

Prices vs. Quantities: Robust Regulation*

Zi Yang Kang[†]

March 2026

Abstract

This paper revisits the classic instrument choice problem in a setting with consumption externalities, through the lens of robust mechanism design. A regulator can implement any incentive-compatible policy but is uncertain about how individual demand is correlated with marginal externalities, and evaluates policies by worst-case welfare. The optimal policy is a quantity control: a floor for positive externalities and a ceiling for negative externalities. If the sign of the correlation is known, a uniform tax or subsidy can be optimal. The framework also applies to regulatory uncertainty and costly screening, providing a welfare-based explanation for the prevalence of non-price policies.

JEL classification: D81, D82, H23, L51

Keywords: quantity regulation, externalities, mechanism design, correlation uncertainty, robustness

* This paper is dedicated to Jeremy Bulow. I thank Gabriel Carroll, Rob McMillan, Paul Milgrom, Andy Skrzypacz, Colin Stewart, Shosh Vasserman, and especially Laura Doval, for helpful discussions. `Refine.ink` was used to check the paper for consistency and clarity.

[†] Department of Economics, University of Toronto; zy.kang@utoronto.ca.

1 Introduction

Policymakers frequently rely on quantity controls to regulate externalities. For example, the recreational use of some opiates and opioids is widely banned, while immunizations are often mandated by schools and employers. Even when outright bans and mandates are not used, quantity restrictions may still apply. Most countries impose compulsory education laws that require children to receive some minimum amount of schooling. One-gun-a-month laws and the Combat Methamphetamine Epidemic Act, respectively, limit the quantity of firearms and pseudoephedrine that any individual may purchase within a given month.

However, many economists believe in the superiority of price regulation through taxes and subsidies. The COVID-19 pandemic provides a stark illustration of this tension between economic prescriptions and policy practice. Several prominent economists, including Robert Litan and Gregory Mankiw, advocated paying individuals to get vaccinated in order to achieve herd immunity (Litan, 2020; Mankiw, 2020). Yet few schools, employers, or governments heeded these calls, choosing instead to impose vaccine mandates.

At first glance, this may look like the familiar prices-versus-quantities problem studied by Weitzman (1974). However, as a large subsequent literature has emphasized, two aspects of that benchmark are ill-suited to many policy settings. First, it places substantial informational demands on the regulator (Baumol and Oates, 1988). Second, it typically restricts attention to simple instruments, even though sufficiently rich price regulation (e.g., nonlinear pricing) can dominate quantity controls (Kaplow and Shavell, 2002). Put differently, Weitzman's approach is most natural in settings with relatively *rich* information but a *limited* set of instruments, whereas many real-world environments feature *limited* information even when a *rich* set of instruments is feasible.

This paper takes a new approach to the instrument choice problem by embedding it in a mechanism-design framework with limited information and a rich set of feasible instruments. This framework comprises two key ingredients: *heterogeneity* and *robustness*. Agents are *heterogeneous* in that they differ both in their consumption preferences for a homogeneous good and in the amount of externality they generate per unit of consumption. In the COVID-19 vaccination example, individuals vary in their willingness to pay for the vaccine as well as in the marginal social benefit their vaccination generates, for instance because they differ in their propensity to wear masks or avoid large crowds. While the regulator

may observe revealed preferences through consumption decisions, she does not observe the amount of externality generated by each agent and may have only limited information about its distribution, such as its mean. This reflects real-world settings in which policymakers know the average benefit of vaccination, but not how marginal benefits are distributed or correlated with willingness to pay. Faced with this uncertainty, the regulator adopts a *robust* approach by evaluating policies according to their worst-case welfare performance. Importantly, the regulator is not restricted to simple instruments *a priori*: she can choose any incentive-compatible policy, including nonlinear price schedules.

The first main result shows that robustness to correlation uncertainty arising from agent heterogeneity favors quantity regulation. Although the regulator can choose any incentive-compatible policy, her best welfare guarantee is achieved by a simple policy: imposing a quantity floor or ceiling when the externality is positive or negative, respectively. In special cases, such as when agents have unit demand, this result takes the form of mandates and bans.

To see why, consider the COVID-19 vaccination example. On the one hand, paying individuals to get vaccinated performs well when those with the highest willingness to pay also generate the largest marginal social benefit. On the other hand, if willingness to pay and marginal social benefit are negatively correlated—for instance, if individuals who are most eager to get vaccinated already take extensive precautions—then subsidies disproportionately attract individuals whose vaccination generates relatively little additional social benefit. In such cases, any given subsidy can fail to realize essentially all of the potential social benefit. By contrast, a mandate may achieve a better welfare guarantee by ensuring full participation, albeit at the cost of potentially vaccinating some individuals whose willingness to pay and marginal social benefit are both low. Depending on how this tradeoff resolves, the optimal policy is either to mandate vaccination or not to intervene at all, but never to subsidize vaccination.

The second main result shows that the optimal regulatory instrument—price or quantity—depends critically on the sign of the correlation between willingness to pay and amount of externality generated, when this sign is known. In particular, knowledge of the sign of the correlation can rule out the worst-case selection patterns that drive the case for quantity regulation under correlation uncertainty. Of course, if willingness to pay and marginal social benefit are known to be negatively correlated, the previous logic continues to apply and quantity regulation remains optimal. However, if willingness to pay and

marginal social benefit are known to be positively correlated—for instance, if individuals who are most eager to get vaccinated are also those at greatest risk of exposure and transmission—then price regulation becomes optimal. In this case, the optimal policy nonetheless remains simple: the regulator optimally sets a uniform subsidy equal to the *average* social benefit generated. While the presence of such positive correlation is unclear in the COVID-19 context, it is more plausible for other vaccines, such as those for human papillomavirus (HPV) or mpox, where individuals at higher risk may also have higher willingness to pay. Similar considerations arise in other externality settings: for example, many countries subsidize electric vehicle adoption, where individuals with higher willingness to pay, such as high-mileage drivers, are expected to generate larger marginal social benefits through greater emissions reductions.

It is worth emphasizing that both heterogeneity and robustness are crucial for the results in this paper. Without heterogeneity, the marginal amount of externality generated by each agent would be the same. In this case, the classical analysis due to [Pigou \(1920\)](#) applies: the regulator can do no better than to set a tax or subsidy equal to that amount of externality. Without robustness, the regulator might instead maximize expected welfare in a standard Bayesian environment with a known joint distribution of willingness to pay and amount of externality generated. This is a substantial informational requirement on the regulator—a point that I revisit below. But putting that aside, given the joint distribution, the regulator would optimally set marginal taxes or subsidies equal to the expected marginal amount of externality generated by agents consuming that marginal unit of the good. In such settings, quantity regulation is optimal only in special cases or when used in conjunction with price regulation, as [Kaplow and Shavell \(2002\)](#) have argued. In this sense, it is the interaction between heterogeneity and robustness that drives the results in this paper.

While the argument above has been framed in terms of externalities, it extends more broadly to other economic frictions. Consider, for instance, environments in which goods are allocated through costly and socially wasteful screening, such as waiting or delay. In many such environments, policymakers instead rely on random or non-price allocation mechanisms that aim to reduce or even eliminate screening—such as arbitration rather than costly litigation, and lotteries rather than ordeals. The results of this paper suggest that these policies can be understood as a robust response to uncertainty about how agents' willingness to pay is correlated with their waiting costs. If individuals with higher willingness to pay also systematically find it less costly to wait, then screening mechanisms can improve allocative efficiency.

However, when policymakers have limited information about this correlation, mechanisms that rely on costly screening may yield poor welfare guarantees. As in the externality setting, uncertainty about this correlation leads to the robust optimality of non-price allocation.

1.1 Related Literature

This paper shares a similar motivation with a positive literature that studies why quantity regulation is so prevalent even when price regulation is theoretically appealing. Key contributions by [Glaeser and Shleifer \(2001\)](#), [Glaeser, Johnson and Shleifer \(2001\)](#), and [Glaeser and Shleifer \(2002\)](#) develop two complementary enforcement-based explanations: quantity restrictions can both strengthen enforcement incentives and reduce enforcement costs by making violations easier to detect. A different strand of the literature, pioneered by [Buchanan and Tullock \(1975\)](#), examines political economy mechanisms. The present paper complements these explanations by developing a welfare-based argument for quantity regulation and by adding nuance: price regulation—also observed in practice—can be explained by the same argument when the sign of the correlation between demand and marginal externality is known.

The argument developed in this paper relates directly to a growing literature that emphasizes the importance of heterogeneity in the regulation of externalities. The seminal work of [Diamond \(1973\)](#) shows that when individuals generate heterogeneous marginal externalities, the optimal uniform corrective tax generally differs from the “naïve” uniform tax based on the unweighted average marginal harm, and instead depends on how marginal externalities covary with individuals’ demand elasticities. A recent empirical literature has documented such heterogeneity and correlation in specific settings where rich data are available, including work by [Knittel and Sandler \(2018\)](#), [Griffith, O’Connell and Smith \(2019\)](#), and [Grummon, Lockwood, Taubinsky and Allcott \(2019\)](#). At the same time, these studies underscore how demanding the informational requirements of optimal taxation are: identifying the relevant correlation structure typically requires detailed micro-level data on both consumption and social costs, data that are unavailable in many policy settings in which such heterogeneity is likely to be quantitatively important, such as those discussed by [Fleischer \(2015\)](#). In this sense, the present paper complements the existing literature by studying robust regulatory policies when such detailed data are unavailable.

Another closely related literature studies the normative question of how to choose between price and quantity instruments in regulation. A key premise of this literature that resonates with the present paper is the difficulty of measuring externalities precisely, as described by [Baumol and Oates \(1971\)](#):

“The proper level of the Pigouvian tax (subsidy) [. . .] is equal to the marginal net damage (benefit) produced by that activity. The difficulty is that it is usually not easy to obtain a reasonable estimate of the money value of this marginal damage.”

Given this difficulty, pioneering contributions by [Weitzman \(1974\)](#), [Roberts and Spence \(1976\)](#), and [Baumol and Oates \(1988\)](#) examine the relative performance of price regulation and quantity regulation, as well as hybrid instruments. Subsequent work has refined these insights and applied them to important policy environments (e.g., [Pizer, 2002](#); [Newell and Pizer, 2003](#)). However, [Kaplow and Shavell \(2002\)](#) have argued that these results rely on restrictions on price regulation—such as linear taxes—and that sufficiently rich price instruments can do better. The present paper departs from this literature in two ways: *(i)* the regulator can choose any incentive-compatible policy without restriction; and *(ii)* the regulator seeks robustness to uncertainty by evaluating policies under a worst-case welfare criterion, with limited information about the distribution of externalities.

The focus on robustness also connects this paper to the literature on mechanism design with worst-case objectives. [Carroll \(2019\)](#) provides a detailed survey of this literature. Recent work has applied worst-case criteria to justify simple policies on the basis of robustness, including [Frankel \(2014\)](#) in delegation; [Garrett \(2014\)](#) in cost-based procurement; [Carroll \(2015\)](#), [Walton and Carroll \(2022\)](#), and [Carroll and Bolte \(2023\)](#) in moral hazard; [Carroll \(2017\)](#) in multidimensional screening; [Guo and Shmaya \(2025\)](#) in monopoly regulation; and [Kang and Vasserman \(2025\)](#) in welfare analysis. Closer to the present paper’s motivation of how worst-case reasoning bears on the comparison between price and quantity regulation, [Mishra, Patil and Pavan \(2025\)](#) study a monopoly regulation problem in which the regulator has a conjectured model but does not fully trust it: the regulator selects mechanisms that maximize a worst-case welfare guarantee over a set of plausible models and chooses among them by their expected performance under the conjectured model.¹ Their analysis compares price and quantity regulations as

¹ [Dworczak and Pavan \(2022\)](#) use a similar selection rule in their study of robust information design.

rich families of direct mechanisms in which the regulator specifies either a price rule or a quantity rule as a function of the monopolist's reported type. Because demand is itself uncertain, these families are not equivalent, and the comparison is driven by how a regulated price maps into different realized quantities across plausible demand models. They show that price and quantity regulations yield the same worst-case welfare guarantee and establish conditions under which one family outperforms the other. By contrast, the present paper studies a different economic environment and source of uncertainty: the regulator wishes to regulate a heterogeneously-generated externality and seeks robustness with respect to the distribution of externalities—in particular, the unknown correlation between demand and externality.² In this setting, although the regulator can choose any incentive-compatible policy, optimal policies take a simple form: price regulation means uniform taxes or subsidies, while quantity regulation means a quantity ceiling or floor *without* other instruments. This delivers a unique robust policy and a clean instrument-choice implication: when the sign of this correlation is unknown, quantity regulation is optimal, whereas when the sign is known, price regulation can be optimal.

2 Framework

This section develops the regulator's problem in terms of regulating a positive consumption externality and discusses modeling assumptions. As Section 4 demonstrates, the model extends straightforwardly to environments with negative externalities, regulatory uncertainty, and costly screening.

2.1 Model

There is a unit mass of risk-neutral agents in a market for a homogeneous good. Consumption of the good generates a positive externality. Agents differ both in their consumption preferences and in the amount of externality they generate from consumption. The good is supplied competitively at a constant marginal cost, $c > 0$.

² This informational structure (where only mean or moment-based information is available) has been considered by other authors, such as Carrasco, Luz, Kos, Messner, Monteiro and Moreira (2018) and Che and Zhong (2025), who build on Bergemann and Schlag's (2011) analysis of monopoly pricing under worst-case objectives when the seller is uncertain about demand.

Heterogeneity in consumption preferences is captured by a *type*, $\theta \in \Theta \subset \mathbb{R}_+$. Each agent is privately informed of θ , which determines the utility $u(q, \theta)$ that he derives from consuming a quantity q of the good. To ensure that consumption is always bounded, I assume that there exists a maximum quantity $A > 0$ that each agent can consume, so that $q \in [0, A]$. As is standard, $u : [0, A] \times \Theta \rightarrow \mathbb{R}$ is assumed to be increasing and strictly concave in quantity and to satisfy the single-crossing property: $u_q > 0$, $u_{qq} < 0$, and $u_{q\theta} > 0$ in the interior of the domain of u . The type space Θ is assumed to be compact and convex (i.e., a closed interval of \mathbb{R}_+), and types have a continuously differentiable cumulative distribution function F that admits a positive density on Θ .

Heterogeneity in externality is captured by $\xi \in \mathbb{R}_+$, which denotes the marginal benefit (measured in dollar terms) that the agent generates per unit of consumption. For simplicity, agents are also assumed to be privately informed about ξ (as shown below, whether agents observe ξ does not affect the analysis).

Each agent's utility depends on both his own consumption of the good and the aggregate externality generated by all agents. For expositional simplicity, I focus on the case where utility is additively separable and linear in the aggregate externality. Denote the joint distribution of (θ, ξ) by H . Thus, a type- θ agent who consumes a quantity $q(\theta, \xi)$ of the good, generates an externality of ξ per unit of consumption, and makes a payment of $t(\theta, \xi)$ obtains a utility of

$$u(q(\theta, \xi), \theta) - t(\theta, \xi) + E, \quad \text{where } E = \iint_{(\hat{\theta}, \hat{\xi}) \in \Theta \times \mathbb{R}_+} \hat{\xi} q(\hat{\theta}, \hat{\xi}) dH(\hat{\theta}, \hat{\xi}).$$

A regulator seeks to regulate the positive externality generated by consumption. To this end, she chooses a *mechanism* (q, t) , which consists of an allocation function $q : \Theta \times \mathbb{R}_+ \rightarrow [0, A]$ and a payment function $t : \Theta \times \mathbb{R}_+ \rightarrow \mathbb{R}$. These respectively determine the quantity of the good that each agent consumes and the payment that he makes. By the revelation principle (Myerson, 1979), it is without loss of generality for the regulator to consider only incentive-compatible mechanisms under which agents report (θ, ξ) truthfully, such that for all $(\theta, \xi) \in \Theta \times \mathbb{R}_+$,

$$(\theta, \xi) \in \arg \max_{(\hat{\theta}, \hat{\xi})} [u(q(\hat{\theta}, \hat{\xi}), \theta) - t(\hat{\theta}, \hat{\xi}) + E]. \quad (\text{IC})$$

Notice that individual reports $(\hat{\theta}, \hat{\xi})$ do not affect the aggregate externality E in the market; therefore, as in the standard Pigouvian analysis, each agent fails to internalize how much his own consumption affects others. Standard arguments from mechanism design imply that (q, t) must be essentially independent of ξ : as can be seen from (IC), each agent's report of ξ affects his payoff only through the mechanism.³ Consequently, no incentive-compatible mechanism can truthfully elicit information about ξ ; instead, the regulator must form beliefs about ξ based only on agents' revealed consumption preferences. Given that incentive-compatible mechanisms can truthfully elicit only information about θ , implementable allocation functions⁴ must be nondecreasing in θ due to an adaptation of Myerson's (1981) lemma:

Lemma 1. *Define*

$$\mathcal{Q} := \{q : \Theta \rightarrow [0, A] \text{ is nondecreasing}\}.$$

Then an allocation function $q : \Theta \times \mathbb{R}_+ \rightarrow [0, A]$ is implementable only if there exists $\hat{q} \in \mathcal{Q}$ such that $q(\theta, \xi) = \hat{q}(\theta)$ for almost every $(\theta, \xi) \in \Theta \times \mathbb{R}_+$.

Lemma 1 is proven formally in Appendix A and shows that the formulation—and thus solution—of the regulator's problem does not depend on whether agents observe the amount of externality they generate. If agents do not observe ξ , the regulator can trivially do no better than to elicit only information about θ . However, even if agents observe ξ , Lemma 1 shows that the regulator can also do no better. Henceforth, \mathcal{Q} defined in Lemma 1 is referred to as the set of implementable allocation functions; and q and t in any incentive-compatible mechanism are written as functions of only θ .

Given any incentive-compatible mechanism (q, t) , the regulator evaluates its performance with respect to the worst-case total surplus. If the joint distribution H of (θ, ξ) were known, then the total surplus would be⁵

$$\int_{\Theta} [u(q(\theta), \theta) + (\mathbf{E}_H[\xi \mid \theta] - c) q(\theta)] dF(\theta).$$

³ See, for example, analogous characterization results by Jehiel and Moldovanu (2001), Che, Dessein and Kartik (2013), and Dworzak @ Kominers @ Akbarpour (2021).

⁴ Following the mechanism design literature, an allocation function $q : \Theta \times \mathbb{R}_+ \rightarrow [0, A]$ is said to be *implementable* if there exists a payment function $t : \Theta \times \mathbb{R}_+ \rightarrow \mathbb{R}$ such that the mechanism (q, t) satisfies (IC).

⁵ As monetary transfers are pure redistributions, the payment function t does not affect total surplus. Given any implementable allocation q , the associated payment function t is pinned down by the envelope theorem up to an additive constant (Milgrom and Segal, 2002). Consequently, the analysis below focuses on the choice of the allocation function q .

However, the regulator does not observe H and instead knows only the unconditional mean $\mathbf{E}_H[\xi] = \mu$. This reflects real-world settings in which policymakers know only the average marginal benefit, but not how marginal benefits are distributed or correlated with consumption. Below, the conditional mean of ξ is denoted by $m(\theta) := \mathbf{E}_H[\xi \mid \theta]$.

To evaluate the worst-case total surplus, the regulator considers three informational benchmarks:

(i) **Unknown correlation.** The regulator allows all conditional mean functions m in the set

$$\mathcal{M}_0 := \{m : \Theta \rightarrow \mathbb{R}_+ \text{ such that } m \text{ is measurable and } \mathbf{E}_F[m(\theta)] = \mu\}.$$

This captures settings where the regulator has no reliable information about how marginal benefits covary with consumption.

(ii) **Positive correlation.** The regulator allows all conditional mean functions m in the set

$$\mathcal{M}_+ := \{m : \Theta \rightarrow \mathbb{R}_+ \text{ such that } m \text{ is nondecreasing and } \mathbf{E}_F[m(\theta)] = \mu\}.$$

Here, higher-demand agents generate larger marginal benefits in expectation.

(iii) **Negative correlation.** The regulator allows all conditional mean functions m in the set

$$\mathcal{M}_- := \{m : \Theta \rightarrow \mathbb{R}_+ \text{ such that } m \text{ is nonincreasing and } \mathbf{E}_F[m(\theta)] = \mu\}.$$

Here, higher-demand agents generate smaller marginal benefits in expectation.

For each $* \in \{0, +, -\}$, the regulator's problem is therefore

$$\sup_{q \in \mathcal{Q}} \inf_{m \in \mathcal{M}_*} \int_{\Theta} [u(q(\theta), \theta) + [m(\theta) - c] q(\theta)] dF(\theta). \quad (\mathbf{P}_*)$$

Problem (\mathbf{P}_*) can be interpreted as a zero-sum game between the regulator and Nature. The regulator first chooses an implementable allocation function $q \in \mathcal{Q}$, after which Nature selects a conditional mean function $m \in \mathcal{M}_*$ that minimizes total surplus, subject to the constraints of the respective benchmark.

2.2 Discussion of Assumptions

To conclude this section, I discuss the key modeling choices underlying the framework and clarify which assumptions are essential for the subsequent analysis and main results.

Information about the externality. A key assumption is that the regulator knows the average marginal externality $\mu = \mathbf{E}_H[\xi]$, but not how the marginal benefit is distributed across agents or how it covaries with demand. This assumption reflects real-world policy settings in which average effects can be estimated from aggregated data, even when individual-level marginal benefit is difficult to observe. For example, policymakers may know the efficacy of a vaccine and the average number of social interactions individuals engage in, but have far less reliable information about how the marginal social benefits from vaccination vary across individuals or how they are correlated with willingness to pay.

Information about demand. The model also assumes that the regulator knows the distribution of types, F . This reflects the fact that consumption behavior is typically easier to observe and measure than external harm: willingness to pay can be estimated from observed consumption choices, whereas marginal social benefits or costs cannot. Nevertheless, Section 3 shows that the regulator’s optimal mechanism for each benchmark requires only a modest amount of information about F : it depends on F at most through a one-dimensional sufficient statistic. This suggests that the regulator’s optimal mechanism is *a fortiori* also robust to demand uncertainty.

Demand and supply specifications. The assumptions of continuous demand, constant marginal cost, and constant marginal externality are made for ease of exposition. As Section 4 shows, the subsequent analysis and main results hold in the case of unit demand (e.g., vaccines) and for more general cost and externality functions. In addition, while I follow standard textbook analyses by assuming that agent utility is additively separable and linear in the aggregate externality, the subsequent analysis and main results readily generalize. Specifically, the form of the optimal regulation under each informational benchmark—and hence the comparison between price and quantity regulation—remains the same when agent utility $u(q, \theta, E)$ increases in a potentially nonlinear way with the aggregate positive externality E .

Intuitively, even when agent utility is neither additively separable nor linear in the aggregate externality, the aggregate externality implies a shadow cost that potentially varies with policy but does not change its optimal form.

3 Main Results

I now solve the regulator’s problem under the three informational benchmarks—unknown, positive, and negative correlation—and show how information about the correlation between demand and marginal externality shapes the optimal allocation.

For ease of notation, let $D(p, \theta)$ denote the demand function of an agent of type θ when the per-unit price of the good is p . The laissez-faire allocation is thus given by $q^{\text{LF}}(\theta) = D(c, \theta)$.

3.1 Unknown Correlation

I begin with the benchmark in which the regulator has no information about how the marginal externality covaries with demand beyond its average level.

Theorem 1. *Under the unknown-correlation benchmark, there exists a quantity floor $\underline{q} \in [0, A]$ such that the regulator’s unique optimal allocation function is*

$$q_0^*(\theta) = \max \left\{ q^{\text{LF}}(\theta), \underline{q} \right\}.$$

Theorem 1 shows that the optimal regulation has two important features that allow it to be interpreted as a form of quantity regulation. First, it imposes a quantity floor: no agent can consume less than \underline{q} units of the good. Second, it imposes no intervention above that floor: agents for whom the quantity floor does not bind simply consume their laissez-faire allocation.

The first feature is intuitive: a quantity floor maximally lower-bounds the amount of external benefit that any single agent can generate, subject to incentive compatibility. Under the unknown correlation benchmark, the regulator must guard against the possibility that marginal benefit is concentrated among low types. At the same time, incentive compatibility implies that any implementable allocation must be

nondecreasing in type (Lemma 1). As such, to maximally expand consumption (and hence externality generated) among the lowest types while respecting incentive compatibility, the regulator can do no better than to flatten the allocation function below some threshold by imposing a quantity floor.

The second feature is perhaps more surprising: any intervention above the quantity floor can hurt the regulator more than it can help. Notably, this feature does not hold in Bayesian models: as [Kaplow and Shavell \(2002\)](#) have observed, a general lesson of the Bayesian mechanism design literature is that nonlinear pricing can—and except in special cases, strictly—improve upon quantity regulations such as a quantity floor. By contrast, under the unknown-correlation benchmark, the regulator must instead guard against the possibility that intervention unnecessarily distorts consumption if high types turn out not to generate much marginal benefit. This force dominates above the quantity floor, meaning that laissez-faire consumption is optimal for these agents.

The proof of Theorem 1 formalizes these intuitions:

Proof. To begin, observe that Nature’s problem is to choose a conditional mean function $m \in \mathcal{M}_0$ to minimize external benefit, given the regulator’s choice of allocation function q :

$$\inf_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)q(\theta)] \geq \mathbf{E}_F[m(\theta)] \cdot \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta).$$

To show that equality is in fact attained, consider the sequence $\{m_n\}_{n=1}^\infty \subset \mathcal{M}_0$ defined by

$$m_n(\theta) := \begin{cases} 0, & \text{if } F(\theta) > 1/n, \\ n\mu, & \text{otherwise.} \end{cases}$$

Given the regulator’s choice of allocation function q , for any $\varepsilon > 0$, $\Pr_F[q(\theta) < \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) + \varepsilon] > 0$ by the definition of the essential infimum. Consequently, there exists $N_\varepsilon \in \mathbb{N}$ such that

$$\Pr_F \left[q(\theta) < \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) + \varepsilon \right] \geq 1/N_\varepsilon.$$

Since q is nondecreasing by Lemma 1, this choice of N_ε yields

$$\mathbf{E}_F[m_{N_\varepsilon}(\theta)q(\theta)] \leq \mu \left[\operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) + \varepsilon \right].$$

In turn, this implies that

$$\lim_{\varepsilon \rightarrow 0} \mathbf{E}_F[m_{N_\varepsilon}(\theta)q(\theta)] = \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) \implies \inf_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)q(\theta)] = \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta). \quad (1)$$

This shows that Nature's best response is to concentrate marginal benefit on low types, in line with the intuition given above.⁶

Importantly, the worst-case external benefit in equation (1) depends on the regulator's choice of q only through its essential infimum, $\operatorname{ess\,inf}_{\theta \in \Theta} q(\theta)$. Thus, the regulator's problem (\mathbf{P}_0) can be rewritten as

$$\begin{aligned} & \sup_{q \in \mathcal{Q}} \left[\int_{\Theta} [u(q(\theta), \theta) - cq(\theta)] \, dF(\theta) + \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) \right] \\ &= \sup_{\underline{q} \in [0, A]} \left[\sup_{q \in \mathcal{Q}} \left\{ \int_{\Theta} [u(q(\theta), \theta) - cq(\theta)] \, dF(\theta) : \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) = \underline{q} \right\} + \mu \underline{q} \right]. \end{aligned} \quad (2)$$

For any choice of $\underline{q} \geq q^{\text{LF}}(\underline{\theta})$, the solution to the inner constrained problem in equation (2) is given by the pointwise maximizer $q^{\text{LF}}(\theta)$ unless the constraint is binding, yielding the desired solution form

$$q_0^*(\theta) = \max \left\{ q^{\text{LF}}(\theta), \underline{q} \right\}.$$

This solution form remains the same if $\underline{q} < q^{\text{LF}}(\underline{\theta})$. This solution is also unique since u is strictly concave in q , which means that the regulator's objective function above is strictly concave in q . \square

The proof of Theorem 1 shows that the quantity floor in the optimal regulation can be characterized by a simple first-order condition whenever it is interior. Specifically, the solution to the outer problem in

⁶ Technically, Nature does not always have a best response as the worst-case external benefit is generally attained only in the limit. However, a best response exists for the optimal allocation function chosen by the regulator.

equation (2) satisfies

$$\mathbf{E}_F \left[\left[c - u_q(\underline{q}, \theta) \right] \mathbf{1}_{q^{\text{LF}}(\theta) < \underline{q}} \right] = \mu. \quad (3)$$

In words, the optimal quantity floor equates the marginal gain in allocative efficiency from lowering the quantity floor—due to agents whose laissez-faire consumption falls below \underline{q} —to the average marginal external benefit.

The key technical insight exploited in the proof is that Nature’s limiting best response results in a worst-case externality that depends on the regulator’s allocation function only through a one-dimensional statistic: its essential infimum. This allows the regulator’s problem to be rewritten as a maximization problem with a single constraint on $\text{ess inf}_{\theta \in \Theta} q(\theta)$, rather than an infinite-dimensional constraint that would arise if Nature’s best response depended on the entire allocation function.

An implication of Theorem 1 is that the optimal regulation places relatively limited informational requirements on the distribution of types. Unlike Bayesian mechanism design models, which typically require estimation of the entire distribution F , the optimal regulation here depends on F only through the condition in equation (3). Thus, implementing the optimal regulation requires only information about a welfare statistic—specifically, how much allocative efficiency is lost by imposing the quantity floor—rather than detailed knowledge of the full demand distribution. This property emerges as a consequence of robustness to correlation uncertainty, rather than an *a priori* restriction imposed on the regulator.

While Theorem 1 uses only information about the average marginal benefit μ , the same analytical techniques extend to ambiguity sets that impose additional structure on the distribution of ξ . To illustrate, suppose the regulator also knows that $\xi \in [0, \bar{\xi}]$ for some $\bar{\xi} \geq \mu$. This restriction implies $0 \leq m(\theta) \leq \bar{\xi}$; since $\mathbf{E}_F[m(\theta)] = \mu$, any feasible m must be positive on a set of F -measure at least $\alpha := \mu/\bar{\xi}$, so Nature can no longer concentrate the entire mean μ on a vanishing set of types. Moreover, to minimize $\mathbf{E}_F[m(\theta)q(\theta)]$ subject to these constraints, Nature assigns $m(\theta) = \bar{\xi}$ on an α -measure set of types with the smallest allocation probabilities and $m(\theta) = 0$ elsewhere. As a result, the essential infimum term in equation (1) is replaced by a lower-tail expected shortfall functional:

$$\inf_{0 \leq m \leq \bar{\xi}} \{ \mathbf{E}_F[m(\theta)q(\theta)] : \mathbf{E}_F[m(\theta)] = \mu \} = \bar{\xi} \int_0^\alpha q(F^{-1}(u)) \, du = \mu \cdot \frac{1}{\alpha} \int_0^\alpha q(F^{-1}(u)) \, du.$$

Accordingly, the objective still depends on q only through a one-dimensional functional, and the resulting mechanism design problem can be analyzed using the same techniques as in the proof of Theorem 1. Similar arguments apply when the regulator has additional moment information (e.g., a variance bound) rather than a support bound on ξ .

Although these extensions are analytically tractable, I view them as more artificial from a policy perspective. In many real-world applications, policymakers lack credible information about the *shape* of the externality distribution—especially its tails—beyond coarse evidence about an average effect. In the exceptional cases where such information *is* available, it is arguably more natural to take a Bayesian approach. Nevertheless, when these additional restrictions are loose—so substantial residual uncertainty about heterogeneity remains (e.g., $\bar{\xi}$ is large or moment bounds are weak)—it can be shown that the optimal policies are well-approximated by the quantity-control characterization in Theorem 1. For these reasons, I do not pursue these extensions here.

3.2 Positive vs. Negative Correlation

I next consider the benchmark in which the regulator knows the sign of the correlation between demand and marginal externality.

Theorem 2.

- (a) *Under the positive-correlation benchmark, the regulator's unique optimal allocation function can be implemented by a per-unit subsidy equal to μ :*

$$q_+^*(\theta) = D(c - \mu, \theta).$$

- (b) *Under the negative-correlation benchmark, there exists a quantity floor $\underline{q} \in [0, A]$ such that the regulator's unique optimal allocation function is*

$$q_-^*(\theta) = q_0^*(\theta) = \max \left\{ q^{\text{LF}}(\theta), \underline{q} \right\}.$$

Theorem 2 shows that the sign of the correlation determines whether price or quantity regulation is optimal. When the correlation is positive, the optimal regulation is a uniform subsidy equal to the average marginal benefit. When the correlation is negative, the optimal regulation consists of a quantity floor.

The negative-correlation benchmark parallels the unknown-correlation benchmark in Theorem 1. In particular, the regulator must guard against the same worst-case possibility—namely, that marginal benefit is concentrated among agents with low demand. Consequently, the optimal allocation coincides with that in Theorem 1 and again takes the form of a quantity floor above which there is no intervention.

The positive-correlation benchmark is qualitatively different. Nature can no longer concentrate high externalities on low-demand agents, and as the proof of Theorem 2 below shows, Nature’s best response is degenerate: $m(\theta) \equiv \mu$, which removes selection altogether. The regulator therefore faces an effectively homogeneous externality of size μ , and the optimal policy collapses to a uniform subsidy.

Perhaps interestingly, the optimal regulation in the positive-correlation benchmark sets a subsidy equal to the average marginal benefit, μ . In Bayesian models, this is *not* the optimal uniform subsidy. In fact, [Diamond \(1973\)](#) and subsequent authors caution against this “naïve” subsidy, showing that policymakers can generally do better by setting a subsidy equal to a weighted average of marginal benefits, with weights given by demand responses. By contrast, Theorem 2 provides a robustness-based argument for why this “naïve” subsidy can be optimal when the sign of the correlation is known but more detailed information about its structure is unavailable.

It should be noted that *no* information about the distribution of types is required under the positive-correlation benchmark. Thus, the optimal regulation in this case is also robust to demand uncertainty.

The proof of Theorem 2 shows why the optimal regulation in the positive-correlation benchmark is simply a uniform subsidy equal to μ :

Proof. To prove the result for the positive-correlation benchmark,⁷ observe that the worst-case external benefit can be bounded more tightly than in Theorem 1. Given the regulator’s choice of allocation function q ,

$$\inf_{m \in \mathcal{M}_+} \mathbf{E}_F[m(\theta)q(\theta)] \geq \mathbf{E}_F[m(\theta)] \cdot \mathbf{E}_F[q(\theta)] = \mu \mathbf{E}_F[q(\theta)].$$

⁷ As noted above, the result for the negative-correlation benchmark follows directly from the proof of Theorem 1.

The inequality above follows by the continuous version of Chebyshev's sum inequality. To show that equality is in fact attained, consider Nature's best response: $m_+^*(\theta) \equiv \mu$. Given this strategy,

$$\mathbf{E}_F[m_+^*(\theta)q(\theta)] = \mu \mathbf{E}_F[q(\theta)] = \inf_{m \in \mathcal{M}_+} \mathbf{E}_F[m(\theta)q(\theta)].$$

This shows that Nature's best response is weakly dominant and distributes marginal benefit equally.

As in the proof of Theorem 1, the worst-case external benefit under the positive-correlation benchmark depends on the regulator's choice of q only through a one-dimensional statistic: in this case, its average, $\mathbf{E}_F[q(\theta)]$. Thus, the regulator's problem (\mathbf{P}_+) can be rewritten as

$$\sup_{q \in \mathcal{Q}} \left[\int_{\Theta} [u(q(\theta), \theta) - cq(\theta)] dF(\theta) + \mu \mathbf{E}_F[q(\theta)] \right] = \sup_{q \in \mathcal{Q}} \int_{\Theta} [u(q(\theta), \theta) - (c - \mu)q(\theta)] dF(\theta).$$

The solution to the regulator's problem is therefore given by the pointwise maximizer $q_+^*(\theta) = D(c - \mu, \theta)$.

As in the proof of Theorem 1, uniqueness of the solution follows from the strict concavity of u . \square

4 Applications

This section applies the main results of Section 3 to four settings: vaccines (as in the motivating example of Section 1), negative externalities, regulatory uncertainty, and costly screening. Proofs are deferred to Appendix A.

4.1 Vaccines

To apply the main results to vaccines, I extend the framework of Section 2 to a setting with unit demand. Each agent demands at most one unit of the vaccine, and willingness to pay is denoted by θ . Interpreting $q(\theta)$ as the probability that an agent of type θ is vaccinated, the regulator chooses an allocation function $q : \Theta \rightarrow [0, 1]$. As Proposition 1 shows below, probabilistic rationing does not arise in the regulator's optimal mechanism; however, allowing for probabilistic allocations implies that the chosen mechanism is also optimal within a broader class.

Agent utility is given by $u(q, \theta) = \theta q$. This specification can be viewed as a limiting special case of the utility functions permitted in Section 2. In particular, although u is not strictly concave in q , it is the pointwise limit of the strictly concave family $u_n(q, \theta) = \theta q - q^2/n$ as $n \rightarrow +\infty$.

Vaccination generates an external benefit $\xi \in \mathbb{R}_+$ that varies across agents. As motivated in Section 1, this may reflect heterogeneity in individuals' propensity to wear masks or avoid large crowds. In turn, ξ may be positively or negatively correlated with θ , depending on whether willingness to pay is associated with greater exposure or greater precaution.

The following proposition extends Theorems 1 and 2 to this setting. As before, the vaccine is supplied competitively at a constant marginal cost c , and the regulator maximizes worst-case total surplus subject to the three informational benchmarks.

Proposition 1.

(a) *Under the unknown-correlation benchmark and the negative-correlation benchmark, the regulator's optimal allocation function is*

$$q^*(\theta) = \begin{cases} 1, & \text{if } \mu \geq \mathbf{E}_F[(c - \theta)_+], \\ q^{\text{LF}}(\theta) = \mathbf{1}_{\theta > c}, & \text{if } \mu < \mathbf{E}_F[(c - \theta)_+]. \end{cases}$$

(b) *Under the positive-correlation benchmark, the regulator's optimal allocation function is*

$$q^*(\theta) = \mathbf{1}_{\theta > c - \mu}.$$

Proposition 1 shows that there are only three candidates for the regulator's optimal policy: a mandate ($q^* \equiv 1$), no intervention ($q^* = q^{\text{LF}}$), and a subsidy equal to the average external benefit μ .⁸ A subsidy is optimal only when willingness to pay is known to be positively correlated with the external benefit—for instance, when individuals at greater risk of exposure also have higher demand for vaccination. This

⁸ Here, the terms “mandate” and “subsidy” refer only to the induced allocation rule in a frictionless total-surplus benchmark. In practice, mandates may raise enforcement and exemption issues, and subsidies may require financing through distortionary taxes. These institutional margins are not modeled here as the goal of the application is to isolate how limited information about heterogeneity in external benefits can rationalize different allocation rules.

direction of selection is plausible for vaccines such as those for HPV and mpox, where individuals may seek vaccination in anticipation of increased infection risk. By contrast, if the direction of selection is negative or ambiguous—such as when individuals with higher demand for vaccination also engage in more precautionary behavior, as may have occurred during the COVID-19 pandemic—then the optimal policy is either a mandate or no intervention.

The regulator’s choice between a mandate and no intervention is determined by a comparison between the average allocative deadweight loss imposed by a mandate and the average external benefit from vaccination. Vaccinating an agent with willingness to pay θ imposes an allocative deadweight loss of $(c - \theta)_+$; thus, $\mathbf{E}_F[(c - \theta)_+]$ represents the average loss in allocative efficiency generated by a mandate. When the average external benefit μ exceeds this loss, mandating vaccination yields a positive welfare guarantee and is optimal. However, when μ is smaller than $\mathbf{E}_F[(c - \theta)_+]$, the regulator optimally refrains from intervention as *any* intervention would generate more allocative inefficiency than external benefit in the worst case. This condition extends the logic of the tradeoff characterized in equation (3).

Proposition 1 implies that mandates are optimal in a wider set of environments than a Bayesian model would suggest. If the true joint distribution H of (θ, ξ) were known, the regulator would be able to exploit information about the conditional average external benefit $m(\theta) = \mathbf{E}_H[\xi | \theta]$. It can be verified that a Bayesian regulator (who maximizes expected total surplus) optimally employs a mandate if and only if for every $\hat{\theta} \in [\underline{\theta}, \bar{\theta}]$,

$$\int_{\underline{\theta}}^{\hat{\theta}} [\theta - c + m(\theta)] dF(\theta) \geq 0. \quad (4)$$

This condition is strictly stronger than that in Proposition 1: taking $\hat{\theta} = c$, equation (4) implies that

$$\mathbf{E}_F[(c - \theta)_+] \leq \mathbf{E}_F[m(\theta)\mathbf{1}_{\theta \leq c}] \leq \mathbf{E}_F[m(\theta)] = \mu.$$

The reverse implication does not generally hold. Consequently, the robust criterion used here rationalizes mandates in environments where a Bayesian model—requiring substantially more information—would not. This may help explain why vaccine mandates have been so widely adopted, including during the COVID-19 pandemic.

Perhaps surprisingly, Proposition 1 can also rationalize no intervention. With a known joint distribution, a Bayesian regulator subsidizes vaccination whenever expanding participation below c yields positive expected social surplus for some types: that is, whenever $\theta - c + m(\theta) > 0$ on a set of positive measure for $\theta \leq c$. By contrast, under the robust criterion considered here, limited information about how external benefits are distributed across agents can make any intervention welfare-reducing in the worst case. In such environments, refraining from intervention maximizes the regulator's welfare guarantee despite the presence of a positive externality.

Finally, the proof of Proposition 1 in Appendix A shows that, despite agent utility no longer being strictly concave as in Section 3, the optimal allocation functions characterized above are unique except in a knife-edge case. Specifically, when $\mu = \mathbf{E}_F[(c - \theta)_+]$, the regulator is indifferent across a continuum of optimal allocations that include both a mandate and no intervention; but in all other cases, the optimal allocation function is unique.

4.2 Negative Externalities

Although I have so far focused on a framework with positive externalities, my analysis extends to the case of negative externalities. Specifically, suppose that a type- θ agent who consumes a quantity $q(\theta, \xi)$ of the good, generates a negative externality of ξ per unit of consumption, and makes a payment of $t(\theta, \xi)$ obtains a utility of

$$u(q(\theta, \xi), \theta) - t(\theta, \xi) - E, \quad \text{where } E = \iint_{(\hat{\theta}, \hat{\xi}) \in \Theta \times \mathbb{R}_+} \hat{\xi} q(\hat{\theta}, \hat{\xi}) dH(\hat{\theta}, \hat{\xi}).$$

Relative to the results of Section 3, quantity ceilings—rather than quantity floors—become optimal when there are negative externalities. The following result extends Theorem 1:

Theorem 1'. *Under the unknown-correlation benchmark for negative externalities, there exists a quantity ceiling $\bar{q} \in [0, A]$ such that the regulator's unique optimal allocation function is*

$$q_0^*(\theta) = \min \{q^{\text{LF}}(\theta), \bar{q}\}.$$

Moreover, the results for positive and negative correlation are reversed. The following result extends Theorem 2 to the case of negative externalities:

Theorem 2’.

- (a) *Under the positive-correlation benchmark for negative externalities, there exists a quantity ceiling $\bar{q} \in [0, A]$ such that the regulator’s unique optimal allocation function is*

$$q_+^*(\theta) = q_0^*(\theta) = \min \{q^{\text{LF}}(\theta), \bar{q}\}.$$

- (b) *Under the negative-correlation benchmark for negative externalities, the regulator’s unique optimal allocation function can be implemented by a per-unit tax equal to μ :*

$$q_-^*(\theta) = D(c + \mu, \theta).$$

Theorems 1’ and 2’ show that with negative externalities, the correlation–instrument mapping reverses relative to Section 3. When the direction of correlation is unknown, and also when willingness to pay is known to be positively correlated with marginal harm, the regulator optimally uses a quantity ceiling, leaving laissez-faire behavior unchanged below the cap. By contrast, when willingness to pay is known to be negatively correlated with marginal harm, a uniform tax equal to the average marginal external cost μ is robustly optimal. Thus, absent *a priori* evidence that high-demand agents generate less marginal harm, the robust criterion favors quantity regulation—quotas, limits, or bans—over price regulation.

4.3 Regulatory Uncertainty

While the analysis so far has focused on heterogeneity across agents, this subsection turns to uncertainty faced by the regulator—a classic problem studied by Weitzman (1974) and a large subsequent literature in environmental economics, law and economics, and other fields. In this setting, my results show that the desire for robustness with respect to uncertainty can help explain the use of quantity regulation.

Consider a regulator who wishes to regulate an externality-producing industry by reducing the amount of the externality it produces. Let q denote the level of externality reduction. For simplicity, suppose that the marginal benefit of externality reduction is constant. This assumption is a reasonable approximation when the industry has a small overall impact on the environment and is commonly imposed in textbook treatments of the problem (see, e.g., Chapter 21 of [Viscusi, Harrington and Vernon, 2005](#)). By contrast, the industry’s marginal cost of externality reduction is increasing, reflecting rising abatement costs.

The regulator faces uncertainty over both the cost and benefit of externality reduction. This uncertainty is captured by the random variables $\theta \in \Theta$ and $\xi \in \mathbb{R}_+$, which represent the aggregate industry’s cost and benefit shifters. The cost function $C : \mathbb{R}_+ \times \Theta \rightarrow \mathbb{R}_+$ is increasing and strictly convex in q ($C_q > 0$ and $C_{qq} > 0$), and higher values of θ correspond to higher marginal costs ($C_{q\theta} > 0$). Likewise, ξ shifts the marginal benefit curve: the marginal benefit of reducing the externality by an additional unit is ξ .

Importantly, cost and benefit shifters may be correlated in an unknown way. For example, a global demand shock can raise both the marginal benefits of abatement (because baseline emissions from other industries are higher) and the marginal costs of abatement (because reductions must be achieved at higher output levels). Conversely, if scarce innovation effort must be allocated between abatement technologies (which lower θ) and mitigation or capture technologies (which lower ξ), then cost and benefit shifters may be negatively correlated. [Stavins \(1996\)](#) discusses other examples of how correlation may arise.

This setup departs from the previous literature along two dimensions: the set of feasible policies and the regulator’s objective.

First, the regulator chooses a potentially nonlinear payment schedule $T : \mathbb{R}_+ \rightarrow \mathbb{R} \cup \{+\infty\}$, under which the industry pays $T(q)$ when it reduces q units of the externality. In particular, the regulator is not restricted to particular instruments (such as constant marginal prices or fixed quantity instruments), as in the models of [Weitzman](#) and [Stavins](#)—and indeed much of the subsequent literature, as critiqued by [Kaplow and Shavell \(2002\)](#).

Second, the regulator does not know the joint distribution of (θ, ξ) and does not treat uncertainty over θ and ξ symmetrically. As [Stavins](#) notes, “more often than not, it is benefit uncertainty that seems to be of substantially greater magnitude.” For example, while policymakers often have access to information about marginal costs through abatement technologies and compliance costs, marginal benefits depend on

more complex processes—such as environmental damages, health impacts, or ecological responses—that are difficult to measure and may vary substantially across states of the world. Accordingly, the regulator is assumed to know only $\mu = \mathbf{E}[\xi]$ and the marginal distribution F of θ , and she evaluates each policy according to its worst-case ex-ante welfare performance.

The game proceeds as follows:

1. The regulator commits to a payment schedule $T : \mathbb{R}_+ \rightarrow \mathbb{R} \cup \{+\infty\}$.
2. The cost and benefit shifters θ and ξ are realized and observed by the industry.
3. The industry chooses the level of externality reduction q and pays $T(q)$.

While this setup differs slightly from the model in Section 2, the same analysis applies. In particular, θ and ξ can equivalently be interpreted as the industry’s private information, and the regulator’s problem can be formulated as one of choosing an incentive-compatible mechanism. Thus, Theorem 1 implies the following result.

Proposition 2. *There exists a quantity floor \underline{q} such that the regulator’s optimal payment schedule is, up to a lump-sum transfer that does not affect total surplus,*

$$T^*(q) = \begin{cases} 0, & \text{if } q \geq \underline{q}, \\ +\infty \text{ (or a sufficiently large payment),} & \text{if } q < \underline{q}. \end{cases}$$

Even though the regulator can, in principle, choose any nonlinear payment schedule, Proposition 2 shows that the optimal policy implements a binding quantity requirement. Under the optimal payment schedule T^* , the industry chooses the minimum level of externality reduction $q = \underline{q}$ in every realization of (θ, ξ) . Because externality reduction is costly, favorable cost realizations reduce the cost of meeting the requirement but do not induce additional abatement beyond \underline{q} . Thus, robustness collapses the optimal policy to a noncontingent quantity rule as considered by Weitzman and Stavins, despite the regulator’s ability to employ a much richer class of instruments.

As in the quantity floor in Theorem 1, the quantity floor \underline{q} in Proposition 2 can be characterized by a simple first-order condition whenever it is interior:

$$\mathbf{E}_F[C_q(\underline{q}, \theta)] = \mu. \quad (5)$$

When the expected marginal cost always exceeds the expected marginal benefit (i.e., $\mathbf{E}_F[C_q(0, \theta)] > \mu$), the optimal policy is no intervention. Conversely, when the expected marginal cost is always less than the expected marginal benefit (i.e., $\lim_{\underline{q} \rightarrow +\infty} \mathbf{E}_F[C_q(\underline{q}, \theta)] < \mu$), the optimal quantity floor is to ban the production of the externality entirely.

Despite the intuitive appeal of the first-order condition (5), it should be emphasized that it does *not* arise in traditional Bayesian models. As noted by Roberts and Spence (1976) and Kaplow and Shavell (2002), for example, the Bayesian-optimal allocation satisfies $C_q(q(\theta), \theta) = \mathbf{E}_F[\xi \mid \theta]$ *pointwise* for each θ , rather than *in expectation* over θ .

The optimal policy characterized in Proposition 2 also admits a natural interpretation as a cap-and-trade policy. The quantity floor \underline{q} on abatement levels equivalently imposes a cap on the total amount of externality produced by the industry, which can be implemented by allocating permits to firms in the industry. While trading permits is not explicitly modeled, it helps allocate abatement efforts more efficiently across firms within the industry, which determines the industry's cost function $C(q, \theta)$. This logic mirrors the operation of real-world cap-and-trade programs such as the U.S. Acid Rain Program for sulfur dioxide, the European Union Emissions Trading System, California's cap-and-trade program, and the Regional Greenhouse Gas Initiative.

While Proposition 2 focuses only on the case in which the correlation between cost and benefit shifters is unknown, Theorem 2 extends to this setting as well. In particular, if the regulator knows that the cost and benefit shifters are negatively correlated—so that states of the world in which abatement is cheaper tend to also be those in which abatement is more valuable—then a per-unit subsidy for externality reduction, equal to the average marginal benefit μ , becomes optimal. As noted previously, such a correlation could arise if scarce innovation effort must be allocated between abatement and capture technologies. However, where the direction of correlation is uncertain (for example, advances in abatement might have positive

spillovers into capture technologies and vice versa), then the analysis favors quantity regulation.

Beyond environmental regulations, Proposition 2 also has implications for a broader set of policy settings in which regulators face uncertainty about costs and benefits. The insights of Weitzman (1974) have been applied to tariff policy in international trade (e.g., Dasgupta and Stiglitz, 1977 and Young and Anderson, 1980), to safety regulation (e.g., Shavell, 1984), and to innovation policy (e.g., Wright, 1983). In these settings, as in the present framework, regulators choose between price and quantity instruments under limited information, and the robustness considerations developed here provide a complementary perspective on why quantity regulations are often observed in practice.

Finally, it is important to emphasize the limited normative scope of Proposition 2. The analysis characterizes the policy a regulator would choose when information about costs and benefits is scarce and the objective is to maximize a worst-case welfare guarantee. While this may well reflect the informational constraints under which some policies are made, it does not imply that policymakers should ignore opportunities to acquire additional information. In practice, policymakers can invest in monitoring and data collection, and should optimally trade off the costs of information acquisition against the benefits of implementing more finely targeted policies.

4.4 Costly Screening

The final application in this section extends the framework of Section 2 to costly screening environments. In such environments, a good is allocated through a costly and socially wasteful action, such as waiting in line. Costly screening has been used to study a wide range of economic settings, including bargaining, lobbying, and redistribution programs. I build on a standard model of costly screening and show that robustness with respect to correlation can rationalize the use of “non-price mechanisms” that rely on random allocation rather than screening.

There is a quantity $Q \in (0, 1)$ of a good to be allocated to a unit mass of risk-neutral agents. Each agent demands at most one unit of the good and is privately informed of his value v . Each agent can also wait for time $t \in \mathbb{R}_+$ and is privately informed of his waiting cost w per unit time. Thus, given a probability $q \in [0, 1]$ of receiving the good after time t , an agent’s payoff is $vq - wt$.

A regulator designs the allocation mechanism under limited information about agents' preferences. Specifically, the regulator observes the distribution F of the marginal rate of substitution $\theta = v/w \in \Theta$, but not the joint distribution H of (v, w) . In addition, she knows the average waiting cost $\mathbf{E}_H[w] = \mu$.

Given this informational constraint, the regulator seeks to maximize worst-case aggregate surplus across all joint distributions consistent with these observables. The regulator chooses a direct mechanism (q, t) , consisting of an allocation function q and a wait-time function t . As Dworzak et al. (2021) show, the analog of Lemma 1 in this environment implies that the mechanism can truthfully elicit only the marginal rate of substitution θ , rather than v and w separately.

Despite the apparent differences in economic environments, Proposition 3 shows that the insights of Theorem 1 carry over to costly screening: non-price allocation is optimal when the regulator wishes to ensure robustness to the correlation between agents' marginal rates of substitution and waiting costs.

Proposition 3. *Under the unknown-correlation benchmark, the mechanism $(q^*, t^*) = (Q, 0)$ is optimal. Moreover, it is the unique optimal mechanism if $\underline{\theta} > 0$.*

Under the mechanism $(q^*, t^*) = (Q, 0)$ in Proposition 3, the regulator randomly allocates the good without agents engaging in costly screening. This can be implemented, for example, via a lottery under which each agent receives the good with probability Q .

Proposition 3 helps rationalize the widespread use of institutional arrangements that deliberately avoid costly screening, even in settings where such screening could improve allocative efficiency under stronger informational assumptions. In the context of dispute resolution, for instance, arbitration and settlement procedures can be understood as responses to uncertainty about how litigation costs correlate with claim strength. Another example is the increasing popularity of lotteries—rather than ordeals (Nichols and Zeckhauser, 1982)—in the allocation of public housing, vouchers, school seats, and other public services. Whereas such lotteries are often justified on equity or administrative grounds, Proposition 3 shows that they are also consistent with maximizing worst-case efficiency.

From a technical perspective, Proposition 3 shows that the max–min logic of Theorem 1 extends to environments in which the main friction affects agents' payoffs directly (here, through waiting costs), rather than entering only through the regulator's objective via the allocation rule. In the application to

costly screening, the regulator effectively chooses a utility schedule through the joint determination of allocation probabilities and waiting times.⁹ As shown in Appendix A, robustness requires applying the earlier max–min arguments to utility rather than to the allocation rule itself.

5 Conclusion

Much of the “prices versus quantities” debate starts from Weitzman’s (1974) classic benchmark in which the regulator compares relatively simple instruments under relatively rich knowledge of primitives. A subsequent literature emphasizes that, once a broader set of instruments is allowed, price-based policies can dominate quantity controls (Kaplow and Shavell, 2002).

This paper departs from that benchmark on both dimensions. The regulator is not restricted to simple instruments: through a mechanism-design formulation of the instrument choice problem, she may choose any incentive-compatible policy, including nonlinear price schedules. At the same time, the regulator has limited information about how willingness to pay covaries with marginal externalities, and thus evaluates policies by their worst-case welfare performance subject only to information about the average externality.

The first main result provides a new explanation for the prevalence of quantity regulation based on heterogeneity and robustness. When the sign of the correlation between willingness to pay and marginal externalities is unknown, any policy that operates through selection—including sophisticated nonlinear price schedules—can be driven to perform poorly by adverse correlation patterns. Quantity regulation instead delivers a welfare guarantee by directly controlling externality-relevant behavior, and under the max–min criterion it is uniquely optimal: a quantity floor for positive externalities and a quantity ceiling for negative externalities (with mandates and bans as unit-demand special cases).

The second main result characterizes when price regulation is justified and why it should be simple. Once the sign of the correlation is known, the adverse selection patterns that drive the worst case are excluded. In the baseline model with positive externalities, if the correlation is positive, price regulation becomes optimal, and the optimal policy takes a starkly simple form: a uniform subsidy equal to the

⁹ Concretely, the worst-case term takes the form $\mathbf{E}_F[m(\theta)q(\theta)]$ in the framework of Section 2, while it takes the form $\mathbf{E}_F[m(\theta)U(\theta)]$ in the present application, where $U(\theta) := \theta q(\theta) - t(\theta)$ is the induced utility schedule.

average marginal externality. If the correlation is negative, the argument for quantity regulation continues to apply. In this sense, the determinant of instrument choice is not whether the regulator can implement rich mechanisms—she can—but whether she can credibly sign the correlation that governs selection.

A natural question is whether the correlation patterns that matter for robustness are economically meaningful. On the one hand, the signs of these correlations do not require extreme assumptions or implausible behavior. In many environments, willingness to pay and marginal social costs are shaped by common underlying factors, so selection is a realistic possibility rather than a theoretical curiosity. For example, individuals who are most eager to get vaccinated may already engage in extensive precautionary behavior, reducing the marginal social benefit of their vaccination; conversely, individuals with the lowest willingness to pay may be those whose behavior generates the largest marginal social benefit. On the other hand, a growing empirical literature suggests that externalities are often driven by tail behavior. For example, [Griffith et al. \(2019\)](#) note that alcohol-related externalities are highly nonlinear in consumption. In the United States, 8% (respectively, 3%) of alcohol drinkers account for 51% (respectively, 28%) of total drinks consumed ([Esser, Sacks, Sherk, Karriker-Jaffe, Greenfield, Pierannunzi and Brewer, 2020](#)). Similarly, [Knittel and Sandler \(2018\)](#) find strong skew in the distribution of tailpipe emissions and observe that the marginal social costs of accidents and congestion are highly skewed as well. These patterns are consistent with the worst-case correlations characterized in this paper.

Finally, a broader implication of this paper is that quantity regulation is a robust welfare benchmark under limited correlation information. Beyond externalities, the same logic extends to environments with costly and socially wasteful screening, where robustness can favor rules that suppress screening—such as lotteries rather than ordeals ([Nichols and Zeckhauser, 1982](#))—and procedures that avoid costly litigation. The framework developed in this paper thus provides a unified way to understand when simple policy instruments are not administrative compromises, but rather justified by robustness concerns in response to informational constraints. Exploring additional applications of this framework, as well as the role of information acquisition and commitment in mitigating robustness concerns, is a promising direction for future research.

References

- BAUMOL, W. J., AND W. E. OATES (1971): “The Use of Standards and Prices for Protection of the Environment,” *Swedish Journal of Economics*, 73, 42–54.
- (1988): *The Theory of Environmental Policy*, Cambridge, UK: Cambridge University Press, 2nd edition.
- BERGEMANN, D., AND K. SCHLAG (2011): “Robust Monopoly Pricing,” *Journal of Economic Theory*, 146, 2527–2543.
- BUCHANAN, J. M., AND G. TULLOCK (1975): “Polluters’ Profits and Political Response: Direct Controls Versus Taxes,” *American Economic Review*, 65, 139–147.
- 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.
- CARROLL, G. (2015): “Robustness and Linear Contracts,” *American Economic Review*, 105, 536–563.
- (2017): “Robustness and Separation in Multidimensional Screening,” *Econometrica*, 85, 453–488.
- (2019): “Robustness in Mechanism Design and Contracting,” *Annual Review of Economics*, 11, 139–166.
- CARROLL, G., AND L. BOLTE (2023): “Robust Contracting Under Double Moral Hazard,” *Theoretical Economics*, 18, 1623–1663.
- CHE, Y.-K., W. DESSEIN, AND N. KARTIK (2013): “Pandering to Persuade,” *American Economic Review*, 103, 47–79.
- CHE, Y.-K., AND W. ZHONG (2025): “Robustly Optimal Mechanisms for Selling Multiple Goods,” *Review of Economic Studies*, 92, 2923–2951.

- DASGUPTA, P., AND J. STIGLITZ (1977): “Tariffs vs. Quotas as Revenue Raising Devices under Uncertainty,” *American Economic Review*, 67, 975–981.
- DIAMOND, P. A. (1973): “Consumption Externalities and Imperfect Corrective Pricing,” *Bell Journal of Economics and Management Science*, 4, 526–538.
- DWORCZAK, P. & S. D. KOMINERS & M. AKBARPOUR (2021): “Redistribution Through Markets,” *Econometrica*, 89, 1665–1698.
- DWORCZAK, P., AND A. PAVAN (2022): “Preparing for the Worst but Hoping for the Best: Robust (Bayesian) Persuasion,” *Econometrica*, 90, 2017–2051.
- ESSER, M. B., J. J. SACKS, A. SHERK, K. J. KARRIKER-JAFFE, T. K. GREENFIELD, C. PIERANNUNZI, AND R. D. BREWER (2020): “Distribution of Drinks Consumed by U.S. Adults by Average Daily Alcohol Consumption: A Comparison of Two Nationwide Surveys,” *American Journal of Preventive Medicine*, 59, 669–677.
- FLEISCHER, V. (2015): “Curb Your Enthusiasm for Pigovian Taxes,” *Vanderbilt Law Review*, 68, 1673–1713.
- FRANKEL, A. (2014): “Aligned Delegation,” *American Economic Review*, 104, 66–83.
- GARRETT, D. F. (2014): “Robustness of Simple Menus of Contracts in Cost-Based Procurement,” *Games and Economic Behavior*, 87, 631–641.
- GLAESER, E. L., S. JOHNSON, AND A. SHLEIFER (2001): “Coase versus the Coasians,” *Quarterly Journal of Economics*, 116, 853–899.
- GLAESER, E. L., AND A. SHLEIFER (2001): “A Reason for Quantity Regulation,” *American Economic Review (Papers and Proceedings)*, 91, 431–435.
- (2002): “Legal Origins,” *Quarterly Journal of Economics*, 117, 1193–1229.
- GRIFFITH, R., M. O’CONNELL, AND K. SMITH (2019): “Tax Design in the Alcohol Market,” *Journal of Public Economics*, 172, 20–35.

- GRUMMON, A. H., B. B. LOCKWOOD, D. TAUBINSKY, AND H. ALLCOTT (2019): “Designing Better Sugary Drink Taxes,” *Science*, 365, 989–990.
- GUO, Y., AND E. SHMAYA (2025): “Robust Monopoly Regulation,” *American Economic Review*, 115, 599–634.
- JEHIEL, P., AND B. MOLDOVANU (2001): “Efficient Design with Interdependent Valuations,” *Econometrica*, 69, 1237–1259.
- KANG, Z. Y., AND S. VASSERMAN (2025): “Robustness Measures for Welfare Analysis,” *American Economic Review*, 115, 2449–2487.
- KAPLOW, L., AND S. SHAVELL (2002): “On the Superiority of Corrective Taxes to Quantity Regulation,” *American Law and Economics Review*, 4, 1–17.
- KNITTEL, C. R., AND R. SANDLER (2018): “The Welfare Impact of Second-Best Uniform-Pigouvian Taxation: Evidence from Transportation,” *American Economic Journal: Economic Policy*, 10, 211–42.
- LITAN, R. E. (2020): “Want Herd Immunity? Pay People to Take the Vaccine,” *Brookings* (August 18, 2020), available at <https://www.brookings.edu/articles/want-herd-immunity-pay-people-to-take-the-vaccine/>.
- MANKIW, N. G. (2020): “Pay People to Get Vaccinated,” *The New York Times* (September 9, 2020), available at <https://www.nytimes.com/2020/09/09/business/pay-people-vaccine-coronavirus.html>.
- MILGROM, P., AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MISHRA, D., S. PATIL, AND A. PAVAN (2025): “Robust Procurement Design,” *Working paper*.
- MYERSON, R. B. (1979): “Incentive Compatibility and the Bargaining Problem,” *Econometrica*, 47, 61–73.

- (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- NEWELL, R. G., AND W. A. PIZER (2003): “Regulating Stock Externalities Under Uncertainty,” *Journal of Environmental Economics and Management*, 45, 416–432.
- NICHOLS, A. L., AND R. J. ZECKHAUSER (1982): “Targeting Transfers through Restrictions on Recipients,” *American Economic Review*, 72, 372–377.
- PIGOU, A. C. (1920): *The Economics of Welfare*, London, UK: Macmillan.
- PIZER, W. A. (2002): “Combining Price and Quantity Controls to Mitigate Global Climate Change,” *Journal of Public Economics*, 85, 409–434.
- ROBERTS, M. J., AND M. SPENCE (1976): “Effluent Charges and Licenses Under Uncertainty,” *Journal of Public Economics*, 5, 193–208.
- SHAVELL, S. (1984): “A Model of the Optimal Use of Liability and Safety Regulation,” *RAND Journal of Economics*, 15, 271–280.
- STAVINS, R. N. (1996): “Correlated Uncertainty and Policy Instrument Choice,” *Journal of Environmental Economics and Management*, 30, 218–232.
- VISCUSI, W. K., J. E. HARRINGTON, JR., AND J. M. VERNON (2005): *Economics of Regulation and Antitrust*, Cambridge, MA: The MIT Press, 4th edition.
- WALTON, D., AND G. CARROLL (2022): “A General Framework for Robust Contracting Models,” *Econometrica*, 90, 2129–2159.
- WEITZMAN, M. L. (1974): “Prices vs. Quantities,” *Review of Economic Studies*, 41, 477–491.
- WRIGHT, B. D. (1983): “The Economics of Invention Incentives: Patents, Prizes, and Research Contracts,” *American Economic Review*, 73, 691–707.
- YOUNG, L., AND J. E. ANDERSON (1980): “The Optimal Policies for Restricting Trade under Uncertainty,” *Review of Economic Studies*, 47, 927–932.

A Omitted Proofs

A.1 Proof of Lemma 1

For any $(\theta, \xi), (\theta', \xi') \in \Theta \times \mathbb{R}_+$, observe that (IC) requires that

$$\begin{cases} u(q(\theta, \xi), \theta) - t(\theta, \xi) & \geq u(q(\theta', \xi'), \theta) - t(\theta', \xi'), \\ u(q(\theta', \xi'), \theta') - t(\theta', \xi') & \geq u(q(\theta, \xi), \theta') - t(\theta, \xi). \end{cases}$$

Adding both inequalities yields, for any $(\theta, \xi), (\theta', \xi') \in \Theta \times \mathbb{R}_+$,

$$\int_{\theta'}^{\theta} \int_{q(\theta', \xi')}^{q(\theta, \xi)} u_{q\theta}(x, s) \, dx \, ds \geq 0.$$

Because u satisfies the single-crossing property $u_{q\theta} > 0$, this implies that for any $(\theta, \xi), (\theta', \xi') \in \Theta \times \mathbb{R}_+$, $q(\theta, \xi) \geq q(\theta', \xi')$ whenever $\theta > \theta'$. Therefore, for each fixed $\xi \in \mathbb{R}_+$, the function $\theta \mapsto q(\theta, \xi)$ is nondecreasing and hence continuous almost everywhere. In turn, this means that for almost every $\theta \in \Theta$ and $\xi, \xi' \in \mathbb{R}_+$,

$$q(\theta, \xi') \leq \lim_{\varepsilon \downarrow 0} q(\theta + \varepsilon, \xi) = q(\theta, \xi).$$

Swapping ξ and ξ' yields the opposite inequality; hence, for almost every $\theta \in \Theta$ and $\xi, \xi' \in \mathbb{R}_+$,

$$q(\theta, \xi) = q(\theta, \xi').$$

This implies that there exists $\hat{q} \in \mathcal{Q}$ such that $q(\theta, \xi) = \hat{q}(\theta)$ for almost every $(\theta, \xi) \in \Theta \times \mathbb{R}_+$, as claimed.

A.2 Proof of Proposition 1

As noted in Section 4.1, Proposition 1 can be obtained as a corollary of Theorems 1 and 2 by taking the pointwise limit of the optimal allocation functions associated with the strictly concave family of utility functions $u_n(q, \theta) = \theta q - q^2/n$ as $n \rightarrow +\infty$. Below, I instead present a direct proof of Proposition 1 using arguments similar to those in the proofs of Theorems 1 and 2.

(a) Under the unknown-correlation benchmark, the worst-case external benefit can be bounded as in Theorem 1. Given the regulator's choice of allocation function q ,

$$\inf_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)q(\theta)] \geq \mathbf{E}_F[m(\theta)] \cdot \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) = \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta).$$

To show that equality is in fact attained, consider the sequence $\{m_n\}_{n=1}^\infty \subset \mathcal{M}_- \subset \mathcal{M}_0$ defined by

$$m_n(\theta) := \begin{cases} 0, & \text{if } F(\theta) > 1/n, \\ n\mu, & \text{otherwise.} \end{cases}$$

Given the regulator's choice of allocation function q , for any $\varepsilon > 0$, $\Pr_F[q(\theta) < \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) + \varepsilon] > 0$ by the definition of the essential infimum. Consequently, there exists $N_\varepsilon \in \mathbb{N}$ such that

$$\Pr_F \left[q(\theta) < \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) + \varepsilon \right] \geq 1/N_\varepsilon.$$

Since q is nondecreasing by Lemma 1, this choice of N_ε yields

$$\mathbf{E}_F[m_{N_\varepsilon}(\theta)q(\theta)] \leq \mu \left[\operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) + \varepsilon \right].$$

In turn, this implies that

$$\lim_{\varepsilon \rightarrow 0} \mathbf{E}_F[m_{N_\varepsilon}(\theta)q(\theta)] = \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) \implies \inf_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)q(\theta)] = \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta). \quad (6)$$

The worst-case external benefit in equation (6) depends on the regulator's choice of q only through its essential infimum, $\operatorname{ess\,inf}_{\theta \in \Theta} q(\theta)$. Let $\mathcal{Q}_{[0,1]} \subset \mathcal{Q}$ denote the set of implementable allocation

functions with image in $[0, 1]$. Thus, the regulator's problem (P_0) can be rewritten as

$$\begin{aligned} & \sup_{q \in \mathcal{Q}_{[0,1]}} \left[\int_{\Theta} (\theta - c) q(\theta) dF(\theta) + \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) \right] \\ &= \sup_{\underline{q} \in [0,1]} \left[\sup_{q \in \mathcal{Q}_{[0,1]}} \left\{ \int_{\Theta} (\theta - c) q(\theta) dF(\theta) : \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) = \underline{q} \right\} + \mu \underline{q} \right]. \end{aligned} \quad (7)$$

For any choice of $\underline{q} \geq q^{\text{LF}}(\underline{\theta})$, the solution to the inner constrained problem in equation (7) is given by the pointwise maximizer $q^{\text{LF}}(\theta) = \mathbf{1}_{\theta > c}$ unless the constraint is binding, yielding the solution to the inner constrained problem,

$$q^*(\theta; \underline{q}) = \max \left\{ q^{\text{LF}}(\theta), \underline{q} \right\}.$$

This solution form remains the same if $\underline{q} < q^{\text{LF}}(\underline{\theta})$. By substitution, the outer maximization problem in equation (7) can be written as

$$\sup_{\underline{q} \in [0,1]} \left[\mathbf{E}_F[(\theta - c)_+] + (\mu - \mathbf{E}_F[(c - \theta)_+]) \underline{q} \right].$$

Clearly, the maximum is obtained with $\underline{q}^* = \mathbf{1}_{\mu \geq \mathbf{E}_F[(c - \theta)_+]}$, yielding the desired optimal allocation function

$$q^*(\theta) = \begin{cases} 1, & \text{if } \mu \geq \mathbf{E}_F[(c - \theta)_+], \\ q^{\text{LF}}(\theta) = \mathbf{1}_{\theta > c}, & \text{if } \mu < \mathbf{E}_F[(c - \theta)_+]. \end{cases}$$

This solution is unique unless $\mu = \mathbf{E}_F[(c - \theta)_+]$, in which case the regulator is indifferent across a continuum of optimal allocations; in particular, a mandate and no intervention are both optimal.

The above solution also holds under the negative-correlation benchmark since $\{m_n\}_{n=1}^{\infty} \subset \mathcal{M}_-$.

- (b) Given the regulator's choice of allocation function q , the continuous version of Chebyshev's sum inequality again implies that

$$\inf_{m \in \mathcal{M}_+} \mathbf{E}_F[m(\theta)q(\theta)] \geq \mathbf{E}_F[m(\theta)] \cdot \mathbf{E}_F[q(\theta)] = \mu \mathbf{E}_F[q(\theta)].$$

Equality is attained by setting $m^*(\theta) \equiv \mu$, in which case

$$\mathbf{E}_F[m^*(\theta)q(\theta)] = \mu \mathbf{E}_F[q(\theta)] = \inf_{m \in \mathcal{M}_+} \mathbf{E}_F[m(\theta)q(\theta)].$$

Using the same notation as part (a) above, the regulator's problem (\mathbf{P}_+) can be rewritten as

$$\sup_{q \in \mathcal{Q}_{[0,1]}} \left[\int_{\Theta} (\theta - c) q(\theta) dF(\theta) + \mu \mathbf{E}_F[q(\theta)] \right] = \sup_{q \in \mathcal{Q}_{[0,1]}} \int_{\Theta} (\theta - c + \mu) q(\theta) dF(\theta).$$

The unique solution is thus given by the pointwise maximizer $q^*(\theta) = \mathbf{1}_{\theta > c - \mu}$.

A.3 Proof of Theorem 1'

In the case of negative externalities, Nature's problem is to choose a conditional mean function $m \in \mathcal{M}_0$ to maximize external harm, given the regulator's choice of allocation function q :

$$\sup_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)q(\theta)] \leq \mathbf{E}_F[m(\theta)] \cdot \operatorname{ess\,sup}_{\theta \in \Theta} q(\theta) = \mu \|q\|_{\infty}.$$

To show that equality is in fact attained, consider the sequence $\{m_n\}_{n=1}^{\infty} \subset \mathcal{M}_0$ defined by

$$m_n(\theta) := \begin{cases} n\mu, & \text{if } F(\theta) > 1 - 1/n, \\ 0, & \text{otherwise.} \end{cases}$$

Given the regulator's choice of allocation function q , for any $\varepsilon > 0$, $\Pr_F[q(\theta) > \|q\|_{\infty} - \varepsilon] > 0$ by the definition of the essential supremum. Consequently, there exists $N_{\varepsilon} \in \mathbb{N}$ such that

$$\Pr_F[q(\theta) > \|q\|_{\infty} - \varepsilon] \geq 1/N_{\varepsilon}.$$

Since q is nondecreasing by Lemma 1, this choice of N_{ε} yields

$$\mathbf{E}_F[m_{N_{\varepsilon}}(\theta)q(\theta)] \geq \mu (\|q\|_{\infty} - \varepsilon).$$

In turn, this implies that

$$\lim_{\varepsilon \rightarrow 0} \mathbf{E}_F[m_{N_\varepsilon}(\theta)q(\theta)] = \mu \|q\|_\infty \implies \sup_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)q(\theta)] = \mu \|q\|_\infty.$$

This shows that Nature's best response is to concentrate marginal harm on high types. Thus, the regulator's problem can be rewritten as

$$\begin{aligned} & \sup_{q \in \mathcal{Q}} \left[\int_{\Theta} [u(q(\theta), \theta) - cq(\theta)] dF(\theta) - \mu \|q\|_\infty \right] \\ &= \sup_{\bar{q} \in [0, A]} \left[\sup_{q \in \mathcal{Q}} \left\{ \int_{\Theta} [u(q(\theta), \theta) - cq(\theta)] dF(\theta) : \|q\|_\infty = \bar{q} \right\} - \mu \bar{q} \right]. \end{aligned} \quad (8)$$

For any choice of $\bar{q} \leq q^{\text{LF}}(\bar{\theta})$, the solution to the inner constrained problem in equation (8) is given by the pointwise maximizer $q^{\text{LF}}(\theta)$ unless the constraint is binding, yielding the desired solution form

$$q_0^*(\theta) = \min \{q^{\text{LF}}(\theta), \bar{q}\}.$$

This solution form remains the same if $\bar{q} > q^{\text{LF}}(\bar{\theta})$. As in Theorem 1, this solution is unique since u is strictly concave in q .

A.4 Proof of Theorem 2'

As the result for the positive-correlation benchmark follows from the proof of Theorem 1', it suffices to prove the result for the negative-correlation benchmark in the case of negative externalities. Given the regulator's choice of allocation function q ,

$$\sup_{m \in \mathcal{M}_-} \mathbf{E}_F[m(\theta)q(\theta)] \leq \mathbf{E}_F[m(\theta)] \cdot \mathbf{E}_F[q(\theta)] = \mu \mathbf{E}_F[q(\theta)].$$

The inequality above follows by the continuous version of Chebyshev's sum inequality. To show that equality is in fact attained, consider Nature's best response: $m_-^*(\theta) \equiv \mu$. Given this strategy,

$$\mathbf{E}_F[m_-^*(\theta)q(\theta)] = \mu \mathbf{E}_F[q(\theta)] = \sup_{m \in \mathcal{M}_-} \mathbf{E}_F[m(\theta)q(\theta)].$$

Thus, the regulator's problem can be rewritten as

$$\sup_{q \in \mathcal{Q}} \left[\int_{\Theta} [u(q(\theta), \theta) - cq(\theta)] dF(\theta) - \mu \mathbf{E}_F[q(\theta)] \right] = \sup_{q \in \mathcal{Q}} \int_{\Theta} [u(q(\theta), \theta) - (c + \mu)q(\theta)] dF(\theta).$$

The solution to the regulator's problem is therefore given by the pointwise maximizer $q_-^*(\theta) = D(c + \mu, \theta)$.

As in the proof of Theorem 2, uniqueness of the solution follows from the strict concavity of u .

A.5 Proof of Proposition 2

First, consider the regulator's choice of a direct mechanism rather than a payment schedule $T : \mathbb{R}_+ \rightarrow \mathbb{R} \cup \{+\infty\}$. Since ξ does not enter the industry's objective function, an analogous argument to Lemma 1 implies that it suffices to restrict attention to nonincreasing allocation functions $q : \Theta \rightarrow \mathbb{R}_+$. Denote

$$\mathcal{Q}^\downarrow := \{q : \Theta \rightarrow \mathbb{R}_+ \text{ is nonincreasing}\}.$$

Given the regulator's choice of allocation function q , a similar argument as in the proof of Theorem 1' shows that

$$\inf_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)q(\theta)] = \mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta).$$

Thus, the regulator's problem can be rewritten as

$$\begin{aligned} & \sup_{q \in \mathcal{Q}^\downarrow} \left[\mu \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) - \int_{\Theta} C(q(\theta), \theta) dF(\theta) \right] \\ &= \sup_{\underline{q} \in \mathbb{R}_+} \left[\mu \underline{q} - \inf_{q \in \mathcal{Q}^\downarrow} \left\{ \int_{\Theta} C(q(\theta), \theta) dF(\theta) : \operatorname{ess\,inf}_{\theta \in \Theta} q(\theta) = \underline{q} \right\} \right]. \end{aligned} \quad (9)$$

For any choice of \underline{q} , because costs are increasing with externality reduction ($C_q > 0$), the solution to the inner constrained problem in equation (9) is simply given by $q(\theta) \equiv \underline{q}$.

Next, to implement such an allocation function, consider the payment schedule defined in the statement of Proposition 2:

$$T^*(q) = \begin{cases} 0, & \text{if } q \geq \underline{q}, \\ +\infty, & \text{if } q < \underline{q}. \end{cases}$$

Under T^* , the industry chooses $q \geq \underline{q}$; moreover, since $C_q > 0$, it chooses $q(\theta) \equiv \underline{q}$ independently of the realization of θ . Thus, T^* implements the regulator's optimal allocation function, as claimed.

A.6 Proof of Proposition 3

Let $m(\theta) := \mathbf{E}_H[w \mid \theta]$. For each $* \in \{0, +, -\}$, observe that the regulator's problem is now

$$\sup_{(q,t) \text{ satisfies (IC)}} \left\{ \inf_{m \in \mathcal{M}_*} \int_{\Theta} m(\theta) [\theta q(\theta) - t(\theta)] dF(\theta) : \int_{\Theta} q(\theta) dF(\theta) \leq Q \right\}. \quad (\mathbf{P}'_*)$$

Denoting $U(\theta) := \theta q(\theta) - t(\theta)$, notice that by the envelope theorem (Milgrom and Segal, 2002),

$$U(\theta) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} q(s) ds.$$

Because q is nonnegative and nondecreasing, U must be nondecreasing and convex. Consequently, given the regulator's choice of allocation function q ,

$$\inf_{m \in \mathcal{M}_0} \mathbf{E}_F[m(\theta)U(\theta)] \geq \mathbf{E}_F[m(\theta)] \cdot \operatorname{ess\,inf}_{\theta \in \Theta} U(\theta) = \mu \operatorname{ess\,inf}_{\theta \in \Theta} U(\theta).$$

The same argument as in the proof of part (a) of Proposition 1 shows that equality is in fact attained in the limit by the sequence $\{m_n\}_{n=1}^{\infty} \subset \mathcal{M}_0$ defined by

$$m_n(\theta) := \begin{cases} 0, & \text{if } F(\theta) > 1/n, \\ n\mu, & \text{otherwise.} \end{cases}$$

Note that $\text{ess inf}_{\theta \in \Theta} U(\theta)$ can be expressed in terms of q : $\text{ess inf}_{\theta \in \Theta} U(\theta) = U(\underline{\theta}) = \underline{\theta}q(\underline{\theta}) - t(\underline{\theta})$.

Thus, the regulator's problem (P'_0) can be rewritten as

$$\sup_{q \in \mathcal{Q}_{[0,1]}} \sup_{t(\underline{\theta}) \in \mathbb{R}_+} \left\{ \mu [\underline{\theta}q(\underline{\theta}) - t(\underline{\theta})] : \int_{\Theta} q(\theta) dF(\theta) \leq Q \right\}.$$

Because wait times must be nonnegative, $t(\underline{\theta}) \geq 0$. Moreover, since q is nondecreasing,

$$q(\underline{\theta}) \leq \int_{\Theta} q(\theta) dF(\theta) \leq Q.$$

Therefore, the value of the regulator's problem (P'_0) is bounded from above by

$$\sup_{q \in \mathcal{Q}_{[0,1]}} \sup_{t(\underline{\theta}) \in \mathbb{R}_+} \left\{ \mu [\underline{\theta}q(\underline{\theta}) - t(\underline{\theta})] : \int_{\Theta} q(\theta) dF(\theta) \leq Q \right\} \leq \mu \underline{\theta} Q.$$

It is easy to see that the mechanism $(q^*, t^*) = (Q, 0)$ is incentive-compatible and attains this value.

Hence, (q^*, t^*) is optimal, as claimed. This solution is unique whenever $\underline{\theta} > 0$.