Markets for Goods with Externalities*

Zi Yang Kang†

This version: April 2020

Abstract

I consider the welfare and profit maximization problems in markets with externalities. I show that when externalities depend generally on allocation, a Pigouvian tax is often suboptimal. Instead, the optimal mechanism has a simple form: a finite menu of rationing options with corresponding prices. I derive sufficient conditions for a single price to be optimal. I show that a monopolist may ration *less* relative to a social planner when externalities are present, in contrast to the standard intuition that non-competitive pricing is indicative of market power. My characterization of optimal mechanisms uses a new methodological tool—the *constrained maximum principle*—which leverages the combined mathematical theorems of Bauer (1958) and Szapiel (1975). This tool generalizes the concavification technique of Aumann and Maschler (1995) and Kamenica and Gentzkow (2011), and has broad applications in economics.

JEL classification: C61, D47, D61, D62, D82

Keywords: externalities, Pigouvian tax, mechanism design, constrained maximum principle, nonconvexities

1 Introduction

Consider a market for a vaccine that is costly to produce. Consumers differ in two ways: their value for the vaccine and the amount of social contact that they maintain. Unvaccinated consumers risk catching a virus, which spreads with a probability that is proportional to their amount of social contact. Healthcare costs, paid for equally by all consumers in the market, increase with the number of consumers that catch the virus.

* Acknowledgements will be added.
† Stanford Graduate School of Business, 655 Knight Way, Stanford, CA 94305; zykang@stanford.edu.

If both the vaccine and the amount of social contact could be jointly priced, the competitive equilibrium is easily seen to be efficient. But the analysis is no longer straightforward when only the vaccine is sold and no market exists for social contact. By choosing to remain unvaccinated, a consumer imposes a negative externality on all other consumers. What then is the efficient outcome? How can it be implemented?

These questions are of fundamental economic importance and arise in numerous other settings. Many environmental challenges that society faces today, such as climate change, pollution and resource depletion, are intimately connected to transportation, energy production, heavy industries, and the negative externalities that they impose on the environment. Internet platforms, education, financial services and public health are but a few more examples of markets for goods with externalities—both positive and negative.

Many economics textbooks emphasize the role of the Pigouvian tax, under which consumers internalize the externalities that their actions impose on others. To see how this might work, suppose that consumers' values for the vaccine are independent of the amount of social contact that they maintain. Then the magnitude of healthcare costs depends only on the number of unvaccinated consumers: vaccination status, under *any* pricing scheme, is completely uninformative about the amount of social contact a consumer maintains. In the Pigouvian tradition, the sum of healthcare costs and the cost of producing the vaccine represent the *social cost*. Crucially, the marginal social cost is, by construction, higher than marginal production cost; and the efficient outcome occurs where marginal social cost intersects demand, which represents the marginal social benefit. Consequently, implementing a tax equal to the difference between marginal social cost and marginal production cost, where the marginal social cost intersects demand, achieves efficiency.

However, in more general settings, a Pigouvian tax fails to induce an efficient outcome. Critically, the analysis above leans on the assumption that the marginal social cost curve depends on allocation through *only* the total number of unvaccinated consumers. This is no longer true when consumers' values for the vaccine can be correlated with their amount of social contact; correlation breaks the optimality of a Pigouvian tax. For instance, if consumers' values and amount of social contact are perfectly *anti-correlated*, any Pigouvian tax would induce the consumers with the highest values for the vaccine (and hence lowest amounts of social contact) to get vaccinated, thereby leaving unvaccinated the most socially costly consumers.

In this paper, I study a large market for an indivisible good with externalities that can depend on allocation in a general way. Each consumer has quasilinear preferences characterized by a multidimensional vector, consisting of the consumer's unidimensional value for the good and a multidimensional vector of values for different externalities. In my analysis, I allow for both *conditional* and *unconditional* externalities. A conditional externality is an externality that the consumer experiences only if they are allocated the good, such as congestion in the market for vehicle ownership.[1] By contrast, an unconditionality externality is experienced by all consumers regardless of their own allocation, such as pollution in the market of vehicle ownership, or healthcare costs in the earlier example of vaccines.

A key assumption of my model is that externalities are *aggregate*: First, changes in allocation for any small set of consumers should not change the magnitude of externalities. Second, externalities depend on allocation through a potentially large, but finite, number of *moments*. These moments can be thought of as "sufficient statistics" of the allocation that determine the externalities.[2] In the vaccine example, the single moment characterizing the externality is the total amount of social contact by unvaccinated consumers. I discuss the implications of this assumption when I introduce the framework in Section 2.

To determine the efficient outcome, I formulate the allocation problem as one of optimal mechanism design. A market designer chooses a mechanism so as to maximize social welfare, subject to individual rationality. Incentive compatibility is also required: while the designer knows the joint distribution of consumer preferences, individual preferences are not observed and have to be inferred through the mechanism. Finally, I allow the designer to subsidize the market up to a predetermined amount, which introduces a budget constraint.

In this general framework, there are two main difficulties. First, unlike typical mechanism design problems, the designer's objective function can depend nonlinearly on allocation due to externalities. Because very few restrictions are placed on precisely how externalities depend on allocation moments, the ironing technique of Myerson (1981) has very little traction here. Second, consumers have multidimensional preferences, which complicate the characterization of incentive compatibility.

Despite these difficulties, and perhaps surprisingly, optimal mechanisms have a simple structure. I show that, in addition to offering a price for the good, the designer also provides a finite set of *rationing options*, each at a different price: consumers pick from these options to receive the good with some probability strictly between 0 and 1. Interestingly, the number of rationing options in the optimal mechanism is bounded above by the number of moments that

[1] I discuss the market for vehicle ownership as an example in Section 6. Other examples are also provided there, as well as at the end of Section 2.

[2] For instance, price theory (see, *e.g.*, Weyl, 2019) typically assumes that total quantity allocated is a sufficient moment that characterizes allocation.

characterize the externalities, rather than the number of externalities. In the vaccine example, this implies that, even without further specification of how consumers' values and their amount of social contact are correlated, or what the social cost function is, the optimal mechanism is characterized by at most one rationing point and a price.

The optimal mechanism suggests the following appealing way of viewing the design problem, which has important empirical implications.[3] To address nonconvexities in the allocation problem, the designer need only contend with the *moments* of allocation that these nonconvexities depend on. The optimal mechanism *prices* these moments. In empirical settings, this implies that counterfactual analysis can be carried out as long as these moments can be accurately measured.

To derive the optimal mechanism, I overcome the two difficulties as follows. First, I employ the simple but useful observation that the designer can be thought of as choosing the magnitude of allocation moments (which determine externalities) before selecting a mechanism.[4] Conditional on the magnitude of allocation moments, the resulting problem is linear in allocation. Second, I show that even though the resulting problem still has multidimensional types, types can be projected onto a single unidimensional "effective type." In principle, the effective type can depend on allocation; but this dependence can only be through the allocation moments, which are held fixed.

While this argument solves the difficulties of nonlinearity and multidimensionality, it also introduces a third difficulty: the resulting problem is *constrained* by the magnitude of allocation moments chosen by the designer. To overcome this last difficulty, I introduce a methodological technique based on the mathematical results of Bauer (1958) and Szapiel (1975). This technique, which I call the "constrained maximum principle," includes as special cases concavification[5] and Lagrangian duality methods developed for constrained information design.[6] Using the constrained maximum principle, the solution to the constrained mechanism design problem can be found by solving a finite-dimensional problem, and can be shown to have the simple form described above.

The constrained maximum principle has valuable applications to various other problems in economics, and is of independent interest. For this reason, I separately present the constrained maximum principle in Section 3. I also include in Section 7 an overview of some results in the literature for which application of the constrained maximum principle greatly simplifies proofs.

Having characterized the optimal mechanism in Section 4, I then examine when the optimal

[3] I am grateful to Shosh Vasserman for this observation.
[4] This implements in a mechanism design framework the price-theoretic trick of simplifying problems via choosing quantities rather than prices as in Weitzman (1974), which was first taught to me by Jeremy Bulow.
[5] See Aumann and Maschler (1995) and Kamenica and Gentzkow (2011).
[6] See, for example, Doval and Skreta (2018) and references therein.

mechanism can be implemented via a Pigouvian tax in Section 5. I find that, when externalities and costs depend on allocation through either total quantity allocated (as in the vaccine example with values independent of social contact) or total value allocated, a Pigouvian tax attains the efficient outcome when the designer has a non-binding budget constraint.[7] By contrast, when externalities and costs depend on allocation through both total quantity and total value allocated (as in the vaccine example with perfect anti-correlation), this does not need to be the case: depending on the magnitude of values relative to costs, even a pure lottery—in which the designer sets only a rationing option, with no (or an arbitrarily high) price—can be optimal.

While I have described results concerning efficiency so far, my analysis also extends to the case of a profit-maximizing monopolist. In a setting with only unconditional externalities (as in the vaccine example), under regularity conditions on demand similar to those in Myerson (1981), the profit-maximizing mechanism sets a single price for the good with no rationing option. Yet the welfare-maximizing mechanism in the same setting could involve as many rationing options as the number of moments through which unconditional externalities depend on allocation. This stands in stark contrast to standard intuition that non-competitive pricing is an indicator of market power in the absence of externalities: here, inefficiencies arise even with perfect competition, precisely because of the *failure* to ration.

The results of this paper have several policy implications. Pigouvian taxation is an important tool in policymaking;[8] this paper extends that tool for a large family of problems. The form of the optimal policy is simple and relies on empirically measurable moments. These have both positive and normative consequences: while the analysis of this paper may help explain the use of price controls and rationing in markets with externalities, it also provides guidance on how empirical measurement can be leveraged for optimal design. I offer a more detailed discussion in Section 6.

The works most closely related to this paper are Condorelli (2013) and Dworczak, Kominers, and Akbarpour (2019). Condorelli (2013) studies an allocation problem for which the designer's objective differs from the agents' willingness to pay. As Condorelli (2013) shows, non-market mechanisms can be optimal due to the non-monotonicity of the "effective demand curve" that takes into account the designer's preferences over allocation. Through the lens of my analysis, the designer's preferences over allocation can be seen as an additional "moment" that the designer chooses when determining the optimal mechanism, thereby introducing the possibility of rationing.

[7] When the budget constraint binds, Samuelson (1984) shows that rationing can be optimal in a setting with adverse selection, even when no externality is present.

[8] For instance, Baumol and Oates (1988) discuss the relevance of Pigouvian taxes in the design of environmental policy.

Dworczak, Kominers, and Akbarpour (2019) analyze a two-sided market for an indivisible good, where agents differ in their values for the good and their marginal utilities of money. They find that a welfare-maximizing designer uses rationing to optimally effect redistribution in the market. While Dworczak, Kominers, and Akbarpour (2019) use concavification techniques in their analysis, the constrained maximum principle demonstrates how their results generalize: every allocation moment that the designer chooses induces the possibility of an additional rationing option.

Finally, this paper relates to extensive literatures on externalities and network goods. Recent works include Sandholm (2005) and Sandholm (2007) on negative externalities; and Rochet and Tirole (2006), Weyl (2010) and Veiga, Weyl, and White (2017) on platform design.[9] The results of this paper are also connected to studies of non-market mechanisms, notably by Weitzman (1977), Bulow and Klemperer (2012) and Che, Gale, and Kim (2013).

2 Framework

2.1 Setup

There is a unit mass of risk-neutral consumers in a market for an indivisible good with externalities. There are two classes of externalities, namely: *conditional* externalities, which a consumer experiences only if they are allocated a good; and *unconditional* externalities, which all consumers experience regardless of their own allocation.

Each consumer has unit demand for the good. An individual consumer is characterized by a multidimensional *type* (θ, η, ζ), consisting of:

(i) a value $\theta \in [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$ for the good itself;

(ii) a k-dimensional vector $\eta = (\eta_1, \ldots, \eta_k) \in [\underline{\eta}_1, \overline{\eta}_1] \times \cdots \times [\underline{\eta}_k, \overline{\eta}_k] \equiv [\underline{\eta}, \overline{\eta}] \subset \mathbb{R}^k$ for the conditional externalities that the good induces; and

(iii) an ℓ-dimensional vector $\zeta = (\zeta_1, \ldots, \zeta_\ell) \in [\underline{\zeta}_1, \overline{\zeta}_1] \times \cdots \times [\underline{\zeta}_\ell, \overline{\zeta}_\ell] \equiv [\underline{\zeta}, \overline{\zeta}] \subset \mathbb{R}^\ell$ for the unconditional externalities that the good induces.

For notational succinctness, I denote the type space by $\mathscr{T} = [\underline{\theta}, \overline{\theta}] \times [\underline{\eta}, \overline{\eta}] \times [\underline{\zeta}, \overline{\zeta}]$.

Consumer preferences are linear in allocation, payment and externalities: conditional on receiving an allocation X of the good for a payment T, and experiencing conditional

[9] In a follow-up paper (Kang and Vasserman, 2020), Shosh Vasserman and I show how the techniques of this paper extend to yield novel insights on platforms and multi-sided markets.

externalities $W = (W_1, \ldots, W_k)$ and unconditional externalities $Z = (Z_1, \ldots, Z_\ell)$, a type-$(\theta, \eta, \zeta)$ consumer realizes a payoff of

$$\theta X + (\eta \cdot W)\, X + \zeta \cdot Z - T.$$

Consumer types are distributed according to an absolutely continuous distribution $F(\theta, \eta, \zeta)$ with positive density $f(\theta, \eta, \zeta)$ everywhere. While the designer does not observe individual realizations of types, the designer knows the distribution $F(\theta, \eta, \zeta)$.

I consider three different objectives for the designer, namely: (i) the *Pigouvian objective* of maximizing welfare subject to a budget constraint; (ii) the *monopolist objective* of maximizing profit; and (iii) the *Ramsey objective* of maximizing welfare subject to a minimum profit constraint. I defer the formalism of these objectives to Section 4.

For any given objective, the designer chooses a *mechanism* (x, t), consisting of an *allocation function* $x : \mathscr{T} \to [0, 1]$ and a *payment function* $t : \mathscr{T} \to \mathbb{R}$. The choice of mechanism determines externalities. To formalize externalities in this setting, I introduce the following definition:

Definition 1. Let $\mathcal{X}$ denote the space of all allocation functions $x : \mathscr{T} \to [0, 1]$ and endow $\mathcal{X}$ with the L^1 topology. A vector-valued function $H : \mathcal{X} \to \mathbb{R}^s$ is an *m-aggregate* function if:

(i) $H(x_1) = H(x_2)$ for any $x_1, x_2 \in \mathcal{X}$ such that $x_1 = x_2$ almost everywhere;[10] and

(ii) there exist continuous affine functions[11] $\phi_1, \ldots, \phi_m : \mathcal{X} \to \mathbb{R}$ and an arbitrary function $\widetilde{H} : \mathbb{R}^m \to \mathbb{R}^s$ such that

$$H(x) = \widetilde{H}(\phi_1(x), \ldots, \phi_m(x)) \quad \text{for any } x \in \mathcal{X}.$$

A function is *aggregate* if it is m-aggregate for some m. For an aggregate function, if there exists a representation such that $\widetilde{H}$ in condition (ii) is upper-semicontinuous (resp. lower-semicontinuous), then the function is said to be *upper-semicontinuously* (resp. *lower-semicontinuously*) *aggregate*. A function that is both upper- and lower-semicontinuously aggregate is *continuously aggregate*.

In the statement of Definition 1, observe that condition (i) requires any aggregate function to remain unchanged by modifications to allocation of measure zero,[12] while condition (ii) requires

[10] Here, and throughout this paper, I will take the underlying measure to be the Lebesgue measure.

[11] Since $\mathcal{X}$ equipped with the L^1 norm is a normed linear (in fact, Banach) space, this assumption is equivalent to the assumption that $\phi_1, \ldots, \phi_m : \mathcal{X} \to \mathbb{R}$ are bounded affine functions.

[12] To see that condition (i) is not implied by condition (ii), note that $H(x) = x(\underline{\theta}, \underline{\eta}, \underline{\zeta})$ is linear in allocation and so satisfies condition (ii), but clearly fails condition (i).

any aggregate function to depend only on a finite number of moments of the allocation; this dependence, however, may be arbitrary.[13] I give a more detailed interpretation of Definition 1 in the next subsection after describing the rest of the model.

Given Definition 1, *externalities* are defined by a pair (w, z) of continuously aggregate vector-valued functions,[14] where $w : \mathcal{X} \to \mathbb{R}^k_+$ is a *conditional externality function* and $z : \mathcal{X} \to \mathbb{R}^\ell_+$ is an *unconditional externality function*.[15]

By the revelation principle, restriction to direct mechanisms—so that each consumer truthfully reports their type—is without loss of generality; hence the mechanism is subject to *incentive compatibility* and *individual rationality* constraints. That is, for any (θ, η, ζ) and $(\widehat{\theta}, \widehat{\eta}, \widehat{\zeta})$:

$$[\theta + \eta \cdot w(x)]\, x(\theta, \eta, \zeta) - t(\theta, \eta, \zeta) \geq [\theta + \eta \cdot w(x)]\, x(\widehat{\theta}, \widehat{\eta}, \widehat{\zeta}) - t(\widehat{\theta}, \widehat{\eta}, \widehat{\zeta}), \tag{IC}$$

$$[\theta + \eta \cdot w(x)]\, x(\theta, \eta, \zeta) - t(\theta, \eta, \zeta) \geq 0. \tag{IR}$$

Because externalities are aggregate, their magnitudes are unaffected by individual misreporting. Consequently, unconditional externalities play no role in (IC). Moreover, note that (IR) requires each consumer to receive no less than what they otherwise would from the unconditional externalities alone. In particular, this prevents the designer from extracting payments from consumers even when they are not allocated the good.

Just as externalities are endogenous to the designer's choice of mechanism, production costs for the good can also depend on the allocation. The *cost function* $c : \mathcal{X} \to \mathbb{R}$ is assumed to be lower-semicontinuously aggregate.[16]

[13] Given that this dependence may be arbitrary and that the assumption of affine dependence is imposed on allocations rather than types, Definition 1 is less restrictive than may first appear. For example, in many economic settings, it is reasonable to consider 1-aggregate functions that depend on the allocation through only the total quantity allocated or the total value of consumers that the allocation achieves.

[14] While continuity guarantees the existence of an optimal mechanism, it is not necessary and is assumed only for expositional ease. A sufficient condition is that the designer's objective is upper semi-continuous in allocation.

[15] Because types are permitted to be negative, restriction of each component of the externality functions to the nonnegative reals is without loss of generality.

[16] This assumption permits, for example, fixed costs and type-dependent marginal costs. Type-dependent marginal costs are important for applications to settings with adverse selection such as health insurance (see, *e.g.*, Einav and Finkelstein, 2011). Lower-semicontinuity allows for capacity constraints in production to be incorporated in the cost function.

2.2 Examples

To my knowledge, the notion of "aggregate" functions described by Definition 1 is new.[17] To illustrate what the notion of "aggregate" in the context of externalities and costs captures, I provide four examples.

Aggregate cost functions

The first two examples focus on familiar models *without* externalities. Both examples, however, have aggregate cost functions, and can be seen as special cases of the framework presented above.

Example 1 (Myerson, 1981; Bulow and Roberts, 1989). Consider a monopolist selling a good to a unit continuum of consumers with unit demand. Consumers have heterogeneous values θ drawn from a distribution F.

- When the monopolist has a constant marginal cost c_0, then the cost function is

$$c(x) = c_0 \int_{\underline{\theta}}^{\overline{\theta}} x(\theta) \, \mathrm{d}\theta.$$

  The total quantity of the good sold is itself a continuous[18] linear function of allocation:

$$\phi_1(x) = \int_{\underline{\theta}}^{\overline{\theta}} x(\theta) \, \mathrm{d}\theta.$$

  The function $\widetilde{H}$ in condition (ii) of Definition 1 can be taken to be "multiplication by c_0" here, with total quantity allocated $\phi_1(x)$ as the single moment. Moreover, $c(x)$ does not change when the allocation $x(\theta)$ changes on a set of measure zero. Hence $c(x)$ is 1-aggregate, and in fact *continuously* 1-aggregate.

- When the monopolist faces a capacity constraint of $C > 0$ with zero marginal cost, then the

---

[17] However, it should be noted that the notion of "aggregate" functions described by Definition 1 shares much in common with the notion of *aggregative* games (surveyed by Jensen, 2018), in which agents' payoffs depend on the actions of other agents through a single statistic of their actions. For example, the Cournot oligopoly game with linear demand is an aggregative game as each firm's payoff depends on only its own choice of quantity and the *sum* of all other firms' quantities.

[18] Continuity follows immediately from the dominated convergence theorem.

cost function is[19]

$$c(x) = \begin{cases} 0 & \text{if } \int_{\underline{\theta}}^{\overline{\theta}} x(\theta)\ \mathrm{d}\theta \leq C, \\ +\infty & \text{otherwise.} \end{cases}$$

Here, $\phi_1(x)$ can be taken as a single moment, while $\widetilde{H}$ in condition (ii) of Definition 1 can be defined by $\widetilde{H}(\widehat{\phi}_1) = 0$ if $\widehat{\phi}_1 \leq C$, and $+\infty$ otherwise. As before, $c(x)$ is 1-aggregate; however, it is only lower-semicontinuously (and not upper-semicontinuously) aggregate.

Example 2 (Samuelson, 1984). Suppose that a monopolist sells insurance to a unit continuum of consumers with demand. There is adverse selection: consumer risks are indexed by θ, distributed according to F; and the marginal cost of serving a consumer of risk θ is $\gamma(\theta)$. In this case, the cost function can be written as

$$c(x) = \int_{\underline{\theta}}^{\overline{\theta}} \gamma(\theta)x(\theta)\ \mathrm{d}\theta.$$

Here, $c(x)$ itself is single moment that is continuous and linear in allocation; and $\widetilde{H}$ in condition (ii) of Definition 1 is the identity function. Indeed, $c(x)$ is continuously 1-aggregate.

Aggregate externality and cost functions

I now introduce two other examples—with aggregate externality and cost functions—that are accommodated by my framework.[20]

Example 3 (social media). Consider a market for a network good (*e.g.*, a social media account) supplied at zero marginal cost. Consumers have an inherent value for the good, benefit (conditional on their own usage) from having high-value consumers on the network, but (unconditional on their usage) experience a loss in emotional connection that depends on the size of the network. Consumer preferences can thus be captured by a vector $(\theta, \eta, \zeta) \in \mathbb{R}^3$, with equilibrium utilities

[19] Technically, $+\infty$ is not permitted as a value of $c(x)$, but it can be replaced by an arbitrarily large cost; the framework can also easily accommodate functions that map to the extended real line at the expense of more unwieldy notation.

[20] I discuss further practical examples and implications for policy in Section 6.

given by

$$\theta x(\theta,\eta,\zeta) + \eta \underbrace{\widetilde{H}_1\left(\int_{\mathscr{T}} \theta' x(\theta',\eta',\zeta')\, \mathrm{d}F(\theta',\eta',\zeta')\right)}_{\text{benefit from high-value consumers on network, } w(x)} x(\theta,\eta,\zeta)$$
$$+\, \zeta \underbrace{\widetilde{H}_2\left(\int_{\mathscr{T}} x(\theta',\eta',\zeta')\, \mathrm{d}F(\theta',\eta',\zeta')\right)}_{\text{loss in emotional connection, } z(x)} - t(\theta,\eta,\zeta).$$

With $\widetilde{H}_1$ and $\widetilde{H}_2$ given by *any* continuous function, the conditional and unconditional externality functions, $w(x)$ and $z(x)$, are each continuously 1-aggregate.

Example 4 (vaccines). To illustrate how the example in the introduction can be formulated in my framework, consider a market for vaccines. Each vaccine is produced at a constant marginal cost, which yields a continuous 1-aggregate cost function (cf. Example 1). Consumers are distinguished by their value for vaccines θ and their amount of social interaction η. Their equilibrium utilities are given by, for some continuous function $\widetilde{H}$,

$$\theta x(\theta,\eta) + \underbrace{\widetilde{H}\left(\int_{\mathscr{T}} \eta' x(\theta',\eta')\, \mathrm{d}F(\theta',\eta')\right)}_{\text{healthcare costs arising from total social interaction}} - t(\theta,\eta).$$

To reconcile this with my framework, I make two observations. First, the conditional externality function $w(x)$ in this example can be thought of as a real-valued function identically equal to zero. Second, consumers can be thought of as having a preference $\zeta \equiv 1$ for the unconditional externality. Here, the unconditional externality is a continuous 1-aggregate function.

3 The constrained maximum principle

In this section, I introduce two theorems due to Bauer (1958) and Szapiel (1975). By combining both results, the *constrained maximum principle* describes how the solution of a constrained maximization problem relates to the solution of the unconstrained maximization problem. The constrained maximum principle has broad applications to other problems in economics. As it is of independent interest, it is presented separately here.

Throughout this section, V will denote a locally convex Hausdorff space, and $K \subset V$ will denote a compact convex subset. An *extreme point* of a convex set $S \subset V$ is any point $x \in S$ that cannot be represented as a strict convex combination $x = \lambda x_1 + (1-\lambda)\,x_2$, $0 < \lambda < 1$, of two distinct points $x_1, x_2 \in S$. The set of extreme points of S is written as $\operatorname{ex} S$.

3.1 Bauer's maximum principle

Theorem (Bauer's maximum principle). Let $\Omega : K \to \mathbb{R}$ be a continuous quasiconvex function.[21] Then Ω has a maximizer on K that is an extreme point of K:

$$\max_{x \in K} \Omega(x) = \max_{x \in \operatorname{ex} K} \Omega(x).$$

The Bauer maximum principle implies that restriction to the extreme points of K is without loss of generality, for the purpose of maximizing a continuous and quasiconvex objective function. The advantage of doing so is that the extreme points of K are often relatively straightforward to characterize, as the result of Szapiel (1975) below shows. A familiar special case of Bauer's maximum principle is the well-known result in linear programming that the extremal values of a linear function over a compact convex set are attained at extreme points.

3.2 Extreme points of the constraint set

Consider a compact convex "constraint set" $C \subset K$, and suppose the set $\operatorname{ex} K$ is known. Szapiel's (1975) theorem provides a way of characterizing each point in $\operatorname{ex} C$ as a convex combination of points in $\operatorname{ex} K$.[22]

Theorem (Szapiel, 1975). Let $\Lambda : K \to \mathbb{R}^n$ be an affine function. Suppose that $C = \Lambda^{-1}(\Sigma)$ for some convex set $\Sigma \subset \operatorname{im} \Lambda$. Then:

(i) C is a convex set; and

(ii) any extreme point $x \in \operatorname{ex} C$ can be expressed as a convex combination of at most $n+1$ extreme points of K.

[21] The original version of the Bauer maximum principle, shown by Bauer (1958), required Ω to be convex. Note that continuity can be relaxed to upper-semicontinuity, but quasiconvexity would have to be strengthened to *explicit* quasiconvexity, so that Ω is quasiconvex and, in addition, $\Omega(x) < \Omega(y)$ implies $\Omega(\lambda x + (1-\lambda)\,y) < \Omega(y)$ for any $\lambda \in (0,1)$ and $x, y \in K$; see Corollary 7.75 in Aliprantis and Border (2006).

[22] I am grateful to a referee of Kang and Vondrák (2019) for the reference to Winkler (1988), which led me to discover the work of Szapiel (1975). While Szapiel (1975) studies compact convex subsets of locally convex topological spaces, Winkler's (1988) generalization shows that Szapiel's (1975) theorem does not, in fact, rely on the underlying topological structure of the space. Nonetheless, in the setting here, the topological structure of the space has to be used for the optimization problem considered in the constrained maximum principle.

Szapiel's (1975) theorem can be interpreted as a dual to Carathéodory's theorem.[23] Carathéodory's theorem states that at most $n+1$ points in S are required to represent any point of a set A as a convex combination, provided that A is the *convex hull* of S. By contrast, Szapiel's (1975) theorem asserts that at most $n+1$ points of $\operatorname{ex} K$ are required to represent any point of $\operatorname{ex} C$ as a convex combination, provided that C is a *section* of K.

It should be noted that the class of constraints allowed by Szapiel's (1975) theorem is fairly general. They include, for example, *equality* constraints of the form $g(x) = \gamma$, where $g : K \to \mathbb{R}$ is affine and $\gamma \in \mathbb{R}$; and *inequality* constraints of the form $h(x) \leq \eta$, where $h : K \to \mathbb{R}$ is affine and $\eta \in \mathbb{R}$.

3.3 Constrained maximum principle

Combined with Szapiel's (1975) theorem, the Bauer maximum principle immediately yields the following *constrained maximum principle*:

Theorem 1 (constrained maximum principle). Let $\Lambda : K \to \mathbb{R}^n$ be an affine function. Suppose that $C = \Lambda^{-1}(\Sigma)$ for some convex set $\Sigma \subset \operatorname{im} \Lambda$. Let $\Omega : K \to \mathbb{R}$ be continuous and quasiconvex. Then Ω has a maximizer on C that is the convex combination of no more than $n+1$ extreme points of K.

The constrained maximum principle relates the solution of the n-dimensional constrained problem, $\max_{x \in C} \Omega(x)$, to the solution of the unconstrained problem, $\max_{x \in K} \Omega(x)$. In particular, the solution of the n-dimensional constrained problem can be viewed as a *randomization* over at most $n+1$ (possibly infeasible) points in K. Crucially, every additional constraint introduces a new level of randomization to the optimal solution.

3.4 Relation to existing methods[24]

The constrained maximum principle has widespread applications, such as in mechanism design, information design and robustness analysis. Economists have developed and extended tools such as ironing, linear programming, optimal control and concavification to solve various problems in these fields. The constrained maximum principle adds to the set of available tools.

[23] See, for example, the discussion in Sections III.9, IV.1 and IV.3 of Barvinok (2002).

[24] I discuss in Section 7.2 how the constrained maximum principle can be applied to problems in mechanism design, information design and robust analysis. Here, I focus on the relationship between the constrained maximum principle to existing methods.

Traditional tools such as ironing (as in Myerson, 1981) and concavification (as in Aumann and Maschler, 1995; and Kamenica and Gentzkow, 2011) were initially developed as tools to solve infinite-dimensional linear programs. In the case of ironing, later works have extended it to convex programs as well, which maximize concave objective functions over convex sets (as in Toikka, 2011). In the case of concavification, later works have extended it to equality and inequality constraints via Lagrangian duality methods (as in Le Treust and Tomala, 2019; and Doval and Skreta, 2018.)

By contrast, the constrained maximum principle applies to the maximization of a *quasiconvex* objective function over a convex set. Consequently, the constrained maximum principle is comparable to ironing and concavification for linear programs.

Yet, even in this case, the approaches differ slightly. Denoting the linear objective function by $\Omega(x) = \langle \alpha, x \rangle$,[25] observe that, through ironing or concavification, solutions to the (possibly constrained) maximization problem are inferred by solving a "relaxed" problem, $\max_x \overline{\Omega}(x) = \langle \overline{\alpha}, x \rangle$. Importantly, ironing and concavification modify α, the dual to x, to infer properties about the maximizer x^*. By contrast, the constrained maximum principle directly asserts properties about the maximizer x^*. In this sense, ironing and concavification can both be viewed as dual methods to the constrained maximum principle.

While the family of objective functions that the constrained maximum principle applies to is much broader, the constrained maximum principle can already be viewed as a generalization of concavification techniques even in the case of linear objective functions: the space of priors (*i.e.*, the simplex of probability measures over a compact set) is generalized to any compact convex subset of a locally convex Hausdorff space, and the class of constraints allowed is more permissive.

Finally, recent remarkable and independent contributions by Arieli, Babichenko, Smorodinsky, and Yamashita (2020) and Kleiner, Moldovanu, and Strack (2020) consider a different version of the constrained maximum principle, where moment constraints are replaced with a majorization constraint. By characterizing the extreme points of the constraint set in that setting, they derive the solution to a class of constrained optimization problems that are complementary to those I consider in this paper. Interestingly, that version of the constrained maximum principle has, as they show, a wide range of economic applications as well.

[25] Here, $\langle \cdot, \cdot \rangle$ is a bilinear form that defines a duality between two vector spaces. In the case of mechanism design with linear utility, this can be defined by $\langle \alpha, x \rangle = \int \alpha(\theta) x(\theta) \; \mathrm{d}F(\theta)$, where F is a distribution over types θ with a fully supported density over the compact convex type space $\Theta \subset \mathbb{R}^d$. In the case of Bayesian persuasion, this can be defined by $\langle \alpha, x \rangle = \int \alpha(\mu) \; \mathrm{d}x(\mu)$, where $\mu \in \Delta(\Omega)$ is a prior over outcomes in a finite set Ω.

4 Optimal mechanisms

I now return to the framework of Section 2 and derive the welfare-maximizing mechanism, subject to a budget constraint. I show that the welfare-maximizing mechanism generally requires *rationing* at different prices; hence a Pigouvian tax typically fails to achieve an efficient outcome.

In addition to determining the efficient outcome, understanding the potential impact of market power on allocation is also important for policymaking. Towards this goal, I complement the analysis above by also deriving the profit-maximizing mechanism.

The key arguments of all results in this section are sketched out under a unified framework in Section 4.3. Technical details of the proofs are left to Appendix A.

4.1 Welfare objective

Under the Pigouvian framework, the designer's goal is to maximize utilitarian welfare, given a budget $B \geq c(0)$.[26] That is, the designer maximizes:

$$\Omega^{\text{welfare}}(x,t) \equiv \int_{\mathscr{T}} \{[\theta + \eta \cdot w(x)]\, x(\theta,\eta,\zeta) + \zeta \cdot z(x)\} \;\mathrm{d}F(\theta,\eta,\zeta) - c(x),$$

subject to the constraints (IC), (IR) and

$$B + \int_{\mathscr{T}} t(\theta,\eta,\zeta)\;\mathrm{d}F(\theta,\eta,\zeta) \geq c(x). \tag{BB}$$

The main result of this section can now be stated:

Theorem 2. Suppose that the *externality-cost function* $H^{\text{welfare}}(x) = (w(x), z(x), c(x))$ is m-aggregate. Then the welfare-maximizing mechanism $(x^{\text{welfare}}, t^{\text{welfare}})$ satisfies

$$\operatorname{im} x^{\text{welfare}} \subset \{0, q_1, \ldots, q_{m+1}, 1\} \quad \text{for some } 0 < q_1 < \cdots < q_m < 1.$$

Moreover, if (BB) does not bind, then

$$\operatorname{im} x^{\text{welfare}} \subset \{0, q_1, \ldots, q_m, 1\} \quad \text{for some } 0 < q_1 < \cdots < q_m < 1.$$

The payment function t^{welfare} can be determined from the allocation function via the Milgrom

[26] This inequality ensures that non-allocation is always feasible for the designer. Note that "0" here denotes the allocation function $x(\theta,\eta,\zeta) = 0$ for any (θ,η,ζ).

and Segal (2002) envelope theorem by assuming that (IR) binds for the lowest type $(\underline{\theta}, \underline{\eta}, \underline{\zeta})$. Therefore, Theorem 2 completely characterizes the welfare-maximizing mechanism.

Theorem 2 demonstrates that the welfare-maximizing mechanism generally entails rationing with different probabilities, each at a distinct price. In other words, externalities are generally only partially internalized. The number of rationing options depend on the number of sufficient statistics (cf. Definition 1) through which externalities and costs depend on allocation. In particular, a Pigouvian tax—which specifies a single price—is generally suboptimal.

Two familiar special cases of Theorem 2 should be noted. First, when $m = 0$ with no budget constraint, Theorem 2 reduces to the optimality of a single price. Second, even in the absence of externalities, the welfare-maximizing mechanism can involve rationing due to the interaction between the cost function and the budget constraint. An important example is the adverse selection setting considered by Samuelson (1984), for which rationing can be welfare-optimal if the budget constraint binds.

4.2 Profit objective

By contrast, a profit-maximizing designer does not face a budget constraint. Instead, the designer maximizes profit, defined by

$$\Omega^{\text{profit}}(x, t) \equiv \int_{\mathscr{T}} t(\theta, \eta, \zeta) \, \mathrm{d}F(\theta, \eta, \zeta) - c(x),$$

subject to the constraints (IC) and (IR). The profit-maximizing mechanism is characterized by:

Theorem 3. Suppose that the *externality-cost function* $H^{\text{profit}}(x) = (w(x), c(x))$ is m-aggregate. Then the profit-maximizing mechanism $(x^{\text{profit}}, t^{\text{profit}})$ satisfies

$$\operatorname{im} x^{\text{profit}} \subset \{0, q_1, \dots, q_m, 1\} \quad \text{for some } 0 < q_1 < \cdots < q_m < 1.$$

The only difference between the profit-maximizing mechanism and the welfare-maximizing mechanism (cf. Theorem 2) is the number of rationing points in the allocation function. Unlike the welfare-maximizing mechanism, the profit-maximizing mechanism fails to even partially internalize unconditional externalities. Rationing arises only because of conditional externalities and costs.

An important special case of Theorem 3 is the classic result of Myerson (1981), whose setting can be interpreted through the lens of Bulow and Roberts (1989) as a capacity-constrained monopolist's problem. Capacity constraints can be incorporated by a 1-aggregate cost function specifying zero cost if total quantity is below capacity, and arbitrarily large cost if total quantity exceeds capacity. In this setting, Theorem 3 implies that rationing can be optimal; this is the case when the maximum capacity is met at an "ironed" portion of the marginal revenue curve.

4.3 Deriving optimal mechanisms

I now explain the key ideas behind Theorems 2 and 3. The generality of the framework in Section 2 presents three challenges. First, the designer's objective functions are nonlinear in allocation due to externalities and costs. Second, heterogeneity in consumer preferences over externalities result in the multidimensionality of types, which makes characterization of (IC) difficult. Third, the designer must account for constraints such as (BB) that generally depend on the mechanism.

To address these three challenges, I outline three key ideas. Throughout this proof sketch, I denote the designer's objective function and externality-cost function by $\Omega \in \{\Omega^{\text{welfare}}, \Omega^{\text{profit}}\}$ and $H \in \{H^{\text{welfare}}, H^{\text{profit}}\}$ respectively. I also include (BB) as a constraint, with the understanding that (BB) does not bind if $\Omega = \Omega^{\text{profit}}$.

To simplify this proof sketch, I will assume that the optimal mechanisms (for both welfare maximization and profit maximization) exist. Existence is verified in Appendix A.

Key idea #1: Formulating externalities and costs as constraints

I begin by addressing nonlinearity of Ω in allocation. Crucially, the dependence of Ω on allocation can be linearized by expressing externalities and costs as *constraints*. Since the externality-cost function $H(x)$ is m-aggregate by assumption, there exist m affine functions $\phi_1, \ldots, \phi_m$ and an arbitrary function $\widetilde{H}$ such that

$$H(x) = \widetilde{H}\left(\phi_1(x), \ldots, \phi_m(x)\right).$$

Then the designer's problem can be written as a nested optimization problem:

$$\begin{aligned}&\max_{(x,t)} \{\Omega(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\}\\ &= \max_{\widehat{\phi}_1,\ldots,\widehat{\phi}_m} \left[\max_{(x,t):(\phi_1(x),\ldots,\phi_m(x))=(\widehat{\phi}_1,\ldots,\widehat{\phi}_m)} \{\Omega(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\} \right].\end{aligned}$$

The inner maximization problem can be viewed as the designer's problem holding fixed externalities and costs—or, more precisely, holding fixed the *moments* of allocation that determine the externalities and costs. Because externalities and costs are fixed, the objective function of the inner maximization problem is, up to a constant, equivalent to the corresponding problem without externalities and production costs, for which dependence on allocation is linear.

If the optimal solution to the inner maximization problem satisfy the properties of Theorems 2 and 3 for *any* choice of $\widehat{\phi}_1, \ldots, \widehat{\phi}_m$, then the properties must also hold for the optimal mechanism. As such, it suffices to restrict attention to the inner maximization problem

$$\max_{(x,t):(\phi_1(x),\ldots,\phi_m(x))=(\widehat{\phi}_1,\ldots,\widehat{\phi}_m)} \{\Omega(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\}. \qquad (*)$$

Key idea #2: Finding an effective type

I now address the issue of multidimensional types. Since externalities and costs are completely determined by the choice of $\widehat{\phi}_1, \ldots, \widehat{\phi}_m$, write $w(x) = w_0$ and $c(x) = c_0$. I show the following lemma in Appendix A:

Lemma 1. Conditional on $(w(x), c(x)) = (w_0, c_0)$, x is implementable[27] if and only if there exists a non-decreasing function $y : \mathbb{R} \to [0,1]$ such that

$$x(\theta, \eta, \zeta) = y(\theta + \eta \cdot w_0) \quad \text{almost everywhere in } \mathscr{T}.$$

Lemma 1 motivates the definition of an *effective type* $\xi \equiv \theta + \eta \cdot w_0$. Intuitively, even though types are multidimensional, the allocation problem is unidimensional. Therefore types can be projected onto a unidimensional effective type space. In principle, the effective type can depend on allocation; but this dependence can only be through $(\phi_1, \ldots, \phi_m)$, which is held fixed at $(\widehat{\phi}_1, \ldots, \widehat{\phi}_m)$.

An immediate consequence of Lemma 1 is that the allocation function must be non-decreasing in the effective type, due to the Myerson (1981) monotonicity lemma. Write $x(\theta, \eta, \zeta) = \widetilde{x}(\xi)$ for $\xi \in [\underline{\xi}, \overline{\xi}] \equiv [\underline{\theta} + \underline{\eta} \cdot w_0, \overline{\theta} + \overline{\eta} \cdot w_0]$, and define[28]

$$K \equiv \left\{ y \in L^1([\underline{\xi}, \overline{\xi}]) : y \overset{\text{a.e.}}{=} \widetilde{y} \text{ for some non-decreasing } \widetilde{y} : [\underline{\xi}, \overline{\xi}] \to [0,1] \right\}.$$

[27] An allocation function x is implementable if and only if there exists a mechanism $(\widetilde{x}, \widetilde{t})$ satisfying (IC) such that $x = \widetilde{x}$.

[28] Technically, note that K is a subset of $L^1([\underline{\xi}, \overline{\xi}])$; hence elements of K are in fact *equivalence classes* of functions (*i.e.*, those that are equal up to a set of measure zero), rather than functions. Nonetheless, I will refer to elements of K as functions in this proof sketch; Appendix A presents the formal treatment.

The extreme points of K are easily characterized:

Lemma 2. $y \in \operatorname{ex} K$ if and only if there exists a non-decreasing function $\widetilde{y} : [\underline{\xi}, \overline{\xi}] \to [0,1]$ such that $\operatorname{im} \widetilde{y} \subset \{0,1\}$ and $y \overset{\text{a.e.}}{=} \widetilde{y}$.

Lemma 2 demonstrates that the solution to the designer's *unconstrained* problem are precisely those implementable by a single price. That is, in the absence of (IR), (BB) and the constraints $(\phi_1(x), \ldots, \phi_m(x)) = (\widehat{\phi}_1, \ldots, \widehat{\phi}_m)$, the designer sets a single price to maximize the objective Ω.

Key idea #3: Applying the constrained maximum principle

Conditional on $(\phi_1(x), \ldots, \phi_m(x)) = (\widehat{\phi}_1, \ldots, \widehat{\phi}_m)$, the payment function $t(\theta, \eta, \zeta)$ is determined up to the equilibrium payoff (gross of unconditional externalities) of the lowest effective type $\underline{\xi}$ using the Milgrom and Segal (2002) envelope theorem:

Lemma 3. Conditional on $(\phi_1(x), \ldots, \phi_m(x)) = (\widehat{\phi}_1, \ldots, \widehat{\phi}_m)$, the payment function $t(\theta, \eta, \zeta)$ can be expressed as $t(\theta, \eta, \zeta) = \widetilde{t}(\xi)$, where $\widetilde{t} : [\underline{\xi}, \overline{\xi}] \to \mathbb{R}$ is defined by

$$\widetilde{t}(\xi) = \widetilde{t}(\underline{\xi}) - \underline{\xi}\widetilde{x}(\underline{\xi}) + \xi \widetilde{x}(\xi) - \int_{\underline{\xi}}^{\xi} \widetilde{x}(\xi') \, \mathrm{d}\xi'.$$

In fact, the payoff of the lowest effective type, $\underline{\xi}\widetilde{x}(\underline{\xi}) - \widetilde{t}(\underline{\xi})$, can be set to zero to satisfy (IR).[29] It follows that

$$\widetilde{t}(\xi) = \xi \widetilde{x}(\xi) - \int_{\underline{\xi}}^{\xi} \widetilde{x}(\xi') \, \mathrm{d}\xi'.$$

Denoting the distribution of ξ under F by $\widetilde{F}$, (BB) can be written as

$$\phi_0(\widetilde{x}) \equiv B + \int_{\underline{\xi}}^{\overline{\xi}} \left[\xi \widetilde{x}(\xi) - \int_{\underline{\xi}}^{\xi} \widetilde{x}(\xi') \, \mathrm{d}\xi' \right] \mathrm{d}\widetilde{F}(\xi) - c_0 \geq 0.$$

Notably, ϕ_0 is affine in $\widetilde{x}$, and can be considered as an additional "moment" that (BB) induces.

[29] More precisely, this equality is necessary for the profit-maximizing objective and if the budget constraint binds for the welfare-maximizing objective; and this equality is without loss of generality if the budget constraint does not bind for the welfare-maximizing objective.

Thus we can define the *constraint function*[30]

$$\Lambda(\widetilde{x}) \equiv (\phi_0(\widetilde{x}), \phi_1(\widetilde{x}), \ldots, \phi_m(\widetilde{x}))\,.$$

Moreover, let

$$\Sigma \equiv (\operatorname{im}\phi_0 \cap \mathbb{R}_+) \times \{(\widehat{\phi}_1, \ldots, \widehat{\phi}_m)\}.$$

Using the fact that the objective function is linear in $\widetilde{x}$ conditional on $(\phi_1(x), \ldots, \phi_m(x)) = (\widehat{\phi}_1, \ldots, \widehat{\phi}_m)$, the constrained maximum principle implies[31] that there is a maximizer of the objective function over $\Lambda^{-1}(\Sigma)$ that can be expressed as the convex combination combination of at most $m+2$ extreme points of K if (BB) binds, and at most $m+1$ extreme points if (BB) does not bind. Any function x^* that is a convex combination of n extreme points of K (which are characterized by Lemma 2) must satisfy

$$\operatorname{im} x^* \subset \{0, q_1, \ldots, q_m, 1\} \quad \text{for some } 0 < q_1 < \cdots < q_n < 1.$$

This concludes the proofs of Theorems 2 and 3.

5 When is a single price sufficient?

While the results of Section 4 establishes that rationing is generally optimal when externalities are present, I now derive conditions under which a single price is sufficient. I separately consider the cases of a welfare-maximizing social planner and a profit-maximizing monopolist.

5.1 When is a Pigouvian tax welfare-optimal?

The possibility of multiple rationing options arises in part because of the generality of the externalities and costs. However, in many instances, externalities and costs may depend on only one moment of allocation. A common example is that of network goods, where—at least up to a first-order approximation—externalities that consumers experience often depend on the *total quantity* of the good sold.[32] For some network goods, such as communication devices or apps, externalities might instead on the *total value* of the good to other consumers, which determines

[30] Here, I abuse notation slightly by writing $\phi_1, \ldots, \phi_m$ as functions of $\widetilde{x}$ rather than x.
[31] I verify that the technical conditions for the constrained maximum principle to apply are met in Appendix A.
[32] See, for example, Katz and Shapiro (1985), Economides (1996) and Farrell and Klemperer (2007).

their frequency of usage.[33]

In each of these two cases, I show that a Pigouvian tax is in fact welfare-optimal:

Theorem 4. Suppose that the externality-cost function $H(x)$ depends on the allocation x only through *total quantity* $Q(x)$ or *total value* $V(x)$, where:

$$\begin{cases} Q(x) & \equiv \int_{\mathcal{T}} x(\theta,\eta,\zeta) \, \mathrm{d}F(\theta,\eta,\zeta), \\ V(x) & \equiv \int_{\mathcal{T}} \theta x(\theta,\eta,\zeta) \, \mathrm{d}F(\theta,\eta,\zeta). \end{cases}$$

Then a Pigouvian tax is welfare-optimal.

The intuition behind the first part of Theorem 4 is clear: When the externalities and costs depend only on allocation through total quantity, social benefit and cost curves can be drawn; hence the traditional price-theoretic approach applies. That is, the social planner can solve the allocation problem by thinking in terms of *quantities*, as in Weitzman (1974); because the quantity of good sold completely determines the externalities that consumers experience, the social planner's choice of quantity alone—which determines the size of the Pigouvian tax—is efficient.

This approach of thinking in terms of quantities carries over to the second part of Theorem 4. Instead of quantities, the social planner now chooses the total value to allocate. The welfare objective depends on allocation in two ways, namely, the value that consumers receive directly from purchasing the good and the externalities that consumers experience; but the latter is completely determined by the former. Conditional on the total value chosen by the social planner, it is easy to see that a Pigouvian tax implements it: simply allocate the good greedily to those who, net of externalities, value the good more.

5.2 Profit maximization versus welfare maximization

Unlike the case of welfare maximization, standard intuition suggests that a profit-maximizing monopolist might choose to ration where a single price might have been efficient. In contrast to the social planner, who intersects an *a priori* downward-sloping demand with marginal cost, the monopolist intersects marginal revenue with marginal cost instead; crucially, marginal revenue need not be downward sloping. Rationing therefore arises from the interaction between the

[33] As might be familiar, the mere *ownership* of an email account is often no guarantee of email usage, which affects how much another consumer might value email.

monopolist's costs and the non-monotonicity of marginal revenue.

As a consequence of this classic result, it may seem intuitive that non-competitive practices such as rationing can be seen as an exercise of market power.[34] This intuition, however, breaks down in the presence of externalities. To make starker the comparison between profit maximization and welfare maximization, I focus on the case of regular distributions:

Definition 2. The distribution F is said to be *regular* if, for any $\theta \in [\underline{\theta}, \overline{\theta}]$,

$$\theta \mapsto \theta - \frac{1 - \int_{\underline{\theta}}^{\theta} \int_{\underline{\eta}}^{\overline{\eta}} \int_{\underline{\zeta}}^{\overline{\zeta}} f(\theta', \eta, \zeta)\, \mathrm{d}\theta'\, \mathrm{d}\eta\, \mathrm{d}\zeta}{\int_{\underline{\eta}}^{\overline{\eta}} \int_{\underline{\zeta}}^{\overline{\zeta}} f(\theta, \eta, \zeta)\, \mathrm{d}\eta\, \mathrm{d}\zeta} \quad \text{is increasing.}$$

Definition 2 extends Myerson's (1981) notion of regularity to this setting with externalities: it requires that the marginal revenue curve, now averaged over the preferences that consumers have about externalities, must be decreasing.

Absent externalities, regularity of F is sufficient to ensure that the monopolist's profit-maximizing mechanism is to set a single price. In fact, the same is true in a large class of models even when externalities are present:

Theorem 5. Suppose that there are are only unconditional externalities, and that the cost function $c(x)$ depends on the allocation x only through total quantity $Q(x)$. If F is regular, then:

(i) the profit-maximizing mechanism can be implemented with a single price; and

(ii) the welfare-maximizing mechanism generally involves rationing.

The assumption that the costs depend only on total quantity is fairly general and accommodate fixed costs and varying returns to scale.[35] By contrast, no additional assumptions are placed on unconditional externalities. The monopolist fails to account for unconditional externalities, and sets the optimal price as if no externalities were present.

Under the welfare-maximizing mechanism, the social planner would have consumers internalize part of the unconditional externalities. If the internalization of externalities is at odds with allocating the goods to the consumers with the highest value, the social benefit *surface* will be

[34] For instance, Loertscher and Muir (2020) show that rationing is consistent with profit maximization whenever the monopolist's revenue function is convex and if resale can be prevented.

[35] However, it does rule out type-dependent marginal costs that are important in adverse selection models. The general model accommodates these costs.

generally non-monotone—note that a *curve* is no longer sufficient to describe social benefit, since it could in principle vary with allocation via multiple moments. This non-monotonicity, like that of marginal revenue curves in the setting of Myerson (1981), induces the social planner to ration.

Theorem 5 therefore shows that the presence of externalities *reverses* the usual intuition of apparently non-competitive practices. Indeed, inefficiencies arise even for the case of perfect competition, precisely because of the failure to ration.

6 Policy implications

Pigouvian taxation is an important tool in policymaking, such as in environmental policy design and public health programs. Indeed, taxes on carbon emissions, tobacco, alcohol and soda stem from the intuition that, by adjusting the prices of goods to reflect their social cost, consumers can be induced to internalize the externalities that their consumption imposes on others. However, as Theorem 4 shows, this intuition is only correct when special assumptions are imposed on externalities and costs. In general, the optimal policy will require rationing *à la* Theorem 2.

How might rationing look in reality? Interpreting allocation as quality,[36] rationing in policy applications could correspond to different *tiers*, such as in the market for public health insurance. Health insurance exchanges in the U.S., for instance, offer four "metal" tiers (Bronze, Silver, Gold and Platinum) associated with different levels of care. Of course, my model abstracts away from many important institutional details of healthcare, such as the availability of employer-provided health insurance and the specifics of insurance contracts in different tiers (*e.g.*, in the form of premiums, deductibles and copay). However, my results help explain why such a tiered system might arise: Consumers differ in various characteristics—such as age and underlying health conditions—that insurance exchanges are legally prohibited from pricing. These characteristics affect the social benefit of incremental coverage that consumers experience, but are also imperfectly correlated with their value for incremental coverage.

Perhaps more importantly for policy design, my results provide guidance as to how these tiers could be chosen in practice. Indeed, the possibility of rationing derives from various empirically measurable *moments* that externalities depend on. For instance, a policymaker might decide that a positive externality arises from covering families with young children—perhaps because young children are more vulnerable to infectious diseases, and could also expose other young children to these diseases in daycare centers. In this case, the coverage of families with young children

[36] I give a more complete discussion as to how this can be done in Section 7.

then becomes a targeted moment. Depending on how this moment interacts with other targeted moments, the policymaker might decide to introduce a new tier of coverage that caters specifically to these families.

While these insights may appear intuitive for the case of health insurance, they could yet inform policy decisions in other markets. Vehicle ownership, for instance, imposes many externalities on others, via pollution, carbon emissions, congestion and accidents, to name a few. There are, however, limitations on the extent to which these externalities can be priced directly: equity concerns may restrict the magnitude of fuel taxes that can be levied, while extensive congestion pricing may be theoretically appealing but politically intractable. In the absence of markets for these externalities, the question of whether vehicle *ownership* itself might be regulated arises.

Singapore provides a fascinating example of how vehicle ownership regulation could work in practice. Singapore maintains a unique vehicle quota system, which has been discussed and analyzed by Koh and Lee (1994). As Koh and Lee (1994) describe, the vehicle quota system was introduced in 1990, only after a mix of regulation and taxes was found to be ineffective at controlling congestion. Under the vehicle quota system, a consumer can purchase a permit to drive the vehicle (known as the "Certificate of Entitlement") under various categories; prices in each category are separately determined via auction. Interestingly, each vehicle does not fall into a unique category. To purchase a permit to drive a big car (defined as having 1601 cc to 2000 cc of horsepower), a consumer can bid in the "big car" category, the "open car" category (in which all consumers are allowed to bid regardless of vehicle), or the "weekend car" category (which would permit the consumer to drive the car only during certain weekday off-peak hours and on the weekends). The different winning probabilities and prices in each categories, along with the option of purchasing a more restrictive permit (presumably at a lower price), demonstrate how rationing can be implemented in public policy.

Finally, the prospect of non-competitive pricing in markets for goods with externalities—both by a profit-maximizing monopolist and by a welfare-maximizing social planner—opens up a host of questions about market power and optimal regulation. Traditional measures of inefficiencies induced by market power, such as deadweight loss and markups, assume that both the monopolist and the social planner choose only prices, and therefore are inadequate in this setting. While I explore some of these issues in the context of platforms in a follow-up paper (Kang and Vasserman, 2020), these are important open questions from a policymaking perspective.

7 Discussion

7.1 Generalizations

Despite the generality of the framework presented in Section 2, it should be noted that further generalizations can be easily accommodated, including the following.

Ramsey objective

In practice, attaining the efficient outcome via the welfare-maximizing mechanism can prove difficult for a variety of reasons, including political economy constraints and implementation challenges within the public sector. As such, the designer may instead contract with a private firm to implement a welfare-maximizing mechanism, subject to a minimum profit guarantee π for the private firm. Formally, the designer maximizes:

$$\Omega^{\text{welfare}}(x,t) \equiv \int_{\mathscr{T}} \{[\theta + \eta \cdot w(x)]\, x(\theta,\eta,\zeta) + \zeta \cdot z(x)\} \; \mathrm{d}F(\theta,\eta,\zeta) - c(x),$$

subject to the constraints (IC), (IR) and

$$\int_{\mathscr{T}} t(\theta,\eta,\zeta) \; \mathrm{d}F(\theta,\eta,\zeta) - c(x) \geq \pi. \tag{MP}$$

Despite the important difference in economic interpretation, (MP) is mathematically equivalent to (BB). As such, the characterization of the optimal Ramsey mechanism follows immediately from Theorem 2:

Theorem 6. Suppose that the *externality-cost function* $H^{\text{Ramsey}}(x) = (w(x), z(x), c(x))$ is m-aggregate. Then the optimal Ramsey mechanism $(x^{\text{Ramsey}}, t^{\text{Ramsey}})$ satisfies

$$\operatorname{im} x^{\text{Ramsey}} \subset \{0, q_1, \ldots, q_{m+1}, 1\} \quad \text{for some } 0 < q_1 < \cdots < q_m < 1.$$

Moreover, if (MP) does not bind, then

$$\operatorname{im} x^{\text{Ramsey}} \subset \{0, q_1, \ldots, q_m, 1\} \quad \text{for some } 0 < q_1 < \cdots < q_m < 1.$$

Other constraints in the designer's objective

While the Ramsey objective includes only the additional constraint (MP), other constraints can be similarly accommodated, as long as the constraints are affine in allocation. Each additional affine constraint adds one additional rationing option in the allocation mechanism. Subject to existence of the optimal mechanism, aggregate constraints can be accommodated too.

Interpreting allocation as quality

While I have formulated consumer utility to depend linearly on allocation, similar results hold when consumer preferences are nonlinear in the following way: conditional on receiving an allocation X of the good for a payment T, and experiencing (unconditional) externalities $Z = (Z_1, \ldots, Z_\ell)$, a type-$(\theta, \zeta)$ consumer realizes a payoff of

$$\theta h(X) + \zeta \cdot Z - T,$$

where $h : \mathbb{R} \to \mathbb{R}$ is an increasing convex function.[37] Similar to the framework presented in Section 2, I assume that externality and cost functions are aggregate; externality functions are continuous and the cost function is lower-semicontinuous.

To interpret this model, I adapt the interpretation of Mussa and Rosen's (1978) model. The good in question can be produced with varying *quality*, which is represented by X. Quality is bounded from below and above; the bounds are normalized to 0 and 1 respectively. Externalities and production costs vary according to quality-based moments.

Importantly, the convexity assumption on h is equivalent to consumers experiencing increasing marginal returns to quality. For instance, reconsider the vaccine example introduced in Section 1. Quality here can be thought of as a measure of vaccine effectiveness; a quality of 0 is tantamount to not receiving the vaccine, while a quality of 1 represents complete immunization against the virus. With a vaccine of modest quality, however, the virus has a chance of mutating into a more resistant strain; this probability drops off sharply as vaccine quality increases. In this case, consumer preferences would be convex in vaccine quality.

Under these assumptions, it can be shown that all the results of the paper extend; in fact, exactly the same proofs apply. This illustrates a key difference between the constrained maximum principle and concavification techniques: while the latter typically require a linear

[37] This is a similar observation to that made by Mussa and Rosen (1978) of their model, with the caveat that their extension requires h to be an increasing *concave* function. With a concave h, convex programming tools can be applied (as in Mussa and Rosen, 1978); the constrained maximum principle applies when h is convex.

objective function, the constrained maximum principle shows that the linearity assumption on the objective function can be relaxed.

7.2 Applications of the constrained maximum principle

To illustrate the broad applicability of the constrained maximum principle, I now discuss a number of other economic settings where it can be applied. In many of these settings, the constrained maximum principle substantially simplifies the proofs of known results.

Mechanism design

In Section 4, I have already discussed how various results of Myerson (1981) and Samuelson (1984) follow from the constrained maximum principle as special cases of this paper's results. The constrained maximum principle can also be applied to the setting of Myerson and Satterthwaite (1983), who study the bilateral trade problem between a seller with an indivisible good and a buyer; each agent has private information about their value for the good. With a budget constraint, the constrained maximum principle implies that, perhaps interestingly, the second-best allocation mechanism x^* satisfies $\operatorname{im} x^* \subset \{0, q, 1\}$ for some $0 < q < 1$. To my knowledge, this result is new.[38]

The Bauer maximum principle has been employed by Manelli and Vincent (2007) to characterize revenue-maximizing mechanisms in a multiple-good setting with a single buyer. They do so by analyzing the problem in utility space. In their model, the buyer has linear utility, so incentive compatibility implies that buyer's expected payoff function is convex; their analysis makes use of the characterization of extreme points in the cone of convex functions.[39] While the constrained maximum principle is not required in the setting that Manelli and Vincent (2007) consider, it suggests that their results can potentially generalize to other environments with constraints.

[38] Using Lagrangian duality methods, Myerson and Satterthwaite (1983) characterize the second-best allocation mechanism assuming that agent distributions are regular. They show that the second-best allocation x^* must then satisfy $\operatorname{im} x^* \subset \{0, 1\}$. By contrast, the result that I state here applies to distributions that need not be regular (but subject to the condition that they have positive density everywhere in each agents' type space).

[39] The relevant compact convex set K in their application of the Bauer maximum principle is the convex cone of convex functions u in the space of continuous functions on the multidimensional type space $[0, 1]^N$, subject to a regularity constraint ($\nabla u \in [0, 1]^N$ a.e.) and an initial value constraint ($u(0) = 0$).

Information design

The constrained maximum principle applies to the setting of Kamenica and Gentzkow (2011). The relevant compact convex set K in their environment is the space of distributions over posteriors for a finite set of N outcomes, the extreme points of which are the Dirac delta distributions on posteriors, which can be identified as posteriors themselves. There are N Bayes-plausibility constraints, one for each outcome; however, one of the Bayes-plausibility constraints is redundant as it is implied by the fact that any distribution over posteriors must have probability mass on each outcome that sum to one. As such, the constrained maximum principle implies that the maximizer of any linear objective function can be expressed as a convex combination of at most N posteriors.

Le Treust and Tomala (2019) consider an additional constraint, which arises from a capacity constraint on information transmission, to the information design problem. The additional constraint is affine in posteriors. The constrained maximum principle implies that one more posterior is generally needed; hence the maximizer of any linear objective function in the information design problem can be expressed as a convex combination of at most $N + 1$ posteriors. Doval and Skreta (2018) generalize Le Treust and Tomala's (2019) result by showing that with K additional affine constraints, the maximizer of any linear objective function in the information design problem can be expressed as a convex combination of at most $N + K$ posteriors; but this is again immediate by appealing to the constrained maximum principle.

Robustness

Carrasco, Luz, Kos, Messner, Monteiro, and Moreira (2018) study a robust version of the classic revenue maximization problem of a seller with an indivisible good and a buyer. In their setting, the seller knows only the first N statistical moments of the buyer's distribution of valuation. A key lemma that they use to characterize optimal selling mechanisms in their setting involves showing that the worst-case distribution can be expressed as the convex combination of $N + 1$ probability masses over buyer types. While they employ Lagrangian duality methods to prove this result, this follows directly from the constrained maximum principle.

Other applications

In an inspiring contribution, the insights of Le Treust and Tomala (2019) and Doval and Skreta (2018) have been applied by Dworczak, Kominers, and Akbarpour (2019) to study redistribution and inequality in a two-sided market. They show that, with two constraints (namely, budget balance and market clearing), at most three prices are needed to represent the optimal mechanism for redistribution. In fact, the constrained maximum principle can be applied directly in their setting to obtain the same result.

The constrained maximum principle can also be used to shed light on economic phenomena that have, so far, proven difficult to analyze, such as cross-subsidization on platforms and in multi-sided markets. While Weyl (2010) presents a price-theoretic framework of viewing platforms and multi-sided markets, his analysis requires that externalities[40] depend only on the number of participants on each side of the market. Kang and Vasserman (2020) employ the constrained maximum principle to extend the techniques of the present paper to platforms and multi-sided markets; doing so significantly generalizes Weyl's (2010) framework and demonstrates that, in general, price discrimination via tiered pricing (rather than a single price, as in Weyl, 2010) is optimal.

7.3 Limitations of the constrained maximum principle

A limitation of the constrained maximum principle is that it establishes only an upper bound on how many variables are needed to represent the solution of the constrained maximization problem; generally, more work (such as in Theorem 4) is required to lower this bound.

However, this limitation must be qualified: The main advantage of the constrained maximum principle is that its application turns an infinite-dimensional optimization problem into a finite-dimensional one. Importantly, this reduction in dimension makes the solution computable.

Although the constrained maximum principle applies to many different settings, it is not always clear *a priori* how it is applied. Crucially, the proof techniques that I have presented in Section 4.3 rely on two other key ideas—namely, formulating externalities and costs as constraints, and projecting multidimensional types onto a unidimensional effective type—in addition to applying the constrained maximum principle. While these techniques extend to other settings (such as in Kang and Vasserman, 2020), it remains to be seen if there are other ways of applying the constrained maximum principle.

Finally, while the constrained maximum principle presented in Section 3 deals with the case

[40] In platforms and multi-sided markets (such as Uber), externalities can either be between members of the same side of the market (*e.g.*, having more drivers on Uber leads to increased competition for a fixed number of riders, and hence decreases the value of a driver to be on Uber), or between members of different sides of the market (*e.g.*, having more riders on Uber leads to faster matching rates for a fixed number of drivers, and hence increases the value of a driver to be on Uber).

of moment constraints, different versions of the constrained maximum principle can accommodate other constraints. This is done in recent independent work by Arieli, Babichenko, Smorodinsky, and Yamashita (2020) and Kleiner, Moldovanu, and Strack (2020), who consider a version of the constrained maximum principle with a majorization constraint rather than moment constraints.

8 Conclusion

In this paper, I have considered a general framework of markets for goods with externalities. The traditional price-theoretic approach of Pigouvian taxes has limited applicability here due to potential sources of nonconvexities in the form of externalities and costs. By contrast, an approach based on mechanism design and the idea that nonconvexities can be mostly abstracted away—by focusing on only a few relevant moments of the allocation function—yields a number of new insights on the problem.

While I have emphasized the applicability of the constrained maximum principle in a wide variety of economic problems throughout this paper, its relevance is really driven by the idea that many nonconvexities can be captured by a few sufficient statistics. This should not be a radical idea in economics: after all, the price system is premised on the philosophy that prices (or, equivalently, quantities) serve as sufficient statistics to describe demand and supply for goods. The results of this paper can be viewed as an extension of the price system to markets where more moments than quantities are needed for a complete description of demand and supply.

Appendix A Omitted proofs

A.1 Proofs of Theorem 2 and 3

While an outline of the argument has been given in Section 4.3, for the sake of completeness, I present a self-contained proof here. I proceed in steps, the enumeration of which follows that of the key ideas sketched in Section 4.3.

I follow the notation of Section 4. Denote the designer's objective function and externality-cost function by $\Omega \in \{\Omega^{\text{welfare}}, \Omega^{\text{profit}}\}$ and $H \in \{H^{\text{welfare}}, H^{\text{profit}}\}$ respectively. I also include (BB) as a constraint, with the understanding that (BB) does not bind if $\Omega = \Omega^{\text{profit}}$.

I proceed under the assumption that the designer's problem admits a solution: that is, an optimal mechanism exists. This is formally verified in Step 4.

Step 1: Formulating externalities and costs as constraints

Since the externality-cost function $H(x)$ is m-aggregate by assumption, there exist m affine functions $\phi_1, \ldots, \phi_m$ and an arbitrary function $\widetilde{H}$ such that

$$H(x) = \widetilde{H}\left(\phi_1(x), \ldots, \phi_m(x)\right).$$

By assumption (and as will be formally verified in Step 4), an optimal mechanism exists; hence the designer's problem can be written as a nested optimization problem:

$$\begin{aligned}
&\max_{(x,t)} \{\Omega(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\} \\
&= \max_{\widehat{\phi}_1,\ldots,\widehat{\phi}_m} \left[\max_{(x,t):(\phi_1(x),\ldots,\phi_m(x))=(\widehat{\phi}_1,\ldots,\widehat{\phi}_m)} \{\Omega(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\} \right].
\end{aligned}$$

Denote the optimal mechanism by (x^*, t^*), and let $(\widehat{\phi}_1^*, \ldots, \widehat{\phi}_m^*) = (\phi_1(x^*), \ldots, \phi_m(x^*))$. I focus on the inner maximization problem:

$$\max_{(x,t):(\phi_1(x),\ldots,\phi_m(x))=(\widehat{\phi}_1^*,\ldots,\widehat{\phi}_m^*)} \{\Omega(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\}. \tag{†}$$

Step 2: Finding an effective type

Since externalities and costs are completely determined by the choice of $\widehat{\phi}_1^*, \ldots, \widehat{\phi}_m^*$, write $w(x) = w^*$ and $c(x) = c^*$.

Lemma 1. Conditional on $(w(x), c(x)) = (w^*, c^*)$, x is implementable if and only if there exists a non-decreasing function $y : \mathbb{R} \to [0, 1]$ such that

$$x(\theta, \eta, \zeta) = y(\theta + \eta \cdot w^*) \quad \text{almost everywhere in } \mathscr{T}.$$

Proof of Lemma 1. Suppose that the mechanism (x, t) satisfies (IC). Define

$$\widetilde{y}(\xi, \eta, \zeta) \equiv x(\xi - \eta \cdot w^*, \eta, \zeta) \quad \text{for } \xi \in [\underline{\theta} + \underline{\eta} \cdot w^*, \overline{\theta} + \overline{\eta} \cdot w^*] \equiv [\underline{\xi}, \overline{\xi}].$$

Since (x, t) satisfies (IC),

$$[(\theta + \eta \cdot w^*) - (\theta' + \eta' \cdot w^*)] \, [x(\theta, \eta, \zeta) - x(\theta', \eta', \zeta')] \geq 0 \quad \text{for any } (\theta, \eta, \zeta), (\theta', \eta', \zeta') \in \mathscr{T}.$$

By the construction of $\widetilde{y}$, this is equivalent to

$$(\xi - \xi') \, [\widetilde{y}(\xi, \eta, \zeta) - \widetilde{y}(\xi', \eta', \zeta')] \geq 0 \quad \text{for any } \xi, \xi' \in [\underline{\xi}, \overline{\xi}]; \; \eta, \eta' \in [\underline{\eta}, \overline{\eta}]; \; \zeta, \zeta' \in [\underline{\zeta}, \overline{\zeta}].$$

In particular, $\widetilde{y}(\cdot, \eta, \zeta)$ is non-decreasing for any (η, ζ); hence it is continuous almost everywhere. However, the inequality above also implies that

$$\widetilde{y}(\xi + \varepsilon, \eta, \zeta) \geq \widetilde{y}(\xi, \eta', \zeta') \quad \text{for any } \varepsilon > 0; \; \xi \in [\underline{\xi}, \overline{\xi}]; \; \eta, \eta' \in [\underline{\eta}, \overline{\eta}]; \; \zeta, \zeta' \in [\underline{\zeta}, \overline{\zeta}].$$

Using the fact that $\widetilde{y}(\cdot, \eta, \zeta)$ is continuous almost everywhere, in the limit where $\varepsilon \to 0$:

$$\widetilde{y}(\xi, \eta, \zeta) \geq \widetilde{y}(\xi, \eta', \zeta') \quad \text{almost everywhere for any } \xi \in [\underline{\xi}, \overline{\xi}]; \; \eta, \eta' \in [\underline{\eta}, \overline{\eta}]; \; \zeta, \zeta' \in [\underline{\zeta}, \overline{\zeta}].$$

So $\widetilde{y}(\xi, \eta, \zeta) = y(\xi)$ is constant in (η, ζ) and non-decreasing in ξ almost everywhere. □

Call $\xi \equiv \theta + \eta \cdot w^*$ the "effective type." Lemma 1 implies that the allocation function must be non-decreasing in the effective type, due to the Myerson (1981) monotonicity lemma. Denote

$$V \equiv L^1([\underline{\xi}, \overline{\xi}]).$$

Write $x(\theta, \eta, \zeta) = y(\xi)$ for $\xi \in [\underline{\xi}, \overline{\xi}] \equiv [\underline{\theta} + \underline{\eta} \cdot w^*, \overline{\theta} + \overline{\eta} \cdot w^*]$, and define

$$K \equiv \left\{ y \in L^1([\underline{\xi}, \overline{\xi}]) : y \overset{\text{a.e.}}{=} \widetilde{y} \text{ for some non-decreasing } \widetilde{y} : [\underline{\xi}, \overline{\xi}] \to [0, 1] \right\}.$$

Denote by $[\widetilde{y}]$ the equivalence class of functions that are almost every everywhere equal to $\widetilde{y}$ on $[\underline{\xi}, \overline{\xi}]$. Clearly, K is convex: if $[\widetilde{y}_1], [\widetilde{y}_2] \in \overline{K}$, where $\widetilde{y}_1, \widetilde{y}_2 : [\underline{\xi}, \overline{\xi}] \to [0,1]$ are non-decreasing, then any convex combination[41] of $\widetilde{y}_1$ and $\widetilde{y}_2$ is non-decreasing and has image in $[0,1]$; hence the convex combination of their equivalence classes is in K.

The following characterizes the extreme points of K:

Lemma 2. $y \in \operatorname{ex} K$ if and only if there exists a non-decreasing function $\widetilde{y} : [\underline{\xi}, \overline{\xi}] \to [0,1]$ such that $\operatorname{im} \widetilde{y} \subset \{0,1\}$ and $y \overset{\text{a.e.}}{=} \widetilde{y}$.

Proof of Lemma 2. See Lemma 4 of Manelli and Vincent (2007). □

Note that restriction to *equivalence classes* of implementable allocation functions is without loss of generality: externalities and costs remain unchanged between any two allocation functions that are equal almost everywhere [cf. condition (i) of Definition 1], and the designer's objective function is unchanged between any two allocation functions that are equal almost everywhere as the distribution of types F is atomless.

Step 3: Applying the constrained maximum principle

Lemma 3. Conditional on $(\phi_1(x), \ldots, \phi_m(x)) = (\widehat{\phi}_1^*, \ldots, \widehat{\phi}_m^*)$, the payment function $t(\theta, \eta, \zeta)$ can be expressed as $t(\theta, \eta, \zeta) = \tau(\xi)$, where $\tau : [\underline{\xi}, \overline{\xi}] \to \mathbb{R}$ is defined by

$$\tau(\xi) = \tau(\underline{\xi}) - \underline{\xi} y(\underline{\xi}) + \xi y(\xi) - \int_{\underline{\xi}}^{\xi} y(\xi') \, \mathrm{d}\xi'.$$

Proof of Lemma 3. (IC) implies that $t(\theta, \eta, \zeta)$ must be a function of ξ; otherwise, different types with the same ξ would report an effective type that results in the lowest payment. Write $t(\theta, \eta, \zeta) = \tau(\xi)$. (IC) also implies that

$$\xi y(\xi) - \tau(\xi) = \sup_{\widehat{\xi} \in [\underline{\xi}, \overline{\xi}]} \left[\xi y(\widehat{\xi}) - \tau(\widehat{\xi}) \right].$$

The desired formula follows by the Milgrom and Segal (2002) envelope theorem. □

If the designer maximizes profit, or if (BB) binds, then the designer optimally sets (IR) to bind for the lowest effective type $\underline{\xi}$. If (BB) does not bind and the designer maximizes welfare, then

[41] Addition and scalar multiplication in $L^1([\underline{\xi}, \overline{\xi}])$ are defined respectively by pointwise addition and scalar multiplication of representative functions from the respective equivalence classes.

setting (IR) to bind for the lowest effective type $\underline{\xi}$ is without loss of generality, since transfers do not affect the welfare objective. It follows that

$$\tau(\xi) = \xi y(\xi) - \int_{\underline{\xi}}^{\xi} y(\xi')\ \mathrm{d}\xi'.$$

Denoting the distribution of ξ under F by $\widetilde{F}$, (BB) can be written as

$$\phi_0(y) \equiv B + \int_{\underline{\xi}}^{\overline{\xi}} \left[\xi y(\xi) - \int_{\underline{\xi}}^{\xi} y(\xi')\ \mathrm{d}\xi'\right]\ \mathrm{d}\widetilde{F}(\xi) - c^* \geq 0.$$

Notably, ϕ_0 is affine. With a slight abuse of notation, define the *constraint function*

$$\Lambda(y) \equiv (\phi_0(y), \phi_1(y), \ldots, \phi_m(y)).$$

Moreover, let

$$\Sigma \equiv (\operatorname{im}\phi_0 \cap \mathbb{R}_+) \times \{(\widehat{\phi}_1^*, \ldots, \widehat{\phi}_m^*)\}.$$

Endow V with the topology induced by the L^1 norm, $||\cdot||_1$; it is well-known that $(V, ||\cdot||_1)$ is a Banach (and hence locally convex Hausdorff) space.

Lemma 4. $K \subset V$ is compact.

Proof of Lemma 4. Since $(V, ||\cdot||_1)$ can be viewed as a metric space, it suffices to show that K is sequentially compact. Let $\{[y_n]\} \subset K$ for non-decreasing $y_n : [\underline{\xi}, \overline{\xi}] \to [0,1]$; by Helly's selection theorem (cf. Exercise 7.13 of Rudin, 1964), y_n may be assumed (*i.e.*, up to passing to a subsequence) to have a pointwise limit, $y^\circ : [\underline{\xi}, \overline{\xi}] \to [0,1]$. By the bounded convergence theorem, y_n converges to y° in the L^1 norm; hence $[y_n] \xrightarrow{L^1} [y^\circ] \in V$. □

Let $\widetilde{G}$ denote the distribution of ζ under F. Define

$$\widetilde{\Omega}(y) = \begin{cases} \displaystyle\int_{\underline{\xi}}^{\overline{\xi}} \xi y(\xi)\ \mathrm{d}\widetilde{F}(\xi) + \int_{\underline{\zeta}}^{\overline{\zeta}} \zeta \cdot z_0\ \mathrm{d}\widetilde{G}(\zeta) - c_0 & \text{for the welfare objective,} \\ \displaystyle\int_{\underline{\xi}}^{\overline{\xi}} \left[\xi y(\xi) - \int_{\underline{\xi}}^{\xi} y(\xi')\ \mathrm{d}\xi'\right]\ \mathrm{d}\widetilde{F}(\xi) - c_0 & \text{for the profit objective.} \end{cases}$$

It is routine to verify that $\widetilde{\Omega}(y)$ is continuous in y (with respect to the L^1 topology); moreover, since it is affine in y, it is also quasiconvex. Thus the constrained maximum principle applies. Consequently, there is a maximizer of the objective function over $\Lambda^{-1}(\Sigma)$ that can be expressed as the convex combination combination of at most $m+2$ extreme points of K if (BB) binds, and at most $m+1$ extreme points of K if (BB) does not bind. By Lemma 2, the optimal allocation function x^* can be written as

$$\operatorname{im} x^* \subset \{0, q_1, \dots, q_k, 1\} \quad \text{for } 0 < q_1 < \cdots < q_k < 1, \text{ where } k = \begin{cases} m+1 & \text{if (BB) binds,} \\ m & \text{if (BB) does not bind.} \end{cases}$$

Step 4: Existence of the optimal mechanism

I now show that an optimal mechanism exists. Since $B \geq c(0)$, there must be at least one feasible mechanism. Suppose $\{(x_n, t_n)\}$ is a maximizing sequence of mechanisms in the designer's optimization problem:

$$\sup_{(x,t)} \{\Omega(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\}.$$

By the results above, it can be assumed without loss of generality that each (x_n, t_n) satisfies

$$\operatorname{im} x_n \subset \{0, q_1^{(n)}, \dots, q_{m+1}^{(n)}, 1\} \quad \text{for } 0 < q_1^{(n)} < \cdots < q_{m+1}^{(n)} < 1.$$

Moreover, (x_n, t_n) can be viewed as functions of $\theta + \eta \cdot w(x_n)$; and x_n can be viewed as a non-decreasing function of $\theta + \eta \cdot w(x_n)$: there exist $\underline{\theta} + \underline{\eta} \cdot w(x_n) \leq \xi_1^{(n)} \leq \cdots \leq \xi_{m+2}^{(n)} \leq \overline{\theta} + \overline{\eta} \cdot w(x_n)$ such that

$$x_n(\theta + \eta \cdot w(x_n)) = \begin{cases} 1 & \text{for } \theta + \eta \cdot w(x_n) \in (\xi_{m+2}^{(n)}, \overline{\theta} + \overline{\eta} \cdot w(x_n)], \\ q_{m+1}^{(n)} & \text{for } \theta + \eta \cdot w(x_n) \in (\xi_{m+1}^{(n)}, \xi_{m+2}^{(n)}], \\ & \vdots \\ q_1^{(n)} & \text{for } \theta + \eta \cdot w(x_n) \in (\xi_1^{(n)}, \xi_2^{(n)}], \\ 0 & \text{for } \theta + \eta \cdot w(x_n) \in [\underline{\theta} + \underline{\eta} \cdot w(x_n), \xi_1^{(n)}]. \end{cases}$$

Therefore, x_n is completely characterized by the finite-dimensional vector:

$$\beta_n = (\phi_1(x_n), \dots, \phi_m(x_n), q_1^{(n)}, \dots, q_{m+1}^{(n)}, \xi_1^{(n)}, \dots, \xi_{m+2}^{(n)}).$$

Note that $\phi_1, \ldots, \phi_m$ are bounded since they are continuous linear functions. By a diagonalization argument and passing to a subsequence if necessary, $\beta_n \to \beta^*$, where

$$\beta^* = (\phi_1^*, \ldots, \phi_m^*, q_1^*, \ldots, q_{m+1}^*, \xi_1^*, \ldots, \xi_{m+2}^*).$$

Let the value of w given $\phi_1^*, \ldots, \phi_m^*$ be w^*. Observe that β^* determines the function

$$x^*(\theta + \eta \cdot w^*) = \begin{cases} 1 & \text{for } \theta + \eta \cdot w^* \in (\xi_{m+2}^*, \overline{\theta} + \overline{\eta} \cdot w^*], \\ q_{m+1}^* & \text{for } \theta + \eta \cdot w^* \in (\xi_{m+1}^*, \xi_{m+2}^*], \\ & \vdots \\ q_1^* & \text{for } \theta + \eta \cdot w^* \in (\xi_1^*, \xi_2^*], \\ 0 & \text{for } \theta + \eta \cdot w^* \in [\underline{\theta} + \underline{\eta} \cdot w^*, \xi_1^*]. \end{cases}$$

Since β^* is the limit of $\{\beta_n\}$, $0 \leq q_1^* \leq \cdots q_{m+1}^* \leq 1$ and $\underline{\theta} + \underline{\eta} \cdot w^* \leq \xi_1^* \leq \cdots \leq \xi_{m+2}^* \leq \overline{\theta} + \overline{\eta} \cdot w^*$; this employs the fact that w is continuously aggregate.

By construction, x_n converges pointwise almost everywhere to x^*. By the dominated convergence theorem, $x_n \to x^*$ in the L^1 norm. Consequently, by continuity of $\phi_1, \ldots, \phi_m$,

$$(\phi_1(x_n), \ldots, \phi_m(x_n)) = (\phi_1^*, \ldots, \phi_m^*).$$

By Lemma 1, x^* is implementable. Let t^* be the payment function induced by x^* using the Milgrom and Segal (2002) envelope theorem, with (IR) binding for the effective type $\underline{\theta} + \underline{\eta} \cdot w^*$:

$$t^*(\theta + \eta \cdot w^*) = (\theta + \eta \cdot w^*)\, x^*(\theta + \eta \cdot w^*) - \int_{\underline{\theta} + \underline{\eta} \cdot w^*}^{\theta + \eta \cdot w^*} x^*(\xi) \, \mathrm{d}\xi.$$

Then (x^*, t^*) satisfies (IC) and (IR).

Denote the distribution of $\theta + \eta \cdot w(x_n)$ under F by $\widetilde{F}_n$, and the distribution of $\theta + \eta \cdot w^*$ under F by $\widetilde{F}^*$. To check that (x^*, t^*) satisfies (BB), observe that the dominated convergence theorem—together with the assumption that c is lower-semicontinuously aggregate implies that

$$\begin{aligned} &B + \int_{\underline{\theta}+\underline{\eta}\cdot w(x^*)}^{\overline{\theta}+\overline{\eta}\cdot w(x^*)} \left[\xi x^*(\xi) - \int_{\underline{\theta}+\underline{\eta}\cdot w(x^*)}^{\theta+\eta\cdot w(x^*)} x^*(\xi')\ \mathrm{d}\xi'\right]\ \mathrm{d}\widetilde{F}^*(\xi) \\ &= \liminf_{x_n\to x^*}\left\{B + \int_{\underline{\theta}+\underline{\eta}\cdot w(x_n)}^{\overline{\theta}+\overline{\eta}\cdot w(x_n)} \left[\xi x_n(\xi) - \int_{\underline{\theta}+\underline{\eta}\cdot w(x_n)}^{\theta+\eta\cdot w(x_n)} x_n(\xi')\ \mathrm{d}\xi'\right]\ \mathrm{d}\widetilde{F}_n(\xi)\right\} \geq \liminf_{x_n\to x^*} c(x_n) \geq c(x^*). \end{aligned}$$

Finally, it is easy to see that

$$\liminf_{x_n\to x^*} \Omega(x_n, t_n) \leq \Omega(x^*, t^*) \leq \lim_{n\to\infty} \Omega(x_n, t_n).$$

Therefore (x^*, t^*) attains the supremum of the designer's problem; hence an optimal mechanism exists. This concludes the proofs of Theorems 2 and 3.

A.2 Proof of Theorem 4

By Step 4 in the proof of Theorem 2, an optimal mechanism x^* exists. Let $\phi \in \{Q, V\}$, and denote $\phi(x^*) = \phi^*$. The designer's problem can be written as a nested optimization problem:

$$\max_{\phi^*}\left[\max_{(x,t):\phi(x^*)=\phi^*}\left\{\Omega^{\text{welfare}}(x,t) : (x,t) \text{ satisfies (IC), (IR) and (BB)}\right\}\right].$$

Given the existence of x^*, it suffices to examine the inner maximization problem. Define an effective type $\xi \equiv \theta + \eta \cdot w(x^*) \in [\underline{\theta} + \underline{\eta} \cdot w(x^*), \overline{\theta} + \overline{\eta} \cdot w(x^*)] \equiv [\underline{\xi}, \overline{\xi}]$; and let the distribution of ξ and ζ under F be $\widetilde{F}$ and $\widetilde{G}$ respectively. Note that Lemmas 1 and 3 apply; so the inner maximization problem can be rewritten as

$$\max_{x:\phi(x^*)=\phi^*}\left\{\int_{\underline{\xi}}^{\overline{\xi}} \xi x(\xi)\ \mathrm{d}\widetilde{F}(\xi) + \int_{\underline{\zeta}}^{\overline{\zeta}} \zeta \cdot z(x^*)\ \mathrm{d}\widetilde{G}(\zeta) - c(x^*) : x(\xi) \text{ is increasing in } \xi\right\}.$$

Note that $x(\xi)$ enters the objective function only through the following integral:

$$J(x) = \int_{\underline{\xi}}^{\overline{\xi}} \xi x(\xi)\ \mathrm{d}\widetilde{F}(\xi).$$

I begin by assuming that (BB) does not bind at the optimal mechanism x^*. By Theorem 2, x^* can be expressed in the following form for some $0 < q < 1$ and $\underline{\xi} \leq \xi_1 \leq \xi_2 \leq \overline{\xi}$:

$$x^*(\xi) = \begin{cases} 1 & \text{for } \xi \in (\xi_2, \overline{\xi}], \\ q & \text{for } \xi \in (\xi_1, \xi_2], \\ 0 & \text{for } \xi \in [\underline{\xi}, \xi_1]. \end{cases}$$

If $\phi = Q$, then I claim that $\xi_1 = \xi_2$. Otherwise, consider the alternative allocation function x^{**} defined by

$$x^{**}(\xi) = \begin{cases} 1 & \text{for } \xi \in (q\xi_1 + (1-q)\,\xi_2, \overline{\xi}], \\ 0 & \text{for } \xi \in [\underline{\xi}, q\xi_1 + (1-q)\,\xi_2)]. \end{cases}$$

By construction, $Q(x^*) = Q(x^{**})$. However, it is clear that $J(x^{**}) > J(x^*)$, contradicting the assumption that x^* is the optimal mechanism.

If $\phi = V$, suppose that $\xi_1 \neq \xi_2$ (otherwise the desired result already holds). Consider the alternative allocation function x^{**} defined by

$$x^{**}(\xi) = \begin{cases} 1 & \text{for } \xi \in (\xi^{**}, \overline{\xi}], \\ 0 & \text{for } \xi \in [\underline{\xi}, \xi^{**}], \end{cases}$$

where $\underline{\xi} \leq \xi^{**} \leq \overline{\xi}$ is chosen so that $J(x^*) = J(x^{**})$. Then the optimal mechanism can be implemented by a Pigouvian tax.

In the argument above, note that the transformation from x^* to x^{**} can only weaken the (BB); hence, if (BB) does not bind at x^*, it cannot bind at x^{**}. The case where (BB) binds at the optimal mechanism x^*, where x^* has 1 or 2 rationing options, is completely analogous: the same argument shows that the number of rationing options can be reduced by 1, such that (BB) does not bind at this new mechanism. Thus the only possibility that remains is when the budget constraint binds and x^* can be implemented by a Pigouvian tax.

A.3 Proof of Theorem 5

Part (ii) of Theorem 5 follows from Theorem 2, so it suffices to show part (i). By Step 4 in the proof of Theorem 3, the profit-maximizing mechanism x^* exists. Write $Q(x^*) = Q^*$, and let $\widetilde{F}$ (with density $\widetilde{f}$) denote the distribution of θ under F. The inner maximization problem for the profit-maximizing designer is

$$\max_{x:Q(x)=Q^*}\left\{\int_{\underline{\theta}}^{\overline{\theta}}\left[\theta x(\theta)-\int_{\underline{\theta}}^{\theta}x(\theta')\ \mathrm{d}\theta'\right]\ \mathrm{d}\widetilde{F}(\theta)-c(x^*):x(\theta)\text{ is increasing in }\theta\right\}.$$

Define

$$J(x)\equiv\int_{\underline{\theta}}^{\overline{\theta}}\left[\theta x(\theta)-\int_{\underline{\theta}}^{\theta}x(\theta')\ \mathrm{d}\theta'\right]\ \mathrm{d}\widetilde{F}(\theta)=\int_{\underline{\theta}}^{\overline{\theta}}\left[\theta-\frac{1-\widetilde{F}(\theta)}{\widetilde{f}(\theta)}\right]x(\theta)\ \mathrm{d}\widetilde{F}(\theta).$$

By Theorem 3, x^* can be expressed in the following form for some $0<q<1$ and $\underline{\xi}\leq\xi_1\leq\xi_2\leq\overline{\xi}$:

$$x^*(\xi)=\begin{cases}1 & \text{for }\xi\in(\xi_2,\overline{\xi}],\\ q & \text{for }\xi\in(\xi_1,\xi_2],\\ 0 & \text{for }\xi\in[\underline{\xi},\xi_1].\end{cases}$$

I claim that $\xi_1=\xi_2$. Otherwise, consider the alternative allocation function x^{**} defined by

$$x^{**}(\xi)=\begin{cases}1 & \text{for }\xi\in(q\xi_1+(1-q)\,\xi_2,\overline{\xi}],\\ 0 & \text{for }\xi\in[\underline{\xi},q\xi_1+(1-q)\,\xi_2)].\end{cases}$$

By construction, $Q(x^*)=Q(x^{**})$, so $c(x^*)=c(x^{**})$. However, it is clear that $J(x^{**})>J(x^*)$ by the assumption that F is regular, contradicting the assumption that x^* is the optimal mechanism. Thus a single price is optimal.

References

ALIPRANTIS, C. D., AND K. C. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 3^{rd} edn.

ARIELI, I., Y. BABICHENKO, R. SMORODINSKY, AND T. YAMASHITA (2020): "Optimal Persuasion via Bi-Pooling," *Working paper*.

AUMANN, R. J., AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*. MIT Press.

BARVINOK, A. (2002): *A Course in Convexity*, vol. 54 of *Graduate Studies in Mathematics*. American Mathematical Society.

BAUER, H. (1958): "Minimalstellen von Funktionen und Extremalpunkte," *Archiv der Mathematik*, 9(4), 389–393.

BAUMOL, W. J., AND W. E. OATES (1988): *The Theory of Environmental Policy*. Cambridge University Press, Cambridge, UK, 2^{nd} edn.

BULOW, J., AND P. KLEMPERER (2012): "Regulated Prices, Rent Seeking, and Consumer Surplus," *Journal of Political Economy*, 120(1), 160–186.

BULOW, J., AND J. ROBERTS (1989): "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, 97(5), 1060–1090.

CARRASCO, V., V. F. LUZ, N. KOS, M. MESSNER, P. MONTEIRO, AND H. MOREIRA (2018): "Optimal Selling Mechanisms under Moment Conditions," *Journal of Economic Theory*, 177, 245–279.

CHE, Y.-K., I. GALE, AND J. KIM (2013): "Assigning Resources to Budget-Constrained Agents," *Review of Economic Studies*, 80(1), 73–107.

CONDORELLI, D. (2013): "Market and Non-Market Mechanisms for the Optimal Allocation of Scarce Resources," *Games and Economic Behavior*, 82, 582–591.

DOVAL, L., AND V. SKRETA (2018): "Constrained Information Design: Toolkit," *Technical note*.

DWORCZAK, P., S. D. KOMINERS, AND M. AKBARPOUR (2019): "Redistribution through Markets," *Working paper*.

Economides, N. (1996): "The Economics of Networks," *International Journal of Industrial Organization*, 14(6), 673–699.

Einav, L., and A. Finkelstein (2011): "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 25(1), 115–38.

Farrell, J., and P. Klemperer (2007): "Coordination and Lock-In: Competition with Switching Costs and Network Effects," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3, pp. 1967–2072. Elsevier.

Jensen, M. K. (2018): "Aggregative Games," in *Handbook of Game Theory and Industrial Organization, Volume I*. Edward Elgar Publishing.

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

Kang, Z. Y., and S. Vasserman (2020): "Platform Market Power," *Working paper*.

Kang, Z. Y., and J. Vondrák (2019): "Fixed-Price Approximations to Optimal Efficiency in Bilateral Trade," *Working paper*.

Katz, M. L., and C. Shapiro (1985): "Network Externalities, Competition, and Compatibility," *American Economic Review*, 75(3), 424–440.

Kleiner, A., B. Moldovanu, and P. Strack (2020): "Extreme Points and Majorization: Economic Applications," *Working paper*.

Koh, W. T., and D. K. Lee (1994): "The Vehicle Quota System in Singapore: An Assessment," *Transportation Research Part A: Policy and Practice*, 28(1), 31–47.

Le Treust, M., and T. Tomala (2019): "Persuasion with Limited Communication Capacity," *Journal of Economic Theory*, 184, 104940.

Loertscher, S., and E. V. Muir (2020): "Monopoly Pricing, Optimal Rationing, and Resale," *Working paper*.

Manelli, A. M., and D. R. Vincent (2007): "Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly," *Journal of Economic Theory*, 137(1), 153–185.

Milgrom, P., and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), 583–601.

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.

Myerson, R. B., and M. A. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29(2), 265–281.

Rochet, J.-C., and J. Tirole (2006): "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, 37(3), 645–667.

Rudin, W. (1964): *Principles of Mathematical Analysis*. McGraw-Hill, New York, 3rd edn.

Samuelson, W. (1984): "Bargaining under Asymmetric Information," *Econometrica*, 54(2), 995–1005.

Sandholm, W. H. (2005): "Negative Externalities and Evolutionary Implementation," *Review of Economic Studies*, 72(3), 885–915.

——— (2007): "Pigouvian Pricing and Stochastic Evolutionary Implementation," *Journal of Economic Theory*, 132(1), 367–382.

Szapiel, W. (1975): "Points Extrémaux dans les Ensembles Convexes (I), Théorie Générale," *Bulletin de l'Académie Polonaise des Sciences. Série des Sciences Mathématiques, Astronomiques et Physiques*, 23(9), 939–945.

Toikka, J. (2011): "Ironing without Control," *Journal of Economic Theory*, 146(6), 2510–2526.

Veiga, A., E. G. Weyl, and A. White (2017): "Multidimensional Platform Design," *American Economic Review*, 107(5), 191–95.

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

——— (1977): "Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need It Most?," *Bell Journal of Economics*, 8(2), 517–524.

Weyl, E. G. (2010): “A Price Theory of Multi-Sided Platforms,” *American Economic Review*, 100(4), 1642–72.

——— (2019): “Price Theory,” *Journal of Economic Literature*, 57(2), 329–84.

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