

Robust Monopoly Regulation[†]

By YINGNI GUO AND ERAN SHMAYA*

We study how to regulate a monopolistic firm using a robust-design, non-Bayesian approach. We derive a policy that minimizes the regulator's worst-case regret, where regret is the difference between the regulator's complete-information payoff and his realized payoff. When the regulator's payoff is consumers' surplus, he caps the firm's average revenue. When his payoff is the total surplus of both consumers and the firm, he offers a piece rate subsidy to the firm while capping the total subsidy. For intermediate cases, the regulator combines these three policy instruments to balance three goals: protecting consumers' surplus, mitigating underproduction, and limiting potential overproduction. (JEL D21, D42, D83, H25, L51)

Regulating monopolies is challenging. A monopolistic firm has the market power to set its price above the price in an oligopolistic or competitive market. For instance, Cooper et al. (2018) show that prices at monopoly hospitals are 12 percent higher than those in markets with 4 or 5 competitors. In order to protect consumers' surplus, a regulator may want to constrain the firm's price. However, a price-constrained firm may fail to obtain enough revenue to cover its fixed cost, so it may end up not producing. The regulator must balance the need to protect consumers' surplus and the need to not distort production.

This challenge could be solved easily if the regulator had complete information about the industry. The regulator could ask the firm to produce the efficient quantity and to set its price equal to the marginal cost. He could then subsidize the firm for all of its other costs. However, the regulator typically has limited information about consumer demand or the production costs of the firm. How should the regulatory policy be designed when the regulator knows much less about the industry than the

*Guo: Duke University and Northwestern University (email: yingni.guo@northwestern.edu). Shmaya: Stony Brook University (email: eran.shmaya@stonybrook.edu). Sylvain Chassang was the coeditor for this article. We are grateful to Debraj Ray, the coeditor who reviewed the first round before leaving the AER and Mark Armstrong for a very helpful discussion at CEPR Virtual IO Seminars. For their valuable feedback, we thank seminar audiences at Northwestern IO seminar, third Columbia Conference on Economic Theory, Notre Dame Mini-Conference on Information Economics and Theory, Theory at Penn State—Information in Markets, University of Bonn, CMU/Pitt Theory Seminar, Boston University, New York University, Yale University, the University of Chicago, University of Michigan, University College London, Seminars in Economic Theory, Queen Mary University of London, London School of Economics, University of Oxford, CUHK-HKU-HKUST, Duke University, CEPR Virtual IO Seminars, European Research Council First Workshop on Verifiability, PUC-Chile, Stanford University, Peking University, Aalto University, University of Cambridge, MIT IO Seminar Series, Virtual Seminar at AMETS, Georgetown University, Cornell SBE Seminar, Princeton Political Economy workshop, University of Washington, UCLA, USC, and the University of Tokyo.

[†]Go to <https://doi.org/10.1257/aer.20191950> to visit the article page for additional materials and author disclosure statement(s).

firm does? If the regulator wants a policy that works “fairly well” in all circumstances, what should this policy look like?

We study this classic problem of monopoly regulation (e.g., Baron and Myerson 1982) using a robust-design, non-Bayesian approach. The regulator’s payoff is a weighted sum of consumers’ surplus and the firm’s profit, with weakly more weight on consumers’ surplus than on the firm’s profit. He can regulate the firm’s quantity and price. He can both subsidize the firm and tax it. Given a policy, the firm chooses a quantity-price pair under its demand curve to maximize its profit. The regulator’s *regret* is, by definition, the difference between what he could have gotten if he had complete information about the industry and what he actually gets. Thus, regret can be interpreted as “money left on the table” due to the regulator’s lack of information. The regulator evaluates a policy by the worst-case regret under this policy, that is, the maximal regret he can incur across all possible demand and cost scenarios. An optimal policy minimizes the worst-case regret.

Our use of the minimax-regret approach to uncertainty is the main way that our approach differs from that taken in the existing literature on monopoly regulation, where the Bayesian approach is the norm. The Bayesian approach would assign a prior to the regulator over all possible demand and cost scenarios and characterize a policy that minimizes the regulator’s expected regret.¹ We instead focus on industries in which information asymmetry is so pronounced that there is no obvious way to formulate a prior, or industries where new sources of uncertainty arise all the time. In response, the regulator looks for a policy that leaves as little “money” on the table as possible in all circumstances. We will show that the minimax-regret approach (i) provides useful insights on the regulation problem and (ii) remains tractable for broad kinds of multidimensional information asymmetry.

To illustrate our optimal policy, we begin with two polar cases of the regulator’s payoff. At one end, he puts no weight on the firm’s profit, so his payoff is consumers’ surplus. At the other end, he puts the same weight on the firm’s profit as on consumers’ surplus, so his payoff is the total surplus of consumers and the firm.

If the regulator’s payoff is consumers’ surplus, we show that it is optimal to just cap the firm’s average revenue. A cap on average revenue bounds how much consumers’ surplus the firm can extract. In this case, consumers benefit from a lower cap. However, the cap may discourage a firm that should have produced from producing. In this case, consumers lose due to the firm’s underproduction. The optimal level of the cap is chosen to balance consumers’ gain from their protected surplus and their loss from the firm’s underproduction. We show that the regulator can also implement this cap on average revenue by imposing a price cap.

If the regulator’s payoff is the total surplus of consumers and the firm, he has no redistribution incentive but simply wants the firm to produce as efficiently as possible. An unregulated, monopolistic firm tends to produce less than the efficient quantity in order to charge a high price. To mitigate such underproduction, the regulator offers a piece rate subsidy: for each unit that the firm sells, he subsidizes the firm for the difference between its price and a target level. This subsidy lifts the firm’s average revenue to the target level, so it encourages the firm to serve more consumers

¹ Minimizing the regulator’s expected regret is the same as maximizing his expected payoff because the regulator’s expected complete-information payoff is constant over all policies.

than just those with high values. On the other hand, the subsidy may induce the firm to produce more than the efficient quantity. To limit the degree of potential overproduction, the regulator imposes a cap on the total subsidy to the firm.

For intermediate cases, the regulator puts some weight on the firm's profit, but less than the weight he puts on consumers' surplus. In these cases, the optimal policy combines the three policy instruments that are used in the two polar cases. First, the firm's average revenue is capped at some target level. Second, if the firm prices below this target level, the regulator offers a piece rate subsidy that lifts the firm's average revenue to this target level. Third, the regulator sets a cap on the total subsidy to the firm. As the weight on the firm's profit increases, the regulator's redistribution incentive weakens, so the cap on the firm's average revenue increases as well.

The minimax-regret approach advances our knowledge of monopoly regulation in two significant ways. First, it delineates the trade-off that the regulator must make among his three goals: protecting consumers' surplus, mitigating underproduction, and limiting potential overproduction. Second, it demonstrates how the regulator can use three different policy instruments to achieve these three goals. Specifically, the cap on average revenue protects consumers' surplus, the piece rate subsidy mitigates underproduction, and the cap on the total subsidy prevents severe overproduction. We show that, among all possible policy instruments that the regulator can choose, these three instruments are sufficient. In Section III, we discuss in detail the policy implications of these results.

Our baseline model analyzes a setting in which the regulator has no knowledge about the firm's demand or cost scenarios except for an upper bound on consumers' values. In some applications, the regulator may know more than this. In Section IV we address how to incorporate the regulator's additional knowledge. For instance, the regulator may know that the firm's demand curve is bounded from above by one function and from below by another. We characterize an optimal policy for any given bounding functions. One special case is that the regulator knows the firm's demand curve as the firm does. In this case, the regulator knows exactly how much total consumer value the firm has created. We show that the optimal policy no longer induces overproduction. Another special case is the one in which the regulator knows both an upper bound and a lower bound on consumers' values. We show that imposing a price cap is optimal when consumers are sufficiently homogeneous or when the weight on the firm's profit is small enough. Our analysis therefore shows that the minimax-regret approach remains tractable for incorporating specific knowledge and continues to provide insights on the problem of monopoly regulation.

Related Literature.—This paper contributes to the literature on monopoly regulation. For an overview of earlier contributions in this field, see Caillaud et al. (1988); Braeutigam (1989); and Laffont and Tirole (1993). Armstrong and Sappington (2007) discuss recent developments. Our paper is closely related to Baron and Myerson (1982); Lewis and Sappington (1988a, b); and Armstrong (1999), all of whom study how to regulate a privately informed firm using the Bayesian approach. Baron and Myerson (1982) assume one-dimensional uncertainty about the cost function. Lewis and Sappington (1988a) assume one-dimensional uncertainty about the demand function. Lewis and Sappington (1988b) and Armstrong (1999) assume two-dimensional uncertainty about both cost and demand scenarios. In our model,

the firm also has private information about both cost and demand scenarios, and this private information is infinite-dimensional. We introduce the minimax-regret approach to the regulation problem. We show that this approach provides new insights on the regulation problem and also broadens the range of multidimensional screening settings that can be studied.

The minimax-regret approach to uncertainty dates back to Wald (1950) and Savage (1954). It has since been applied broadly in many areas, including game theory (e.g., Hannan 1957; Hart and Mas-Colell 2000), machine learning (e.g., Bubeck and Cesa-Bianchi 2012), designing treatment rules (e.g., Manski 2004; Stoye 2009), forecast aggregation (e.g., Arieli, Babichenko, and Smorodinsky 2018; Babichenko and Garber 2021), and algorithm design (e.g., Vera and Banerjee 2021; Feng et al. 2024). Our paper contributes especially to the literature on mechanism design with the minimax-regret approach. To start with, Hurwicz and Shapiro (1978) examine a moral hazard problem. Bergemann and Schlag (2008, 2011) examine monopoly pricing. Renou and Schlag (2011) apply the solution concept of ε -minimax-regret to the problem of implementing social-choice correspondences. Caldentey, Liu, and Lobel (2017) characterize the dynamic-pricing rule that minimizes the seller's worst-case regret. Beviá and Corchón (2022) examine the contest that minimizes the designer's worst-case regret. Malladi (2022) studies the optimal approval rules for innovation; Bergemann, Gan, and Li (2023) study robust persuasion when the receiver is uncertain about the sender's preferences and the set of feasible signal structures; and Guo and Shmaya (2023) study the optimal mechanism for project choice within organizations.

More broadly, we contribute to the growing literature on mechanism design with worst-case objectives. See, for instance, Chassang (2013); Carroll (2015); and Carroll (2019) for a survey on robustness in mechanism design. In regard to the monopoly regulation environment, one related paper in this literature is Garrett (2014), which considers the cost-based procurement problem (Laffont and Tirole 1986). More recently, Mishra and Patil (2024) apply the prior-free criterion of undomination to Baron and Myerson (1982) and show that undominated mechanisms feature a quantity floor.

Our work also contributes to the delegation literature (e.g., Holmström 1984). Alonso and Matouschek (2008); Kolotilin and Zapechelnyuk (forthcoming); and Amador and Bagwell (2022) characterize conditions under which price cap regulation is optimal under the restriction that transfers are infeasible. In our environment, the regulator and the firm can make transfers to each other. We characterize conditions under which price cap regulation is optimal. To our knowledge, we are the first to show that a contract that does not use transfers is optimal even in a contracting environment in which both parties can make transfers to each other.² From this aspect, our paper also relates to Drexler and Kleiner (2018), which shows that a transfer-free, qualified majority voting rule maximizes expected utilitarian welfare among strategyproof and anonymous rules.

² Armstrong and Vickers (2010) show that it can be optimal not to use transfers in an environment in which the agent must receive nonnegative transfers from the principal.

I. Environment

There is a monopolistic firm and a mass one of consumers. Let $P: [0, 1] \rightarrow [0, \bar{v}]$ be a decreasing upper-semicontinuous *inverse-demand function*. Let $C: [0, 1] \rightarrow \mathbb{R}_+$ with $C(0) = 0$ be an increasing lower-semicontinuous *cost function*. The total consumer value of quantity q is the area under the inverse-demand function, given by $\int_0^q P(z)dz$. The total surplus of quantity q is thus the total consumer value minus the cost, given by $\int_0^q P(z)dz - C(q)$. The *maximal total surplus* is given by

$$(1) \quad OPT(P, C) = \max_{q \in [0, 1]} \left\{ \int_0^q P(z)dz - C(q) \right\}.$$

If the firm produces and sells q units, then *distortion* is the maximal total surplus minus the actual total surplus, given by

$$DSTR(P, C, q) = OPT(P, C) - \left[\int_0^q P(z)dz - C(q) \right].$$

To simplify notation, we sometimes omit the dependence of OPT on (P, C) and the dependence of $DSTR$ on (P, C, q) when no confusion arises. We will do the same for other terms.

EXAMPLE 1: Suppose that $P(q) = 1 - q$ and $C(q) = q/2$. It is efficient to produce and sell $q^* = 1/2$ units. The maximal total surplus is $\int_0^{q^*} (1 - z)dz - q^*/2 = 1/8$. If the firm chooses $q < q^*$, we say that it underproduces. If the firm chooses $q > q^*$, we say that it overproduces. In both cases, distortion is strictly positive.

Regulatory Policies, Timing of the Game, and Payoffs.—The firm knows (P, C) , but the regulator does not know (P, C) . The firm chooses both the quantity q to sell and the price p paid by consumers. A quantity-price pair (q, p) is *feasible* if and only if it is under the inverse-demand function, that is, if and only if $p \leq P(q)$. The firm can choose any feasible quantity-price pair. A consumer chooses to buy if the price p is below their value. If the firm’s choice (q, p) cannot clear the market, then the product is rationed.³

The regulator observes the firm’s choice (q, p) . A regulatory *policy* is given by an upper-semicontinuous function $\rho: [0, 1] \times [0, \bar{v}] \rightarrow \mathbb{R} \cup \{-\infty\}$. If the firm chooses (q, p) , it receives the revenue $\rho(q, p)$. This revenue is the sum of the market revenue qp and any tax or subsidy, $\rho(q, p) - qp$, imposed by the regulator. We assume that $\rho(0, 0) \geq 0$, so the firm can stay out of business without suffering a negative profit. This is the firm’s participation constraint.

³The firm’s choice (q, p) cannot clear the market if $\sup\{q' \in [0, 1] : P(q') > p\} > q$.

The timing of the game is as follows: (i) the regulator publicly chooses and commits to a policy ρ ; (ii) the firm privately observes (P, C) , publicly chooses (q, p) , and obtains the market revenue qp ; (iii) the regulator transfers $\rho(q, p) - qp$ to the firm. (If $\rho(q, p) - qp < 0$, it represents a tax on the firm.) The firm's choice (q, p) provides evidence to the regulator that $P(q) \geq p$. This evidence aspect of the firm's choice (q, p) was not present in Baron and Myerson (1982) because they assume that the regulator knows P , or in Lewis and Sappington (1988a, b) and Armstrong (1999) because they assume that the regulator cannot observe the firm's sales q .

There are many policy instruments that the regulator can use. To illustrate, we give four examples of policies:

- The regulator can give the firm a lump-sum subsidy $s > 0$ if it sells more than a certain quantity \tilde{q} . The policy is $\rho(q, p) = qp$ if $q < \tilde{q}$ and $\rho(q, p) = qp + s$ if $q \geq \tilde{q}$.
- The regulator can charge a proportional tax by setting $\rho(q, p) = (1 - \tau)qp$ for some $\tau \in (0, 1)$.
- The regulator can impose a price cap at k , such that the firm cannot price above k . It gets the market revenue qp if it prices below k . The policy is $\rho(q, p) = qp$ if $p \leq k$ and $\rho(q, p) = -\infty$ if $p > k$.
- The regulator can require that the firm get no more than k per unit by setting $\rho(q, p) = \min\{qk, qp\}$. If the firm prices above k , it pays a tax of $q(p - k)$ to the regulator.

The regulator could ask the firm to report its inverse-demand and cost functions, and then determine the firm's quantity, price, and revenue as a function of its report. For any such deterministic direct-revelation mechanism, there exists a revenue function $\rho(q, p)$ that induces the same outcome. This is referred to as the Taxation Principle (Rochet 1986; Guesnerie 1998). Hence, it is without loss of generality to work directly with the revenue function $\rho(q, p)$.

Fix a policy ρ and a pair of (P, C) functions. If the firm sells q units at price p , then *consumers' surplus* and *the firm's profit* are given by

$$(2) \quad CS(\rho, P, q, p) = \int_0^q P(z) dz - \rho(q, p), \quad \text{and}$$

$$FP(\rho, P, C, q, p) = \rho(q, p) - C(q).$$

In the definition of consumers' surplus, we make two assumptions. First, we assume that any subsidy to the firm is paid by consumers through their taxes and that any tax paid by the firm is passed on to consumers; moreover, there is no social cost of public funds (e.g., Baron and Myerson 1982). Second, we assume efficient rationing, so when the firm's choice (q, p) cannot clear the market, consumers with the highest values are served.⁴ In Section IIC, we discuss whether and when the firm wants to clear the market so rationing is not needed.

⁴Suppose that the firm's choice (q, p) cannot clear the market. Let $q'' = \sup\{q' \in [0, 1] : P(q') > p\}$. Without efficient rationing, consumers' surplus is at least $\int_{q''-q}^{q''} P(z) dz - \rho(q, p)$ and at most $\int_0^q P(z) dz - \rho(q, p)$.

The firm's profit does not depend directly on the inverse-demand function P . However, since P determines the set of feasible quantity-price pairs for the firm, we include P as an argument in the firm's profit. We say that (q, p) is a firm's best response to (P, C) under policy ρ if it maximizes the firm's profit over all feasible quantity-price pairs. The firm may have multiple best responses. Its participation constraint implies that $FP(\rho, P, C, q, p) \geq 0$ for every best response (q, p) .

The regulator's payoff is a weighted sum, $CS + \alpha FP$, of consumers' surplus and the firm's profit. The parameter $\alpha \in [0, 1]$ is the welfare weight the regulator puts on the firm's profit.

The Regulator's Complete-Information Payoff.—Fix a pair of (P, C) functions. We let $CIP(P, C)$ denote the regulator's complete-information payoff. This is what he could achieve if he knew (P, C) and thus chose a policy ρ according to (P, C) . Formally,

$$CIP(P, C) = \max_{\rho, q, p} \{CS(\rho, P, q, p) + \alpha FP(\rho, P, C, q, p)\},$$

where the maximum is over all ρ and all of the firm's best responses (q, p) to (P, C) under ρ .

Claim 1 states that the regulator's complete-information payoff equals the maximal total surplus. The regulator would ask the firm to sell the efficient quantity and to price at the marginal consumer's value at this quantity. He would then make sure that the firm's revenue equals its cost. This policy ensures that the maximal total surplus is generated and that all of this surplus goes to consumers. Hence, although the regulator's payoff is a function of the weight α on the firm's profit, his complete-information payoff does not depend on α .

CLAIM 1: For any pair of (P, C) functions, $CIP(P, C) = OPT(P, C)$.

PROOF:

First, the regulator's complete-information payoff is at most $OPT(P, C)$. Indeed,

$$\begin{aligned} CS(\rho, P, q, p) + \alpha FP(\rho, P, C, q, p) &\leq CS(\rho, P, q, p) + FP(\rho, P, C, q, p) \\ &\leq OPT(P, C), \end{aligned}$$

for every policy ρ and every best response (q, p) of the firm to (P, C) under ρ . Here, the first inequality follows from $\alpha \leq 1$ and the fact that $FP(\rho, P, C, q, p) \geq 0$, and the second inequality follows from the definitions of OPT , CS , and FP in (1) and (2).

Second, let q^* denote an efficient quantity that achieves the maximal total surplus. The regulator can achieve $OPT(P, C)$ by setting $\rho(q, p) = C(q^*)$ if $(q, p) = (q^*, P(q^*))$, and $\rho(q, p) = 0$ otherwise. Choosing $(q, p) = (q^*, P(q^*))$ is a

firm's best response to (P, C) under ρ . Since $CS(\rho, P, q, p) = OPT(P, C)$ and $FP(\rho, P, C, q, p) = 0$, it follows that the regulator's payoff is $OPT(P, C)$. ■

Regret.—When the regulator does not know (P, C) , no policy can guarantee that the regulator gets his complete-information payoff. Given a policy ρ , a pair of (P, C) functions, and a firm's best response (q, p) to (P, C) under ρ , the regulator's *regret* is the difference between his complete-information payoff and his actual payoff:

$$RGRT(\rho, P, C, q, p) = CIP(P, C) - [CS(\rho, P, q, p) + \alpha FP(\rho, P, C, q, p)].$$

Our next result shows that regret is a weighted sum of distortion and the firm's profit.

CLAIM 2: *Given a policy ρ , a pair of (P, C) functions, and a firm's best response (q, p) to (P, C) under ρ , we have*

$$RGRT(\rho, P, C, q, p) = DSTR(P, C, q) + (1 - \alpha)FP(\rho, P, C, q, p).$$

PROOF:

Suppressing the dependence on (ρ, P, C, q, p) , we have

$$\begin{aligned} RGRT &= CIP - (CS + \alpha FP) \\ &= OPT - (CS + \alpha FP) \\ &= OPT - (CS + FP) + (1 - \alpha)FP \\ &= DSTR + (1 - \alpha)FP. \end{aligned}$$

Here, the first equality is the definition of regret, the second equality follows from Claim 1 that $CIP = OPT$, and the last equality follows from the definition of distortion. ■

Regret has a natural interpretation in our setting. $DSTR$ represents the loss in the regulator's efficiency objective since he wishes the firm to produce as efficiently as possible. $(1 - \alpha)FP$ represents the loss in his redistribution objective since the regulator wants to give more surplus to consumers rather than to the firm. The less weight α the regulator puts on the firm's profit, the more he cares about redistribution, and the higher his regret is.

The Regulator's Problem.—The regulator chooses a policy that minimizes his worst-case regret. Thus, *the regulator's problem* is

$$\underset{\rho}{\text{minimize}} \quad \max_{P, C, q, p} \quad RGRT(\rho, P, C, q, p),$$

where the minimization is over all ρ and the maximum is over all (P, C) and all of the firm's best responses (q, p) to (P, C) under ρ .

Formulating the regulator’s problem as a minimax-regret problem is our main departure from the prior literature on monopoly regulation (e.g., Armstrong and Sappington 2007). If we assigned a prior to the regulator over the possible inverse-demand and cost functions, then choosing a policy that maximizes the regulator’s expected payoff would be the same as choosing one that minimizes his expected regret.⁵ Instead, we consider environments in which information asymmetry is so pronounced that there is no obvious way to formulate a prior. The regulator looks for a policy that works fairly well in all circumstances.

Remark 1: In the definition of the regulator’s complete-information payoff, we assume that the firm breaks ties in favor of the regulator, whereas in the definition of the regulator’s problem, we assume that the firm breaks ties against the regulator. These assumptions are for convenience only and do not affect the value of the regulator’s complete-information payoff in Claim 1 or the solution to the regulator’s problem in Theorems 1 to 3.⁶

II. Main Result

In this section we first establish a lower bound on the regulator’s worst-case regret under any policy. We do this in Section IIA by showing that there is a nontrivial trade-off among (i) protecting consumers’ surplus, (ii) mitigating underproduction, and (iii) limiting potential overproduction. We then show in Section IIB that the worst-case regret under our policy is at most this lower bound, so it is optimal. In Section IIC, we discuss when the firm wants to clear the market.

A. Lower Bound on Worst-Case Regret

We begin with an example in Figure 1 to illustrate the trade-off between (i) protecting consumers’ surplus and (ii) mitigating underproduction. Suppose that the regulator constrains how much consumers’ surplus the firm can extract by imposing a price cap at k . This price cap has opposing implications for the two market scenarios in Figure 1.

In the left panel, every consumer has the highest value \bar{v} , and the cost is zero. The firm will price at k and serve all consumers. There is no distortion since all consumers are served, as they should be. The firm’s profit is k , so regret is $(1 - \alpha)k$. The lower the price cap k , the lower this regret from the firm’s profit. In the right panel, every consumer still has the highest value \bar{v} , but now the firm has a fixed cost of k . It is a firm’s best response not to produce. The firm’s profit is zero, but distortion is $(\bar{v} - k)$, which is the maximal total surplus that could have been generated. Thus,

⁵This is because the expected complete-information payoff of the regulator is the same across all policies.

⁶If in the definition of $CIP(P, C)$, we assumed that the firm breaks ties against the regulator, we would define

$$CIP(P, C) = \sup_{\rho} \min_{q,p} \{CS(\rho, P, q, p) + \alpha FP(\rho, P, C, q, p)\},$$

where the minimum is over all of the firm’s best responses (q, p) to (P, C) under ρ . Then the supremum may not be achieved, but the value of $CIP(P, C)$ would be the same. Similarly, if we assumed that the firm breaks ties in favor of the regulator in the regulator’s problem, then the “worst-case” pair (P, C) may not exist, but the solution to the regulator’s problem would remain the same.

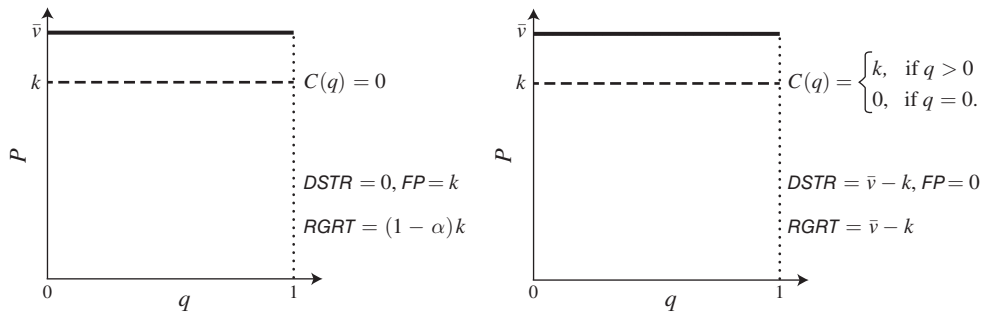


FIGURE 1. PROTECTING CONSUMERS' SURPLUS VERSUS MITIGATING UNDERPRODUCTION

regret is $(\bar{v} - k)$. The lower the price cap k , the higher this regret from distortion. Combining these two market scenarios, we conclude that the worst-case regret given the price cap k is at least $\max\{(1 - \alpha)k, \bar{v} - k\}$.

This example assumes a price cap policy, yet the trade-off in the example applies to any policy. If a policy offers very high revenue to the firm, it may give the firm too much profit, causing the regulator high regret due to the firm's profit. However, if the policy constrains the firm's revenue too much, it may worsen the problem of underproduction, causing high regret from distortion. We now provide a lower bound on the worst-case regret based on this trade-off.

CLAIM 3: Let $f_\alpha = 1/(2 - \alpha)$ be such that $(1 - \alpha)f_\alpha = 1 - f_\alpha$. Then the worst-case regret under any policy is at least $(1 - f_\alpha)\bar{v} = (1 - \alpha)\bar{v}/(2 - \alpha)$.

PROOF:

Fix a policy ρ . Let $k = \max_{(q,p) \in [0,1] \times [0,\bar{v}]} \rho(q,p)$ be the highest revenue under ρ . If the firm's (P, C) is given by the left panel of Figure 1, it chooses (q, p) that maximizes its revenue $\rho(q, p)$. Its profit is k , so regret is at least $(1 - \alpha)k$. If the firm's (P, C) is given by the right panel of Figure 1, it is a firm's best response not to produce, so regret is $(\bar{v} - k)$. Hence, the worst-case regret under ρ is at least $\max\{(1 - \alpha)k, \bar{v} - k\}$, which is at least $(1 - f_\alpha)\bar{v}$. ■

Building on Claim 3, our next theorem incorporates overproduction and the trade-off between mitigating underproduction and limiting potential overproduction. An unregulated, monopolistic firm tends to underproduce, especially when consumers are sufficiently heterogeneous in their values. In such cases, the firm tends to serve only consumers with high values in order to charge a high price. If the regulator does not encourage the firm to produce more and thereby also serve consumers with moderate values, he suffers regret from underproduction. This will be reflected in the first term of (3) in Theorem 1 and in Cases 1 and 2 of the proof. However, if the regulator subsidizes the firm for producing more, his subsidy may cause overproduction because the firm may choose an inefficiently large quantity due to the subsidy. In such cases, the regulator suffers regret from overproduction. This will be reflected by the second term in (3) and Case 3 in the Proof of Theorem 1.

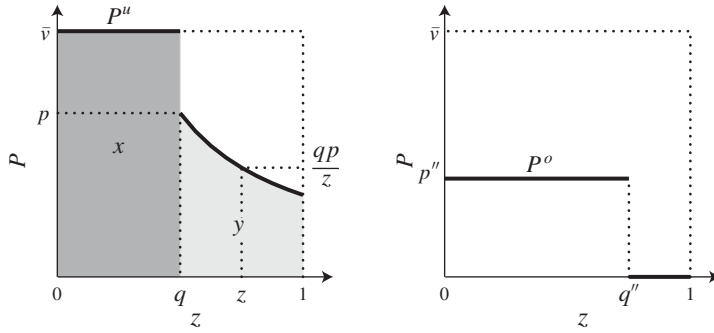


FIGURE 2. INVERSE-DEMAND FUNCTIONS P^u AND P^o IN THE PROOF OF THEOREM 1

Theorem 1 refines the lower bound on the worst-case regret in Claim 3. We later show in Section IIB that the lower bound in this theorem is tight.

THEOREM 1 (Lower Bound on Worst-Case Regret): *Let*

$$(3) \quad r_\alpha = \max_{(q,p) \in [0,1] \times [0, f_\alpha \bar{v}]} \min \{ q(1 - f_\alpha) \bar{v} - qp \log q, q(f_\alpha \bar{v} - p) \}.$$

The worst-case regret under any policy is at least r_α .

PROOF:

Fix a policy ρ . Fix a pair $(q, p) \in [0, 1] \times [0, f_\alpha \bar{v}]$. Let P^u denote the inverse-demand function in the left panel of Figure 2: $P^u(z) = \bar{v}$ if $z \leq q$, and $P^u(z) = qp/z$ if $z > q$.⁷ Let $x = \max_{q' \leq q} \rho(q', p')$ be the highest revenue among all (q', p') in the dark-gray area and $y = \max_{q' \geq q, q'p' \leq qp} \rho(q', p')$ be the highest revenue among all (q', p') in the light-gray area.

We discuss three cases, which vary depending on the values of x and y .

Case 1: $\max\{x, y\} \leq qf_\alpha \bar{v}$. If the firm has a fixed cost of $qf_\alpha \bar{v}$ and its inverse-demand function is P^u , then it is a firm’s best response not to produce. Regret equals distortion, which equals the maximal total surplus:

$$RGRT = DSTR = q\bar{v} + \int_q^1 \frac{qp}{z} dz - qf_\alpha \bar{v} = q(1 - f_\alpha) \bar{v} - qp \log q.$$

Case 2: $\max\{x, y\} > qf_\alpha \bar{v}$ and $x > y$. If the firm’s inverse-demand function is P^u and its cost is zero, then the firm produces at most q and its profit is $x \geq qf_\alpha \bar{v}$. Regret is

$$\begin{aligned} RGRT &\geq (1 - \alpha)qf_\alpha \bar{v} + DSTR \\ &\geq (1 - \alpha)qf_\alpha \bar{v} + \int_q^1 \frac{qp}{z} dz = q(1 - f_\alpha) \bar{v} - qp \log q. \end{aligned}$$

The last step follows from the definition of f_α such that $(1 - \alpha)f_\alpha = 1 - f_\alpha$.

⁷The inverse-demand function P^u is characterized by being the highest inverse-demand function where the market revenue does not exceed qp if the firm chooses any quantity $z > q$.

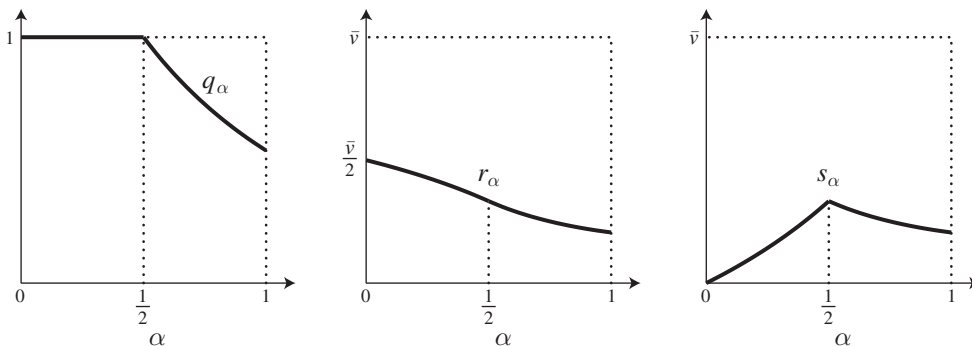


FIGURE 3. VALUES OF q_α , r_α , AND s_α

Case 3: $\max\{x, y\} > qf_\alpha \bar{v}$ and $x \leq y$. Let (q'', p'') be the maximizer in the definition of $y = \max_{q' \geq q, q' p' \leq qp} \rho(q', p')$. If the firm’s cost function is $C(q) = y$ for $q > 0$ and its inverse-demand function is P^o as given by the right panel of Figure 2 (i.e., $P^o(z) = p''$ if $z \leq q''$, and $P^o(z) = 0$ if $z > q''$), then it is a firm’s best response to choose (q'', p'') . Since the fixed cost y exceeds the total consumer value $q'' p''$ that the firm has created, the firm overproduces. Regret equals distortion from overproduction:

$$RGRT = DSTR = y - q'' p'' \geq q f_\alpha \bar{v} - qp = q(f_\alpha \bar{v} - p).$$

For any $(q, p) \in [0, 1] \times [0, f_\alpha \bar{v}]$, one of these cases occurs. It follows that for any such (q, p) , the worst-case regret is at least $\min\{q(1 - f_\alpha) \bar{v} - qp \log q, q(f_\alpha \bar{v} - p)\}$. Therefore, the worst-case regret is at least r_α . ■

Let q_α achieve the maximum in the definition of r_α in (3). The explicit values of q_α and r_α are given by

$$q_\alpha = \begin{cases} 1, & \text{if } \alpha \leq \frac{1}{2}, \\ e^{1 - \frac{\alpha + \sqrt{\alpha(\alpha+4)}}{2}}, & \text{if } \alpha > \frac{1}{2}; \end{cases}$$

$$r_\alpha = \begin{cases} \bar{v} \left(\frac{1 - \alpha}{2 - \alpha} \right), & \text{if } \alpha \leq \frac{1}{2}, \\ \frac{[2 + \alpha - \sqrt{\alpha(\alpha+4)}] \bar{v} e^{1 - \frac{\alpha + \sqrt{\alpha(\alpha+4)}}{2}}}{2(2 - \alpha)}, & \text{if } \alpha > \frac{1}{2}. \end{cases}$$

We depict the values of q_α and r_α in the left and middle panels of Figure 3.⁸

⁸For detailed calculations, see the Supplemental Appendix.

When the weight α on the firm's profit is small enough (i.e., $\alpha \leq 1/2$), q_α equals one, so r_α equals the lower bound $(1 - f_\alpha)\bar{v}$ in Claim 3. This shows that, when the regulator's redistribution objective is strong enough, his worst-case regret is pinned down by the trade-off between protecting consumers' surplus and mitigating underproduction. In contrast, when α is large enough (i.e., $\alpha > 1/2$), q_α is smaller than one and r_α is greater than the lower bound in Claim 3. This shows that, when the redistribution objective is not so strong or, equivalently, when the regulator's main concern is efficiency, the additional trade-off between under- and overproduction leads to a greater r_α than the lower bound in Claim 3.

B. Optimal Policy

To illustrate our optimal policy and the intuition behind it, we begin with two polar cases of the weight α : $\alpha = 0$ and $\alpha = 1$. We then extend the intuition from these two cases to any weight between zero and one.

Optimal Policy for $\alpha = 0$ and That for $\alpha = 1$.—We first consider the case that $\alpha = 0$. One unit of the firm's profit translates into one unit of regret. Compared to the cases in which $\alpha > 0$, the regulator with $\alpha = 0$ has the strongest redistribution objective and thus the strongest incentive to constrain the firm's revenue. We next show that it is optimal to cap the firm's average revenue at $f_0\bar{v} = \bar{v}/2$.

CLAIM 4 (Optimal Policy for $\alpha = 0$): *Assume that $\alpha = 0$. The policy*

$$(4) \quad \rho(q, p) = \min\{qf_0\bar{v}, qp\} = \min\{q\bar{v}/2, qp\}$$

achieves the worst-case regret $r_0 = \bar{v}/2$.

Both Claim 4 and Claim 5 follow from Theorem 2, so we omit their proofs. When the market price is p , the regulator learns that among consumers that are served, each values the firm's product at least p . The regulator uses this information by setting the firm's average revenue to be $\min\{\bar{v}/2, p\}$. Under this policy, the firm's inverse-demand function is effectively $\min\{\bar{v}/2, P(q)\}$. The firm chooses (q, p) to maximize its profit subject to this inverse-demand function and its cost function C . Since the firm's marginal revenue at any q is at most $P(q)$, it will not overproduce. In contrast, capping the firm's average revenue may induce underproduction, as is evident from the right panel of Figure 1. The optimal cap is $\bar{v}/2$ because distortion from underproduction and the firm's profit are equally undesirable for the regulator with $\alpha = 0$.

Policy (4) can be implemented by imposing a price cap at $\bar{v}/2$. The corresponding price cap policy is

$$(5) \quad \rho(q, p) = \begin{cases} qp, & \text{if } p \leq \frac{\bar{v}}{2}, \\ -\infty, & \text{if } p > \frac{\bar{v}}{2}. \end{cases}$$

These two policies are equivalent in the sense that (i) (q, p) is a best response under policy (4) if and only if $(q, \min\{p, \bar{v}/2\})$ is a best response under policy (5) and (ii) these two responses induce the same profit for the firm and the same consumers' surplus under the two respective policies. The price cap policy is especially appealing for applications where it is easier to regulate the price than to tax the firm. We discuss more the price cap implementation after Theorem 2 and in Section IIC.

We next consider the case that $\alpha = 1$. The firm's profit causes no regret, so the regulator's sole concern is efficiency. If the regulator knew the inverse-demand function P , he would equate the firm's revenue with the total consumer value that it has created (e.g., Baron and Myerson 1982). This policy would perfectly align the firm's incentives with the regulator's, so no distortion or regret would occur. However, when the regulator does not know P , he is uncertain how much total consumer value the firm has created; he knows only that this value is at least qp and at most $q\bar{v}$ if the firm chooses (q, p) . Offering less revenue than the total consumer value that the firm has created may cause underproduction, but offering more may cause overproduction. For small enough q , the regulator offers the upper bound $q\bar{v}$ since he is more concerned about underproduction for small quantities. He then imposes a cap on the total subsidy, which limits potential overproduction.

CLAIM 5 (Optimal Policy for $\alpha = 1$): Assume that $\alpha = 1$. The policy

$$\rho(q, p) = \min\{qf_1\bar{v}, qp + r_1\} = \min\{q\bar{v}, qp + r_1\}$$

achieves the worst-case regret r_1 .

Under this policy, given an inverse-demand function P , there exists a threshold quantity $\hat{q}(P) \in (0, 1]$ such that the cap on the total subsidy starts binding for $q > \hat{q}(P)$. Formally, $\hat{q}(P) = \max\{q \in [0, 1] : q\bar{v} \leq qP(q) + r_1\}$. The firm's revenue is $q\bar{v}$ if it chooses $q \leq \hat{q}(P)$ and is $qP(q) + r_1$ if it chooses $q > \hat{q}(P)$. The higher the inverse-demand function P is, the larger the threshold quantity $\hat{q}(P)$ is.

How does this policy constrain regret from both under- and overproduction? The piece rate subsidy lifts the firm's marginal revenue to \bar{v} for any $q \leq \hat{q}(P)$, which is weakly higher than the consumer value $P(q)$ at q . Therefore, the subsidy encourages production if it is efficient to produce more, at least until the cap on the total subsidy starts binding. On the other hand, the policy caps the total subsidy to the firm at r_1 . Since distortion from overproduction is at most the total subsidy, regret from overproduction is also at most r_1 .

Optimal Policy for $\alpha \in [0, 1]$.—We now characterize an optimal policy for any $\alpha \in [0, 1]$.

THEOREM 2 (Optimal Policy for $0 \leq \alpha \leq 1$): Let

$$(6) \quad s_\alpha = \left(\sup\{q(f_\alpha\bar{v} - p) : q(1 - f_\alpha)\bar{v} - qp \log q > r_\alpha, (q, p) \in [0, 1] \times [0, f_\alpha\bar{v}]\} \right)^+.$$

Then $s_\alpha \leq r_\alpha$, and for every $s \in [s_\alpha, r_\alpha]$, the policy

$$(7) \quad \rho(q, p) = \min\{qf_\alpha\bar{v}, qp + s\}$$

achieves the worst-case regret r_α .

Policy (7) uses the same instruments as in the two polar cases for the same intuition. First, the firm’s average revenue is capped at $f_\alpha\bar{v}$. This caps how much consumers’ surplus the firm can extract and therefore caps the regulator’s regret from the firm’s profit. Second, if the firm prices below $f_\alpha\bar{v}$, the policy offers a piece rate subsidy that lifts the firm’s average revenue to $f_\alpha\bar{v}$ for small enough quantities. The subsidy encourages the firm to serve more consumers than just those with high values, so it prevents severe underproduction. Third, the total subsidy to the firm is capped at some level s , so regret from overproduction is also at most s .

As the weight α on the firm’s profit increases from zero to one, the value of $f_\alpha\bar{v}$ increases from $\bar{v}/2$ to \bar{v} . This shows that, as the regulator’s redistribution objective weakens, he is less willing to constrain the firm’s average revenue.

According to (6), s_α is the maximal regret from overproduction that the regulator must face if he wants to keep regret from underproduction and the firm’s profit under r_α . Hence, the cap s on the total subsidy is at least s_α . On the other hand, the cap s cannot exceed r_α since otherwise, regret from overproduction may exceed r_α . We depict the value of s_α in the right panel of Figure 3. For $\alpha < 1/2$, $s_\alpha = \alpha f_\alpha\bar{v} < r_\alpha$, and for $\alpha \geq 1/2$, $s_\alpha = r_\alpha$. Thus, there is freedom in choosing the cap s for $\alpha < 1/2$ but no freedom for $\alpha \geq 1/2$. This contrast reflects again the idea that for small α , the worst-case regret r_α is pinned down only by the trade-off between protecting consumers’ surplus and mitigating underproduction, while for large α , the additional trade-off between under- and overproduction is crucial in pinning down r_α .

Policy (7) caps the firm’s average revenue at $f_\alpha\bar{v}$. As in the case of $\alpha = 0$, this cap on average revenue can be implemented by imposing a price cap at $f_\alpha\bar{v}$. The price cap implementation corresponds to the following policy:

$$(8) \quad \rho(q, p) = \begin{cases} \min\{qf_\alpha\bar{v}, qp + s\}, \text{ with } s \in [s_\alpha, r_\alpha], & \text{if } p \leq f_\alpha\bar{v}, \\ -\infty, & \text{if } p > f_\alpha\bar{v}. \end{cases}$$

The optimal policy in Theorem 2 features three properties. First, the firm’s average revenue is capped at $f_\alpha\bar{v}$. Second, if $s_\alpha > 0$, then for some quantity-price pair, the total subsidy to the firm is at least s_α . Third, the total subsidy to the firm is at most r_α . Not every optimal policy has the same form as policy (7) does, but Theorem 3 asserts that every optimal policy has similar properties. Recall that q_α achieves the maximum in the definition of r_α in (3).

THEOREM 3 (Properties of Any Optimal Policy): *Let ρ be an optimal policy. Then,*

1. (Cap on Average Revenue): $\rho(q, p) \leq qf_\alpha\bar{v}$ for every $q \leq q_\alpha$.

2. (*Subsidy*): If $s_\alpha > 0$, then there exists some (q, p) such that $\rho(q, p) \geq qp + s_\alpha$.
3. (*Cap on Total Subsidy*): $\rho(q, p) \leq qp + r_\alpha$ for every (q, p) .

In particular, since $q_\alpha = 1$ for $\alpha \leq 1/2$, it follows from Theorem 3 that for $\alpha \leq 1/2$, the firm's average revenue is capped at $f_\alpha \bar{v}$ for every quantity $q \in [0, 1]$.

C. When Does the Firm Clear the Market?

In this subsection, we discuss whether and when the firm wants to clear the market.

First, the optimal policy (7) is weakly increasing in price p , so if (q, p) is a best response of the firm under this policy, then $(q, P(q))$ is also a best response. This shows that under policy (7), the firm always has a best response that clears the market.

Second, we have shown that the cap on average revenue can be implemented by imposing a price cap. The price cap implementation (8) is no longer weakly increasing in p , so under this policy the firm may not have a best response that clears the market. For instance, suppose that $\alpha = 0$ and that the regulator imposes a price cap at $\bar{v}/2$. Suppose also that the firm's inverse-demand function is $P(q) = \bar{v}$ and its cost function is $C(q) = \bar{v}q^2/2$. The firm's best response is $(q, p) = (1/2, \bar{v}/2)$, which cannot clear the market.

Third, a common assumption in the monopoly-regulation literature is that the firm has decreasing average cost. We next show that, if we assume decreasing average cost, then even under the price cap implementation (8), the firm has a best response that clears the market.

PROPOSITION 1: *Assume that $C(q)/q$ decreases in q for $q > 0$. Then under policy (8), the firm has a best response that clears the market.*

PROOF:

Let (q, p) be a firm's best response to its (P, C) under policy (8). Since policy (8) is increasing in p for $p \leq f_\alpha \bar{v}$, we can assume without loss that $p = P(q)$ or $p = f_\alpha \bar{v}$ (since otherwise, we can replace p by $\min\{P(q), f_\alpha \bar{v}\}$). In the former case, $(q, P(q))$ clears the market. In the latter case, let $q' \geq q$ be the maximal quantity that the firm can sell at price $f_\alpha \bar{v}$. Then since $\rho(q', f_\alpha \bar{v}) = q' f_\alpha \bar{v}$ and $\rho(q, f_\alpha \bar{v}) = q f_\alpha \bar{v}$, we get that

$$\begin{aligned} FP(\rho, P, C, q', f_\alpha \bar{v}) &= q' \left(f_\alpha \bar{v} - \frac{C(q')}{q'} \right) \geq q' \left(f_\alpha \bar{v} - \frac{C(q)}{q} \right) \\ &= \frac{q'}{q} FP(\rho, P, C, q, f_\alpha \bar{v}) \geq FP(\rho, P, C, q, f_\alpha \bar{v}), \end{aligned}$$

where the first inequality follows from decreasing average cost. Therefore, $(q', f_\alpha \bar{v})$ is also a best response and clears the market by the definition of q' . ■

III. Regulatory Policies in Practice

Our results offer a number of insights and implications for the practical design of regulatory policies.

First, we show that the minimax-regret criterion selects a price cap policy as optimal if the regulator cares only about consumers' surplus (i.e., if $\alpha = 0$). In Proposition 2 of Section IVA, we extend this result by showing that a price cap policy is optimal if consumers are sufficiently homogeneous or if the weight α on the firm's profit is small enough. Price cap regulation was first developed in the 1980s and has since been applied in many industries and locations (e.g., Cowan 2002; Sappington and Weisman 2010). A usual rationale is that the policy is simple to implement and provides strong incentives for the firm to cut costs (e.g., Armstrong and Sappington 2007). By formalizing the regret-minimizing property of the price cap policy, we provide an additional explanation for why this policy is widespread in practice.

Second, we recommend a piece rate subsidy up to a target level to mitigate underproduction in less competitive markets. One such policy is the feed-in tariff in the renewable energy sector, which takes the form of a guaranteed price above the market price. The purpose is to encourage investment that otherwise may not take place. In 2008, a report by the European Commission concluded that "well-adapted feed-in tariff regimes are generally the most efficient and effective support schemes for promoting renewable electricity" (European Union 2008). By early 2010, at least 50 countries had feed-in tariffs (REN21 2010). Our analysis provides a theoretical foundation for feed-in tariffs while also highlighting the risk of overproduction due to such a policy. This risk was exemplified by Spain's solar-market experience around 2007–2009. Its feed-in tariff invited a flood of investment and a tariff payment as high as \$26.4 billion, which led to taxpayer backlash and suspension of the tariffs (Voosen 2009). Our result suggests that a cap on the total subsidy should have been imposed to prevent excessive overproduction. Such a cap is similar in spirit to Germany's practice "to reduce tariff rates if its capacity targets were exceeded" (Voosen 2009).

Third, although we assume a monopolistic firm in a product market, our analysis can be useful for other markets in which considerable monopoly or monopsony power is present (due to market concentration or other frictions like search cost). One possible application is the labor market, as a growing body of empirical studies has shown that labor market monopsony is pervasive.⁹ In Appendix C, we show how to recast our model to model labor market monopsony. The optimal policy in Theorem 2 can be interpreted as a combination of a minimum wage and a wage subsidy, or the Earned Income Tax Credit (EITC). This interpretation is elaborated in Appendix C.2.¹⁰ In particular, wage subsidies or EITC can make up the difference between a worker's reservation wage and the minimum wage paid by the firm, so they can encourage more employment than at the monopsony level. Our intuition resonates with the observation in Naidu and Posner (2022) that the minimum wage can "restrain firms' wage-setting power" and that wage subsidies can "rectify

⁹Most recent contributions include Dube, Giuliano, and Leonard (2019); Dube et al. (2020); Dube, Manning, and Naidu (2020); Bassier, Dube, and Naidu (2022).

¹⁰We thank Suresh Naidu for suggesting this labor market interpretation of our optimal policy.

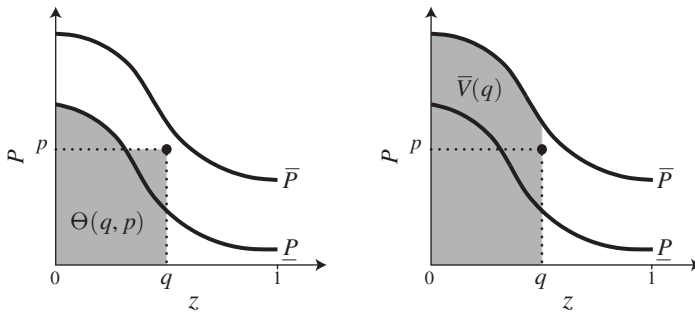


FIGURE 4. TOTAL CONSUMER VALUE THE FIRM HAS CREATED IS ABOVE $\Theta(q, p)$ AND BELOW $\bar{V}(q)$

inefficiencies” (p. S285). Our results thus provide a theoretical foundation for these policy tools.

IV. Incorporating Additional Knowledge

In our baseline model in Section I, the only assumptions we made on the inverse-demand and cost functions were monotonicity, semicontinuity, and an upper bound \bar{v} on the consumers’ values. We view this minimally informed regulator as a natural starting point. In some applications, the regulator may know more than this. We can incorporate the regulator’s knowledge by restricting the set of inverse-demand and cost functions in his problem.

Let \mathcal{E} be the set of possible inverse-demand and cost functions. The regulator chooses a policy ρ that minimizes the worst-case regret over all $(P, C) \in \mathcal{E}$ and all of the firm’s best responses (q, p) to (P, C) under ρ :

$$\text{minimize}_{\rho} \max_{(P, C) \in \mathcal{E}(q, p)} \text{RGRT}(\rho, P, C, q, p).$$

We now solve the regulator’s problem for some specific \mathcal{E} and thus demonstrate the adaptability of the minimax-regret approach.

A. Additional Knowledge about Demand

The regulator may know more specific upper and lower bounds on the inverse-demand function P than in the baseline model. Let $\underline{P}, \bar{P}: [0, 1] \rightarrow \mathbb{R}_+$ be two decreasing upper-semicontinuous functions such that $\underline{P}(z) \leq \bar{P}(z)$ for every $z \in [0, 1]$. Suppose that the regulator knows that P is bounded from below and above by \underline{P} and \bar{P} , respectively. Hence, $\mathcal{E} = \{(P, C) : \underline{P}(z) \leq P(z) \leq \bar{P}(z) \text{ for } z \in [0, 1]\}$. We let $\underline{V}(q)$ and $\bar{V}(q)$ denote the lower and upper bounds on the total consumer value of quantity q :

$$\underline{V}(q) = \int_0^q \underline{P}(z) dz, \quad \bar{V}(q) = \int_0^q \bar{P}(z) dz.$$

In this subsection, we characterize an optimal policy for any given bounds (\underline{P}, \bar{P}) on P . Figure 4 gives an example of such bounds. Our baseline model corresponds to the

case in which $\underline{P} \equiv 0$ and $\bar{P} \equiv \bar{v} > 0$. In Baron and Myerson (1982), the regulator knows P as the firm does, which corresponds to the case in which $\underline{P} = \bar{P} = P$.

If the firm chooses (q, p) , it provides evidence to the regulator that $P(q) \geq p$. The regulator also knows that $P \geq \underline{P}$, so the total consumer value that the firm has created is *at least*

$$\Theta(q, p) = \int_0^q \max\{\underline{P}(z), p\} dz.$$

On the other hand, the regulator knows that $P \leq \bar{P}$, so the total consumer value that the firm has created is at most $\bar{V}(q)$. We illustrate the value of $\Theta(q, p)$ in the left panel of Figure 4 and that of $\bar{V}(q)$ in the right panel. In our baseline model, $\Theta(q, p)$ equals the market revenue qp since $\underline{P} \equiv 0$, and $\bar{V}(q)$ equals $q\bar{v}$ since $\bar{P} \equiv \bar{v}$.

We now adjust policy (7) in Theorem 2 to incorporate the bounds (\underline{P}, \bar{P}) on P . We replace $q\bar{v}$ and qp with the more general terms of $\bar{V}(q)$ and $\Theta(q, p)$, respectively.

THEOREM 4 (Optimal Policy Given Bounds (\underline{P}, \bar{P})): *Assume that $\underline{P} \leq P \leq \bar{P}$. Let $\underline{d}(q) = \max_{0 \leq z \leq q} \{V(z) - f_\alpha \bar{V}(z)\}$, so $\underline{d}(q)$ is an increasing function and $\underline{d}(0) = 0$. There exists $\bar{s}_\alpha \geq \underline{s}_\alpha \geq 0$ such that for every $s \in [\underline{s}_\alpha, \bar{s}_\alpha]$, the policy*

$$(9) \quad \rho(q, p) = \min\{f_\alpha \bar{V}(q), \Theta(q, p) + s - \underline{d}(q)\}$$

achieves the minimal worst-case regret.

Theorem 4 follows from Theorem A.1 in Appendix A, which further characterizes the value of the minimal worst-case regret as well as the values of \underline{s}_α and \bar{s}_α . The proof is also relegated to Appendix A.

Policy (9) has a form similar to that of policy (7) for a similar intuition. The firm's revenue is at most a fraction f_α of the maximal total consumer value $\bar{V}(q)$ that it could have created. This cap on the firm's revenue keeps the regret from the firm's profit under control. Subject to this cap, the firm's revenue is the total consumer value $\Theta(q, p)$ that the firm proves it has created, plus an amount $(s - \underline{d}(q))$. The amount $(s - \underline{d}(q))$ is designed to keep the regret from both underproduction and overproduction under control. This amount weakly decreases in q , reflecting the idea that the regulator is more concerned about underproduction for small q and more concerned about overproduction for large q .

Policy (9) differs from policy (7) in that it includes an additional term $(- \underline{d}(q))$. This term equals zero in our baseline model since $\underline{P} \equiv 0$. The term has bite in markets where $\underline{P}(z)$ is very close to $\bar{P}(z)$ for small z 's but much smaller than $\bar{P}(z)$ for large z 's. Recall that, for any $z \in [0, q]$, the z th consumer contributes at least $\underline{P}(z)$ to $\Theta(q, p)$. When $\underline{P}(z)$ is very close to $\bar{P}(z)$ for small z 's, $\underline{P}(z)$ is higher than $f_\alpha \bar{P}(z)$. Under policy (9), the firm cannot use such high values of $\underline{P}(z)$ as a justification for increasing its marginal revenue for unit $q > z$, even though these high values of $\underline{P}(z)$ boost the value of $\Theta(q, p)$ for $q > z$.¹¹

¹¹ In the Supplemental Appendix, we provide an example to illustrate this role of $(- \underline{d}(q))$.

Next, we apply Theorem 4 to two special cases. In the first case, the regulator knows the inverse-demand function P . In the second, both \underline{P} and \bar{P} are constant functions.

The Regulator Knows the Inverse-Demand Function: $\underline{P} = \bar{P} = P$.—The assumption that $\underline{P} = \bar{P}$ implies the equality that $\underline{V}(q) = \bar{V}(q)$ for any $q \in [0, 1]$. Since $\underline{V}(q) \leq \Theta(q, p) \leq \bar{V}(q)$, this equality further implies that $\Theta(q, p) = \bar{V}(q)$. It also implies that $\underline{d}(q) = \max_{0 \leq z \leq q} \{ \underline{V}(z) - f_\alpha \bar{V}(z) \} = (1 - f_\alpha) \bar{V}(q)$. Substituting $\Theta(q, p)$ and $\underline{d}(q)$ into the optimal policy (9), we can reduce the policy to

$$(10) \quad \begin{aligned} \rho(q, p) &= \min \{ f_\alpha \bar{V}(q), \Theta(q, p) + s - \underline{d}(q) \} \\ &= \min \{ f_\alpha \bar{V}(q), f_\alpha \bar{V}(q) + s \} = f_\alpha \bar{V}(q). \end{aligned}$$

The firm's revenue is a fixed fraction f_α of the total consumer value $\bar{V}(q)$ that it has created. Because the regulator knows P already, the firm's price p provides no new information about P and thus no new information about the total consumer value that the firm has created. Hence, unlike in our baseline model, the firm's revenue does not depend on p . As the weight α goes from 0 to 1, the fraction f_α goes from 1/2 to 1. If f_α equals one, the firm's revenue equals the total consumer value that it has created, so it will choose an efficient quantity.

Next, we argue that the firm never overproduces under the optimal policy (10). Suppose otherwise that the firm chooses q even though it is more efficient to choose a lower quantity $q^* < q$. Since q^* is more efficient than q , the total consumer value in $[q^*, q]$ is strictly lower than the additional cost from q^* to q . Under policy (10), the firm's additional revenue from q^* to q is just a fraction f_α of the total consumer value in $[q^*, q]$, so it is strictly lower than the additional cost from q^* to q . This contradicts the presumption that choosing q is more profitable than choosing q^* . Hence, the firm never overproduces.

This result stands in contrast to the possibility of overproduction in our baseline model, suggesting that overproduction occurs only when the regulator is uncertain about the demand. Intuitively, when the regulator is uncertain about P , he is uncertain about how much total consumer value the firm has created. Thus, the optimal policy may offer more revenue to the firm than the total consumer value and thereby cause overproduction. Our results are consistent with what Baron and Myerson (1982) and Armstrong (1999) find using Bayesian models. Baron and Myerson (1982) assume that the regulator knows P . They also find that the firm never overproduces under the optimal policy. Armstrong (1999) assumes that the regulator is uncertain about both P and C . Using a four-type example, he shows that both over- and underproduction occur under the optimal policy. Our results generalize these findings to settings in which the firm's private information is infinite-dimensional.

Both \bar{P} and \underline{P} Are Constant Functions.—Let $\bar{P} \equiv \bar{v} > 0$ and $\underline{P} \equiv \underline{v} \in [0, \bar{v}]$. If $\underline{v} = 0$, we are back to the baseline model. If $\underline{v} = \bar{v}$, the regulator knows that every consumer has value \bar{v} . The ratio \underline{v}/\bar{v} quantifies the demand uncertainty faced by the regulator: the higher \underline{v}/\bar{v} is, the smaller the uncertainty. It can also be viewed

as a measure of how homogeneous consumers are. We next show that a price cap policy is optimal when consumers are sufficiently homogeneous.

PROPOSITION 2 (Price Cap Optimality): *If $\underline{v}/\bar{v} \geq f_\alpha = 1/(2 - \alpha)$, it is optimal to impose a price cap at $f_\alpha\bar{v}$.*

PROOF:

Given that $\bar{P} \equiv \bar{v}$ and $\underline{P} \equiv \underline{v}$, it follows that $\bar{V}(q) = q\bar{v}$ and $\Theta(q, p) = q \max\{\underline{v}, p\}$. Given that $\underline{v} \geq f_\alpha\bar{v}$, it follows that $\underline{d}(q) = q(\underline{v} - f_\alpha\bar{v})$. The optimal policy (9) reduces to

$$(11) \quad \rho(q, p) = \min\{f_\alpha q\bar{v}, q \max\{\underline{v}, p\} + s - q(\underline{v} - f_\alpha\bar{v})\} = qf_\alpha\bar{v},$$

where the last equality follows from the fact that $q \max\{\underline{v}, p\} - q(\underline{v} - f_\alpha\bar{v}) \geq q\underline{v} - q(\underline{v} - f_\alpha\bar{v}) = qf_\alpha\bar{v}$ and that $s \geq 0$. We next show that policy (11) can be implemented by the price cap policy $\rho^*(q, p)$: $\rho^*(q, p) = qp$ if $p \leq f_\alpha\bar{v}$, and $\rho^*(q, p) = -\infty$ if $p > f_\alpha\bar{v}$. First, a best response $(q, f_\alpha\bar{v})$ under ρ^* is also a best response under policy (11), and this $(q, f_\alpha\bar{v})$ induces the same profit for the firm and the same consumers' surplus under both policies. Second, if (q, p) is a best response under policy (11), then $(q, f_\alpha\bar{v})$ is a best response under ρ^* , and $(q, f_\alpha\bar{v})$ induces the same profit and the same consumers' surplus under ρ^* as (q, p) does under policy (11). ■

When \underline{v} is sufficiently close to \bar{v} , consumers with the highest values do not value the product much more than consumers with the lowest values. In such a case, the regulator is not concerned that the firm will serve just a small group of consumers with very high values. Hence, he has no incentive to offer a subsidy to encourage more production. Similarly, when α is small enough, the regulator's redistribution objective is sufficiently strong that he is not willing to subsidize the firm.

B. Additional Knowledge about Cost

The regulator may know that the firm has a fixed cost and a constant marginal cost but does not know the cost levels. Hence, $\mathcal{E} = \{(P, C) : P(q) \in [0, \bar{v}], C(q) = aq + b \text{ for } q > 0 \text{ for some } a, b \geq 0\}$. This is the type of cost function used most frequently in studies of monopoly regulation.

In our Proof of Theorem 1, we show that the worst-case regret under any policy is at least r_α using only fixed-cost functions. This means that Theorem 1 remains true for every set of cost functions that includes the set of all fixed-cost functions. In particular, Theorem 1 remains true for $\mathcal{E} = \{(P, C) : P(q) \in [0, \bar{v}], C(q) = aq + b \text{ for } q > 0 \text{ for some } a, b \geq 0\}$. Therefore, policy (7) in Theorem 2 remains optimal since the worst-case regret under policy (7) is at most r_α for any $(P, C) \in \mathcal{E}$.

Another natural restriction on the cost function is that the firm has decreasing average cost, so $\mathcal{E} = \{(P, C) : P(q) \in [0, \bar{v}], C(q)/q \text{ decreases in } q \text{ for } q > 0\}$. Since this set also includes the set of all fixed-cost functions, Theorems 1 and 2 remain true for this \mathcal{E} as well.

V. Conclusion

How to regulate a monopolistic firm when information asymmetries are pronounced and multidimensional is an important and thus far underexplored topic (Armstrong and Sappington 2007). We study such a setting in which the regulator has very little information about the firm's demand and cost scenarios. The regulator looks for a policy that works "fairly well" in all circumstances by minimizing his worst-case regret. We show that the optimal policy features three instruments: (i) a cap on the firm's average revenue to protect consumers' surplus, (ii) a piece rate subsidy up to a target level to mitigate underproduction, and (iii) a cap on the total subsidy to limit potential overproduction. When the consumers are sufficiently homogeneous or the regulator's redistributive objective is sufficiently strong, a cap on the firm's average revenue is sufficient, which can be implemented by a price cap policy. Our result thus provides another explanation for why price cap policies are popular. By selecting a piece rate subsidy up to a target level as optimal, we provide a theoretical foundation for subsidy programs like feed-in tariffs. Our analysis also highlights the risk of overproduction under such programs and thus emphasizes the importance of capping the total subsidy to the firm.

The minimax-regret approach used in this paper could be useful for studying other variants of the monopoly regulation problem. Here are a few examples. First, we have assumed that there is no social cost of public funds. This assumption is not needed when a price cap alone is optimal but is needed when the optimal policy must involve a subsidy. It will be useful to examine systematically how the cost of public funds shapes the optimal policy. Second, we have taken the first steps in exploring how the regulator's additional knowledge affects the optimal policy. There might be markets in which the regulator knows that the maximal total surplus is positive so production is efficient; there might be markets in which the regulator has knowledge about the levels of the firm's fixed or marginal costs. It will be useful to explore how these types of knowledge shape the optimal policy. Lastly, we have focused on settings in which the firm chooses a uniform price. There are other settings in which the firm can engage in price discrimination. Our analysis already provides a solution if the firm can engage in first-degree price discrimination and charge each consumer their value. In this case, the firm's market revenue is exactly the total consumer value it has created. This is as if the regulator knew the firm's inverse-demand function P , so policy (10) in Section IVA is optimal. Our model can be modified to study the regret-minimizing two-part tariffs. A natural starting point is that each consumer has a (possibly different) inverse-demand function and that the regulator observes the firm's entry fee, its per unit charge, and the total quantity sold. We leave these questions for future research.

APPENDIX A. BOUNDS (\underline{P}, \bar{P}) ON THE INVERSE-DEMAND FUNCTION P

Theorem A.1 characterizes the value of the minimal worst-case regret and an optimal policy for the environment $\mathcal{E} = \{(P, C) : \underline{P}(z) \leq P(z) \leq \bar{P}(z) \text{ for } z \in [0, 1]\}$ in Section IVA. Theorems 1 and 2 follow as a special case.

For every q and $v \geq \underline{V}(q)$, let

$$\psi_v(q) = \max\{\underline{P}(q) \leq p \leq \bar{P}(q) : \Theta(q, p) \leq v\}$$

be such that if the firm sells q units at price $p \leq \psi_v(q)$, then the total consumer value that the firm has created might be below v .

THEOREM A.1 (General bounds (\underline{P}, \bar{P}) on P): Assume that $\underline{P}(z) \leq P(z) \leq \bar{P}(z)$, $\forall z \in [0, 1]$. Let $\underline{D} = \max_{0 \leq z \leq 1} \{\underline{V}(z) - f_\alpha \bar{V}(z)\}$ and $\bar{D} = \max_{0 \leq z \leq 1} \{\underline{V}(z) - \alpha f_\alpha \bar{V}(z)\}$. For (q, v, c, q'') such that $0 \leq q \leq q'' \leq 1$, $v \geq \underline{V}(q'')$, and $c \geq 0$, let

$$r_u(q, v, c, q'') = (1 - f_\alpha) \bar{V}(q) + \int_q^{q''} \psi_v(z) dz - c, \quad \text{and}$$

$$r_o(q, v, c, q'') = f_\alpha \bar{V}(q) - v + c + \bar{D}.$$

1. Let R_α denote the optimal value of the following problem (12):

$$(12) \quad \text{maximize}_{(q, v, c, q'')} \min \{r_u(q, v, c, q''), r_o(q, v, c, q'')\}$$

subject to

$$0 \leq q \leq q'' \leq 1, \quad v \geq \underline{V}(q''), \quad \text{and} \quad c \geq 0.$$

The worst-case regret under any policy is at least R_α .

2. Let

$$\underline{s}_\alpha = \left(\sup \{f_\alpha \bar{V}(q) - v + c + \underline{D} : r_u(q, v, c, q'') > R_\alpha\} \right)^+, \quad \text{and}$$

$$\bar{s}_\alpha = R_\alpha + \underline{D} - \bar{D}.$$

Then $\underline{s}_\alpha \leq \bar{s}_\alpha$ and for every $s \in [\underline{s}_\alpha, \bar{s}_\alpha]$, the policy

$$(13) \quad \rho(q, p) = \min\{f_\alpha \bar{V}(q), \Theta(q, p) + s - \underline{d}(q)\}$$

achieves the worst-case regret R_α .

Roughly speaking, the term $r_u(q, v, c, q'')$ represents regret from underproduction and the firm's profit, and $r_o(q, v, c, q'')$ represents regret from overproduction. In the expression of r_u , q is the firm's quantity, v is the total consumer value that the firm proves it has created, q'' is an efficient quantity in the case of underproduction, and c is the additional cost from q to q'' .

Problem (12) is a finite constrained optimization problem. We next show that problem (12) admits an optimal solution with $c = 0$ and $q'' = 1$.

CLAIM 6:

1. The optimal value R_α of problem (12) is at least $(1 - f_\alpha)\bar{V}(1)$.
2. There exists an optimal solution to problem (12) such that $c = 0$ and $q'' = 1$.

PROOF:

Denote a feasible point of problem (12) by (q, v, c, q'') .

1. Let $\mu = (1, \underline{V}(1), 0, 1)$, which is a feasible point. Then $r_u(\mu) = (1 - f_\alpha)\bar{V}(1)$ and $r_o(\mu) = f_\alpha\bar{V}(1) - \underline{V}(1) + \bar{D} \geq (1 - \alpha)f_\alpha\bar{V}(1) = (1 - f_\alpha)\bar{V}(1)$ since $\bar{D} \geq \underline{V}(1) - \alpha f_\alpha\bar{V}(1)$. Therefore, $R_\alpha \geq \min\{r_u(\mu), r_o(\mu)\} = (1 - f_\alpha)\bar{V}(1)$, as desired.
2. Let $\nu = (q, v, c, q'')$ be a feasible point of problem (12). Then the point $\tilde{\nu} = (q, \max\{v, \underline{V}(1)\}, c + \underline{V}(1) - \underline{V}(q''), 1)$ is also feasible. Moreover, $r_u(\nu) \leq r_u(\tilde{\nu})$ and $r_o(\nu) \leq r_o(\tilde{\nu})$. Therefore, we can assume that the optimal solution has $q'' = 1$. To show that there exists an optimal solution with $c = 0$, we start with some optimal solution $\nu = (q, v, c, 1)$ and let $\tilde{q} = \max\{q \leq z \leq 1 : c \geq f_\alpha(\bar{V}(z) - \bar{V}(q))\}$ and $\tilde{c} = c - f_\alpha(\bar{V}(\tilde{q}) - \bar{V}(q))$. Then $\tilde{\nu} = (\tilde{q}, v, \tilde{c}, 1)$ is another feasible point with $r_u(\tilde{\nu}) \geq r_u(\nu)$ and $r_o(\tilde{\nu}) = r_o(\nu)$. Therefore, $\tilde{\nu}$ is also an optimal solution with either $\tilde{c} = 0$, which is what we wanted, or $\tilde{q} = 1$. If $\tilde{q} = 1$, then the optimal value R_α is at most $(1 - f_\alpha)\bar{V}(1)$. This, combined with part 1, implies that $R_\alpha = (1 - f_\alpha)\bar{V}(1)$, so $\mu = (1, \underline{V}(1), 0, 1)$ is also an optimal solution. ■

A.1 The Case That $\underline{P} \equiv 0$ and $\bar{P} \equiv \bar{v}$

We now show that Theorems 1 and 2 follow from Theorem A.1 for the case that $\underline{P} \equiv 0$ and $\bar{P} \equiv \bar{v}$. Indeed, in this case, (i) $\Theta(q, p)$ is the market revenue qp , and $\bar{V}(q)$ is $q\bar{v}$; (ii) $\underline{V} \equiv 0$, $\underline{D} = \bar{D} = 0$, and $\underline{d} \equiv 0$; and (iii) the function $\psi_v(z)$ is

$$\psi_v(z) = \begin{cases} \bar{v}, & \text{if } z \leq \frac{v}{\bar{v}}, \\ \frac{v}{z}, & \text{if } z > \frac{v}{\bar{v}}. \end{cases}$$

It follows that for $q \in [0, 1]$ and $v \in [0, qf_\alpha\bar{v}]$,

$$r_u(q, v, 0, 1) = (1 - f_\alpha)q\bar{v} + \int_q^1 \psi_v(z)dz = q(1 - f_\alpha)\bar{v} - v \log q, \quad \text{and}$$

$$r_o(q, v, 0, 1) = qf_\alpha\bar{v} - v.$$

Also, any $(q, v, 0, 1)$ with $v > qf_\alpha \bar{v}$ cannot be an optimal solution to problem (12) because $r_o(q, v, 0, 1) = qf_\alpha \bar{v} - v < 0$ for such $(q, v, 0, 1)$. Therefore, by Claim 6, problem (12) reduces to the lower bound in Theorem 1 with $v = qp$, so $R_\alpha = r_\alpha$. Moreover, the optimal policy (13) reduces to (7). The value of \bar{s}_α is R_α , and (using the arguments in Claim 6 again) the value of \underline{s}_α reduces to the value of s_α in Theorem 2.

A.2 Proof of Lower Bound

In this section we show that the worst-case regret under any policy is at least the optimal value R_α of problem (12). The argument is the same as in the Proof of Theorem 1, with an additional complication that did not appear in the case in which $\underline{P} \equiv 0$. We will show that the function ψ_v plays the same role as the function P^u does in the Proof of Theorem 1. Claim 7 is the counterpart of Claim 3 for general bounds (\bar{P}, \underline{P}) on P . For a policy ρ , let $WCR(\rho)$ be the worst-case regret under ρ .

CLAIM 7: For every policy ρ , it holds that $WCR(\rho) \geq (1 - f_\alpha)\bar{V}(1)$.

PROOF:

Fix a policy ρ . Let $k = \max_{(q,p)} \rho(q,p)$ be the highest revenue that the firm can get under ρ . We separate to two cases.

1. If $k \geq f_\alpha \bar{V}(1)$, then, if the inverse-demand function is \bar{P} and the cost is 0, the firm's profit is k and $RGRT \geq (1 - \alpha)k \geq (1 - \alpha)f_\alpha \bar{V}(1) = (1 - f_\alpha)\bar{V}(1)$.
2. If $k < f_\alpha \bar{V}(1)$, then, if the inverse-demand function is \bar{P} and there is a fixed cost of $f_\alpha \bar{V}(1)$, the firm does not produce. In this case, $RGRT = DSTR = (1 - f_\alpha)\bar{V}(1)$. ■

Let $(q, v, c, q'') = (q, v, 0, 1)$ be an optimal solution to (12) as in Claim 6. Abusing notations, let

$$r_u = r_u(q, v, 0, 1) = (1 - f_\alpha)\bar{V}(q) + \int_q^1 \psi_v(z) dz, \quad \text{and}$$

$$r_o = r_o(q, v, 0, 1) = f_\alpha \bar{V}(q) - v + \bar{D}.$$

Fix a policy ρ . We need to show that $R_\alpha = \min\{r_u, r_o\} \leq WCR(\rho)$.

Let q' be such that $\bar{D} = \underline{V}(q') - \alpha f_\alpha \bar{V}(q')$. Assume first that $q \leq q'$. Then

$$\begin{aligned} r_o &\leq f_\alpha \bar{V}(q) - \underline{V}(1) + \underline{V}(q') - \alpha f_\alpha \bar{V}(q') \\ &\leq (1 - \alpha) f_\alpha \bar{V}(q') \\ &\leq (1 - f_\alpha) \bar{V}(1) \\ &\leq WCR(\rho), \end{aligned}$$

where the first inequality follows from $v \geq \underline{V}(1)$, the second from $q' \geq q$ and $q' \leq 1$, the third from $q' \leq 1$ and $(1 - \alpha)f_\alpha = 1 - f_\alpha$, and the fourth from Claim 7.

We now assume $q' < q$. Let $U_{q,v}$ be the inverse-demand function such that $U_{q,v}(z) = \bar{P}(z)$ for $z \leq q$ and $U_{q,v}(z) = \psi_v(z)$ for $q < z \leq 1$. Let

$$x = \max_{z \leq q, p \leq \bar{P}(z)} \rho(z, p), \quad y = \max_{z \geq q, \Theta(z, p) \leq v} \rho(z, p), \quad \text{and } w = \max_{z \leq q', p \leq \bar{P}(z)} \rho(z, p).$$

Assume the inverse-demand function is $U_{q,v}$. Then x is the highest revenue that the firm can get by producing at most q , y is the highest revenue that it can get by producing at least q , and w the highest revenue that it can get by producing at most q' .

We separate to four (comprehensive, nonexclusive) cases:

1. If $x, y \leq f_\alpha \bar{V}(q)$, then consider the inverse-demand function $U_{q,v}$ with fixed cost $f_\alpha \bar{V}(q)$. The firm will not produce, with $FP = 0$ and $RGRT = DSTR = (1 - f_\alpha) \bar{V}(q) + \int_q^1 \psi_v(z) dz = r_u$ from underproduction.
2. If $x \geq f_\alpha \bar{V}(q)$ and $y \leq x$, then consider the inverse-demand function $U_{q,v}$ with cost 0. The firm will produce at most q units, with $FP = x$ and $DSTR \geq \int_q^1 \psi_v(z) dz$ from underproduction. Therefore, again, $RGRT \geq r_u$.
3. If $y \geq \max\{f_\alpha \bar{V}(q'), w\} + f_\alpha(\bar{V}(q) - \bar{V}(q'))$ and $y \geq x$. Let (\tilde{q}, \tilde{p}) achieve the maximum in the definition of y , so that $\Theta(\tilde{q}, \tilde{p}) \leq v$ and $\rho(\tilde{q}, \tilde{p}) = y$. Consider the inverse-demand function W that is given by

$$W(z) = \begin{cases} \bar{P}(z), & \text{if } z \leq q' \\ \max\{\underline{P}(z), \tilde{p}\}, & \text{if } q' < z \leq \tilde{q} \\ \underline{P}(z), & \text{if } z > \tilde{q}, \end{cases}$$

and the cost function such that producing q' units is costless and producing additional units incurs a fixed cost of $y - w$. If the firm produces q' units, it creates the total surplus of $\bar{V}(q')$. However, the highest profit the firm can get by producing at most q' units is w , and therefore, the firm will produce \tilde{q} units and sell them at price \tilde{p} with $FP = w$. The additional value to consumers created by producing the extra units is given by

$$\begin{aligned} \int_{q'}^{\tilde{q}} W(z) dz &= \int_{q'}^{\tilde{q}} \max\{\underline{P}(z), \tilde{p}\} dz \\ &= \Theta(\tilde{q}, \tilde{p}) - \Theta(q', \tilde{p}) \\ &\leq \Theta(\tilde{q}, \tilde{p}) - \underline{V}(q') \\ &\leq v - \underline{V}(q'). \end{aligned}$$

Let Δ be the difference between the additional cost and this additional value to consumers. Then

$$(14) \quad \Delta \geq y - w - v + \underline{V}(q') \geq r_o - (1 - \alpha)f_\alpha \bar{V}(q') + (f_\alpha \bar{V}(q') - w)^+.$$

We now separate to two cases:

(3.1) If $\Delta \leq 0$, then $r_o \leq (1 - \alpha)f_\alpha \bar{V}(q') = (1 - f_\alpha)\bar{V}(q') \leq WCR(\rho)$, where the first inequality follows from (13) and the second from Claim 7.

(3.2) If $\Delta > 0$, then the firm overproduces, with $DSTR = \Delta$ and $FP = w$, and therefore, from (13), $RGRT = (1 - \alpha)w + \Delta \geq r_o + (1 - \alpha)(w - f_\alpha \bar{V}(q')) + (f_\alpha \bar{V}(q') - w)^+ \geq r_o$.

Once we exclude these three cases, we are left with the case that $f_\alpha \bar{V}(q) \leq y \leq \max\{f_\alpha \bar{V}(q'), w\} + f_\alpha(\bar{V}(q) - \bar{V}(q'))$ and $y \geq x$, which is a subset of the following case.

4. If $w \geq f_\alpha \bar{V}(q')$ and $x, y \leq w + f_\alpha(\bar{V}(q) - \bar{V}(q'))$, then consider the inverse-demand function $U_{q,v}$ with cost such that producing q' units is costless and producing additional units incurs a fixed cost of $f_\alpha(\bar{V}(q) - \bar{V}(q'))$. The firm produces at most q' units, with $FP = w$, and $DSTR \geq (1 - f_\alpha)[\bar{V}(q) - \bar{V}(q')] + \int_q^1 \psi_v(z) dz$ from underproduction. $RGRT \geq (1 - \alpha)f_\alpha \bar{V}(q') + (1 - f_\alpha)[\bar{V}(q) - \bar{V}(q')] + \int_q^1 \psi_v(z) dz = r_u$.

A.3 Proof of Optimality

In this section we prove the second part of Theorem A.1. We first show that $\underline{s}_\alpha \leq \bar{s}_\alpha$. The inequality follows from the following two observations:

- From the definitions of \underline{D} and \bar{D} in Theorem A.1, it follows that

$$\begin{aligned} \bar{D} - \underline{D} &\leq \max_{0 \leq z \leq 1} \{f_\alpha \bar{V}(z) - \alpha f_\alpha \bar{V}(z)\} \\ &= \max_{0 \leq z \leq 1} (1 - \alpha)f_\alpha \bar{V}(z) \\ &= (1 - \alpha)f_\alpha \bar{V}(1) \\ &\leq R_\alpha, \end{aligned}$$

where the last inequality follows from Claim 6. Therefore, $\bar{s}_\alpha = R_\alpha + \underline{D} - \bar{D} \geq 0$.

- For every (q, v, c, q'') such that $r_u(q, v, c, q'') > R_\alpha$, it follows from the definition of R_α in Theorem A.1 that $f_\alpha \bar{V}(q) - v + c + \bar{D} = r_o(q, v, c, q'') \leq R_\alpha$. Therefore, $\bar{s}_\alpha = R_\alpha + \underline{D} - \bar{D} \geq f_\alpha \bar{V}(q) - v + c + \underline{D}$.

These two observations imply that

$$\bar{s}_\alpha \geq \left(\sup \{f_\alpha \bar{V}(q) - v + c + \underline{D} : r_u(q, v, c, q'') > R_\alpha\} \right)^+ = \underline{s}_\alpha.$$

Let $s \in [\underline{s}_\alpha, \bar{s}_\alpha]$ and let ρ be the policy (13). Fix an inverse-demand function P and a cost function C . We need to show that $RGRT \leq R_\alpha$, where R_α is the optimal

value of problem (12). Since $\rho(q, p)$ increases in p , we can assume that the firm sets its price at $P(q)$ and receives revenue $\tilde{\rho}(q) = \min\{f_\alpha \bar{V}(q), \Theta(q, P(q)) + s - \underline{d}(q)\}$.

Efficient Production.—If the firm's quantity q is efficient, then

$$RGRT = (1 - \alpha)FP \leq (1 - \alpha)\tilde{\rho}(q) \leq (1 - \alpha)f_\alpha \bar{V}(q) \leq (1 - \alpha)f_\alpha \bar{V}(1) \leq R_\alpha,$$

where the last inequality follows from Claim 6.

Underproduction.—Assume that the firm underproduces at quantity q , and let $q'' > q$ be the minimal efficient quantity above q . Let $q^* = \max\{z \in [q, q''] : \tilde{\rho}(z) = f_\alpha \bar{V}(z)\}$ with the maximum defined to be q if this set is empty. Let $c = C(q'') - C(q)$ be the additional cost from q to q'' . Then $DSTR = \int_q^{q''} P(z) dz - c \leq \bar{V}(q^*) - \bar{V}(q) + \int_{q^*}^{q''} P(z) dz - c$. Since $FP \leq \tilde{\rho}(q) \leq f_\alpha \bar{V}(q)$, it follows that

$$(15) \quad RGRT = (1 - \alpha)FP + DSTR \leq (1 - f_\alpha)\bar{V}(q^*) - c^* + \int_{q^*}^{q''} P(z) dz,$$

where $c^* = c - f_\alpha(\bar{V}(q^*) - \bar{V}(q))$. We argue that $c^* \geq 0$. Indeed, if $q^* = q$, then $c^* = c \geq 0$; if $q^* > q$, since the firm weakly prefers to choose q with revenue at most $f_\alpha \bar{V}(q)$ over choosing q^* with revenue $f_\alpha \bar{V}(q^*)$ and paying an additional cost at most c , it follows that $c \geq f_\alpha(\bar{V}(q^*) - \bar{V}(q))$.

We separate to two cases:

- If $q^* = q''$, then $RGRT \leq (1 - f_\alpha)\bar{V}(q^*) - c^* \leq (1 - f_\alpha)\bar{V}(1) \leq R_\alpha$ by Claim 6.
- If $q^* < q''$, let $v = \sup\{\Theta(z, P(z)) : z \in (q^*, q'']\}$. Then $P(z) \leq \psi_v(z)$ for every $z \in (q^*, q'']$ by the definition of ψ_v . Therefore, from (15),

$$RGRT \leq (1 - f_\alpha)\bar{V}(q^*) - c^* + \int_{q^*}^{q''} \psi_v(z) dz = r_u(q^*, v, c^*, q'').$$

For every $z \in (q^*, q'']$ such that $\Theta(z, P(z)) \geq \underline{V}(q'')$, the firm could have chosen quantity z with revenue $\tilde{\rho}(z) = \Theta(z, P(z)) + s - \underline{d}(z)$ but chooses not to. It follows that $f_\alpha \bar{V}(q) \geq \Theta(z, P(z)) + s - \underline{d}(z) - c$, which implies that

$$\begin{aligned} s &\leq f_\alpha \bar{V}(q) - \Theta(z, P(z)) + c + \underline{d}(z) \\ &\leq f_\alpha \bar{V}(q) - \Theta(z, P(z)) + c + \underline{D} \\ &= f_\alpha \bar{V}(q^*) - \Theta(z, P(z)) + c^* + \underline{D}. \end{aligned}$$

For every $q^* < k < z$, we must have $\bar{V}(q^*) < \bar{V}(k)$, otherwise, q^* will also be efficient, contradicting $q^* < q''$. Therefore, $s < f_\alpha \bar{V}(k) - \Theta(z, P(z)) + c^* + \underline{D}$. Since $s \geq \underline{s}_\alpha$, this implies, by the definition of \underline{s}_α , that $r_u(k, \Theta(z, P(z)), c^*, q'') \leq R_\alpha$. Taking the limit with $k \downarrow q^*$, it follows that $r_u(q^*, \Theta(z, P(z)), c^*, q'') \leq R_\alpha$.

Since $v = \sup\{\Theta(z, P(z)) : z \in (q^*, q'')\}$, this implies that $RGRT \leq r_u(q^*, v, c^*, q'') \leq R_\alpha$, as desired.

Overproduction.—Assume that the firm overproduces at quantity q , and let $q'' < q$ be the maximal efficient quantity below q . Let $p = P(q)$. Let $c = C(q) - C(q'')$ be the additional cost from q'' to q . Since q maximizes the firm's profit, it follows that $c \leq \tilde{\rho}(q) - \tilde{\rho}(q'')$. Then

$$DSTR = c - \int_{q''}^q P(z)dz \leq \tilde{\rho}(q) - \tilde{\rho}(q'') - \int_{q''}^q P(z)dz.$$

Since we assume that $DSTR > 0$, it follows that

$$\begin{aligned} \tilde{\rho}(q'') &< \tilde{\rho}(q) - \int_{q''}^q P(z)dz \leq \Theta(q, p) + s - \underline{d}(q) - \int_{q''}^q P(z)dz \\ &\leq \Theta(q'', p) + s - \underline{d}(q'') \end{aligned}$$

because $\underline{d}(q'') \leq \underline{d}(q)$ and $\Theta(q, p) \leq \Theta(q'', p) + \int_{q''}^q P(z)dz$ as $P(z) \geq \max\{p, \underline{P}(z)\}$ for $z \in [q'', q]$. Therefore, by the definition of $\tilde{\rho}$, it must be the case that $\tilde{\rho}(q'') = f_\alpha \bar{V}(q'')$.

The firm's revenue is $\tilde{\rho}(q)$, and its cost is at least c , so its profit is at most $\tilde{\rho}(q) - c$. Therefore,

$$\begin{aligned} (16) \quad RGRT &= (1 - \alpha)FP + DSTR \\ &\leq (1 - \alpha)[\tilde{\rho}(q) - c] + c - \int_{q''}^q P(z)dz \\ &\leq \tilde{\rho}(q) - \alpha\tilde{\rho}(q'') - \int_{q''}^q P(z)dz \\ &= \tilde{\rho}(q) - \alpha f_\alpha \bar{V}(q'') - \int_{q''}^q P(z)dz, \end{aligned}$$

where the second inequality follows from $c \leq \tilde{\rho}(q) - \tilde{\rho}(q'')$.

Let $q^* = \max\{0 \leq z \leq q'' : \underline{P}(z) \geq p\}$ with the maximum defined to be 0 if $\underline{P}(0) < p$. We separate to two cases:

- If $p(q'' - q^*) \leq \alpha f_\alpha (\bar{V}(q'') - \bar{V}(q^*))$, then since

$$\begin{aligned} \tilde{\rho}(q) &\leq s + \Theta(q, p) - \underline{d}(q) \\ &= s + \underline{V}(q^*) + p(q'' - q^*) + \int_{q''}^q \max\{p, \underline{P}(z)\}dz - \underline{d}(q) \\ &\leq s + \underline{V}(q^*) + p(q'' - q^*) - \underline{d}(q) + \int_{q''}^q P(z)dz, \end{aligned}$$

we get from (16) and using the assumption $p(q'' - q^*) \leq \alpha f_\alpha (\bar{V}(q'') - \bar{V}(q^*))$ that

$$(17) \quad RGRT \leq s - \alpha f_\alpha \bar{V}(q^*) + \underline{V}(q^*) - \underline{d}(q).$$

Let $\bar{d}(q) = \max_{0 \leq z \leq q} \{V(z) - \alpha f_\alpha \bar{V}(z)\}$. Then $\bar{d} - \underline{d}$ is increasing and $\bar{d}(1) = \bar{D}$. From (17), we get

$$RGRT \leq s + \bar{d}(q) - \underline{d}(q) \leq s + \bar{D} - \underline{D} \leq R_\alpha$$

since $s \leq \bar{s}_\alpha$.

- If $p(q'' - q^*) > \alpha f_\alpha (\bar{V}(q'') - \bar{V}(q^*))$, then it must be the case that $q^* < q''$. Because \bar{V} is concave (since \bar{P} is decreasing), we have that $\int_{q^*}^{q''} P(z) dz \geq p(q - q'') > \alpha f_\alpha (\bar{V}(q) - \bar{V}(q''))$. Therefore, from (16) and from $\tilde{\rho}(q) \leq f_\alpha \bar{V}(q)$, it follows that

$$RGRT \leq (1 - \alpha) f_\alpha \bar{V}(q) = (1 - f_\alpha) \bar{V}(q) \leq R_\alpha,$$

from Claim 6.

APPENDIX B. PROOF OF THEOREM 3

The proof is based on a refined argument of the Proof of Theorem 1. For every (q, p) , we define two inverse-demand functions $P_{q,p}^u$ and $P_{q,p}^o$: (i) $P_{q,p}^u(z) = \bar{v}$ if $z \leq q$ and $P_{q,p}^u(z) = qp/z$ if $z > q$, and (ii) $P_{q,p}^o(z) = p$ if $z \leq q$ and $P_{q,p}^o(z) = 0$ if $z > q$. We depict $P_{q,p}^u$ and $P_{q,p}^o$ in Figure B1. The function $P_{q,p}^u$ exhibits unitary price elasticity when the price is in the range of $[qp, p]$. Similar functions have appeared in Roesler and Szentes (2017) and Condorelli and Szentes (2020, 2022).

Fix a policy ρ . Let $\bar{\rho}(q) = \max_{q' \leq q} \rho(q', p')$ be the highest revenue the firm can get from producing at most q . Let $\hat{\rho}(q, p) = \max_{q' \geq q, q' p' \leq qp} \rho(q', p')$ be the highest revenue if the firm produces at least q and its market revenue is at most qp . As shown in the left panel of Figure B1, $\bar{\rho}(q)$ is the maximum of ρ in the light-gray area, and $\hat{\rho}(q, p)$ is the maximum of ρ in the dark-gray area.

Claim 8 shows that $WCR(\rho)$ is at least the subsidy that policy ρ offers.

CLAIM 8: Fix a policy ρ . Then $WCR(\rho) \geq \rho(q, p) - qp$ for every (q, p) .

PROOF:

If $\rho(q, p) \leq qp$, the assertion follows from the fact that regret is nonnegative. Now suppose that $\rho(q, p) > qp$. If the firm has a fixed cost of $\rho(q, p)$ and its inverse-demand is $P_{q,p}^o$, it is a firm's best response to produce, so the firm incurs the fixed cost of $\rho(q, p)$. Distortion is at least $\rho(q, p) - qp$ due to overproduction, so regret is at least $\rho(q, p) - qp$. ■

Claim 9 shows that, if the firm does not receive enough additional revenue from producing more, there is sizable regret due to underproduction.

CLAIM 9: Fix a policy ρ . Let $\underline{q} \leq q \in [0, 1]$ and $p \in [0, \bar{v}]$. If $\hat{\rho}(q, p) \leq \bar{\rho}(\underline{q}) + (q - \underline{q}) f_\alpha \bar{v}$, then

$$WCR(\rho) \geq (1 - \alpha) [\bar{\rho}(\underline{q}) + (q - \underline{q}) f_\alpha \bar{v}] - qp \log q.$$

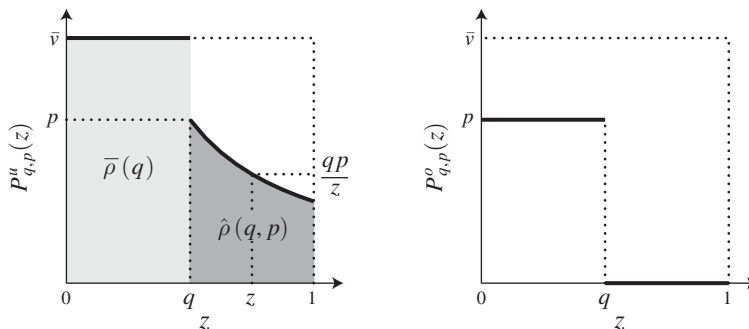


FIGURE B1. INVERSE-DEMAND FUNCTIONS $P_{q,p}^u$ AND $P_{q,p}^o$ AND DEFINITIONS OF $\bar{\rho}(q)$ AND $\hat{\rho}(q,p)$

PROOF:

1. If $\bar{\rho}(q) - \bar{\rho}(\underline{q}) \leq (q - \underline{q})f_\alpha \bar{v}$, then consider the inverse-demand function $P_{q,p}^u$ and a cost function such that producing \underline{q} or fewer units is costless and producing additional units incurs a fixed cost of $(q - \underline{q})f_\alpha \bar{v}$. The firm will produce at most \underline{q} units, with $FP = \bar{\rho}(\underline{q})$ and

$$DSTR \geq (q - \underline{q})(\bar{v} - f_\alpha \bar{v}) - qp \log q = (1 - \alpha)(q - \underline{q})f_\alpha \bar{v} - qp \log q,$$

because of underproduction. Therefore,

$$RGRT = (1 - \alpha)FP + DSTR \geq (1 - \alpha)[\bar{\rho}(\underline{q}) + (q - \underline{q})f_\alpha \bar{v}] - qp \log q.$$

2. If $\bar{\rho}(q) - \bar{\rho}(\underline{q}) \geq (q - \underline{q})f_\alpha \bar{v}$, then consider the inverse-demand function $P_{q,p}^u$ and zero cost. The firm will produce at most q units, with $FP = \bar{\rho}(q) \geq \bar{\rho}(\underline{q}) + (q - \underline{q})f_\alpha \bar{v}$, and $DSTR \geq -qp \log q$ because of underproduction. Therefore,

$$RGRT = (1 - \alpha)FP + DSTR \geq (1 - \alpha)[\bar{\rho}(\underline{q}) + (q - \underline{q})f_\alpha \bar{v}] - qp \log q. \blacksquare$$

Combining Claims 8 and 9, we show that the regulator suffers sizable regret from either underproduction or overproduction.

CLAIM 10: Fix a policy ρ . Let $\underline{q} \leq q \in [0, 1]$ and $p \in [0, \bar{v}]$. Then

$$WCR(\rho) \geq \min \left\{ (1 - \alpha)[\bar{\rho}(\underline{q}) + (q - \underline{q})f_\alpha \bar{v}] - qp \log q, \right. \\ \left. \bar{\rho}(\underline{q}) + (q - \underline{q})f_\alpha \bar{v} - qp \right\}.$$

PROOF:

If $\hat{\rho}(q, p) \geq \bar{\rho}(q) + (q - \underline{q})f_\alpha \bar{v}$, then let (q', p') be such that $q' \geq q$, $q'p' \leq qp$, and $\rho(q', p') = \hat{\rho}(q, p)$. By Claim 8, $WCR(\rho) \geq \rho(q', p') - q'p' \geq \bar{\rho}(q) + (q - \underline{q})f_\alpha \bar{v} - qp$.

If $\hat{\rho}(q, p) < \bar{\rho}(q) + (q - \underline{q})f_\alpha \bar{v}$, then $WCR(\rho) \geq (1 - \alpha)[\bar{\rho}(q) + (q - \underline{q})f_\alpha \bar{v}] - qp \log q$ by Claim 9. ■

PROOF OF THEOREM 3:

1. Let (q_α, p_α) achieve the maximum in the definition of r_α in (3). Assume that $\rho(\underline{q}, p) > \underline{q}f_\alpha \bar{v}$ for some $\underline{q} \leq q_\alpha$ and some p . Then $\bar{\rho}(\underline{q}) > \underline{q}f_\alpha \bar{v}$, and therefore, $\bar{\rho}(\underline{q}) + (q_\alpha - \underline{q})f_\alpha \bar{v} > q_\alpha f_\alpha \bar{v}$. Therefore, by Claim 10 with $q = q_\alpha$ and $p = p_\alpha$, it follows that $WCR(\rho)$ is at least

$$\begin{aligned} & \min\left\{(1 - \alpha)[\bar{\rho}(\underline{q}) + (q_\alpha - \underline{q})f_\alpha \bar{v}] - q_\alpha p_\alpha \log q_\alpha, \right. \\ & \left. \bar{\rho}(\underline{q}) + (q_\alpha - \underline{q})f_\alpha \bar{v} - q_\alpha p_\alpha\right\} \\ & > \min\left\{(1 - \alpha)q_\alpha f_\alpha \bar{v} - q_\alpha p_\alpha \log q_\alpha, q_\alpha(f_\alpha \bar{v} - p_\alpha)\right\} = r_\alpha. \end{aligned}$$

2. Suppose that $\rho(q, p) < qp + s_\alpha$ for every q, p . This implies that $\max_{q', p'}\{\rho(q', p') - p'q'\} < s_\alpha$. Given the definition of s_α in (6), there exists some $q \in [0, 1], p \in [0, f_\alpha \bar{v}]$ such that $(1 - \alpha)qf_\alpha \bar{v} - qp \log q > r_\alpha$ and $q(f_\alpha \bar{v} - p) > \max_{q', p'}\{\rho(q', p') - p'q'\}$. This second inequality, combined with the relation that $\max_{q', p'}\{\rho(q', p') - p'q'\} \geq \max_{q' \geq q, q'p' \leq qp}\{\rho(q', p') - p'q'\} \geq \hat{\rho}(q, p) - qp$, implies that $\hat{\rho}(q, p) < qf_\alpha \bar{v}$. By Claim 9 with $\underline{q} = 0$, we get that $WCR(\rho) > r_\alpha$.
3. Suppose that $\rho(q, p) > qp + r_\alpha$ for some q, p . Then $WCR(\rho) > r_\alpha$ by Claim 8. ■

APPENDIX C. LABOR MARKET MONOPSONY

Let $R: [0, 1] \rightarrow \mathbb{R}$ with $R(0) = 0$ be the employer’s prewage profit function, so her profit from hiring q workers before she pays the wage is $R(q)$. Let $L: [0, 1] \rightarrow \mathbb{R}$ be an increasing labor supply curve, with $L(q)$ being the reservation wage of the q th worker. The employer knows (L, R) , but the regulator does not know (L, R) . The employer chooses both the number of workers to hire and the wage paid to those hired. An employment-wage pair (q, w) is *feasible* if and only if $w \geq L(q)$.

According to a regulatory policy η , if the employer chooses (q, w) , then she pays the total amount of $\eta(q, w)$. If $\eta(q, w) > qw$, then the employer pays a tax of $\eta(q, w) - qw$, which goes to the entire population of the workers, not just those who are hired. If $\eta(q, w) < qw$, then the employer receives a subsidy of $qw - \eta(q, w)$, which is paid for by the entire population of the workers. The timing is as follows. The regulator publicly chooses and commits to a policy η . The employer privately observes (L, R) . She then publicly chooses (q, w) and pays the total amount of $\eta(q, w)$.

Fix a policy η and a pair of (L, R) functions. If the employer chooses (q, w) , then *workers' surplus* and *the employer's profit* are given by

$$WS(\eta, L, q, w) = \eta(q, w) - \int_0^q L(z) dz, \quad \text{and}$$

$$EP(\eta, L, R, q, w) = R(q) - \eta(q, w).$$

The regulator knows that $\underline{L}(q) \leq L(q) \leq \bar{L}(q)$ and that each worker contributes at most \hat{R} to the employer's prewage profit (i.e., $R(q + \varepsilon) - R(q) \leq \hat{R}\varepsilon$ for every q and every $\varepsilon \geq 0$). The regulator's payoff is $WS + \alpha EP$.

C.1 Reduction to the Monopoly Environment

Let $C(q) = \hat{R}q - R(q)$. Then $C(q)$ is an arbitrary increasing and nonnegative function with $C(0) = 0$. It satisfies decreasing average cost if $R(q)/q$ is increasing for $q > 0$.

Let $P(q) = \hat{R} - L(q)$. Then $P(q)$ is decreasing. The regulator knows that $\underline{P}(q) \leq P(q) \leq \bar{P}(q)$ with $\underline{P}(q) = \hat{R} - \bar{L}(q)$ and $\bar{P}(q) = \hat{R} - \underline{L}(q)$. A price p in the monopoly environment corresponds to a wage $w = \hat{R} - p$ in the monopsony environment. A policy ρ in the monopoly environment corresponds to a policy η in the monopsony environment by the relation that $\rho(q, p) = \hat{R}q - \eta(q, w)$.

Under these definitions, (i) (q, p) is feasible in the monopoly environment if and only if (q, w) is feasible in the monopsony environment, (ii) $CS(\rho, P, q, p) = WS(\eta, L, q, w)$, and (iii) $FP(\rho, P, C, q, p) = EP(\eta, L, R, q, w)$. Therefore, the regulator's problems in the monopoly and monopsony environments are equivalent, so we can apply Theorem A.1 to solve for the optimal $\eta(q, w)$.

C.2 Reduction to the Baseline Model in Section I

The case that \bar{L} and \underline{L} are constant functions corresponds to the case that \bar{P} and \underline{P} are constant functions. In particular, $\bar{L}(q) \equiv \bar{\ell}$ and $\underline{L}(q) \equiv \underline{\ell}$ correspond to $\bar{P} \equiv \hat{R} - \underline{\ell}$ and $\underline{P} \equiv \hat{R} - \bar{\ell}$. If we let $\bar{\ell} = \hat{R}$ and $\underline{\ell} = \hat{R} - \bar{v}$, then $\bar{P} \equiv \bar{v}$ and $\underline{P} \equiv 0$, so we can derive an optimal η from policy (7). For every $s \in [s_\alpha, r_\alpha]$, the following policy is optimal:

$$\begin{aligned} \eta(q, w) &= \hat{R}q - \rho(q, p) \\ &= \hat{R}q - \min\{qf_\alpha\bar{v}, qp + s\} \\ &= \max\{\hat{R}q - qf_\alpha\bar{v}, \hat{R}q - qp - s\} \\ &= \max\{q((1 - f_\alpha)\hat{R} + f_\alpha\underline{\ell}), qw - s\}. \end{aligned}$$

This policy uses three instruments. First, the employer's average wage payment, $\eta(q, w)/q$, is at least $((1 - f_\alpha)\hat{R} + f_\alpha\underline{\ell})$. This caps how much workers' surplus the employer can extract. Second, if the wage is above $((1 - f_\alpha)\hat{R} + f_\alpha\underline{\ell})$, the policy

offers a piece rate subsidy that lowers the average wage payment to $((1 - f_\alpha)\hat{R} + f_\alpha\underline{\ell})$ for sufficiently low employment level q . This subsidy prevents severe underemployment. Third, the total subsidy is at most s , so regret from overemployment is at most s .

For $\alpha = 0$, $s_\alpha = 0$, so it is optimal to choose $\eta(q, w) = \max\{q(\hat{R} + \underline{\ell})/2, qw\}$. This policy can be implemented by imposing a minimum wage at $(\hat{R} + \underline{\ell})/2$. The intuition is similar to that in the monopoly environment: \hat{R} is the highest possible value a worker can generate for the employer, and $\underline{\ell}$ is the lowest possible reservation wage. The minimum wage is halfway between \hat{R} and $\underline{\ell}$ in order to balance regret from inefficient underemployment and that from extracted workers' surplus.

REFERENCES

- Alonso, Ricardo, and Niko Matouschek. 2008. "Optimal Delegation." *Review of Economic Studies* 75 (1): 259–93.
- Amador, Manuel, and Kyle Bagwell. 2022. "Regulating a Monopolist with Uncertain Costs without Transfers." *Theoretical Economics* 17 (4): 1719–60.
- Arieli, Itai, Yakov Babichenko, and Rann Smorodinsky. 2018. "Robust Forecast Aggregation." *Proceedings of the National Academy of Sciences* 115 (52): E12135–43.
- Armstrong, Mark. 1999. "Optimal Regulation with Unknown Demand and Cost Functions." *Journal of Economic Theory* 84 (2): 196–215.
- Armstrong, Mark, and David E. M. Sappington. 2007. "Recent Developments in the Theory of Regulation." In *Handbook of Industrial Organization*, Vol. 3, edited by Mark Armstrong and Robert H. Porter, 1557–1700. Amsterdam: Elsevier.
- Armstrong, Mark, and John Vickers. 2010. "A Model of Delegated Project Choice." *Econometrica* 78 (1): 213–44.
- Babichenko, Yakov, and Dan Garber. 2021. "Learning Optimal Forecast Aggregation in Partial Evidence Environments." *Mathematics of Operations Research* 46 (2): 628–41.
- Baron, David P., and Roger B. Myerson. 1982. "Regulating a Monopolist with Unknown Costs." *Econometrica* 50 (4): 911–30.
- Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu. 2022. "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies." *Journal of Human Resources* 57 (S): S50–86.
- Bergemann, Dirk, Tan Gan, and Yingkai Li. 2023. "Managing Persuasion Robustly: The Optimality of Quota Rules." Unpublished.
- Bergemann, Dirk, and Karl H. Schlag. 2008. "Pricing without Priors." *Journal of the European Economic Association* 6 (2/3): 560–69.
- Bergemann, Dirk, and Karl Schlag. 2011. "Robust Monopoly Pricing." *Journal of Economic Theory* 146 (6): 2527–43.
- Beviá, Carmen, and Luis Corchón. 2022. "Contests with Dominant Strategies." *Economic Theory* 74 (4): 1163–81.
- Braeutigam, Ronald R. 1989. "Optimal Policies for Natural Monopolies." In *Handbook of Industrial Organization*, Vol. 2, edited by Richard Schmalensee and Robert Willig, 1289–1346. Amsterdam: Elsevier.
- Bubeck, Sébastien, and Nicolo Cesa-Bianchi. 2012. "Regret Analysis of Stochastic and Nonstochastic Multi-Armed Bandit Problems." *Foundations and Trends in Machine Learning* 5 (1): 1–122.
- Caillaud, B., R. Guesnerie, P. Rey, and J. Tirole. 1988. "Government Intervention in Production and Incentives Theory: A Review of Recent Contributions." *RAND Journal of Economics* 19 (1): 1–26.
- Caldentey, René, Ying Liu, and Ilan Lobel. 2017. "Intertemporal Pricing Under Minimax Regret." *Operations Research* 65 (1): 104–29.
- Carroll, Gabriel. 2015. "Robustness and Linear Contracts." *American Economic Review* 105 (2): 536–63.
- Carroll, Gabriel. 2019. "Robustness in Mechanism Design and Contracting." *Annual Review of Economics* 11 (1): 139–66.
- Chassang, Sylvain. 2013. "Calibrated Incentive Contracts." *Econometrica* 81 (5): 1935–71.

- Condorelli, Daniele, and Balázs Szentes.** 2020. "Information Design in the Holdup Problem." *Journal of Political Economy* 128 (2): 681–709.
- Condorelli, Daniele, and Balázs Szentes.** 2022. "Surplus Bounds in Cournot Oligopoly." *Theoretical Economics* 17 (3): 955–75.
- Cooper, Zack, Stuart V. Craig, Martin Gaynor, and John Van Reenen.** 2018. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." *Quarterly Journal of Economics* 134 (1): 51–107.
- Cowan, Simon.** 2002. "Price-Cap Regulation." *Swedish Economic Policy Review* 9 (2): 167–88.
- Drexler, Moritz, and Andreas Kleiner.** 2018. "Why Voting? A Welfare Analysis." *American Economic Journal: Microeconomics* 10 (3): 253–71.
- Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard.** 2019. "Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior." *American Economic Review* 109 (2): 620–63.
- Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri.** 2020. "Monopsony in Online Labor Markets." *American Economic Review: Insights* 2 (1): 33–46.
- Dube, Arindrajit, Alan Manning, and Suresh Naidu.** 2020. "Monopsony and Employer Mis-optimization Explain Why Wages Bunch at Round Numbers." Unpublished.
- European Union: European Commission.** 2008. *Commission Staff Working Document: Accompanying Document to the Proposal for a Directive of the European Parliament and of the Council on the Promotion of the Use of Energy from Renewable Sources*. Brussels: European Union.
- Feng, Yiding, Vahideh Manshadi, Rad Niazadeh, and Saba Neyshabouri.** 2024. "Robust Dynamic Staffing with Predictions." Unpublished.
- Garrett, Daniel F.** 2014. "Robustness of Simple Menus of Contracts in Cost-Based Procurement." *Games and Economic Behavior* 87: 631–41.
- Guesnerie, Roger.** 1998. *A Contribution to the Pure Theory of Taxation*. Cambridge: Cambridge University Press.
- Guo, Yingni, and Eran Shmaya.** 2023. "Regret-Minimizing Project Choice." *Econometrica* 91 (5): 1567–93.
- Hannan, James.** 1957. "Approximation to Bayes Risk in Repeated Play." In *Contributions to the Theory of Games*, Vol. 3, edited by M. Dresher, A. W. Tucker, and P. Wolfe, 97–139. Princeton, NJ: Princeton University Press.
- Hart, Sergiu, and Andreu Mas-Colell.** 2000. "A Simple Adaptive Procedure Leading to Correlated Equilibrium." *Econometrica* 68 (5): 1127–50.
- Holmström, Bengt.** 1984. "On the Theory of Delegation." In *Bayesian Models in Economic Theory*, edited by M. Boyer and R. Kihlstrom, 115–41. Amsterdam: Elsevier.
- Hurwicz, Leonid, and Leonard Shapiro.** 1978. "Incentive Structures Maximizing Residual Gain under Incomplete Information." *Bell Journal of Economics* 9 (1): 180–91.
- Kolotilin, Anton, and Andriy Zapechelyuk.** Forthcoming. "Persuasion meets Delegation." *Econometrica*.
- Laffont, Jean-Jacques, and Jean Tirole.** 1986. "Using Cost Observation to Regulate Firms." *Journal of Political Economy* 94 (3): 614–41.
- Laffont, Jean-Jacques, and Jean Tirole.** 1993. *A Theory of Incentives in Procurement and Regulation*. Cambridge, MA: MIT Press.
- Lewis, Tracy R., and David E. M. Sappington.** 1988a. "Regulating a Monopolist with Unknown Demand." *American Economic Review* 78 (5): 986–98.
- Lewis, Tracy R., and David E. M. Sappington.** 1988b. "Regulating a Monopolist with Unknown Demand and Cost Functions." *RAND Journal of Economics* 19 (3): 438–57.
- Malladi, Suraj.** 2022. "Delegated Costly Screening." Unpublished.
- Manski, Charles F.** 2004. "Statistical Treatment Rules for Heterogeneous Populations." *Econometrica* 72 (4): 1221–46.
- Mishra, Debasis, and Sanket Patil.** 2024. "Undominated Monopoly Regulation." Unpublished.
- Naidu, Suresh, and Eric A. Posner.** 2022. "Labor Monopsony and the Limits of the Law." *Journal of Human Resources* 57 (S): S284–S323.
- REN21.** 2010. *Renewables 2010 Global Status Report*. Paris: REN21.
- Renou, Ludovic, and Karl H. Schlag.** 2011. "Implementation in Minimax Regret Equilibrium." *Games and Economic Behavior* 71 (2): 527–33.
- Rochet, Jean-Charles.** 1986. "Le contrôle des équations aux dérivées partielles issues de la théorie des incitations." PhD diss. Paris IX.
- Roesler, Anne-Katrin, and Balázs Szentes.** 2017. "Buyer-Optimal Learning and Monopoly Pricing." *American Economic Review* 107 (7): 2072–80.

- Sappington, David E. M., and Dennis L. Weisman.** 2010. "Price Cap Regulation: What Have We Learned from 25 Years of Experience in the Telecommunications Industry?" *Journal of Regulatory Economics* 38 (3): 227–57.
- Savage, Leonard J.** 1954. *The Foundations of Statistics*. Hoboken, NJ: Wiley Publications in Statistics.
- Stoye, Jörg.** 2009. "Minimax Regret Treatment Choice with Finite Samples." *Journal of Econometrics* 151 (1): 70–81.
- Vera, Alberto, and Siddhartha Banerjee.** 2021. "The Bayesian Prophet: A Low-Regret Framework for Online Decision Making." *Management Science* 67 (3): 1368–91.
- Voosen, Paul.** 2009. "Spain's Solar Market Crash Offers a Cautionary Tale About Feed-In Tariffs." *New York Times*, August 8. <https://archive.nytimes.com/www.nytimes.com/gwire/2009/08/18/18green-wire-spains-solar-market-crash-offers-a-cautionary-88308.html?pagewanted=2>.
- Wald, Abraham.** 1950. *Statistical Decision Functions*. Hoboken, NJ: Wiley.