$$

for all $i \leq n - 1$, where Γ_i denotes the partial derivative of Γ in the ith coordinate. Then, t

solves:

$$\min_{\hat{t} \preceq_{MPS} t} \Gamma(\hat{t}).$$

Finally, combining Lemma 7 and Theorem 3 gives my main result. Theorem 4 shows that enactability and no local contractions are sufficient for a mechanism to be robustly enactable, and that the converse holds if that mechanism is strictly monotone.

Theorem 4. *A deterministic mechanism, (q, t), is robustly enactable if:*

1. *It is enactable.*

2. *The mechanism satisfies the no local contractions condition,*

$$\Gamma_i(t) - \Gamma_{i+1}(t) \geq 0,$$

for all $i < n$.

Further, if t is strictly increasing, then (q, t) is robustly enactable only if conditions 1 and 2 are satisfied.

Combined with Lemma 6, Theorem 4 gives a sharp (in the case of strictly monotone mechanisms), easily checkable answer to the question of robust enactability. Lemma 6 shows that it is without loss for the regulator to attempt to enact a deterministic mechanism, and Theorem 4 characterizes the enactable deterministic mechanisms.

6 Monopolist randomization: continuous types

Let us now return to our baseline model with continuous types. Recall that the CDF of the type distribution, F, and the monopolist's cost function, c, are both twice differentiable.

I will state two definitions, in order to present the (conjectured) partial characterization of robustly enactable mechanisms in this setting.

Definition 14. *Given a price schedule, t, and some positive integer n, let*

$$t^{[n]} := (t(F^{-1}(1/n)), t(F^{-1}(2/n)), \ldots, t(F^{-1}(1)))$$

denote the (equal mass) n-discretization of t.

Additionally, let us define $\Gamma^{[n]}$ as in Section 5 by:

$$\Gamma^{[n]}(t) := \frac{1}{n} \sum_{i=1}^{n} c \left(\frac{t_i}{\theta_i} + \sum_{j=1}^{i-1} \left(\frac{1}{\theta_j} - \frac{1}{\theta_{j+1}} \right) t_j \right).$$

The following conjecture is the analog of Theorem 4 in the continuous setting.

Conjecture 1. *A deterministic mechanism, (q, t), is robustly enactable if:*

1. *It is enactable.*
2. *The mechanism satisfies the no local contractions condition,*

$$\lim_{n \to \infty} n^2 \left(\Gamma^{[n]}_{\lceil nx \rceil}(t) - \Gamma^{[n]}_{\lceil nx \rceil + 1}(t) \right) \geq 0,$$

for all $x \in [0, 1]$.

Further, if t is strictly increasing, then (q, t) is robustly enactable only if conditions 1 and 2 are satisfied.

Let me briefly describe some challenges for establishing this conjecture. Fix some t, $\hat{t}$, with $\hat{t} \preceq_{MPS} t$. We would like to show that we can find a sequence of (equal mass) discretizations, $t^{[n^*]}$, $\hat{t}^{[n^*]}$ (not necessarily $t^{[n]}$ and $\hat{t}^{[n]}$) that converge in distribution to t and $\hat{t}$, and maintain $\hat{t}^{[n^*]} \preceq_{MPS} t^{[n^*]}$ along the sequence. Second, we would like to show that, if we have such a sequence, the total length of contractions obtained from running the algorithm in Lemma 9 to produce $\hat{t}^{[n^*]}$ from $t^{[n^*]}$ grows with n^2. I have proved (not in this draft) this second claim.

The following corollary of Conjecture 1 restates the continuous no local contractions condition. Intuitively, when n is large, we can compute the impact of a small local contraction by taking a second derivative of Γ at t. Please check back in the coming days for a derivation of this corollary.

Corollary 3. *A deterministic mechanism, (q, t), is robustly enactable if:*

1. *It is enactable.*

2. *The mechanism satisfies the differential no local contractions condition,*

$$\left(\frac{2}{\theta} + \frac{f'(\theta)}{f(\theta)}\right) \int_{\theta}^{\overline{\theta}} c'(q(x))dF(x) \geq f(\theta)[\theta c''(q(\theta))q'(\theta) - 2c'(q(\theta))]$$

for all $\theta \in \Theta$.

Further, if t is strictly increasing, then (q, t) is robustly enactable only if conditions 1 and 2 are satisfied.

Finally, recall that the efficient quality schedule, q^{eff}, satisfies $c'(q^{\text{eff}}(\theta)) = \theta$. In this case, observing that we can write $c''(q(\theta))q'(\theta) = (c'(q(\theta)))'$, the differential no local contractions condition reduces to:

$$\left(\frac{2}{\theta} + \frac{f'(\theta)}{f(\theta)}\right) \int_{\theta}^{\overline{\theta}} x dF(x) + \theta f(\theta) \geq 0, \text{ for all } \theta \in \Theta. \tag{13}$$

Notice that, like Myerson regularity, (13) is a condition which concerns only the type distribution, and is satisfied for many commonly used distributions, such as the uniform distribution. Myerson regularity does not imply (13), nor vice versa, but, informally speaking, the two conditions seem to mostly pass and fail on the same distributions.

The regulator can robustly enact the efficient mechanism whenever (13) is satisfied; one might wonder whether it is "easier" or "harder" for her to robustly enact a more consumer surplus-favoring mechanism – say, the regulator's optimal mechanism with $\alpha > 1$ in Section

4.1. Informally, it is easier (though many caveats are required in the formal statement). All of the numerical examples of the regulator's optimal enactable schedules I presented in Section 4 were also robustly enactable.

7 Conclusion

In this paper, I studied price distribution regulation of Mussa-Rosen monopolists. Price distribution regulation is the most powerful regulation which only uses data on transacted prices. As such, any regulatory frictions present under price distribution regulation will also be found in any simpler regulation which is restricted to this data, such as a price-dependent subsidy.

I identified two important regulatory frictions. First, even when the monopolist cannot use a randomized mechanism, the regulator cannot give low-type consumers positive surplus. Notably, this had the consequence of making a more equitable regulator less hawkish on regulation, since she is not impressed by regulations which inefficiently transfer surplus from the government to (primarily) high-type consumers.

Second, when the monopolist can randomize, he can effectively deviate to any deterministic mechanism which induces a price distribution which is a mean-preserving contraction of the regulated price distribution. I gave an easy-to-check condition to determine whether a mechanism is "robustly enactable," meaning that the regulator can induce the monopolist to use it, even with the possibility of such randomization.

For many "reasonable" triples of a regulator objective, a consumer type distribution, and a cost of quality function, the regulator's best enactable mechanism is also robustly enactable. In this case, the regulator's optimal regulation is not affected by the possibility of monopolist randomization. It remains an interesting open question what the regulator should mandate when her best enactable mechanism is not robustly enactable.

References

Akbarpour, M., P. Dworczak, and S. D. Kominers (2023). Redistributive Allocation Mechanisms. *Journal of Political Economy (forthcoming https//doi.erg/10.1086/728111)*.

Armstrong, M., S. Cowan, and J. Vickers (1995). Nonlinear Pricing and Price Cap Regulation. *Journal of Public Economics 58*(1), 33–55.

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

Besanko, D., S. Donnenfeld, and L. J. White (1987). Monopoly and Quality Distortion: Effects and Remedies. *The Quarterly Journal of Economics 102*(4), 743–767.

Börgers, T. (2015). *An Introduction to the Theory of Mechanism Design*. Oxford University Press, USA.

Bulow, J. and J. Roberts (1989). The Simple Economics of Optimal Auctions. *Journal of Political Economy 97*(5), 1060–1090.

Frankel, A. (2014). Aligned Delegation. *American Economic Review 104*(1), 66–83.

Frankena, M. (1975). Alternative Models of Rent Control. *Urban Studies 12*(3), 303–308.

Goldman, M. B., H. E. Leland, and D. S. Sibley (1984). Optimal Nonuniform Prices. *The Review of Economic Studies 51*(2), 305–319.

Hardy, G. H., J. E. Littlewood, and G. Pólya (1929). Some Simple Inequalities Satisfied by Convex Functions. *Messenger of Mathematics 58*, 145–152.

Hart, O., A. Shleifer, and R. W. Vishny (1997). The Proper Scope of Government: Theory and an Application to Prisons. *The Quarterly Journal of Economics 112*(4), 1127–1161.

Jackson, M. O. and H. F. Sonnenschein (2007). Overcoming Incentive Constraints by Linking Decisions. *Econometrica 75*(1), 241–257.

Kihlstrom, R. E. and D. Levhari (1977). Quality, Regulation and Efficiency. *Kyklos 30*(2), 214–234.

Kim, J.-C. and C.-Y. Jung (1995). Regulating a Multi-Product Monopolist. *Journal of Regulatory Economics 8*, 299–307.

Kleiner, A., B. Moldovanu, and P. Strack (2021). Extreme Points and Majorization: Economic Applications. *Econometrica 89*(4), 1557–1593.

Krishna, K. (1990). Protection and the Product Line: Monopoly and Product Quality. *International Economic Review*, 87–102.

Lee, S.-H. (1997). A Note on Regulating a Multiproduct Monopolist. *Journal of Regulatory Economics 12*, 311–317.

Leshno, M., H. Levy, and Y. Spector (1997). A Comment on Rothschild and Stiglitz's "Increasing Risk: I. A Definition". *Journal of Economic Theory 77*(1), 223–228.

Lin, X. and C. Liu (2022). Credible Persuasion. *arXiv preprint arXiv:2205.03495*.

Loertscher, S. and E. V. Muir (2022). Monopoly Pricing, Optimal Randomization, and Resale. *Journal of Political Economy 130*(3), 566–635.

Loertscher, S. and E. V. Muir (2023). Wage Dispersion, Involuntary Unemployment and Minimum Wages in the Presence of Market Power. *Working paper. Available at ellenmuir.net*.

Luo, Y., I. Perrigne, and Q. Vuong (2018). Structural Analysis of Nonlinear Pricing. *Journal of Political Economy 126*(6), 2523–2568.

Maskin, E. and J. Riley (1984). Monopoly with Incomplete Information. *The RAND Journal of Economics 15*(2), 171–196.

Mulligan, C. B. and K. K. Tsui (2016). The Upside-Down Economics of Regulated and Otherwise Rigid Prices. Technical report, National Bureau of Economic Research.

Mussa, M. and S. Rosen (1978). Monopoly and Product Quality. *Journal of Economic Theory 18*(2), 301–317.

Myerson, R. B. (1981). Optimal Auction Design. *Mathematics of Operations Research 6*(1), 58–73.

Nikzad, A. (2022). Constrained Majorization: Applications in Mechanism Design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 330–331.

Riley, J. and R. Zeckhauser (1983). Optimal Selling Strategies: When to Haggle, When to Hold Firm. *The Quarterly Journal of Economics 98*(2), 267–289.

Roberts, K. W. (1979). Welfare Considerations of Nonlinear Pricing. *The Economic Journal 89*(353), 66–83.

Rothschild, M. and J. E. Stiglitz (1971). Increasing Risk: I. A Definition. *Journal of Economic Theory (September 1970)*, 225–243.

Sappington, D. E. and D. S. Sibley (1988). Regulating without Cost Information: The Incremental Surplus Subsidy Scheme. *International Economic Review*, 297–306.

Spence, M. (1977). Nonlinear Prices and Welfare. *Journal of Public Economics 8*(1), 1–18.

Strausz, R. (2006). Deterministic versus Stochastic Mechanisms in Principal–Agent Models. *Journal of Economic Theory 128*(1), 306–314.

Tebaldi, P., A. Torgovitsky, and H. Yang (2023). Nonparametric Estimates of Demand in the California Health Insurance Exchange. *Econometrica 91*(1), 107–146.

Weitzman, M. L. (1974). Prices vs. Quantities. *The Review of Economic Studies 41*(4), 477–491.

Wilson, R. B. (1993). *Nonlinear Pricing*. Oxford University Press, USA.

Wong, A. C. L. (2012). Menu Implementability and Practical Pricing Schemes. *Working paper. Available at SSRN 1915752*.

A Deterministic revelation principle

This section is modeled after the revelation principle in (Börgers, 2015, Chapter 2).

Definition 15. *A **deterministic mechanism** is an extensive form game between the monopolist and the consumer, in which the outcome at each terminal history is a quality,* $q \in \mathbb{R}^+$ *and a transfer,* $t \in \mathbb{R}^+$.

We will be interested in pure-strategy Bayes-Nash equilibria (BNE) of deterministic mechanisms. Such an equilibrium, σ, of some mechanism, induces a quality outcome function and a transfer outcome function, $q_\sigma : \Theta \to \mathbb{R}^+, t_\sigma : \Theta \to \mathbb{R}^+$, which map a type to the outcomes that type obtains under σ. (In the notation, I have suppressed the dependence of these outcomes on the mechanism.) Say that a pure-strategy BNE, σ, is measurable if t_σ is (Lebesgue) measurable.

Definition 16. *Say that a measurable pure-strategy BNE,* σ*, of a deterministic mechanism **generates** a price distribution,* G*, if:*

$$G(\hat{t}) = \mu_F(\theta : t_\sigma(\theta) \leq \hat{t}) \text{ for all } \hat{t} \in \mathbb{R}.$$

Definition 17. *Say that a deterministic mechanism is **direct** if it consists of two functions,* $q : \Theta \to \mathbb{R}^+$ *and* $t : \Theta \to \mathbb{R}^+$*, which map reported types to outcomes.*

A pure strategy in a direct deterministic mechanism can be described by a map, $\sigma : \Theta \to \Theta$, which translates the consumer's type into their reported type.

Theorem 5 (Deterministic revelation principle)**.** *Given any deterministic mechanism,* Γ*, and any measurable pure-strategy BNE,* σ*, of* Γ*, which generates price distribution* G*, there exists a direct deterministic mechanism,* Γ' *with a measurable pure-strategy BNE,* σ'*, satisfying the following:*

1. *(Incentive compatibility)* $\sigma'(\theta) = \theta$ *for all* $\theta \in \Theta$

2. *(Equivalence)* σ' *induces the same outcome functions as* σ

3. *(Legality)* σ' *generates* G

Proof. We will show that the pair $\{\Gamma' = (q_\sigma, t_\sigma), \sigma'(\theta) = \theta\}$ satisfies the desired conditions. We must first show that this σ' is a measurable pure-strategy BNE of Γ'. Immediately, σ' inherits measurability from σ. Now, suppose that σ' is not a BNE of Γ'. Then, some type θ can strictly improve by reporting θ'. Then, it must be that

$$\theta q_\sigma(\theta') - t_\sigma(\theta') > \theta q_\sigma(\theta) - t_\sigma(\theta),$$

and so θ would strictly benefit by deviating by mimicing θ' at all histories under (Γ, σ), and therefore obtaining $(q_\sigma(\theta'), t_\sigma(\theta'))$, contradicting the fact that σ is a BNE. Then, the three properties follow immediately. In particular, legality follows from equivalence and the fact that σ generates G. $\square$

B Monopolist private information

So far, I have assumed the regulator has perfect information about the monopolist's operating environment. She has known his costs, c, and the consumer demand, F. The only "wedge" preventing a first-best policy has been the unobservability of product quality. In this section, I relax this assumption, by allowing the monopolist's costs to be unknown to the regulator, in the spirit of Baron and Myerson (1982). I will prove a preliminary lemma, which shows that the regulator's problem separates into a weighted surplus maximization problem resembling the one studied in Section 4.1, and a screening problem resembling the one in Baron-Myerson. This lemma suggests that the regulator's problem with imperfect information is tractable, though I leave a complete analysis for future work.

Formally, let us enrich the model in two ways. First, the monopolist's cost will be given by $\omega c(q)$, (maintaining our previous assumptions about c). This cost type, ω, will be drawn from a distribution H, which is supported on $\Omega := [\underline{\omega}, \overline{\omega}]$, with $\underline{\omega} \geq 0$; it will be privately known to the monopolist. Second, we will allow transfers between the regulator and the monopolist. Specifically (with a revelation principle in mind), the regulator will choose a direct mechanism, $(G(\cdot), \tau(\cdot))$, which consists of two maps, $G : \Omega \to \Delta(\mathbb{R}^+)$, and $\tau : \Omega \to \mathbb{R}$, which map a reported cost type to, respectively, a mandated price distribution, and a (possibly negative) tax. The regulator will maximize a weighted sum of the consumer's surplus, the monopolist's surplus, and the government's surplus, with weights $\alpha_c, \alpha_m, \alpha_g$, respectively. I will assume that the monopolist uses a deterministic screening mechanism; as before, this is justified for many consumer type distributions by Section 5.

For any function q on Θ, let us define the quantities,

$$\begin{aligned} CS(q) &:= \int [\theta q(\theta) - t(\theta)] dF(\theta) \qquad PS(q, \omega) := \int [t(\theta) - \omega c(q(\theta))] dF(\theta) \\ C(q) &:= \int c(q(\theta)) dF(\theta), \end{aligned}$$

where $t(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(x) dF(x)$. Using Theorem 1, the regulator can optimize over quality schedules, q, rather than distributions, G. Let $\mathfrak{q}$ denote a map from Ω to quality schedules, and denote the quality schedule associated with ω by $\mathfrak{q}_\omega$. Then, the regulator solves:

$$\max_{\mathfrak{q}, \tau} \int \alpha_c CS(\mathfrak{q}_\omega) + \alpha_m [PS(\mathfrak{q}_\omega, \omega) - \tau(\omega)] + \alpha_g \tau(\omega) dH(\omega) \tag{14}$$

$$\begin{aligned} \text{s.t. } & PS(\mathfrak{q}_\omega, \omega) - \tau(\omega) \geq PS(\mathfrak{q}_{\hat{\omega}}, \omega) - \tau(\hat{\omega}) \text{ for all } \omega, \hat{\omega} \in \Omega && \text{(monopolist IC)} \\ & PS(\mathfrak{q}_\omega, \omega) - \tau(\omega) \geq o \text{ for all } \omega \in \Omega && \text{(monopolist IR)} \\ & \mathfrak{q}_\omega \text{ is non-decreasing for all } \omega \in \Omega && \text{(implementability)} \end{aligned}$$

The analysis of (14) is greatly simplified by the following result.

Lemma 8. *Define CS, PS and C as above, and suppose* (14) *has a solution. Then, there* *exists a function, $\overline{c} : \Omega \to \mathbb{R}$, such that for any $\mathfrak{q}$ where the quality schedule $\mathfrak{q}_\omega$ solves:*

$$\max_{q\ incr.} \alpha_c CS(q) + \alpha_g PS(q, \omega) \tag{15}$$
$$s.t.\ C(q) = \overline{c}(\omega),$$

for each $\omega \in \Omega$, there exists a τ such that $(\mathfrak{q}, \tau)$ is a solution to (14).

Lemma 8 says that we can find a solution to (14) with a very specific structure, in which, for each ω, the quality schedule $\mathfrak{q}_\omega$ solves a cost-constrained surplus optimization problem. The surplus being maximized is a weighted sum of consumer surplus and producer surplus, where the weight on producer surplus is taken from the weight on the government in the original program.

The proof contains two ideas. First, each monopolist cost type only cares about the "standardized" cost, $C(q)$, and the total income, $\int t(\theta)dF(\theta) - \tau$, that each bundle he can choose from entails. Thus, the regulator may as well optimize surplus once she has chosen some target cost and income for each type. Second, since utility is perfectly transferable between the government and the monopolist, once the regulator has chosen some target income for a type, decreases to that type's revenue come directly out of the government's pocket. Thus, trade-offs from the choice of q only concern the consumers and the government.

Proof. Consider some solution, $(\mathfrak{q}, \tau)$ to (14), and let $\overline{c}(\omega) = C(\mathfrak{q}_\omega)$. Consider any modified schedule, where for each ω, the modified quality schedule is given by a $\tilde{\mathfrak{q}}_\omega$ satisfying:

$$\tilde{\mathfrak{q}}_\omega \in \arg\max_{q \text{ incr.}} \alpha_c CS(q) + \alpha_g PS(q, \omega) \tag{16}$$
$$\text{s.t. } C(q) = \overline{c}(\omega),$$

and the modified tax is given by:

$$\tilde{\tau}(\omega) = \tau(\omega) + PS(\tilde{\mathfrak{q}}_\omega, \omega) - PS(\mathfrak{q}_\omega, \omega). \tag{17}$$

I claim that, for any $\omega, \hat{\omega} \in \Omega$, we have:

$$PS(\mathfrak{q}_{\hat{\omega}}, \omega) - \tau(\hat{\omega}) = PS(\tilde{\mathfrak{q}}_{\hat{\omega}}, \omega) - \tilde{\tau}(\hat{\omega}). \tag{18}$$

To see this, define the revenue function, $R(q) := \int t(\theta) dF(\theta)$, so that $PS(q, \omega) = R(q) - \omega C(q)$. Rearranging equation (17) applied to $\hat{\omega}$, we have:

$$R(\mathfrak{q}_{\hat{\omega}}) - \hat{\omega} C(\mathfrak{q}_{\hat{\omega}}) - \tau(\hat{\omega}) = R(\tilde{\mathfrak{q}}_{\hat{\omega}}) - \hat{\omega} C(\tilde{\mathfrak{q}}_{\hat{\omega}}) - \tilde{\tau}(\hat{\omega}). \tag{19}$$

Since $C(\tilde{\mathfrak{q}}_{\hat{\omega}}) = C(\mathfrak{q}_{\hat{\omega}})$, let us add $(\omega - \hat{\omega}) C(\mathfrak{q}_{\hat{\omega}})$ to the left-hand side and $(\omega - \hat{\omega}) C(\tilde{\mathfrak{q}}_{\hat{\omega}})$ to the right-hand side of (19); this establishes (18). Now, (18) implies that $(\tilde{\mathfrak{q}}, \tilde{\tau})$ inherits monopolist IC from $(\mathfrak{q}, \tau)$. Additionally, since, by construction $(\tilde{\mathfrak{q}}, \tilde{\tau})$ holds the each monopolist type's surplus (producer surplus less tax) constant, $(\tilde{\mathfrak{q}}, \tilde{\tau})$ inherits monopolist IR. Finally, $\tilde{\mathfrak{q}}$ is non-decreasing by construction. Thus, $(\tilde{\mathfrak{q}}, \tilde{\tau})$ is feasible for (14).

Further, $(\tilde{\mathfrak{q}}, \tilde{\tau})$ weakly increases the total weighted surplus at each ω. Since it holds the monopolist's surplus constant we need only consider changes to consumer and government surplus. The weighted sum of consumer and government surplus under $(\tilde{\mathfrak{q}}, \tilde{\tau})$ at ω is given by:

$$\alpha_c CS(\tilde{\mathfrak{q}}_\omega) + \alpha_g \tilde{\tau}(\omega) = \overbrace{\alpha_c CS(\tilde{\mathfrak{q}}_\omega) + \alpha_g PS(\tilde{\mathfrak{q}}_\omega, \omega)} + \alpha_g (\tau(\omega) - PS(\mathfrak{q}_\omega, \omega)) \qquad \text{(by (17))},$$

while under $(\mathfrak{q}, \tau)$ it is given by:

$$\alpha_c CS(\mathfrak{q}_\omega) + \alpha_g \tau(\omega) = \overbrace{\alpha_c CS(\mathfrak{q}_\omega) + \alpha_g PS(\mathfrak{q}_\omega, \omega)} + \alpha_g (\tau(\omega) - PS(\mathfrak{q}_\omega, \omega)).$$

Since the two expressions differ only by the bracketed terms, by (16), $(\tilde{\mathfrak{q}}, \tilde{\tau})$ weakly increases total surplus at ω. Since this holds for all $\omega \in \Omega$, the modified schedule $(\tilde{\mathfrak{q}}, \tilde{\tau})$ is a solution to (14). $\square$

Lemma 8 allows the regulator to optimize over a single-dimensional cost constraint function, $\overline{c}$, rather than the high-dimensional map $\mathfrak{q}$. If $\alpha_g \geq \alpha_m$, I conjecture that the solution to the remaining problem of choosing $\overline{c}$ optimally has the "standard" screening structure: cost throttling, with no distortion at the bottom. More precisely, let $\overline{c}^{\text{fb}}(\omega)$ be the cost associated with the schedule:

$$\underset{q \text{ incr.}}{\arg\max}\ \alpha_c CS(q) + \alpha_g PS(q, \omega)$$

Then, I conjecture that there exists an optimal $\overline{c}$ such that $\overline{c}(\omega) \leq \overline{c}^{\text{fb}}(\omega)$, for all $\omega \in \Omega$, and $\overline{c}(\underline{\omega}) = c^{\text{fb}}(\underline{\omega})$.

C Mean price regulation

Suppose in this section that the regulator can only control the mean transacted price. I will show that the regulator optimally sets the mean price *above* the mean price charged by the unregulated monopolist. For simplicity, maintain the following assumptions:

Assumption 1 (Myerson regularity). *$\phi(\theta)$ is a non-decreasing function of θ.*

Assumption 2 (Full coverage). *$\phi(\underline{\theta}) \geq c'(0)$.*

Assumption 3 (Strict profitability absent regulation).

$$\int [\phi(\theta) q^{mon}(\theta) - c(q^{mon}(\theta))] dF(\theta) > o,$$

where $q^{mon}(\theta) := c'^{-1}(\phi(\theta))$.

The full-coverage assumption is likely to be satisfied in practical applications if the regulator's knowledge of model primitives comes only from observation of the monopolized market. In this case, she knows little about the unserved population, and may therefore find it imprudent to regulate the monopolist's conduct towards this population. See Luo, Perrigne, and Vuong (2018, p. 2550, 2553) for further discussion of this issue.

Under such a regulation, a monopolist facing a regulated mean price of m solves the following problem, if he operates:

$$\max_{q \text{ incr.}} \int [t(\theta) - c(q(\theta))] dF(\theta) \tag{20}$$
$$\text{s.t.} \quad \int t(\theta) dF(\theta) = m,$$

where $t(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(x)dx$ is the transfer function induced by q. As before, the regulatory constraint specifies the mean transacted price over the whole population, including non-serviced consumers as 0s; this can be interpreted as a mean price constraint, coupled with a quantity target.

The constraint in (20) can be incorporated with Lagrange multiplier $(1 - \lambda)$, so that the monopolist solves:

$$\max_{q \text{ incr.}} \int [\lambda t(\theta) - c(q(\theta))] dF(\theta), \tag{21}$$

for the appropriate λ; finally, using Myerson's Lemma, (21) reduces to the virtual surplus maximization problem,

$$\max_{q \text{ incr.}} \int [\lambda \phi(\theta) q(\theta) - c(q(\theta))] dF(\theta). \tag{22}$$

If the monopolist chooses to operate, he solves the relaxation of (22) in which q is not constrained to be non-decreasing, by choosing

$$q_\lambda(\theta) := c'^{-1}(\lambda \phi(\theta)).$$

Happily, by Myerson regularity, q_λ is in fact non-decreasing, and thus solves (22).

Now, by choosing m, the regulator can induce any λ for which q_λ is not identically 0. This is because the map $\lambda \mapsto \int q_\lambda(\theta)\phi(\theta)dF(\theta)$, which maps a multiplier, λ, to the associated mean price, is strictly monotone for such λ (and 0 for other λ). Since the regulator can also induce the $q_\lambda = 0$ map (for instance by setting $m = 0$), the regulator's problem is equivalent to a choice of λ, so that she can be thought of as solving:

$$\max_{\lambda} \int [\theta q_\lambda(\theta) - c(q_\lambda(\theta))] dF(\theta). \tag{Mean reg}$$

We will show the following result.

Theorem 6. *Any maximizer, λ^*, of* (Mean reg) *satisfies* $\lambda^* > 1$.

Proof. Let $\tilde{q}(\lambda, \theta) := c'^{-1}(\lambda\phi(\theta))$, and $obj(\lambda) := \int_{\underline{\theta}}^{\overline{\theta}} \theta\tilde{q}(\lambda, \theta) - c(\tilde{q}(\lambda, \theta))dF(\theta)$, for $\lambda \geq \frac{c'(0)}{\phi(\underline{\theta})}$ (so that $\tilde{q}$ is defined on the integral's domain). For such λ, we have:

$$obj'(\lambda) = \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial\tilde{q}}{\partial\lambda}(\lambda, \theta) \left[(1-\lambda)\theta + \frac{1-F(\theta)}{f(\theta)}\right] dF(\theta).$$

Assumption 2, the full coverage assumption, guarantees that obj is defined at $\lambda = 1$, and that $q_\lambda(\theta) = \tilde{q}(\lambda, \theta)$ for $\lambda \geq 1$, so that obj' agrees with the right derivative of the regulator's objective at $\lambda = 1$. Let us therefore evaluate $obj'(1)$. The bracketed term reduces to $\frac{1-F(\theta)}{f(\theta)}$, which is strictly positive, except at $\overline{\theta}$. By strict convexity of c, the partial derivative term, $\frac{\partial\tilde{q}}{\partial\lambda}(\lambda, \theta)$, is strictly positive, except where $\phi(\theta) = 0$, which cannot happen on a set of full measure since the unregulated monopolist is strictly profitable. Thus, $obj'(1) > 0$. The regulator can therefore improve on a choice of $\lambda = 1$ by slightly increasing λ. A sufficiently small increase will not affect the monopolist's decision to operate, since his profits are continuous in λ, and he found operation strictly profitable absent regulation. □

Because $q_{\lambda=1} = q^{\text{mon}}$ and the induced mean price is strictly monotone in λ, Theorem 6 implies that the regulator optimally sets m strictly higher than the unregulated monopolist's mean price.

Theorem 6 is intuitive when viewed as a consequence of the Mussa-Rosen downward quality distortion. Since q_λ is increasing in λ, the regulator optimally increases λ above 1 to move q closer (on aggregate) to its efficient level. On the other hand, it is surprising when compared with our usual intuition that monopolists charge excessively high prices, which regulation should decrease. This discrepancy is explained by noticing that mean price regulation, like general price distribution regulation, is able to control quantity in a de-facto manner, by counting non-service as 0 price. This quantity targeting accomplishes the market expansion we usually associate with downward price controls, leaving the mean price regulation free to correct the Mussa-Rosen distortion by demanding price increases.

D Proof of Theorem 3 (no local contractions)

Recall that n is the number of consumer types in the finite model.

Definition 18. *A **contraction** at i, denoted by $\mathcal{C}_i$, is the n-vector*

$$\mathcal{C}_i := (0, \ldots, 1, -1, \ldots, 0),$$

where the positive 1 appears in the i^{th} coordinate.

Definition 19. *A **long contraction** from i to j, denoted by $\mathcal{C}_{i,j}$, is the n-vector*

$$\mathcal{C}_{i,j} := (0, \ldots, 1, \ldots, -1, \ldots, 0),$$

where the positive 1 appears in the i^{th} coordinate, and the negative 1 appears in the j^{th} coordinate, with $j > i$.

The next lemma shows that we can write a majorizing vector as the sum of the majorized vector and a sequence of contractions. This is closely related to a result of Rothschild and Stiglitz (1971) and Leshno, Levy, and Spector (1997). The differences are that their discrete distributions are equally *spaced*, while mine are equally *weighted*, and that their decomposition is into long contractions rather than contractions.

Lemma 9. *Suppose s and t are two n-vectors with $s \preceq_{MPS} t$. Then, there exists a sequence* *of non-negative numbers, $\delta_i \geq 0$, $i \in \{1, \ldots, n\}$, such that:*

$$s = t + \sum_{i=1}^{n} \delta_i \mathcal{C}_i.$$

Proof. I will describe an algorithm to find these δ_i's, by constructing a sequence of contractions from t to s.

<u>Algorithm</u>

Initialize $\hat{t} = t$, and $\hat{\delta}_i = 0$, for all i. Then, repeat the following steps:

1. For each $\hat{t}_i$, color it red (for deficit) if $\hat{t}_i < s_i$, and green (for surplus) if $\hat{t}_i > s_i$. (It receives no color if $\hat{t}_i = s_i$.)

2. Find the indices, r and g, of the lowest-index red and green points in $\hat{t}$. I show below that we must have $r < g$.

3. Let $\delta := \min\{s_r - \hat{t}_r, \hat{t}_g - s_g\}$, and form the long contraction, $\mathcal{C}_{r,g}$.

4. Notice that $\mathcal{C}_{r,g}$ telescopes as: $\mathcal{C}_{r,g} = \sum_{i=r}^{g-1} \mathcal{C}_i$. For $i \in \{r, \ldots, g-1\}$, add δ to $\hat{\delta}_i$.

5. Update $\hat{t}$ to equal $\hat{t} + \delta \mathcal{C}_{r,g}$, and repeat steps 1-5, until $\hat{t} = s$. When the algorithm terminates, set $\delta_i = \hat{\delta}_i$ for all i.

By construction, if the algorithm terminates, we are left with the desired sequence of $\delta_i \geq 0$ satisfying $s = t + \sum_{i=1}^{n} \delta_i \mathcal{C}_i$. Further, provided that it is always possible to form the long contraction $\mathcal{C}_{r,g}$ in Step 3, the algorithm must conclude after at most n repetitions. This is because each repetition causes at least one point in $\hat{t}$ to lose color (become equal to its counterpart in s), and uncolored points never gain color.

It remains only to show that it is always possible to form the long contraction, $\mathcal{C}_{r,g}$, in Step 3. This could fail in two ways. First, in Step 2, it could be that there is a red but not a green point, or vice versa (there must be some colored point, since otherwise the algorithm would have ended). But this is not possible, because it would imply $\sum \hat{t}_i \neq \sum s_i$, whereas the algorithm maintains $\sum \hat{t}_i = \sum s_i$, (because $\hat{t}$ only changes via long contractions).

Second, it could be that, in Step 2, $r > g$; but this is also not possible. When some r and g are selected, it must be that all initially red points with index $< r$ have already fully contracted, having been paired with initially green points with index $\leq g$. The total distance that needs to be contracted (over the course of the entire algorithm, via long contractions) by initially green points up to g therefore satisfies:

$$\sum_{j \leq g : t_j > s_j} (t_j - s_j) \geq \sum_{j \leq r-1 : t_j < s_j} (s_j - t_j) + \hat{t}_g - s_g, \tag{23}$$

since the first term on the RHS measures the distance already contracted by red points these green points were paired with, and the second term measures the distance point $\hat{t}_g$ must still contract. If $r > g$, we have:

$$\begin{aligned}
\sum_{j \leq r-1} (s_j - t_j) &= \sum_{j \leq r-1 : t_j < s_j} (s_j - t_j) + \sum_{j \leq r-1 : t_j > s_j} (s_j - t_j) \\
&\leq \sum_{j \leq r-1 : t_j < s_j} (s_j - t_j) + \sum_{j \leq g : t_j > s_j} (s_j - t_j) && \text{(since } g < r\text{)} \\
&\leq -(\hat{t}_g - s_g) && \text{(by (23))} \\
&< 0,
\end{aligned}$$

contradicting the fact that $s \preceq_{MPS} t$. $\square$

The next lemma proves the main result when the no local contractions condition is strictly satisfied. The idea is that, given some "large" contraction, $\hat{t} \preceq_{MPS} t$, we will form an ε-interpolation, $t_\varepsilon := \varepsilon \hat{t} + (1 - \varepsilon) t$ near t. Then, using Lemma 9 we will decompose t_ε into t plus a sum of "small, local contractions." Then, taking a first-order Taylor expansion, the no local contractions condition will allow us to show that $\Gamma(t_\varepsilon) > \Gamma(t)$; convexity of Γ then implies $\Gamma(\hat{t}) > \Gamma(t)$.

Lemma 10. *Suppose $\Gamma : \mathbb{R}^n \to \mathbb{R}$ is a convex, continuously differentiable function, and suppose that $t \in \mathbb{R}^n$ is non-decreasing and satisfies the strict no local contractions condition,*

$$\Gamma_i(t) - \Gamma_{i+1}(t) > 0,$$

for all $i < n$, where Γ_i denotes the partial derivative of Γ in the ith coordinate. Then, t solves:

$$\min_{\hat{t} \preceq_{MPS} t} \Gamma(\hat{t}).$$

Proof. Consider any $\hat{t} \preceq_{\text{MPS}} t$, with $\hat{t} \neq t$. Using Lemma 9, write

$$\hat{t} = t + \sum_k \delta_k \mathcal{C}_k.$$

Then, let $b := \min_i(\Gamma_i(t) - \Gamma_{i+1}(t))$, and $d := \sum_k \delta_k$, and notice that $b > 0$ by assumption, and $d > 0$, since $\hat{t} \neq t$. Take $\varepsilon > 0$ and define $t_\varepsilon := \varepsilon\hat{t} + (1-\varepsilon)t$. From Taylor's theorem, we have:

$$\Gamma(t_\varepsilon) = \Gamma(t) + \varepsilon \sum_k \delta_k(\Gamma_k(t) - \Gamma_{k+1}(t)) + o(\varepsilon) \geq \Gamma(t) + bd\varepsilon + o(\varepsilon);$$

we will specifically choose ε small enough that $\Gamma(t_\varepsilon) > \Gamma(t)$. Then, by convexity of Γ, we must also have $\Gamma(\hat{t}) > \Gamma(t)$. □

Finally, we use a continuity argument to extend Lemma 10 to the case where the no local contractions condition is weakly satisfied.

Theorem 3. *Suppose $\Gamma : \mathbb{R}^n \to \mathbb{R}$ is a convex, continuously differentiable function, and suppose that $t \in \mathbb{R}^n$ is non-decreasing and satisfies the no local contractions condition,*

$$\Gamma_i(t) - \Gamma_{i+1}(t) \geq 0,$$

for all $i < n$. Then, t solves:

$$\min_{\hat{t} \preceq_{MPS} t} \Gamma(\hat{t}).$$

Proof. For any $\varepsilon > 0$, define the function $\Gamma^\varepsilon(x) := \varepsilon \sum_{i=1}^n (n-i)x_i + \Gamma(x)$. Notice that each Γ^ε is convex and continuously differentiable, and satisfies $\Gamma_i^\varepsilon(t) - \Gamma_{i+1}^\varepsilon(t) > 0$, for all i, because $\Gamma_i(t) - \Gamma_{i+1}(t) \geq 0$ for all i, and $\Gamma_i^\varepsilon(t) - \Gamma_{i+1}^\varepsilon(t) = \Gamma_i(t) - \Gamma_{i+1}(t) + \varepsilon$.

Then, consider any $\hat{t}$ such that $\hat{t} \preceq_{MPS} t$, and consider a sequence $\varepsilon_k \to 0$. By Lemma 10, we have $\Gamma^{\varepsilon_k}(t) \geq \Gamma^{\varepsilon_k}(\hat{t})$ for all k, and so, since $\Gamma^{\varepsilon_k}(t) \to \Gamma(t)$ and $\Gamma^{\varepsilon_k}(\hat{t}) \to \Gamma(\hat{t})$, we have $\Gamma(t) \geq \Gamma(\hat{t})$. □

E Omitted proofs

E.1 Section 3 proofs

Proof of Lemma 2.

Proof. Since t is non-decreasing, (G reg) implies:

$$G(t(\theta)) = F(\sup\{\hat{\theta} : t(\hat{\theta}) \leq t(\theta)\}), \tag{24}$$

for all $\theta \in \Theta$. Consider any $\theta \in \Theta$. We will split the analysis into three cases.

Case 1: There is a unique $\hat{t}$ such that $G(\hat{t}) = F(\theta)$.

In this case, we will show that $t(\theta) = \hat{t}$.

First, observe that (24) implies that $G(t(\theta)) \geq F(\theta)$. Since $\hat{t}$ is unique, we must have $G(\tilde{t}) < F(\theta)$ for any $\tilde{t} < \hat{t}$. Thus, it cannot be that $t(\theta) < \hat{t}$.

Next, since $\hat{t}$ is unique, we have $G(\tilde{t}) > G(\hat{t})$ for any $\tilde{t} > \hat{t}$. For any such $\tilde{t}$, we have:

$$G(\tilde{t}) = F(\sup\{\hat{\theta} : t(\hat{\theta}) \leq \tilde{t}\}) > F(\theta),$$

which implies, since F is strictly increasing:

$$\sup\{\hat{\theta} : t(\hat{\theta}) \leq \tilde{t}\} > \theta,$$

and so, since t is non-decreasing:

$$t(\theta) \leq \tilde{t}. \tag{25}$$

Since (25) holds for all $\tilde{t} > \hat{t}$, we must have $t(\theta) \leq \hat{t}$, completing the argument.

<u>Case 2</u>: There is no $\hat{t}$ such that $G(\hat{t}) = F(\theta)$.

Let $t^+ := \min\{t : G(t) \geq F(\theta)\}$. (Notice that t^+ is well-defined because G is right-continuous.) We will show that $t(\theta) = t^+$ by essentially the same argument as above.

First, suppose that $t(\theta) < t^+$. Then, by construction and uniqueness of t^+, we must have $G(t(\theta)) < F(\theta)$. But (24) implies that $G(t(\theta)) \geq F(\theta)$, so this is not possible.

Next, by the Case 2 condition, we must have $G(t^+) > F(\theta)$. We have:

$$G(t^+) = F(\sup\{\hat{\theta} : t(\hat{\theta}) \leq t^+\}) > F(\theta),$$

which implies as above:

$$\sup\{\hat{\theta} : t(\hat{\theta}) \leq t^+\} > \theta,$$

and so $t(\theta) \leq t^+$.

<u>Case 3</u>: There are multiple values $\hat{t}$ satisfying $G(\hat{t}) = F(\theta)$.

This case occurs for at most countably many values of θ by a standard argument (the $\hat{t}$ interval associated with each θ contains a rational number; there is no surjective mapping from a subset of the rationals to an uncountable set.)

Notice that in Cases 1 and 2, we have $t(\theta) = G^{-1}(F(\theta))$, proving the result. $\square$

Proof of Theorem 1.

Proof. If: The price schedule, t, induces some price distribution, G_t. If the regulator mandates G_t, then, since t satisfies (G_t reg) (and because it must be non-decreasing by implementability), by Lemma 2, we must have $t = G_t^{-1}(F(\theta))$ except possibly on a measure 0 set. Since the same holds for any $\tilde{t}$ which satisfies (G_t reg), any feasible price schedule must equal t almost everywhere.

Then, as discussed in the text, the monopolist solves (3) by choosing the corresponding implementable quality schedule, with $B = 0$. Since (q, t) is q-implementable with $B = 0$, this schedule is q. Finally, by the third condition, the monopolist chooses to produce.

Only if: If the schedule (q, t) is enactable, then there must exist some G that induces the monopolist to produce, and choose it. Thus, it must immediately satisfy the monopolist participation condition. Further, as discussed in the text, the monopolist solves (3) by choosing a t-implementable (q, t) with $B = 0$. Thus, by Lemma 1, (q, t) must be q-implementable with $B = 0$, giving the first and second conditions. $\square$

E.2 Section 4 proofs

Proof of Proposition 2.

Proof. First, let us define the consumer surplus function,

$$CS(q) := \int q(\theta) \frac{1 - F(\theta)}{f(\theta)} dF(\theta).$$

Then, let $\hat{q}$ be a solution to (6) with consumer surplus weight $\tilde{\alpha}$, and define the function $q_{\hat{\gamma}} : \Theta \to \mathbb{R}$ as:

$$q_{\hat{\gamma}}(\theta) := (1 - \hat{\gamma}) \min\{q^{\alpha}(\theta), \hat{q}(\theta)\} + \hat{\gamma}\hat{q}(\theta),$$

for any $\hat{\gamma} \in [0, 1]$.

An intermediate value theorem argument shows that there exists some $\gamma \in [0, 1]$ such that

$PS(q_\gamma) = PS(q^\alpha)$. First, observe that $PS(\hat{q}) \leq PS(q^\alpha)$ (since α-optimality gives $CS(q^\alpha) - CS(\hat{q}) \geq \frac{PS(\hat{q})-PS(q^\alpha)}{\alpha}$, while $\tilde{\alpha}$-optimality gives $CS(q^\alpha) - CS(\hat{q}) \leq \frac{PS(\hat{q})-PS(q^\alpha)}{\tilde{\alpha}}$, which are incompatible if $PS(\hat{q}) > PS(q^\alpha)$.) Then, since $q_1 = \hat{q}$, we have:

$$PS(q_1) \leq PS(q^\alpha). \tag{26}$$

Second, define $\overline{q}(\theta) := \max\{\hat{q}(\theta), q^\alpha(\theta)\}$. Clearly, $\overline{q}$ is pointwise larger than $\hat{q}$, and so $CS(\overline{q}) \geq CS(\hat{q})$ (using $\dot{u}(\theta) = q(\theta)$). Also, since PS is separable in θ, and using the definition of $q_{\hat{\gamma}}$, we have:

$$PS(\overline{q}) - PS(\hat{q}) = PS(q^\alpha) - PS(q_0). \tag{27}$$

Therefore, we must have:

$$PS(q_0) \geq PS(q^\alpha), \tag{28}$$

since otherwise, using (27), we would have $PS(\overline{q}) > PS(\hat{q})$ and $CS(\overline{q}) \geq CS(\hat{q})$, contradicting the $\tilde{\alpha}$-optimality of $\hat{q}$. Then, using (26) and (28), by continuity of PS and the intermediate value theorem, we must be able to find $\gamma \in [0, 1]$ such that $PS(q_\gamma) = PS(q^\alpha)$.

Since q_γ is a convex combination of two non-decreasing functions, it is non-decreasing, and was therefore feasible for (6) with weight α, but not chosen. We therefore have:

$$\alpha CS(q^\alpha) + PS(q^\alpha) \geq \alpha CS(q_\gamma) + PS(q_\gamma). \tag{29}$$

Define the perturbation

$$h(\theta) := q^\alpha(\theta) - q_\gamma(\theta).$$

Since $PS(q^\alpha) = PS(q_\gamma)$ by construction, we must have $CS(q^\alpha) \geq CS(q_\gamma)$, and so

$$\int h(\theta) \frac{1 - F(\theta)}{f(\theta)} dF(\theta) \geq 0. \tag{30}$$

Similarly, $PS(q^\alpha) = PS(q_\gamma)$ gives:

$$\int h(\theta)\phi(\theta)dF(\theta) + \int c(q^\alpha(\theta) - h(\theta)) - c(q^\alpha(\theta))dF(\theta) = 0. \tag{31}$$

Let us compare CS and PS at $\hat{q}$ and $\hat{q}+h$. We have $CS(\hat{q}+h) - CS(\hat{q}) = \int h(\theta)\frac{1-F(\theta)}{f(\theta)}dF(\theta)$, and so, by (30),

$$CS(\hat{q}+h) \geq CS(\hat{q}).$$

Similarly, we have

$$PS(\hat{q}+h) - PS(\hat{q}) = \int h(\theta)\phi(\theta)dF(\theta) + \int c(\hat{q}(\theta)) - c(\hat{q}(\theta) + h(\theta))dF(\theta). \tag{32}$$

Notice that

$$\int c(\hat{q}(\theta)) - c(\hat{q}(\theta) + h(\theta))dF(\theta) \geq \int c(q^\alpha(\theta) - h(\theta)) - c(q^\alpha(\theta))dF(\theta), \tag{33}$$

because the LHS integrand weakly exceeds the RHS integrand for all θ. (If $h(\theta) \geq 0$, the integrands are equal by definition of h, while if $h(\theta) < 0$, the LHS integrand weakly exceeds the RHS integrand, since $\gamma \leq 1$ and c is convex.) Therefore, combining (31), (32) and (33) we have:

$$PS(\hat{q}+h) \geq PS(\hat{q}).$$

Thus, $\hat{q}+h$ dominates $\hat{q}$. Additionally, it is non-decreasing, since

$$\hat{q} + h = (1-\gamma)\max\{q^\alpha, \hat{q}\} + \gamma q^\alpha.$$

Thus it is feasible, and so must also be optimal for the regulator with consumer surplus weight $\tilde{\alpha}$. Taking $q^{\tilde{\alpha}} = \hat{q} + h$ then proves the theorem. □

Proof of Lemma 4.

Proof.

$$\int \rho(\theta)[\theta q(\theta) - t(\theta)]dF(\theta) = \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\theta} \rho(\theta)f(\theta)q(x)dxd\theta = \int_{\underline{\theta}}^{\overline{\theta}} \int_{x}^{\overline{\theta}} \rho(\theta)f(\theta)q(x)d\theta dx$$
$$= \int_{\underline{\theta}}^{\overline{\theta}} q(x) \left(\int_{x}^{\overline{\theta}} \rho(\theta)f(\theta)d\theta \right) dx = \int_{\underline{\theta}}^{\overline{\theta}} q(x)(1 - F_{\rho}(x))dx = \int q(\theta)\frac{1 - F_{\rho}(\theta)}{f(\theta)}dF(\theta),$$

where, for the fourth equality, we have used the fact that $\int \rho(\theta)dF(\theta) = 1$. □

Proof of Proposition 3.

Proof. Let $\hat{q}$ be an optimal solution to (9) with weighting function $\tilde{\rho}$, and define the function $h : \Theta \to R, \theta \mapsto \max\{\hat{q}(\theta) - q^{\rho}(\theta), 0\}$. Since the monopolist's payoff is continuous in q, there exists some $\varepsilon \in (0, 1)$ such that the monopolist's participation constraint is satisfied by $q_{\varepsilon} := q^{\rho} + \varepsilon h$. Now, q_{ε} is non-decreasing, since it is a convex combination of two non-decreasing functions:

$$q_{\varepsilon} = (1 - \varepsilon)q^{\rho} + \varepsilon \max\{q^{\rho}, \hat{q}\}.$$

Thus, q_{ε} was a feasible solution for (9), but was not chosen. We therefore have:

$$\int \alpha(q^{\rho}(\theta) + \varepsilon h(\theta))(\theta)\frac{1 - F_{\rho}(\theta)}{f(\theta)} + (q^{\rho}(\theta) + \varepsilon h(\theta))\phi(\theta) - c(q^{\rho}(\theta) + \varepsilon h(\theta))dF(\theta) \quad (34)$$
$$\leq \int \alpha q^{\rho}(\theta)\frac{1 - F_{\rho}(\theta)}{f(\theta)} + q^{\rho}(\theta)\phi(\theta) - c(q^{\rho}(\theta))dF(\theta).$$

Define the functions ξ and C as:

$$\xi(\rho) := \int \alpha h(\theta)\frac{1 - F_{\rho}(\theta)}{f(\theta)} + h(\theta)\phi(\theta)dF(\theta)$$
$$C(k) := \int c(\hat{q}(\theta)) - c(\hat{q}(\theta) - kh(\theta))dF(\theta).$$

The function ξ represents the weighted surplus increase, ignoring costs, from adding h to any quality schedule. The function C represents the cost saving, relative to the schedule $\hat{q}$, of using the schedule $\hat{q} - kh$. This cost saving function satisfies $C(0) = 0$, and is concave, due to the convexity of c. Using these new functions, we rewrite (34) as:

$$C(1) - C(1-\varepsilon) - \varepsilon\xi(\rho) \geq 0. \tag{35}$$

Since $C(0) = 0$, the concavity of C implies:

$$C(1) - C(1-\varepsilon) \leq \varepsilon C(1), \tag{36}$$

and so, from (35), we have $\varepsilon C(1) - \varepsilon\xi(\rho) \geq 0$. Multiplying both sides by $\frac{1-\varepsilon}{\varepsilon}$, and adding to (35), we obtain:

$$C(1) - C(1-\varepsilon) + (1-\varepsilon)C(1) - \xi(\rho) \geq 0.$$

Finally, again applying (36), we obtain:

$$C(1) - \xi(\rho) \geq 0.$$

Since $\tilde{\rho}$ is weakly more equitable than ρ, we have $\frac{1-F_{\tilde{\rho}}(\theta)}{f(\theta)} \leq \frac{1-F_{\rho}(\theta)}{f(\theta)}$, for all $\theta \in \Theta$. Therefore, $\xi(\tilde{\rho}) \leq \xi(\rho)$, and so:

$$C(1) - \xi(\tilde{\rho}) \geq 0. \tag{37}$$

Now, the regulator's surplus increase, with weighting function $\tilde{\rho}$, from $\hat{q}$ to $\hat{q} - h$ is given by $C(1) - \xi(\tilde{\rho})$, which exceeds 0. Further, $\hat{q} - h$ is non-decreasing, and therefore feasible, because $\hat{q} - h = \min\{q^{\rho}, \hat{q}\}$. Therefore, $\hat{q} - h$ must also be optimal for a regulator with weighting function $\tilde{\rho}$; taking $q^{\tilde{\rho}} = \hat{q} - h$ proves the theorem. $\square$

E.3 Section 5 proofs

Proof of Lemma 5.

Proof. If:

Suppose $G_t \preceq_{MPS} G$, and find random variables X_t, X_G and Z, such that $\mathrm{E}[Z|X_t] = 0$ and

$X_G = X_t + Z$. Then, define $T \in \Delta(\mathbb{R}^+)^n$ so that T_i is the conditional distribution of X_G, given $X_t = t_i$. By construction, for any $\hat{t}$, we have (summing over θ_i): $G(\hat{t}) = \frac{1}{n}\sum_i T_i(\hat{t}) = G_T(\hat{t})$. Further, we have, for each i:

$$\mathrm{E}[T_i] = \mathrm{E}[X_G | X_t = t_i] = \mathrm{E}[X_t + Z | X_t = t_i] = t_i,$$

since $\mathrm{E}[Z|X_t] = 0$. Thus, t is compatible with G.

Only if:

Suppose t is compatible with G. Find the stochastic price schedule, T, such that $\overline{T} = t$ and $G_T = G$. Then, let θ denote any random variable with distribution $(\frac{1}{n} \circ \theta_1, \dots, \frac{1}{n} \circ \theta_n)$, and let X_t be the random variable which takes value t_i whenever θ takes value θ_i. Let X_T be any random variable with value drawn from T_i whenever θ takes value θ_i. Let $Z := X_T - X_t$. We have $\mathrm{E}[Z|X_t] = 0$ because $\overline{T} = t$, so $G_t \preceq_{MPS} G_T = G$. □

Proof of Lemma 6.

Proof. I claim that the distribution $G_{\overline{T}}$ enacts $(q, \overline{T})$. Suppose not. Observe that $\overline{T}$ is feasible for (10) under $G_{\overline{T}}$, because it inherits (IC) and (IR) from the enactability of (q, T) and satisfies ($G_{\overline{T}}$ Reg) by construction. Thus, there must exist a deterministic mechanism, $(\tilde{q}, \tilde{t})$, that satisfies (IC) and (IR), and such that $\tilde{t}$ is compatible with $G_{\overline{T}}$, that is strictly more profitable for the monopolist than $(q, \overline{T})$. By Lemma 5, we must have $G_{\tilde{t}} \preceq_{MPS} G_{\overline{T}}$. Thus, $G_{\tilde{t}} \preceq G_T$, and so $\tilde{t}$ is compatible with G_T. Thus, $(\tilde{q}, \tilde{t})$ is feasible for (10) under G_T, but is strictly more profitable, contradicting the enactability of (q, T). □

Proof of Lemma 7.

Proof. Clearly, (q, t) is robustly enactable if and only if it is robustly enacted by the distribution G_t. Let us examine let us examine (12) under G_t, which has the same value to the monopolist as (10) under G_t:

$$\max_{\hat{q},\hat{t}} \sum_i \hat{t}_i - c(\hat{q}_i) \tag{38}$$

$$\text{s.t. } \theta_i\hat{q}_i - \hat{t}_i \geq \theta_i\hat{q}_j - \hat{t}_j, \text{ for all } \theta_i, \theta_j \in \Theta \qquad \text{(IC)}$$

$$\theta_i\hat{q}_i - \hat{t}_i \geq 0 \text{ for all } \theta_i \in \Theta \qquad \text{(IR)}$$

$$G_{\hat{t}} \preceq_{MPS} G_t. \qquad (G_t \text{ MPC})$$

Notice that any feasible solution for (38) generates the same revenue ($\mathrm{E}_{G_t}[t]$), due to (G_t MPC), and so we can replace the objective with $\min_{\hat{q},\hat{t}} \sum c(\hat{q}_i)$. Second, notice that any $(\hat{q}, \hat{t})$ which solves (38) also solves the optimization with (G_t MPC) repaced by the equality constraint, $G_{\hat{t}} = G_t$. Thus, $(\hat{q}, \hat{t})$ must be enactable. Using Theorem 2, we can further rewrite the monopolist's optimization, facing distribution G_t, as:

$$\min_{\hat{t} \text{ incr.}} \Gamma(\hat{t})$$
$$\text{s.t. } G_{\hat{t}} \preceq_{MPS} G_t,$$

which is equivalent to:

$$\min_{\hat{t} \preceq_{MPS} t} \Gamma(\hat{t}). \tag{39}$$

Now, if t solves (39), we know (q, t) solves (38), and therefore must solve (10) ((q, t) is feasible for (10) since it is enactable, and (38) and (10) have the same value for the monopolist). Conversely, if t does not solve (39), there must exist some feasible $\tilde{t}$ the monopolist strictly prefers, and hence, a strictly better $(\tilde{q}, \tilde{t})$ for (38). Since (38) has the same value as (10), (q, t) cannot be optimal for (10). □

Proof of Theorem 4.

Proof. If (q, t) is enactable, then by Lemma 7 it is robustly enactable if t solves

$$\min_{\hat{t} \preceq_{MPS} t} \Gamma(\hat{t}). \tag{40}$$

By Theorem 3, this occurs if the no local contractions condition is satisfied.

Let us now prove the partial converse, when t is strictly increasing. Clearly, if (q, t) is not enactable, it is not robustly enactable (facing G_t, the monopolist can improve on (q, t) with another deterministic schedule).

This leaves the case where (q, t) is enactable, but $\Gamma_i(t) - \Gamma_{i+1}(t) < 0$, for some i. In this case, using Taylor's theorem, choose $\varepsilon > 0$ sufficiently small that:

$$\Gamma(t_1, \ldots, t_i + \varepsilon, t_{i+1} - \varepsilon, \ldots, t_n) > \Gamma(t),$$

and $t_i + \varepsilon \leq t_{i+1} - \varepsilon$ (we can do the latter since t is strictly increasing). Then, taking $\tilde{t} = (t_1, \ldots, t_i + \varepsilon, t_{i+1} - \varepsilon, \ldots, t_n)$, we have $\tilde{t} \preceq_{MPS} t$ and $\Gamma(\tilde{t}) < \Gamma(t)$, and so by Lemma 7, (q, t) is not robustly enactable. $\square$

CHAPTER 2

SHARPENING WINKLER'S EXTREME POINT THEOREM: ECONOMIC APPLICATIONS

Sharpening Winkler's Extreme Point Theorem: Economic Applications

Christoph Schlom*

Winkler's extreme point theorem (Winkler, 1988) states that the extreme points of a convex set, X, constrained by n linear constraints, are convex combinations of at most $n+1$ extreme points of X. Winkler's condition is always necessary for a point to be extreme in the constrained set; it provides a characterization of those extreme points only when X is a simplex and the constraints are hyperplane equality constraints. I provide a characterization of the extreme points for general linear constraints, and use it to analyze two problems in economic theory.

1 Introduction and related literature

Extreme point techniques are widely used in economic theory. They have perhaps received particular interest in recent years due to the use of closely related concavification methods in finite-dimensional Bayesian persuasion, starting with Kamenica and Gentzkow (2011). I will present a useful theorem due to Winkler, relating the extreme points of a linearly constrained set to the extreme points of the corresponding unconstrained set. Then, I will present two "sharpenings" of that theorem, which characterize the extreme points of the constrained set in settings where Winkler's condition is necessary but not sufficient. I will discuss the relevance of these sharpenings to economic problems via two simple examples.

*University of Chicago, Booth School of Business. I thank Phil Reny, Lars Stole and Emir Kamenica for helpful feedback.

To my knowledge, the most closely related paper to this one is Doval and Skreta (2023). They analyze constrained extreme point techniques in information design, and their Corollary 2, which provides a sharper version of Winkler's condition, is the analogue of my Theorem 3. Relative to that paper, my contribution is (1) to situate that sharper condition in a general constrained optimization setting (2) to show the sense in which that sharper condition is in fact a characterization (3) to state and prove Theorem 2, which applies to a more general form of linear constraints.

2 Winkler's Theorem and sharpenings

I will first state two definitions, which will be useful in the statement of Winkler's theorem. More context can be found in Winkler (1988) and in Barvinok (2002), where the result appears as Theorem 9.2.

Definition 1. *An **extreme point** of a set, K, in a real vector space is a point $e \in K$ such that there do not exist $p_1, p_2 \in K$, with $p_1, p_2 \neq e$ and $\alpha \in [0,1]$, such that $e = \alpha p_1 + (1-\alpha)p_2$.*

Notice that this definition does not require K to be convex. I will use $\text{ext}(K)$ to denote the set of extreme points of a set K.

Next, an informal definition (a formal definition can be found in Winkler). By a **simplex**, I mean a Choquet simplex, which is a generalization of a finite-dimensional simplex. Its key property is that each point in the simplex can be uniquely expressed as a convex combination of the simplex's extreme points.

We can now state Winkler's original result. The version I present is a specialization of Winkler's Proposition 2.1 to the case where X is a simplex.[1]

Theorem 1 (Winkler, 1988)**.** *Let X be a simplex in a real vector space. Let $\Phi : X \to \mathbb{R}^n$*

[1] Winkler further shows that the only if direction holds if X is any convex, linearly compact set.

be a linear map, and let Σ be a subset of $\mathbb{R}^n$. Define:

$$K := \Phi^{-1}(\Sigma).$$

Then a point, $e \in K$, is an extreme point of K only if it can be written as $e = \sum_1^m \alpha_i x_i$, with $\sum \alpha_i = 1$; $\alpha_i > 0$, $x_i \in \operatorname{ext}(X)$ for all i; and satisfying:

(AI) $\Phi(x_1), \ldots, \Phi(x_m)$ *are affinely independent.*

Further, if Σ is a single point, then any such e is an extreme point of K.

Next, I will give a sharper condition, which characterizes the extreme points of constrained simplices when the linear constraints take on Winkler's more general form. The key additional condition requires the candidate extreme point to have an image which is extremal in the "constrained space" ($\mathbb{R}^n$) in an appropriate sense. I call this the barycentric condition.

Theorem 2. *Let X be a simplex in a real vector space. Let $\Phi : X \to \mathbb{R}^n$ be a linear map, and let Σ be a subset of $\mathbb{R}^n$. Define:*

$$K := \Phi^{-1}(\Sigma).$$

Then a point, $e \in K$, is an extreme point of K if and only if it can be written as $e = \sum_1^m \alpha_i x_i$, with $\sum \alpha_i = 1$; $\alpha_i > 0$, $x_i \in \operatorname{ext}(X)$ for all i; and satisfying:

(AI) $\Phi(x_1), \ldots, \Phi(x_m)$ *are affinely independent.*

(B) $\Phi(e) \in \operatorname{ext}\left(\Sigma \cap \operatorname{co}(\Phi(x_1), \ldots, \Phi(x_m))\right).$

Proof. The necessity of the affine independence property was proved in Winkler. To prove the necessity of the barycentric property, we will show that if a point has *any* affinely independent representation which does not satisfy the barycentric property, it cannot be an extreme point.

Choose a point, $p \in K$, and a representation, $p = \sum_1^m \alpha_i x_i$, satisfying (AI) but not (B). Since $\Phi(x_1), \ldots, \Phi(x_m)$ are affinely independent, Φ is one-to-one on the set $\mathrm{co}(x_1, \ldots, x_m)$. Thus, Φ^{-1} is well-defined, linear and one-to-one on its image, $\mathrm{co}(\Phi(x_1), \ldots, \Phi(x_m))$. Since (B) is not satisfied, write:

$$\Phi(p) = 1/2(s' + s''), \tag{1}$$

with $s', s'' \in \Sigma \cap \mathrm{co}(\Phi(x_1), \ldots, \Phi(x_m))$ and not equal to $\Phi(p)$. Using the aforementioned properties of Φ^{-1}, (1) implies: $p = 1/2(\Phi^{-1}(s') + \Phi^{-1}(s''))$, with $\Phi^{-1}(s'), \Phi^{-1}(s'') \in K$ and not equal to p. Thus, p is not an extreme point of K.

To prove sufficiency, choose $e \in K$ with $e = \sum_1^m \alpha_i x_i$ satisfying (AI) and (B). Following Winkler, again observe that, since $\Phi(x_1), \ldots, \Phi(x_m)$ are affinely independent, Φ is one-to-one on the set $S = \mathrm{co}(x_1, \ldots, x_m)$. Suppose $k', k'' \in K$ are such that $e = 1/2(k' + k'')$. Winkler's Lemma 2.3 shows that, since X is a simplex, S is a face, so that $k', k'' \in S$. By linearity, we have $\Phi(e) = 1/2(\Phi(k') + \Phi(k''))$; since $k', k'' \in K \cap S$, we have $\Phi(k'), \Phi(k'') \in \Sigma \cap \mathrm{co}(\Phi(x_1), \ldots, \Phi(x_m))$. Therefore, by (B), $\Phi(k') = \Phi(k'') = \Phi(e)$. Finally, since Φ is one-to-one on S, we have $k' = k'' = e$, so that e is an extreme point of K. $\square$

Finally, I give a specialized characterization when the constraints are a mix of hyperplane equality and inequality constraints. Recall the definitions of a hyperplane and a halfspace:

Definition 2. *A* ***hyperplane*** *in a real vector space, V, is a set, H, corresponding to a non-zero linear map $\phi : V \to \mathbb{R}$ and a constant $c \in \mathbb{R}$, such that:*

$$H = \{v \in V : \phi(v) = c.\}$$

The corresponding ***closed halfspace*** *is the set:*

$$\overline{H} = \{v \in V : \phi(v) \leq c.\}$$

Key to this characterization is a sharp affine independence condition, which resembles Winkler's condition, except that we drop non-binding constraints. It is therefore analogous to a Kuhn-Tucker condition.

Theorem 3. *Let X be a simplex in a real vector space. Consider a collection of n sets, $\hat{H}_1, \ldots, \hat{H}_n$, each of which is either a hyperplane or a closed halfspace. Let H_i and ϕ_i be, respectively, the hyperplane and the linear functional corresponding to $\hat{H}_i$. Define:*

$$K := X \cap \hat{H}_1 \cap \cdots \cap \hat{H}_n.$$

Given some subset of indices, $A = \{a_1, \ldots, a_\ell\} \subset \{1, \ldots, n\}$, define the constraint function restricted to those indices as:

$$\Phi_A : X \to \mathbb{R}^\ell$$
$$x \mapsto (\phi_{a_1}(x), \ldots, \phi_{a_\ell}(x)).$$

Then a point, $e \in K$, is an extreme point of K if and only if it can be written as $e = \sum_1^m \alpha_i x_i$, with $\sum \alpha_i = 1$; $\alpha_i > 0$, $x_i \in \operatorname{ext}(X)$ for all i; and satisfying:

(SAI) Let A be the set of indices of the hyperplanes H_i that e lies on. Then, $\Phi_A(x_1), \ldots, \Phi_A(x_m)$ are affinely independent.

Proof. To prove necessity, suppose $p \in K$ is a point which cannot be represented in the desired form, and let $A = \{a_1, \ldots, a_\ell\}$ be the set of indices of hyperplanes it lies on. Let c_i be the constant corresponding to hyperplane H_i,[2] and define $s := (c_{a_1}, \ldots, c_{a_\ell})$. Define $K^* := \Phi_A^{-1}(s)$. Applying Theorem 1 to K^* lets us write $p = 1/2(a^* + b^*)$, with $a^*, b^* \in K^*$ and not equal to p. For $\varepsilon \in [0, 1]$, define $a(\varepsilon) := \varepsilon a^* + (1 - \varepsilon)p$, and $b(\varepsilon) := \varepsilon b^* + (1 - \varepsilon)p$. By linearity of ϕ_i, we can choose $\varepsilon > 0$ small enough that $\phi_i(a(\varepsilon)), \phi_i(b(\varepsilon)) < c_i$ for all i

[2] That is, $H_i = \{v : \phi_i(v) = c_i\}$.

for which p does not lie on H_i. Choosing such an ε, we have $p = 1/2(a(\varepsilon) + b(\varepsilon))$, with $a(\varepsilon), b(\varepsilon) \in K$ and not equal to p, so that p is not an extreme point of K.

To prove sufficiency, let $e = \sum_1^m \alpha_i x_i$ be a point with the desired representation, and define $K^* := \Phi_A^{-1}(s)$ as in the necessity argument.[3] Suppose $e = 1/2(a + b)$, with $a, b \in K$. For any $i \in \{1, \ldots, \ell\}$, we have $c_{a_i} = \phi_{a_i}(e) = 1/2(\phi_{a_i}(a) + \phi_{a_i}(b))$, and so $\phi_{a_i}(a) = \phi_{a_i}(b) = c_{a_i}$, since $\phi_{a_i}(a), \phi_{a_i}(b) \leq c_{a_i}$. Thus, $a, b \in K^*$. By Theorem 1, e is an extreme point of K^*; thus, $a = b = e$, so e is an extreme point of K. □

3 Economic examples

I will now illustrate, via two examples, the type of characterizing results in economic theory that we can prove using Winkler's theorem and its sharpenings.

3.1 One-dimensional Bayesian persuasion

As a warm-up exercise, we will establish a familiar result, which requires only Winkler's original theorem.

Following Kamenica and Gentzkow (2011) (KG), there is a state, $\omega \in \{0, 1\}$, which is payoff-relevant to a sender and a receiver. Both parties have a common prior that $\omega = 1$ with probability μ_0. The sender chooses an arbitrary public experiment about the state, and, after observing the outcome of the experiment, the receiver takes an action that is payoff-relevant to both parties. KG show that the sender's experiment is able to induce exactly those distributions of posterior beliefs which have expectation μ_0. Then, writing the sender's interim expected payoff from inducing a belief of μ as $V(\mu)$, the sender's problem

[3]That is, $A = \{a_1, \ldots, a_\ell\}$ is the set of indices of hyperplaes e lies on, and $s = (c_{a_1}, \ldots, c_{a_\ell})$, where c_i is the constant for hyperplane H_i.

becomes:

$$\max_F \int V(\mu)dF(\mu) \tag{2}$$
$$\text{s.t.} \quad \int \mu dF(\mu) = \mu_0.$$

It will be useful to understand the extreme points of the set, K, of probability measures on $[0,1]$ (the possible posteriors) which have mean μ_0. Consider first the set X of probability measures on $[0,1]$, as a subset of the vector space of signed measures on $[0,1]$. The set X is a simplex, and has as extreme points the degenerate measures which put unit mass on some point in $[0,1]$. Then, the constraint $\int \mu dF(\mu) = \mu_0$ is of the required form, with $\Phi : X \to \mathbb{R}, F \mapsto \int \mu dF(\mu)$, and $\Sigma = \mu_0$, a single point. Theorem 1 therefore tells us that the extreme points of the constrained set are exactly the measures which put mass on at most two posteriors, and have expectation μ_0.[4]

Further, each of these extreme points is exposed (the extreme point which puts mass on posteriors μ_1 and μ_2 is uniquely optimal if V is 1 at μ_1 and μ_2 and 0 elsewhere). Since the objective is linear in F, by Bauer's maximum principle we have the following result:

Proposition 1. *Fix μ_0 in* (2)*. Then, for any V, there is an optimal solution, F, to* (2) *that puts mass on at most two points in* $[0,1]$ *and has expectation μ_0. Further, for any such F, there exists a V that makes F uniquely optimal.*

When there are n states, we obtain an analogous result, where the optimal solution puts mass on at most n posteriors.[5] Such a result is familiar, and was first observed in Proposition 9 of Kamenica and Gentzkow (2009). The emphasis in this paper is on the sense in which solutions of this form *characterize* the possible solutions as we vary the model primitives. As we shall see, to obtain such a characterization with more complex

[4]This follows from the affine independence condition, once we observe that Φ is one-to-one on the degenerate measures on $[0,1]$.

[5]In this case, the affine independence condition has more bite – for instance, with three states, it rules out measures which put mass on three co-linear points.

constraints, we will need a sharper version of Winkler's result.

3.2 Selling under a price cap

A revenue-maximizing monopolist wishes to sell an indivisible item, subject to a regulatory price cap, which states that the object may not be sold at a price exceeding c. The monopolist sells to a single buyer, who has private valuation, v, drawn from a distribution F, which is known to the monopolist and supported on $[\underline{v}, \overline{v}]$. The buyer has quasi-linear utility for money, and all agents are risk-neutral. Such a constraint seems economically relevant because a price cap is a common prescription against monopoly power. Though I interpret c as a price cap, it can also be interpreted as a known buyer budget, as in Che and Gale (2000).[6]

We will model the monopolist as a mechanism designer, who can commit to an arbitrary extensive form game (potentially involving random moves by nature), between herself and the buyer, each outcome of which is either no sale, or a deterministic sale of the good at a price less than or equal to c. Further, the mechanism must respect an ex-post participation constraint for the buyer. For example, the two-price mechanism $\langle(1/2, 2), (1, 4)\rangle$, in which the buyer selects between a $1/2$ chance of being allowed to buy the item at a price of 2 (otherwise no sale) and a guarantee to be allowed to buy the item at a price of 4, is a legal mechanism under a price cap of 4.

Any such mechanism induces an ex-interim probability that each buyer type will obtain the good, and an expected payment for each buyer type. I will call these maps $p : [\underline{v}, \overline{v}] \to [0, 1]$ and $t : [\underline{v}, \overline{v}] \to \mathbb{R}$. For such maps to be implementable by a legal mechanism, it is necessary that they satisfy the usual IC and IR conditions, as well as the interim price cap constraint, that $t(\overline{v}) \leq c$. Further, in any optimal mechanism, either $t(v) = c$ for all types, or the lowest type gets 0 surplus.[7]

[6] Che and Gale give a regularity condition (their Assumption 1) which rules out the optimality of the two-price mechanisms I describe in this section. This exercise can be seen as a relaxation of that condition.

[7] If neither of these is true, the monopolist can improve by slightly increasing t for the lowest type, and

Thus, in any optimal mechanism for the monopolist, either $t(v) = c$ for all types, or the following hold: (1) p is non-decreasing and (2) p and t satisfy the envelope formula where $\underline{v}$ receives 0 surplus,

$$t(v) = vp(v) - \int_{\underline{v}}^{v} p(x)dx, \text{ for all } v.$$

In the latter case, let us consider the following relaxed optimization over p:

$$\max_{p \text{ incr.}} \int \left[vp(v) - \int_{\underline{v}}^{v} p(x)dx \right] dF(v) \tag{3}$$
$$\text{s.t. } \overline{v}p(\overline{v}) - \int_{\underline{v}}^{\overline{v}} p(x)dx \leq c,$$

where "incr." is shorthand for "non-decreasing."

As in the Bayesian persuasion example, it will be useful to understand the extreme points of K, the set of non-decreasing functions, $p : [\underline{v}, \overline{v}] \to [0, 1]$, which satisfy the price-cap constraint in (3). Following Chapter 2 of Börgers (2015), first consider the vector space of bounded non-decreasing functions on $[\underline{v}, \overline{v}]$, and let X denote the subset of those functions which map into $[0, 1]$. Observe that X is a simplex, and that its extreme points are the 0-1 step functions (that is, non-decreasing functions which are everywhere 0 or 1). Now, since the price-cap constraint is a hyperplane inequality constraint, Theorem 3 tells us that the extreme points of K are exactly (1) the 0-1 step functions and (2) the convex combinations of two 0-1 step functions which satisfy the price-cap constraint with equality.[8] I will refer to these convex combinations of two 0-1 step functions as "0-1 step functions with two jumps."

Further, any such p is uniquely optimal against some distribution F, and implementable via a legal mechanism. Any p of the first type corresponds to a posted price mechanism with price $t_1 \leq c$. It is uniquely optimal against an F which puts unit mass on $v = t_1$.

allowing him to choose his favorite menu item from the modified mechanism.

[8] As before, this follows from the affine independence condition once we observe that $\overline{v}p(\overline{v}) - \int_{\underline{v}}^{\overline{v}} p(x)dx$ is one-to-one on the 0-1 step functions.

Any p of the second type is a convex combination, $\alpha p_1 + (1-\alpha)p_2$, between two posted price mechanisms, with prices $t_1 < c$ and $t_2 > c$. It is uniquely optimal against an F which puts mass only on $v = t_1$ and $v = t_2$, and such that the mass, m, it puts on t_2 satisfies $m > t_1/t_2$. Finally, it is legally implemented by the mechanism $\langle(\alpha, t_1), (1, c)\rangle$.[9]

Finally, notice that any IC and legal mechanism in which $t(v) = c$ for all types must have $p(v) = 1$ for all types, and so be a posted price mechanism with price c, and that this mechanism is already included by the first type of extreme point. Therefore, since the objective in (3) is linear in p, we can again use Bauer's maximum principle to obtain the following characterization:

Proposition 2. *Any optimal mechanism for a price-capped monopoly problem with cap c induces interim probability assignment and transfer functions, p and t, linked via $t(v) = vp(v) - \int_{\underline{v}}^{v} p(x)dx$, for all v, and such that either:*

1. *$p(v)$ is a 0-1 step function with a jump at $t_1 \leq c$*

2. *$p(v)$ is a 0-1 step function with two jumps, and $t(\overline{v}) = c$*

Further, for any such p and t, there exists a consumer value distribution, F, against which any optimal mechanism induces those p and t almost everywhere.

[9]For example, if $c = 4$ and F puts mass 1/2 on 2 and 1/2 on 6, then the 0-1 step function with two jumps,

$$p(v) = \begin{cases} 0 & v < 2 \\ \frac{1}{2} & v \in [2, 6) \\ 1 & v \geq 6, \end{cases}$$

is uniquely optimal for (3), and is legally implemented by the two-price mechanism, $\langle(1/2, 2), (1, 4)\rangle$.

References

Barvinok, A. (2002). *A Course in Convexity*, Volume 54. American Mathematical Soc.

Börgers, T. (2015). *An Introduction to the Theory of Mechanism Design*. Oxford University Press, USA.

Che, Y.-K. and I. Gale (2000). The Optimal Mechanism for Selling to a Budget-Constrained Buyer. *Journal of Economic theory 92*(2), 198–233.

Doval, L. and V. Skreta (2023). Constrained Information Design. *Mathematics of Operations Research*.

Kamenica, E. and M. Gentzkow (2009). Bayesian Persuasion. Working Paper 15540, National Bureau of Economic Research.

Kamenica, E. and M. Gentzkow (2011). Bayesian Persuasion. *American Economic Review 101*(6), 2590–2615.

Winkler, G. (1988). Extreme Points of Moment Sets. *Mathematics of Operations Research 13*(4), 581–587.

CHAPTER 3

CONCAVIFICATION BOUNDS AND MECHANISM SIMPLICITY

Concavification Bounds and Mechanism Simplicity

Christoph Schlom*

I develop a new type of bound, relating to mechanism simplicity, in a class of economic models: the linear functional optimization models. In these models, which are frequently used in economic design contexts, the optimal function is piecewise constant with at most n jumps, where n is a finite number related to the number of constraints in the problem. My bound quantifies, over all model primitives, the maximum shortfall a designer could have from using the optimal piecewise constant function with at most m jumps, with $m < n$. I illustrate my results in the capacity-constrained selling problem of Bulow and Roberts.

1 Introduction

This paper develops a new type of bound (a "concavification bound"), which applies to a particular class of economic optimization problems. That class is the (monotone) linear functional optimization problems, in which the aim is to choose a non-decreasing function to maximize a linear functional, subject to a finite number of linear constraints. These problems often arise in economic design contexts. If there are n linear constraints, a well-known argument, using extreme-point techniques, shows that the optimization problem is solved by some piecewise-constant function with at most $n+1$ jumps. The bounds I will develop answer the question: "How much worse off could the designer be, if, instead of using the optimal function, she used a simpler function: namely, the best piecewise-constant function with $m < n+1$ jumps?"

*University of Chicago, Booth School of Business.

These bounds help us to understand the potential value of mechanism complexity, which is especially useful to do in settings where such complexity comes with an unmodeled cost. For example, in unpublished work (Schlom, 2024), I extend a model of Fryer and Loury (2013) to study optimal affirmative action via randomized school admissions policies. In that model, knowing the test score distributions of advantaged and disadvantaged students, a school sets a monotone admissions policy (a non-decreasing map from test scores to admission probabilities), to maximize the total test score of admitted students, subject to an equity constraint that a certain mass of those admits be disadvantaged. Since the objectives and the constraints (a total capacity constraint and the equity constraint) are both linear in the admissions policy, I show that the optimal policy is a piecewise-constant function with at most 3 jumps. That is, the optimal policy partitions test scores into (no more than) 4 categories: reject, maybe admit, probably admit, definitely admit.

Consider a conversation between an economist (E) who is familiar with this analysis and a school (S).

E: I don't have access to your data on student test scores by group, but I can tell you that the optimal admissions policy, once you look at that, won't be any more complicated than specifying 3 cutoff test scores (and 4 admissions categories).

S: Even so, that's kind of complicated. We feel that a randomized admissions policy would be easier to explain to parents and students if it just had 3 admissions categories: reject, maybe admit, definitely admit.

At this point, the economist would like to be able to advise the school based on bounds of the type I will present in this paper. She might say:

E: That sounds reasonable – without looking at the data, I can tell you that you'd never give up more than 10% of aggregate test scores by using that simpler policy.

Or she might say:

E: I think you should hold off on that decision until you look at the data. In the worst case, you could be sacrificing 50% of aggregate test scores by using that simpler policy.

In this draft of the paper, I will derive such a bound in the capacity-constrained monopoly problem proposed by Bulow and Roberts (1989) and studied further by Loertscher and Muir (2022). In that model, the monopolist optimally uses a 2-jump mechanism; I will derive a bound on the shortfall if the monopolist instead uses a 1-jump (posted price) mechanism. This "2-1" bound represents the simplest non-trivial example of a concavification bound.

2 Concavification bounds in the Bulow-Roberts model

2.1 Model

A capacity-constrained monopolist seller sells a good to a continuum of buyers. The buyers have mass 1, and (private) valuations, v, drawn iid from a distribution with CDF F, which is known to the seller, and is supported on the interval $V := [\underline{v}, \overline{v}]$. The seller has capacity $c < 1$.

Using standard revelation principle arguments, we can think of the seller's design problem as choosing a monotone assignment function, $p : V \to [0,1]$, and then using the direct mechanism comprised of p and the corresponding transfer function, $t(v) := vp(v) - \int_{\underline{v}}^{v} t(x)dx$. The seller optimally chooses p to maximize his revenue, subject to the capacity constraint. He therefore solves the optimization problem:

$$\max_{p} \int [vp(v) - \int_{\underline{v}}^{v} p(x)dx]dF(v) \tag{1}$$
$$\text{s.t.} \int p(v)dF(v) \leq c.$$

2.2 The concavification bound

I will now state the bound we are after. Bulow and Roberts show that the seller's optimal assignment function features at most 2 jumps; we will be interested in the maximum possible (over all WTP distributions, F) revenue ratio between the optimal mechanism and the optimal posted-price mechanism.

Let us begin by defining the class of posted-price assignment functions.

Definition 1. *The (feasible)* ***posted-price assignment function*** *associated with price v^* is the assignment function:*

$$p_{v^*}(v) := \begin{cases} 0 & v < v^* \\ a & v \geq v^*, \end{cases}$$

where $a = \max\{1, c/(1 - F(v^*))\}$.

A posted price mechanism sells the good at a posted price, and uses rationing if the demand at that price exceeds c.

Define the revenue associated with an assignment function, p, under distribution F as:

$$rev(p, F) := \int [vp(v) - \int_{\underline{v}}^{v} p(x)dx] dF(v).$$

Finally, given some distribution, F, define the optimal assignment function, p_F^*, to be any which solves (1). Then, define the **revenue ratio** at that F as:

$$R(F; c) := \frac{rev(p_F^*, F)}{\max_{v^*} rev(p_{v^*}, F)}. \tag{2}$$

The **concavification bound** is given by:

$$B(c) := \max_F R(F; c) = \max_F \min_{v^*} \frac{rev(p_F^*, F)}{rev(p_{v^*}, F)}. \tag{3}$$

Notice that, since the worst-case is taken only over F, the bound depends on the capacity, c. This reflects the idea that the seller knows his capacity, but is (at the time the bound is assessed) uncertain about the distribution of buyers he will face. We will call any distribution, $\hat{F}$, which achieves the maximum in (3), a "worst-case" distribution.

2.3 Computing the concavification bound

A major simplification of our problem comes from observing that there is a worst-case distribution which puts mass on exactly two points. This is a consequence of the familiar "at most 2 jumps" result.

Lemma 1. *There exists a worst-case distribution which puts mass on exactly two points.*

Proof. We will show that, given any distribution, F, we can find some 2-point distribution, $\tilde{F}$ which improves on F, in the sense that $R(\tilde{F}, c) \geq R(F, c)$. Recall that the F-optimal assignment function, p_F^*, has either 1 or 2 jumps, and can, in particular, be written as:

$$p_F^*(v) := \begin{cases} 0 & v < v_1 \\ a & v_1 \leq v < v_2 \\ 1 & v \geq v_2, \end{cases}$$

for some a, v_1, v_2, with $v_1 < v_2$. Now, define $\hat{\tilde{F}} := (m_0 \circ 0, m_1 \circ v_1, m_2 \circ v_2)$, where $m_0 := P_{X\sim F}(X < v_1)$, $m_1 := P_{X\sim F}(v_1 \leq X < v_2)$ and $m_2 := P_{X\sim F}(X \geq v_2)$. That is, $\hat{\tilde{F}}$ moves each unit of probability mass under F to the bottom of its p_F^* "bin." Notice that $rev(p_F^*, \hat{\tilde{F}}) = rev(p_F^*, F)$, and so we must have:

$$rev(p^*_{\hat{\tilde{F}}}, \hat{\tilde{F}}) \geq rev(p_F^*, F). \tag{4}$$

On the other hand, since $F \succeq_{FOSD} \hat{\tilde{F}}$, any fixed posted-price assignment function performs weakly worse against $\hat{\tilde{F}}$ than against F (weakly less buyers pay the posted price under $\hat{\tilde{F}}$). Thus,

$$\max_{v^*} rev(p_{v^*}, \hat{\tilde{F}}) \leq \max_{v^*} rev(p_{v^*}, F), \tag{5}$$

and so, combining inequalities (4) and (5), we find that $R(\hat{\tilde{F}}, c) \geq R(F, c)$. Finally, let $\tilde{F} := (\frac{m_1}{m_1+m_2} \circ v_1, \frac{m_2}{m_1+m_2} \circ v_2)$, and observe that this effectively grows the population

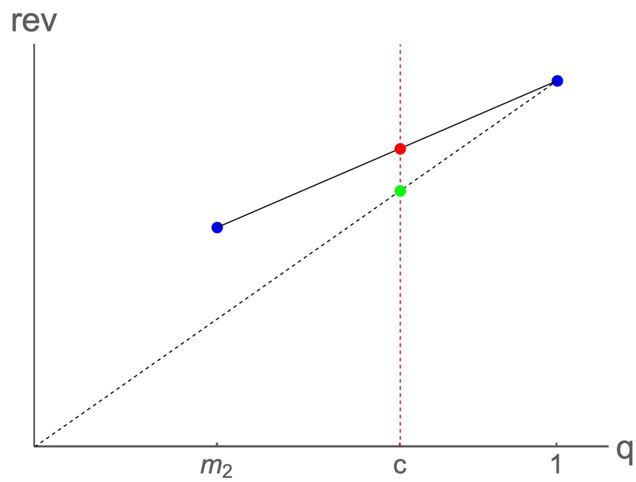

Figure 1: The optimal mechanism and the optimal posted-price mechanism under an arbitrarily chosen 2-point distribution.

from $\hat{\tilde{F}}$ by a factor of $\frac{1}{m_1+m_2}$. This means that both the revenue from the optimal posted-price mechanism and the optimal 2-price mechanism both grow by that factor, and so $R(\tilde{F}, c) = R(\hat{\tilde{F}}, c) \geq R(F, c)$, proving the result.

□

With the simplification from Lemma 1, the problem lends itself to a simple geometric analysis. Let us begin by reviewing the concavification approach that can be used to find the optimal 2-price mechanism.

Figure 1 shows a typical example of the monopolist's problem under a 2-point distribution, $F = (m_1 \circ v_1, m_2 \circ v_2)$. To compute the revenue-maximizing mechanism, we begin by plotting the posted-price mechanisms (ignoring the constraints) in quantity-revenue space. These are the blue points: the left point corresponds to a posted price of v_2, at which a quantity m_2 transacts; the right point corresponds to a posted price of v_1, at which a quantity 1 would transact, absent the capacity constraint. We then incorporate the constraint by drawing the red dotted line at quantity c. Finally, we form the concave envelope of the blue points, and choose the maximal feasible (i.e., weakly left of the red line) point in this envelope (the red point). This point represents the optimal selling mechanism: its y-coordinate represents the optimal revenue, while its assignment function is given by taking the appropriate convex combination of the assignment functions corresponding to the blue points. Since those assignment functions are 0-1 step functions, the optimal mechanism's assignment function has 2 jumps.

This approach works by capturing the linearity in the problem. Any non-decreasing assignment function is a convex combination of 0-1 step functions. Further, for any assignment function, its relevant properties (its revenue and its capacity) are, by linearity, that same convex combination of those properties of those step functions.

Figure 1 also shows the revenue from feasible posted-price mechanisms. There are two reasonable candidates for the seller: the left blue point represents a posted price of v_2, while the green point represents a posted price of v_1 (with rationing). In this example, the green point is higher than the blue point, and so the revenue ratio, R, is given by the ratio of the height of the red point to the height of the green point.

How should we choose F to maximize the revenue ratio? First, notice that we have not yet specified units for the price. Since these units do not matter for the revenue ratio, we may set $v_1 = 1$ without loss, so that the worst-case distribution is given by $(m_1 \circ 1, m_2 \circ v_2)$. Next, notice that, ordinally, the worst-case distribution must resemble the distribution in Figure 1, in that we must have $m_2 < c < 1$ and $m_2 v_2 < 1$; otherwise, the revenue ratio will be 1. Thus, we can view our problem geometrically, in quantity-revenue space, as fixing the right blue point at $(1, 1)$, and therefore the green point at (c, c), and moving around the left blue point, within the region to the left of c.

First, notice that, keeping fixed the height of that left blue point, it is optimal to move it all the way to the left. This increases the height of the red point, without improving either of the candidate posted-price mechanisms. Next, notice that, once the blue point is all the way on the y-axis, it should be no lower than the green point. Indeed, if it is strictly below the green point, then raising it slightly increases the height of the red point, without improving the best candidate posted-price mechanism (the green point). Finally, it is easy to confirm algebraically that it should also be moved *no higher* than the green point. Such a movement raises the revenue from both the

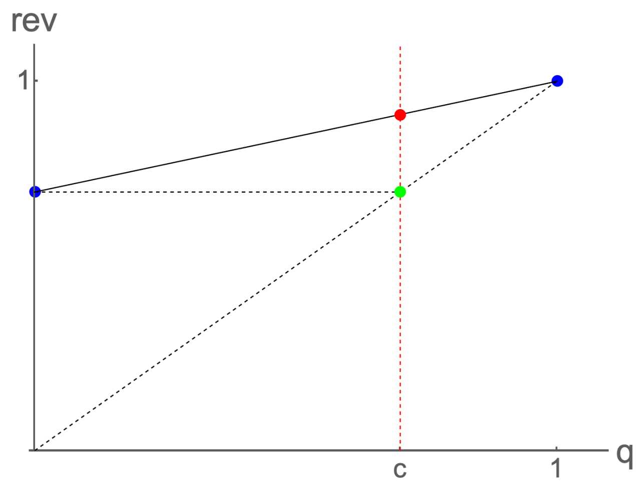

Figure 2: The optimal mechanism and the optimal posted-price mechanisms under the worst-case distribution.

optimal mechanism (red) and the optimal posted price mechanism (blue), but lowers the revenue ratio overall.

These arguments pin down the worst-case distribution, shown in Figure 2. Concretely, the worst-case distribution is given by: $(1 - \frac{1}{B} \circ 1, \frac{1}{B} \circ cB)$, where B is some very large number (infinity, if you like). Against this distribution, both candidate posted price mechanisms yield revenue c, while the optimal mechanism yields $(1-c) \times c + c \times 1 = 2c - c^2$, giving the following result:

Proposition 1. *The concavification bound in the Bulow-Roberts model is given by* $B(c) = 2 - c$.

References

Bulow, J. and J. Roberts (1989). The simple economics of optimal auctions. *Journal of Political Economy 97*(5), 1060–1090.

Fryer, R. G. and G. C. Loury (2013). Valuing Diversity. *Journal of Political Economy 121*(4), 747–774.

Loertscher, S. and E. V. Muir (2022). Monopoly Pricing, Optimal Randomization, and Resale. *Journal of Political Economy 130*(3), 566–635.

Schlom, C. (2024). The Concavification Method, Applied to Affirmative Action via Random Admissions. *In preparation*.