

Robust Monopoly Regulation

Yingni Guo* Eran Shmaya†

December 2019

Abstract

We study the regulation of a monopolistic firm using a non-Bayesian approach. We derive the 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 imposes a price cap. When his payoff is the total surplus of both consumers and the firm, he offers a capped piece-rate subsidy. For intermediate cases, the regulator uses both a price cap and a capped piece-rate subsidy. The optimal policy balances three goals: giving more surplus to consumers, mitigating underproduction, and mitigating overproduction.

JEL: D81, D82, D86

Keywords: monopoly regulation, non-Bayesian, regret, price cap, piece-rate subsidy

*Department of Economics, Northwestern University, yingni.guo@northwestern.edu.

†MEDS, Kellogg School of Management, Northwestern University, and Department of Economics, Stony Brook University, e-shmaya@kellogg.northwestern.edu.

1 Introduction

Regulating monopolies is challenging. A monopolistic firm has the market power to set its price above that 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 four or five 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 at the efficient level 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 technological capacity of the firm. How should the regulatory policy be designed when the regulator knows considerably less about the industry than the 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 non-Bayesian approach. The regulator's payoff is a weighted sum of consumers' surplus and the firm's profit. He can regulate the firm's price and/or quantity. He can give a subsidy to the firm or impose a tax on it. Given a policy, the firm chooses its price and quantity to maximize its profit. The regret of the regulator is, by definition, the difference between what he could have gotten if he had complete information about the industry and what he actually gets. We can think of regret as "money left on the table" due to the regulator's lack of information. The regulator evaluates a policy by his worst-case regret, i.e., the maximal regret he can incur across all possible demand and cost scenarios under this policy. The optimal policy minimizes worst-case regret.

The worst-case regret approach to uncertainty is our most significant difference from Baron and Myerson (1982) and the literature on monopoly regulation in general. Baron and Myerson (1982) take a Bayesian approach to uncertainty by assigning a prior to the regulator over the demand and cost scenarios and characterizing the policy that minimizes the expected regret. (Minimizing the expected regret is the same as maximizing the expected payoff, since the regulator's expected complete-information payoff is constant.) 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. (See Hayek (1945), Weitzman (1974) and Carroll (2019), for instance, for elaboration of these points.) In response, the regulator looks for a policy that works fairly well in all circumstances.

To illustrate our solution, we begin with two extreme cases of the regulator's payoff. If the regulator puts no weight on the firm's profit, so that his payoff is only consumers' surplus, then it is optimal to impose a price cap. A price cap bounds how much consumers' surplus the firm can extract. Consumers benefit from a lower price. However, a price cap might discourage a firm which should have produced from producing. Consumers lose in this case due to the firm's underproduction. The optimal level of the price cap balances consumers' gain from a lower price and their loss from the firm's underproduction.

If the regulator puts the same weight on the firm's profit as he does on consumers' surplus, so that his payoff is the total surplus of both consumers and the firm, then the regulator simply wants the firm to produce as efficiently as possible. Given that an unregulated monopolistic firm tends to supply less than the efficient level, the regulator wants to encourage more production by subsidizing the firm. However, a subsidy might incentivize production above the efficient level. The optimal design of a subsidy must balance the loss from underproduction and that from overproduction. The optimal policy has the following form: The regulator will have a target price and a subsidy cap. For each unit that the firm sells, he

subsidizes the firm for the difference between its price and the target price, subject to the constraint that the total subsidy not exceed the subsidy cap. This piece-rate subsidy up to the target price effectively lifts the firm's selling price, incentivizing the firm to serve more consumers than just those with high values. On the other hand, the cap on the firm's total subsidy makes sure that the regulator doesn't lose too much from potential overproduction.

For intermediate cases, the regulator puts some weight on the firm's profit, but less than the weight he puts on consumers' surplus. He must balance three goals simultaneously: giving more surplus to consumers, mitigating underproduction, and mitigating overproduction. It is optimal to combine the policies described above for the extreme cases, which leads to a regulatory policy with three distinctive features. First, the regulator imposes a price cap so the firm can't get more per unit than the price cap. As the regulator puts more weight on the firm's profit, the price cap increases. Second, the firm gets a piece-rate subsidy. Third, the regulator imposes a cap on the total subsidy that the firm will get.

The worst-case regret approach advances our knowledge of monopoly regulation. First, it highlights the tradeoff among the three goals of the regulator: giving more surplus to consumers, mitigating underproduction, and mitigating overproduction. Second, it highlights the roles of three policy instruments which are common in practice in achieving these three goals. The price cap bounds how much consumers' surplus the firm can extract. The piece-rate subsidy encourages the firm to serve more consumers than just those with high values. Hence, it deals with a monopolistic firm's intrinsic tendency to underproduce. The cap on the total subsidy makes sure that the potential overproduction induced by the subsidy is also under control. We show that, among all the policy instruments that the regulator can choose, these three policy instruments are sufficient.

We address how to incorporate the regulator's additional knowledge about the industry in Subsection 4.1. The worst-case regret approach remains tractable. In particular, we show that price-cap regulation is optimal when the range of consumers' values is small or the

weight the regulator puts on the firm's profit is small.

Related literature

This paper contributes to the literature on monopoly regulation. Caillaud et al. (1988), Braeutigam (1989), and Laffont and Tirole (1993) provide an overview of earlier contributions in this field. Armstrong and Sappington (2007) discuss recent developments. Our paper is closely related to Baron and Myerson (1982), Lewis and Sappington (1988*a,b*), and Armstrong (1999). The most significant difference is our approach to uncertainty. These papers take a Bayesian approach to uncertainty. Armstrong and Sappington (2007) emphasize two limitations of the Bayesian approach. First, since the relevant information asymmetries can be difficult to characterize precisely, it is not clear how to formulate a prior. Second, since multidimensional screening problems are difficult to solve, the form of optimal regulatory policies is generally not known. We take a worst-case regret approach. It does not require a prior and allows us to derive an optimal policy.

A second difference between our paper and the previous papers is the scope of uncertainty. Baron and Myerson (1982) and Lewis and Sappington (1988*a*) assume one-dimensional uncertainty about cost or demand scenarios, respectively. Lewis and Sappington (1988*b*) 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.

Our paper also contributes to the literature on mechanism design where the designer minimizes his worst-case regret. Hurwicz and Shapiro (1978) examine a moral hazard problem and show that a fifty-fifty split is the minimax-regret solution. Bergemann and Schlag (2008, 2011) examine monopoly pricing and argue that minimizing worst-case regret is more relevant than maximizing worst-case payoff, since the criterion of maximizing worst-case payoff suggests pricing to the lowest-value buyer. Caldentey, Liu and Lobel (2017) characterize the

dynamic-pricing rule that minimizes the seller's worst-case regret. Renou and Schlag (2011) apply the solution concept of ε -minimax regret to the problem of implementing social choice correspondences. Beviá and Corchón (2019) characterize contests in which contestants have dominant strategies; within this class they find the contest which minimizes the designer's worst-case regret. More broadly, we contribute to the growing literature of mechanism design with worst-case objectives. See for instance Chassang (2013), Carroll (2015), and Carroll (2019) for a survey. In terms of the monopoly regulation environment, a related paper in this literature is Garrett (2014) which considers the cost-based procurement problem (Laffont and Tirole (1986)).

Minimizing worst-case regret is a more relevant criterion than maximizing worst-case payoff in our setting for two reasons. First, regret in our setting has a natural interpretation: it is the weighted sum of distortion in production and the firm's profit. Second, the regulator's worst-case payoff is zero or less under any policy, since consumers' values might be too low relative to the cost. In this case, there is no surplus even under complete information. When there is no surplus, there is nothing the regulator can do. We argue that the regulator's goal should instead be to protect surplus in situations where there is some surplus to protect. The notion of regret catches this idea.

We also contribute to public policy design under the minimax-regret criterion. Manski (2006) examines the optimal way to search for evidence of crime. Manski (2011) reviews optimal treatment choice for a population in environments in which the policy outcome is only partially identified due to the unobservability of counterfactual policy outcomes.

The worst-case regret approach goes back at least to Savage (1954). Under this approach, when a decision maker has to choose some action while facing uncertainty, he chooses the action that minimizes his worst-case regret across all possible realizations of the uncertainty. Regret is defined as the difference between what the decision maker could achieve given the realization, and what he achieves under this action. In our case, the regulator is uncertain

about demand and cost scenarios and he has to choose a policy. Savage also puts forward an interpretation of the worst-case regret approach in the context of group decision-making, which is relevant for our policy design context. Consider a group of people who must jointly choose a policy. They have the same payoffs but different probability judgments. Under the policy that minimizes worst-case regret, no member of the group faces great regret, so no member will feel that the policy is a serious mistake. Seminal game-theory papers in which players minimize worst-case regret include Hannan (1957) and Hart and Mas-Colell (2000). Minimizing worst-case regret is also the leading approach in online learning, and in particular in multi-armed bandit problems (see Bubeck, Cesa-Bianchi et al. (2012) for a survey). It has also been used in designing treatment rules (e.g., Manski (2004) and Stoye (2009)) and in forecast aggregation (e.g., Areili, Babichenko and Smorodinsky (2017) and Babichenko and Garber (2018)).

Our work also contributes to the delegation literature (e.g., Holmström (1977, 1984)). Alonso and Matouschek (2008), Amador and Bagwell (2019), and Kolotilin and Zapechelnyuk (2019) 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 doesn't use transfers is optimal in a contracting environment in which both parties can make transfers to each other.¹

2 Environment

There is a monopolistic firm and a mass one of consumers. Let $V : [0, 1] \rightarrow [0, \bar{v}]$ be a decreasing upper-semicontinuous *inverse-demand function*. A quantity-price pair $(q, p) \in$

¹Armstrong 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.

$[0, 1] \times [0, \bar{v}]$ is *feasible* if and only if it is below the inverse-demand function, i.e., $p \leq V(q)$. The firm can choose any feasible quantity-price pair. The total value to consumers of quantity q is the area under the inverse-demand function, given by $\int_0^q V(z) dz$.

Let $C : [0, 1] \rightarrow \mathbf{R}_+$ with $C(0) = 0$ be an increasing lower-semicontinuous *cost function*. The *optimal total surplus* is given by:

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

If the firm produces q units, then the (*market*) *distortion* is given by:

$$\text{DSTR}(V, C, q) = \text{OPT}(V, C) - \left(\int_0^q V(z) dz - C(q) \right). \quad (2)$$

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

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

Regulatory policies

A *policy* is given by an upper-semicontinuous function $\rho : [0, 1] \times [0, \bar{v}] \rightarrow \mathbf{R}$. If the firm sells q units at price p , then it receives revenue $\rho(q, p)$. This revenue is the sum of the revenue qp from the marketplace, and any tax or subsidy, $\rho(q, p) - qp$, imposed by the regulator. We also assume that $\rho(0, 0) \geq 0$, so the firm is allowed to stay out of business without suffering a negative profit. This is the participation constraint.

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

1. 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 $q \geq \tilde{q}$.
2. The regulator can charge a proportional tax by setting $\rho(q, p) = (1 - \tau)qp$ for some $\tau \in (0, 1)$.
3. The regulator can require that the firm get no more than k per unit by imposing $\rho(q, p) = \min(qp, qk)$. If the firm prices above k , it pays a tax of $q(p - k)$ to the regulator. This policy effectively creates a price cap at k .
4. If the regulator decides not to intervene, then he chooses $\rho(q, p) = qp$, so the firm's revenue $\rho(q, p)$ equals its revenue from the marketplace.

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 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) and Guesnerie (1998)). Hence, it is without loss of generality to work with the revenue function $\rho(q, p)$ directly.

Fix a policy ρ , an inverse-demand function V , and a cost function C . If the firm produces q units at price p , then *consumers' surplus* and the *firm's profit* are given by:

$$\text{CS}(V, \rho, q, p) = \int_0^q V(z) \, dz - \rho(q, p), \text{ and } \text{FP}(V, C, \rho, q, p) = \rho(q, p) - C(q). \quad (3)$$

The definition of consumers' surplus incorporates the fact that any subsidy to the firm is paid by consumers through their taxes and that any tax imposed on the firm is passed on to consumers.

The firm's profit doesn't depend directly on V , but since V determines which quantity-price pairs are feasible, we include V as an argument in the firm's profit.

We say that (q, p) is a *firm's best response to (V, C) under the policy ρ* if it maximizes the firm's profit over all feasible (q, p) . The firm might have multiple best responses. The participation constraint implies that $\text{FP}(V, C, \rho, q, p) \geq 0$ for every best response (q, p) of the firm.

The regulator's payoff is a weighted sum, $\text{CS} + \alpha\text{FP}$, 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 an inverse-demand function V and a cost function C . We let $\text{CIP}(V, C)$ denote the regulator's complete-information payoff. This is what the regulator would achieve if he could tailor his policy for these inverse-demand and cost functions. Formally,

$$\text{CIP}(V, C) = \max_{\rho, q, p} (\text{CS}(V, \rho, q, p) + \alpha\text{FP}(V, C, \rho, q, p)), \quad (4)$$

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

Claim 1 below shows that the regulator's complete-information payoff equals the optimal total surplus. The regulator would ask the firm to produce the efficient quantity and to set a price equal to the marginal consumer's value at this efficient quantity. He would then give the firm a revenue equal to its cost. Although the regulator's payoff is a function of α , his complete-information payoff doesn't depend on α . This is because the optimal total surplus is generated and all of this surplus goes to consumers.

Claim 1. For any inverse-demand function V and cost function C ,

$$\text{CIP}(V, C) = \text{OPT}(V, C).$$

Proof. First, the regulator's complete-information payoff is at most $\text{OPT}(V, C)$. Indeed,

$$\text{CS}(V, \rho, q, p) + \alpha \text{FP}(V, C, \rho, q, p) \leq \text{CS}(V, \rho, q, p) + \text{FP}(V, C, \rho, q, p) \leq \text{OPT}(V, C),$$

for every policy ρ and every best response (q, p) of the firm to (V, C) under ρ . Here the first inequality follows from $\alpha \leq 1$ and the participation constraint $\text{FP}(V, C, \rho, q, p) \geq 0$, and the second inequality follows from the definitions of OPT , CS , FP in (1) and (3).

Second, let q^* denote a quantity that achieves the optimal total surplus. The regulator can achieve $\text{OPT}(V, C)$ by setting

$$\rho(q, p) = \begin{cases} C(q^*) & \text{if } (q, p) = (q^*, V(q^*)) \\ 0 & \text{otherwise.} \end{cases}$$

Choosing $(q, p) = (q^*, V(q^*))$ is a firm's best response to (V, C) under ρ . Since $\text{CS}(V, \rho, q, p) = \text{OPT}(V, C)$ and $\text{FP}(V, C, \rho, q, p) = 0$, it follows that $\text{CS}(V, \rho, q, p) + \alpha \text{FP}(V, C, \rho, q, p) = \text{OPT}(V, C)$. \square

Regret

When the regulator does not know (V, C) , a policy will usually not give the regulator his complete-information payoff. Given a policy ρ , an inverse-demand function V , and a cost function C , the regulator's *regret* is the difference between what he could have gotten under

complete information and what he actually gets:

$$\text{RGRT}(V, C, \rho, q, p) = \text{CIP}(V, C) - (\text{CS}(V, \rho, q, p) + \alpha \text{FP}(V, C, \rho, q, p)).$$

The following claim shows that the regret is a weighted sum of distortion and the firm's profit.

Claim 2. *For every pair (V, C) of inverse-demand and cost functions and for every policy ρ ,*

$$\text{RGRT}(V, C, \rho, q, p) = \text{DSTR}(V, C, q) + (1 - \alpha) \text{FP}(V, C, \rho, q, p).$$

Proof. Suppressing the dependence on V, C, ρ, q, p , we have

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

Here, the first equality is the definition of regret, the second is from Claim 1 that $\text{CIP} = \text{OPT}$, and the last is from the definition of distortion. \square

Thus, 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. $(1 - \alpha) \text{FP}$ represents the loss in his redistribution objective, since the regulator wants more surplus to go 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

We look for a policy that minimizes worst-case regret. Thus *the regulator's problem* is

$$\text{minimize}_{\rho} \max_{V, C, q, p} \text{RGRT}(V, C, \rho, q, p),$$

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

Formulating the regulator's problem as a minimax-regret problem is our main departure from the literature on monopoly regulation. If we assigned a Bayesian prior to the regulator over the demand and cost scenarios, minimizing the expected regret would be the same as maximizing the expected payoff as in Baron and Myerson (1982). Instead we consider environments where 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. We focus on deterministic policies. If the regulator can randomize but the adversary chooses the worst-case inverse-demand and cost functions after seeing the realized policy, then deterministic policies are without loss. \square

Remark 2. In the definition of consumers' surplus, we made the efficient-rationing assumption, so consumers with the highest values are served when the firm chooses not to clear the market. In Subsection 4.2, we argue that the firm will clear the market under our optimal policy if its average cost is decreasing in quantity q , so our result does not depend on this assumption. \square

Remark 3. In the definition of the regulator's complete-information payoff, we assumed that the firm breaks ties in favor of the regulator, whereas in the definition of the regulator's problem, we assumed 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 3.1 to 3.3.² \square

3 Main result

We begin with a lower bound on the worst-case regret of any policy. We then show that our policy indeed achieves this lower bound, so it is optimal. Both the lower-bound and the upper-bound discussions will center on the tradeoff between giving more surplus to consumers, mitigating underproduction, and mitigating overproduction.

3.1 Lower bound on worst-case regret

We first illustrate that there is a nontrivial tradeoff between giving more surplus to consumers and mitigating underproduction.

Suppose that the regulator constrains how much consumers’ surplus the firm can extract by imposing a price cap k . This price-cap policy has opposing implications for the two market scenarios in Figure 1.

On the left-hand side, every consumer has the highest value \bar{v} , and the firm’s 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 the regret is $(1 - \alpha)k$. The lower the price cap k , the lower this regret. On the right-hand side, every consumer still has the highest value \bar{v} , but now the firm has a fixed cost of k . It is a best response of the firm not to produce. The firm’s profit is zero, but the distortion is $\bar{v} - k$, which is the surplus that

²If in the definition of $CIP(V, C)$ we assumed that the firm breaks ties against the regulator, we would define

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

where the minimum ranges over all of the firm’s best responses (q, p) to (V, C) under ρ . Then the supremum may not be achieved, but the value of $CIP(V, C)$ 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 (V, C) may not exist, but the solution to the regulator’s problem would remain the same.

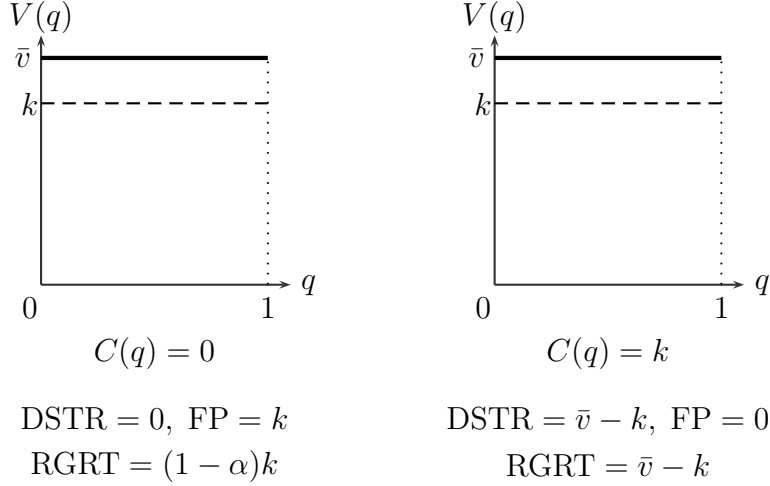


Figure 1: Impact of the price-cap policy on two scenarios

could have been generated. The regret equals this distortion. The lower the price cap k , the higher this regret.

The contrast between the two scenarios shows that a lower price cap may advance the regulator's redistribution objective on the one hand, but may worsen the problem of underproduction on the other. It also shows that the regulator's regret is at least $\max((1 - \alpha)k, \bar{v} - k)$. We let k_α be the price cap that minimizes $\max((1 - \alpha)k, \bar{v} - k)$, so

$$k_\alpha = \frac{\bar{v}}{2 - \alpha}$$

as depicted in the left panel of Figure 2. Not surprisingly, the more weight α the regulator puts on the firm's profit, the higher this price cap k_α is.

Our next claim follows directly from the discussion above. Even if the regulator had the knowledge that he were facing one of the two scenarios in Figure 1, he would have nonnegligible worst-case regret.

Claim 3. *The worst-case regret under any policy is at least $(1 - \alpha)k_\alpha = \frac{1 - \alpha}{2 - \alpha}\bar{v}$.*

Proof. Fix a policy ρ and let $k = \max_{q \in [0, 1], p \in [0, \bar{v}]} \rho(q, p)$ be the highest revenue the firm can




Figure 2: Values of k_α , r_α , q_α , and s_α

get under ρ . If the inverse-demand and cost functions are given by the left-hand side of Figure 1, the regret is at least $(1 - \alpha)k$ since the firm's profit is k . If the inverse-demand and cost functions are given by the right-hand side of Figure 1, it is a firm's best response not to produce and so the regret is $\bar{v} - k$. Hence, the worst-case regret under ρ is at least $\max((1 - \alpha)k, \bar{v} - k)$, which is weakly greater than $(1 - \alpha)k_\alpha$. \square

With this k_α balancing the tradeoff between giving more surplus to consumers and mitigating underproduction, we are ready to establish a tight lower bound on the worst-case regret.

Theorem 3.1 (Lower bound on worst-case regret). *Let*

$$r_\alpha = \max_{q \in [0, 1], p \in [0, k_\alpha]} \min(q(1 - \alpha)k_\alpha - qp \log q, q(k_\alpha - p)). \quad (5)$$

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

For any (q, p) , we argue that the worst-case regret is at least the minimum of two terms. Roughly speaking, the first term, $q(1 - \alpha)k_\alpha - qp \log q$, is the possible regret from underproduction if the revenue to the firm is too low. This occurs when the inverse-demand function is $U_{q,p}$ given in Figure 4. The second term, $q(k_\alpha - p)$, is the possible regret from overproduction if the revenue is too high. This occurs when the inverse-demand function is $W_{q,p}$ given

in Figure 4. No matter how the policy is designed, the regulator has to suffer from one of these two terms. Since the worst-case regret is at least the minimum of these two terms for every (q, p) , it is at least the maximum over any $q \in [0, 1]$ and $p \in [0, k_\alpha]$.

Let q_α achieve the maximum in the definition of r_α in (5). When $\alpha \leq 1/2$, q_α equals one. When $\alpha > 1/2$, q_α is interior. The explicit values of r_α and q_α are given by:

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

The middle and right panels of Figure 2 depict the values of r_α and q_α .

3.2 Optimal policy

Our next theorem shows that our policy guarantees that the worst-case regret is at most r_α , so it is optimal.

Theorem 3.2 (Optimal policy). *Let*

$$s_\alpha = \sup\{q(k_\alpha - p) : q \in [0, 1], p \in [0, k_\alpha], q(1 - \alpha)k_\alpha - qp \log q > r_\alpha\}.$$

The policy

$$\rho(q, p) = \min(qk_\alpha, qp + s) \tag{6}$$

with $s_\alpha \leq s \leq r_\alpha$ achieves the worst-case regret r_α .

We first provide intuition as to how a policy of the form (6) simultaneously addresses the three goals of giving more surplus to consumers, mitigating underproduction, and mitigating overproduction. First, the firm can't get more than k_α for each unit it sells. This caps how much consumers' surplus the firm can extract. Second, a monopolistic firm has the tendency

to serve just those consumers with very high values. In order to incentivize the firm to produce more, the regulator subsidizes the firm for the difference between its price and k_α . This piece-rate subsidy effectively increases the firm's selling price to k_α . Third, the firm's total subsidy is capped by s , so the potential overproduction induced by the subsidy is also under control.

Depending on the welfare weight α that the regulator puts on the firm's profit, the regulator puts different weights on these three goals, and hence varies k_α and s as α varies. The explicit value of s_α is given below, and is depicted as the dashed line in the middle panel of Figure 2. When $\alpha \leq 1/2$, the cap s on the total subsidy to the firm can take any value between s_α and r_α .

$$s_\alpha = \begin{cases} \bar{v} \frac{\alpha}{2-\alpha} & \text{if } \alpha \leq \frac{1}{2} \\ r_\alpha & \text{if } \alpha > \frac{1}{2}. \end{cases}$$

Note that $s_\alpha = 0$ when $\alpha = 0$. Hence, the policy $\rho(q, p) = q \min(\bar{v}/2, p)$ is optimal when $\alpha = 0$. The regulator doesn't subsidize the firm, and simply imposes a price cap at $\bar{v}/2$.

The optimal policy in Theorem 3.2 features three properties. The first property is that $\rho(q, p) \leq qk_\alpha$ for every q implies a price cap: the firm cannot get more than k_α per unit sold. The second property is that for some quantity-price pairs, the total subsidy to the firm is at least s_α . The third property is that the total subsidy to the firm is at most r_α . Not every optimal policy has the same form as in (6), but Theorem 3.3 asserts that every optimal policy has similar properties. Recall that q_α achieves the maximum in the definition of r_α in (5).

Theorem 3.3. *Let ρ be an optimal policy. Then*

1. (*Price cap*): $\rho(q, p) \leq qk_\alpha$ for every $q \leq q_\alpha$.
2. (*Subsidy*): There exists some (q, p) such that $\rho(q, p) \geq qp + s_\alpha$.

3. (*Subsidy cap*): $\rho(q, p) \leq qp + r_\alpha$ for every (q, p) .

In particular, since $q_\alpha = 1$ for $\alpha \leq 1/2$, it follows from Theorem 3.3 that for $\alpha \leq 1/2$ a price cap at k_α is necessary for every level of production.

4 Discussions

4.1 Incorporating additional knowledge

In our model we made no assumptions on the inverse-demand or cost functions except for monotonicity, semicontinuity, and the upper bound of consumers' values (which is \bar{v}). We view this minimally-informed regulator as a natural starting point. The regulator may know more than this. We can extend our framework in an obvious way to incorporate the regulator's knowledge by restricting the set of inverse-demand and cost functions in the regulator's problem.

Let \mathcal{E} be the set of possible inverse-demand and cost functions. The regulator chooses a policy that minimizes the worst-case regret across the elements of \mathcal{E} :

$$\underset{\rho}{\text{minimize}} \quad \max_{(V, C) \in \mathcal{E}, q, p} \text{RGRT}(V, C, \rho, q, p),$$

where the maximum ranges over all $(V, C) \in \mathcal{E}$ and all of the firm's best responses (q, p) to (V, C) under ρ . We illustrate how to solve the regulator's problem for some explicit \mathcal{E} , and demonstrate the adaptability of the worst-case regret approach.

4.1.1 Additional knowledge about cost

The regulator may know that the firm has a constant marginal cost together with a fixed cost, but doesn't know these cost levels. In this case, $\mathcal{E} = \{(V, C) : C(q) = aq + b \text{ for some } a, b \geq 0\}$. This is the type of cost function used most frequently in studies of monopoly regulation.

In our proof of Theorem 3.1, we establish a lower bound on the worst-case regret of any policy using only fixed-cost functions, i.e., $C(q) = b$ for some $b \geq 0$. (See Remark 4 for details.) This means that Theorem 3.1 remains true for every set of cost functions that includes the set of all fixed-cost functions. In particular, Theorem 3.1 remains true for $\mathcal{E} = \{(V, C) : C(q) = aq + b \text{ for some } a, b \geq 0\}$. Once we know that Theorem 3.1 remains true, we know that our policy in Theorem 3.2 remains optimal, since the worst-case regret under our policy is at most r_α across all inverse-demand and cost functions.

4.1.2 Additional knowledge about demand

The regulator may know not only an upper bound but also a lower bound on consumers' values. Let $\bar{v} > 0$ and $\underline{v} \in [0, \bar{v}]$ be the upper and lower bounds on consumers' values, so $\mathcal{E} = \{(V, C) : \underline{v} \leq V(q) \leq \bar{v}\}$. If $\underline{v} = 0$, we are back to our baseline model. If $\underline{v} = \bar{v}$, the regulator knows the inverse-demand function exactly since every consumer has value \bar{v} .

We show in Theorem 5.1 that, for any $\underline{v} \in [0, \bar{v}]$, a policy of the form $\rho(q, p) = \min(qk, qp + s)$ remains optimal. The price cap k is still k_α for any $\underline{v} \in [0, \bar{v}]$, but the cap s on the total subsidy becomes smaller as \underline{v} increases. Hence, as consumers become more homogeneous in their values, the regulator is less willing to subsidize the firm. In particular, we characterize the condition under which it is optimal to choose $s = 0$ so that a price-cap policy is optimal.

Proposition 4.1 (Price cap optimality). *If $\underline{v} \geq \frac{1}{2-\alpha}\bar{v}$, it is optimal to impose a price cap k_α .*

Both this proposition and Proposition 4.2 follow directly from Theorem 5.1. A price-cap policy is optimal when \underline{v} is sufficiently close to \bar{v} or when the welfare weight α on the firm's profit is sufficiently low. When the range of consumers' values is small, consumers with the highest values don't 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 piece-rate subsidy to encourage more production. Similarly, the redistribution objective is more salient when α is small than when α is large, so the regulator is less willing to subsidize when α is small.

It is interesting to derive the comparative statics with respect to \underline{v} . A natural conjecture is that the regulator's worst-case regret decreases as \underline{v} increases. We show that the answer depends on the welfare weight α that the regulator puts on the firm's profit.

Proposition 4.2. *For $\alpha \leq 1/2$, the worst-case regret is constantly $r_\alpha = \frac{1-\alpha}{2-\alpha}\bar{v}$ for any $\underline{v} \in [0, \bar{v}]$. In contrast, for $\alpha > 1/2$, the worst-case regret goes down from r_α to $\frac{1-\alpha}{2-\alpha}\bar{v}$ as \underline{v} goes from 0 to \bar{v} .*

When α is small, the main tradeoff is that between giving more surplus to consumers and mitigating underproduction. Even if all consumers have the highest value \bar{v} , the uncertainty about the cost would alone imply significant regret due to this tradeoff. When α is large, the regulator's main task is to promote efficiency. With a higher \underline{v} , consumers become homogenous in their values, which makes it easier to mitigate underproduction.

Figure 3 shows the worst-case regret as a function of $\underline{v} \in [0, \bar{v}]$ for $\alpha = 1/2, 3/4$, and 1.




Figure 3: Worst-case regret as a function of \underline{v}

4.2 The efficient-rationing assumption

In our model we allow the firm not to clear the market, and we assume that if this happens, then the consumers who are being served are the ones with higher values. Indeed, absent some additional assumptions on the cost function, even a firm which operates under a price cap may choose not to clear the market.

A common assumption in the monopoly regulation literature is that the firm has decreasing average cost, i.e., the average cost $C(q)/q$ is decreasing for $q > 0$. Since the set of all fixed-cost functions satisfies the decreasing average-cost assumption, in accordance with the discussion in Subsection 4.1.1, Theorem 3.1 remains correct under this decreasing average-cost assumption, and our policy in Theorem 3.2 is optimal. Moreover, if the cost function satisfies this assumption, then a firm which operates under our policy will want to clear the market.

5 Proofs

Fix $\bar{v} > 0$ and $\underline{v} \in [0, \bar{v}]$. Throughout this section, we assume that $(V, C) \in \{(V, C) : \underline{v} \leq V(q) \leq \bar{v}\}$. Theorems 3.1 and 3.2 follow from Theorem 5.1 when we assume $\underline{v} = 0$.

Theorem 5.1. *The worst-case regret under any policy is at least*

$$R_\alpha(\underline{v}) = \max \left((1 - \alpha)k_\alpha, \max_{qp \geq \underline{v}} \min (q(1 - \alpha)k_\alpha - qp \log q, q(k_\alpha - p)) \right).$$

Let

$$S_\alpha(\underline{v}) = (\sup\{q(k_\alpha - p) : q \in [0, 1], p \in [0, \bar{v}], qp \geq \underline{v}, q(1 - \alpha)k_\alpha - qp \log q > R_\alpha(\underline{v})\})^+.$$

The policy

$$\rho(q, p) = \min(qk_\alpha, qp + s)$$

with $S_\alpha(\underline{v}) \leq s \leq R_\alpha(\underline{v})$ achieves the worst-case regret $R_\alpha(\underline{v})$.

5.1 Preliminaries

For every q, p , we let $W_{q,p}$ and $U_{q,p}$ be the inverse-demand functions given by:

$$W_{q,p}(z) = \begin{cases} p & \text{if } z \leq q \\ 0 & \text{if } z > q, \end{cases} \quad \text{and} \quad U_{q,p}(z) = \begin{cases} \bar{v} & \text{if } z \leq q \\ \frac{qp}{z} & \text{if } z > q, \end{cases}$$

as shown in Figure 4.




Figure 4: $W_{q,p}$ and $U_{q,p}$ inverse-demand functions

The inverse-demand function $W_{q,p}$ has the property that, among all inverse-demand functions under which (q, p) is feasible, $W_{q,p}$ generates the least total value to consumers. That is, $W_{q,p}$ solves

$$\min_{V(\cdot)} \int_0^1 V(z) dz, \text{ subject to } (q, p) \text{ is feasible.}$$

The inverse-demand function $U_{q,p}$ exhibits unitary price elasticity when the price is in the range of $[qp, p]$.

To understand the role of $U_{q,p}$ in our argument, consider an unregulated firm (i.e., a firm

which operates under the policy $\rho(q, p) = qp$. If the inverse-demand function is $U_{q,p}$ and the cost is zero, then selling q units at price \bar{v} is a best response of the firm. This response causes a distortion of $\int_q^1 U_{q,p}(z) dz = -qp \log q$ due to underproduction. The following lemma shows that this is the worst distortion that can happen when the firm is unregulated.

Lemma 5.2. *Assume that an unregulated firm sells $\bar{q} < 1$ units at a price \bar{p} such that $\bar{p} \geq \sup_{z > \bar{q}} V(z)$. Let*

$$\text{OPT}_{\bar{q}} = \max_{q \geq \bar{q}} \left(\int_{\bar{q}}^q V(z) dz - (C(q) - C(\bar{q})) \right)$$

be the maximal additional surplus to society if the firm has produced \bar{q} units, and let

$$\text{FP}_{\bar{q}, \bar{p}} = \max_{q \geq \bar{q}} \left(q \min(\bar{p}, V(q)) - \bar{q}\bar{p} - (C(q) - C(\bar{q})) \right)$$

be the maximal additional profit to the firm if it has produced \bar{q} units and commits to price at most \bar{p} . Then

$$\text{OPT}_{\bar{q}} \leq \text{FP}_{\bar{q}, \bar{p}} + D(\bar{q}, \bar{p}),$$

where $D(q, p) = \max\{-q'\bar{p} \log q' : q' \geq \bar{q}, q'\bar{p} \geq \underline{v}\}$.

The lemma does not assume that selling \bar{q} or more units is optimal for the firm. Therefore, the assertion in the lemma still holds even if the best response for an unregulated firm is to sell fewer than \bar{q} units at a possibly higher price than \bar{p} .

Proof of Lemma 5.2. We can assume that $\text{FP}_{\bar{q}, \bar{p}} = 0$. If $\text{FP}_{\bar{q}, \bar{p}} > 0$, we can replace C with \bar{C} such that $\bar{C}(z) = C(z)$ if $z \leq \bar{q}$, and $\bar{C}(z) = C(z) + \text{FP}_{\bar{q}, \bar{p}}$ if $z > \bar{q}$. This replacement doesn't change $\text{OPT}_{\bar{q}} - \text{FP}_{\bar{q}, \bar{p}}$ but makes sure that the firm's maximal additional profit is zero.

Let $q^* \in \arg \max_{q \geq \bar{q}} \left(\int_{\bar{q}}^q V(z) dz - (C(q) - C(\bar{q})) \right)$.

If $q^* = \bar{q}$, then $\text{OPT}_{\bar{q}} = 0$ and the assertion in the lemma holds since $D(\bar{q}, \bar{p}) \geq 0$.

Therefore, we assume from now on that $q^* > \bar{q}$. Let $c^* = C(q^*) - C(\bar{q})$.

Since $V(z) \leq \bar{p}$ for $z > \bar{q}$, it follows from the definition of q^* that $c^* \leq (q^* - \bar{q})\bar{p}$. Let q'' solve $(q'' - \bar{q})\bar{p} = c^*$, so $q'' \leq q^*$.

Since the firm does not want to produce more, it follows that $zV(z) - C(z) \leq \bar{q}\bar{p} - C(\bar{q})$ for every $z > \bar{q}$, so that

$$V(z) \leq \frac{\bar{q}\bar{p} + C(z) - C(\bar{q})}{z} \leq \frac{\bar{q}\bar{p} + c^*}{z} = \frac{q''\bar{p}}{z}, \text{ for } \bar{q} < z \leq q^*. \quad (7)$$

Especially, for $z = q^*$, and since $V(q^*) \geq \underline{v}$, it follows that $q^* \leq q''\bar{p}/\underline{v}$. Then

$$\begin{aligned} \text{OPT}_{\bar{q}} &= \int_{\bar{q}}^{q^*} V(z) \, dz - c^* \leq (q'' - \bar{q})\bar{p} + \int_{q''}^{q^*} \frac{q''\bar{p}}{z} \, dz - c^* = q''\bar{p} \log \frac{q^*}{q''} \leq \\ &\quad - q''\bar{p} \log \max(q'', \underline{v}/\bar{p}) \leq - \max(q'', \underline{v}/\bar{p})\bar{p} \log \max(q'', \underline{v}/\bar{p}) \leq D(\bar{q}, \bar{p}), \end{aligned}$$

where the second step uses the fact that $V(z) \leq \bar{p}$ for $\bar{q} < z \leq q''$ and (7), the third step follows from $\bar{q}\bar{p} + c^* = q''\bar{p}$, and the fourth step follows from $q^* \leq 1$ and $q^* \leq q''\bar{p}/\underline{v}$. \square

For $\bar{q} = 0$, $\bar{p} = \bar{v}$, $\underline{v} = 0$, Lemma 5.2 has the following corollary which is interesting for its own sake. It bounds from below an unregulated firm's profit in a market with a high optimal total surplus. We are unaware of previous statements of this corollary, but similar arguments to those in the proof of Lemma 5.2 with zero cost have appeared in Roesler and Szentes (2017) and Condorelli and Szentes (2019a,b).

Corollary 5.3. *Suppose that $\underline{v} = 0$ and that the firm is unregulated. Then*

$$\text{FP} \geq \text{OPT} - \frac{\bar{v}}{e}.$$

5.2 Lower bounds on worst-case regret

For a policy ρ , let

$$\text{WCR}(\rho) = \max_{V, C, q, p} \text{RGRT},$$

where the maximum ranges over all $(V, C) \in \{(V, C) : \underline{v} \leq V(q) \leq \bar{v}\}$ and over all of the firm's best responses (q, p) to (V, C) under ρ .

For a policy ρ , let $\bar{\rho}(q) = \max_{q' \leq q} \rho(q', p')$ be the maximal revenue the firm can get under ρ from selling q or fewer units, and let $\hat{\rho}(q, p) = \max_{q' \geq q, q'p' \leq qp} \rho(q', p')$ be the maximal revenue under ρ if the firm sells at least q units and the revenue from the marketplace is at most qp . As shown in Figure 5, $\bar{\rho}(q)$ 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.




Figure 5: Definitions of $\bar{\rho}(q)$ and $\hat{\rho}(q, p)$

Claim 4 shows that the worst-case regret under a policy is significant if this policy offers a significant subsidy.

Claim 4. *Fix a policy ρ . Then $\text{WCR}(\rho) \geq \rho(q, p) - q \max(p, \underline{v})$ for every q, p .*

Proof. If $\rho(q, p) \leq q \max(p, \underline{v})$, the assertion follows from the fact that regret is nonnegative. Assume that $\rho(q, p) \geq q \max(p, \underline{v})$ and consider the inverse-demand function $W_{q, \max(p, \underline{v})}$ and a fixed cost $\rho(q, p)$. The firm will produce and incur the fixed cost, while the total value to

consumers is at most $q \max(p, \underline{v})$. The regret is at least:

$$\text{RGRT} \geq \text{DSTR} \geq \rho(q, p) - q \max(p, \underline{v}),$$

because of overproduction. □

Claim 5 shows that, if the firm doesn't receive enough additional revenue from producing more, there is sizable regret due to underproduction.

Claim 5. *Fix a policy ρ . Let $\underline{q} \leq q \in [0, 1]$ and let $p \in [0, \bar{v}]$ such that $qp \geq \underline{v}$. If $\hat{\rho}(q, p) \leq \bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha$, then*

$$\text{WCR}(\rho) \geq (1 - \alpha)(\bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha) - qp \log q.$$

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

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

because of underproduction. Therefore,

$$\text{RGRT} = (1 - \alpha)\text{FP} + \text{DSTR} \geq (1 - \alpha)(\bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha) - qp \log q.$$

2. If $\bar{\rho}(q) - \bar{\rho}(\underline{q}) \geq (q - \underline{q}) k_\alpha$, then consider the inverse-demand function $U_{q,p}$ and zero cost. The firm will produce at most q units, with FP = $\bar{\rho}(q) \geq \bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha$, and

DSTR $\geq -qp \log q$ because of underproduction. Therefore,

$$\text{RGRT} = (1 - \alpha)\text{FP} + \text{DSTR} \geq (1 - \alpha) (\bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha) - qp \log q.$$

□

Combining Claims 4 and 5, we show that the regulator suffers sizable regret from either underproduction or overproduction.

Claim 6. Fix a policy ρ . Let $\underline{q} \leq q \in [0, 1]$ and let $p \in [0, \bar{v}]$ such that $qp \geq \underline{v}$. Then

$$\text{WCR}(\rho) \geq \min((1 - \alpha)(\bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha) - qp \log q, \bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha - qp).$$

Proof. If $\hat{\rho}(q, p) \geq \bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha$, then let q', p' be such that $q' \geq q, q'p' \leq qp$ and $\rho(q', p') = \hat{\rho}(q, p)$. Since $q'\underline{v} \leq \underline{v} \leq qp$ and $q'p' \leq qp$, we have $q' \max(p', \underline{v}) \leq qp$. By Claim 4

$$\text{WCR}(\rho) \geq \rho(q', p') - q' \max(p', \underline{v}) \geq \bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha - qp.$$

If $\hat{\rho}(q, p) < \bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha$, then $\text{WCR}(\rho) \geq (1 - \alpha)(\bar{\rho}(\underline{q}) + (q - \underline{q}) k_\alpha) - qp \log q$ by Claim 5. □

5.3 Upper bounds on worst-case regret

We consider a policy of this form: for some $k \in [0, \bar{v}]$ and $s \geq 0$:

$$\rho(q, p) = \min(qk, qp + s). \tag{8}$$

We bound the regret from (8) separately for the case of overproduction and the case of underproduction.

Claim 7. *The regret from overproduction under (8) is at most*

$$\max((1 - \alpha)k, s).$$

Proof. Let q^* be an efficient quantity, let $p^* = V(q^*)$, and assume that the firm chooses (q, p) with $q \geq q^*$ and $p \leq V(q) \leq p^*$. Let $\bar{c} = C(q) - C(q^*)$. Then

$$\text{DSTR} = \bar{c} - \int_{q^*}^q V(z) \, dz \leq \bar{c} - (q - q^*)p, \quad (9)$$

and

$$\bar{c} \leq \rho(q, p) - \rho(q^*, p^*) \quad (10)$$

since (q, p) is a best response. Therefore,

$$\begin{aligned} \text{RGRT} &= (1 - \alpha)\text{FP} + \text{DSTR} \leq (1 - \alpha)(\rho(q, p) - \bar{c}) + \bar{c} - (q - q^*)p \leq \\ &(1 - \alpha)\rho(q^*, p^*) + \rho(q, p) - \rho(q^*, p^*) - (q - q^*)p \leq (1 - \alpha)\rho(q^*, p^*) + \rho(q, p) - \rho(q^*, p) - (q - q^*)p \\ &\leq (1 - \alpha)\rho(q^*, p^*) + (q - q^*)(\rho(q, p)/q - p) \leq (1 - \alpha)q^*k + (1 - q^*/q)s \leq \\ &(1 - \alpha)q^*k + (1 - q^*)s \leq \max((1 - \alpha)k, s). \end{aligned}$$

where the first inequality follows from the definition of FP in (3), the fact that $\bar{c} \leq C(q)$, and (9); the second inequality follows from (10); the third inequality follows from the fact that $p \leq p^*$ and the fact that $p \mapsto \rho(q, p)$ is monotone increasing; the fourth inequality, $\rho(q, p) - \rho(q^*, p) \leq (q - q^*)\rho(q, p)/q$, follows from the fact that $q^* \leq q$ and the fact that $q \mapsto \rho(q, p)/q$ is decreasing; the fifth inequality follows from $\rho(q^*, p^*) \leq q^*k$, $\rho(q, p) \leq qp + s$; and the sixth inequality follows from $q \leq 1$.

□

Claim 8. *The regret from underproduction under (8) is at most*

$$\max \{q \max((1 - \alpha)k, \bar{v} - k) - (qk - s) \log q : qk - s \geq \underline{v} \text{ or } q = 1\}.$$

Proof. Let q^* be an efficient quantity and assume that the firm chooses (q, p) with $q \leq q^*$.

If $q^*(k - V(q^*)) \leq s$, then $\rho(q^*, V(q^*)) = q^*k$ and $\rho(q, p) = qk$. Therefore, since the firm prefers to produce q over q^* , it follows that $C(q^*) - C(q) \geq (q^* - q)k$, which implies that $\text{DSTR} \leq (q^* - q)(\bar{v} - k)$ and

$$\text{RGRT} \leq (1 - \alpha)\rho(q, p) + \text{DSTR} \leq (1 - \alpha)qk + (q^* - q)(\bar{v} - k) \leq \max((1 - \alpha)k, \bar{v} - k).$$

If $q^*(k - V(q^*)) > s$, then let $\bar{q} \in [q, q^*)$ be such that $z(k - V(z)) \leq s$ for $q < z < \bar{q}$, and $z(k - V(z)) > s$ for $z > \bar{q}$. (\bar{q} is the point at which the subsidy is used up, except that if it is already used up before q , then $\bar{q} = q$). Let $\bar{p} = k - s/\bar{q}$. Then it follows from the definition of \bar{q} that $\bar{p} \geq \sup_{z > \bar{q}} V(z)$.

By Lemma 5.2 there exists some $z^* \in [\bar{q}, q^*]$ such that

$$\int_{\bar{q}}^{q^*} V(z) dz - (C(q^*) - C(\bar{q})) \leq z^* \underline{p} - \bar{q}\bar{p} - (C(z^*) - C(\bar{q})) + D(\bar{q}, \bar{p}),$$

with $\underline{p} = \min(\bar{p}, V(z^*))$. Since $z^* \geq \bar{q}$ and $\underline{p} \leq \bar{p}$, it follows from the definition of ρ that

$$\rho(z^*, \underline{p}) - z^* \underline{p} \geq \rho(\bar{q}, \bar{p}) - \bar{q}\bar{p} = \bar{q}(k - \bar{p}).$$

Since the firm prefers to produce q over z^* , it follows that

$$\rho(z^*, \underline{p}) \leq \rho(q, p) + (C(z^*) - C(q)).$$

The last three inequalities and $C(\bar{q}) \leq C(q^*)$ imply that

$$\int_{\bar{q}}^{q^*} V(z) dz \leq C(q^*) - C(q) + \rho(q, p) - \bar{q}k + D(\bar{q}, \bar{p}). \quad (11)$$

Therefore,

$$\begin{aligned} \text{DSTR} &= \int_q^{q^*} V(z) dz - (C(q^*) - C(q)) = \int_q^{\bar{q}} V(z) dz + \int_{\bar{q}}^{q^*} V(z) dz - (C(q^*) - C(q)) \\ &\leq (\bar{q} - q)\bar{v} - \bar{q}k + \rho(q, p) + D(\bar{q}, \bar{p}) \leq (\bar{q} - q)(\bar{v} - k) + D(\bar{q}, \bar{p}), \end{aligned}$$

where the first inequality follows from (11) and $V(z) \leq \bar{v}$, and the second from $\rho(q, p) \leq qk$.

It follows that

$$\begin{aligned} \text{RGRT} &\leq (1 - \alpha)\rho(q, p) + \text{DSTR} \leq (1 - \alpha)qk + (\bar{q} - q)(\bar{v} - k) + D(\bar{q}, \bar{p}) \leq \\ &\bar{q} \max((1 - \alpha)k, \bar{v} - k) - q'\bar{p} \log q' \leq q' \max((1 - \alpha)k, \bar{v} - k) - (q'k - s) \log q', \end{aligned}$$

for some $\bar{q} \leq q' \leq 1$ such that $q'\bar{p} \geq \underline{v}$. Here the last inequality follows from the fact that $\bar{p} = k - s/\bar{q} \leq k - s/q'$. \square

5.4 Proof of Theorem 5.1

We need to show that $\text{WCR}(\rho) \geq \min((1 - \alpha)qk_\alpha - qp \log q, q(k_\alpha - p))$ for every q, p such that $qp \geq \underline{v}$. This follows from Claim 6 with $\underline{q} = 0$. We need to show that $\text{WCR}(\rho) \geq (1 - \alpha)k_\alpha$. This follows from Claim 3.

Consider the policy (8) with $k = k_\alpha$ and $S_\alpha(\underline{v}) \leq s \leq R_\alpha(\underline{v})$. Since $s \leq R_\alpha(\underline{v})$ and $(1 - \alpha)k_\alpha \leq R_\alpha(\underline{v})$, it follows from Claim 7 that the regret from overproduction is at most $R_\alpha(\underline{v})$.

For the regret from underproduction, by Claim 8 we need to show that $q(1 - \alpha)k_\alpha -$

$(qk_\alpha - s) \log q \leq R_\alpha(\underline{v})$ for $q = 1$ and for every $q \in [0, 1]$ such that $qk_\alpha - s \geq \underline{v}$. For $q = 1$, this follows from the fact that $(1 - \alpha)k_\alpha \leq R_\alpha(\underline{v})$. Let $q < 1$ and let $p = k_\alpha - s/q$. Then $qp \geq \underline{v}$. Let $q' > q$. Then $q'(k_\alpha - p) > s$ and $q'p \geq \underline{v}$, and by the assumption on s , this implies that $(1 - \alpha)q'k_\alpha - q'p \log q' \leq R_\alpha(\underline{v})$. Since this is true for every $q' > q$, it follows by continuity that $(1 - \alpha)qk_\alpha - qp \log q \leq R_\alpha(\underline{v})$, as desired.

Remark 4. The proof of Claim 6 for the case of $\underline{q} = 0$ and the proof of Claim 3 rely only on fixed-cost functions, as do the proof of the lower bound on the worst-case regret in Theorem 5.1 and the proof of Theorem 3.1. \square

5.5 Proof of Theorem 3.3

For this theorem, we assume that $\underline{v} = 0$. Let (q_α, p_α) achieve the maximum in the definition of r_α in (5).

1. Assume that $\rho(\underline{q}, p) > \underline{q}k_\alpha$ for some $\underline{q} \leq q_\alpha$ and some p . Then $\bar{\rho}(\underline{q}) > \underline{q}k_\alpha$, and therefore $\bar{\rho}(\underline{q}) + (q_\alpha - \underline{q})k_\alpha > q_\alpha k_\alpha$. Therefore, by Claim 6 with $q = q_\alpha$ and $p = p_\alpha$, it follows that

$$\begin{aligned} \text{WCR}(\rho) &\geq \min((1 - \alpha)(\bar{\rho}(\underline{q}) + (q_\alpha - \underline{q})k_\alpha) - q_\alpha p_\alpha \log q_\alpha, \bar{\rho}(\underline{q}) + (q_\alpha - \underline{q})k_\alpha - q_\alpha p_\alpha) \\ &> \min((1 - \alpha)q_\alpha k_\alpha - q_\alpha p_\alpha \log q_\alpha, q_\alpha(k_\alpha - p_\alpha)) = r_\alpha. \end{aligned}$$

2. Suppose that $\rho(q, p) < qp + s_\alpha$ for every q, p . This implies that $\max_{q', p'}(\rho(q', p') - p'q') < s_\alpha$. There exists some $q \in [0, 1], p \in [0, k_\alpha]$ such that $(1 - \alpha)qk_\alpha - qp \log q > r_\alpha$ and $q(k_\alpha - p) > \max_{q', p'}(\rho(q', p') - p'q') \geq \hat{\rho}(q, p) - qp$, which implies that $\hat{\rho}(q, p) < qk_\alpha$. By Claim 5 with $\underline{q} = 0$, we get that $\text{WCR}(\rho) > r_\alpha$.
3. Suppose that $\rho(q, p) > qp + r_\alpha$ for some q, p . Then $\text{WCR}(\rho) > r_\alpha$ by Claim 4.

References

- Alonso, Ricardo, and Niko Matouschek.** 2008. “Optimal Delegation.” *The Review of Economic Studies*, 75(1): 259–293.
- Amador, Manuel, and Kyle Bagwell.** 2019. “Regulating a Monopolist With Uncertain Costs Without Transfers.” *Working paper*.
- Areili, Itai, Yakov Babichenko, and Rann Smorodinsky.** 2017. “Robust forecast aggregation.” *arXiv preprint arXiv:1710.02838*.
- Armstrong, Mark.** 1999. “Optimal Regulation with Unknown Demand and Cost Functions.” *Journal of Economic Theory*, 84(2): 196 – 215.
- Armstrong, Mark, and David EM Sappington.** 2007. “Recent developments in the theory of regulation.” *Handbook of industrial organization*, 3: 1557–1700.
- Armstrong, Mark, and John Vickers.** 2010. “A Model of Delegated Project Choice.” *Econometrica*, 78(1): 213–244.
- Babichenko, Yakov, and Dan Garber.** 2018. “Learning of optimal forecast aggregation in partial evidence environments.” *arXiv preprint arXiv:1802.07107*.
- Baron, David P., and Roger B. Myerson.** 1982. “Regulating a Monopolist with Unknown Costs.” *Econometrica*, 50(4): 911–930.
- Bergemann, Dirk, and Karl H. Schlag.** 2008. “Pricing without Priors.” *Journal of the European Economic Association*, 6(2/3): 560–569.
- Bergemann, Dirk, and Karl Schlag.** 2011. “Robust monopoly pricing.” *Journal of Economic Theory*, 146(6): 2527–2543.

- Beviá, Carmen, and Luis Corchón.** 2019. “Contests with dominant strategies.” *Economic Theory*.
- Braeutigam, Ronald R.** 1989. “Optimal policies for natural monopolies.” *Handbook of industrial organization*, 2: 1289–1346.
- Bubeck, Sébastien, Nicolo Cesa-Bianchi, et al.** 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.” *The 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): 104–129.
- 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): 139–166.
- Chassang, Sylvain.** 2013. “Calibrated incentive contracts.” *Econometrica*, 81(5): 1935–1971.
- Condorelli, Daniele, and Balázs Szentes.** 2019*a*. “Information Design in the Hold-up Problem.” *Journal of Political Economy*, *forthcoming*.
- Condorelli, Daniele, and Balázs Szentes.** 2019*b*. “Surplus Bounds in Cournot Competition.” *Working paper*.

- 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.” *The Quarterly Journal of Economics*, 134(1): 51–107.
- Garrett, Daniel F.** 2014. “Robustness of simple menus of contracts in cost-based procurement.” *Games and Economic Behavior*, 87: 631–641.
- Guesnerie, Roger.** 1998. *A contribution to the pure theory of taxation*. Cambridge University Press.
- Hannan, James.** 1957. “Approximation to Bayes risk in repeated play.” *Contributions to the Theory of Games*, 3: 97–139.
- Hart, Sergiu, and Andreu Mas-Colell.** 2000. “A Simple Adaptive Procedure Leading to Correlated Equilibrium.” *Econometrica*, 68(5): 1127–1150.
- Hayek, F. A.** 1945. “The Use of Knowledge in Society.” *The American Economic Review*, 35(4): 519–530.
- Holmström, Bengt.** 1977. “On Incentives and Control in Organizations.” *PhD diss., Stanford University*.
- Holmström, Bengt.** 1984. “On the Theory of Delegation.” *in: Bayesian Models in Economic Theory*. Ed. by M. Boyer, and R. Kihlstrom. North-Holland, Amsterdam, 115–141.
- Hurwicz, Leonid, and Leonard Shapiro.** 1978. “Incentive Structures Maximizing Residual Gain under Incomplete Information.” *The Bell Journal of Economics*, 9(1): 180–191.
- Kolotilin, Anton, and Andriy Zapechelnuk.** 2019. “Persuasion meets delegation.” *arXiv preprint arXiv:1902.02628*.

- Laffont, Jean-Jacques, and Jean Tirole.** 1986. “Using Cost Observation to Regulate Firms.” *Journal of Political Economy*, 94(3): 614–641.
- Laffont, Jean-Jacques, and Jean Tirole.** 1993. *A theory of incentives in procurement and regulation*. MIT press.
- Lewis, Tracy R., and David E. M. Sappington.** 1988a. “Regulating a Monopolist with Unknown Demand.” *The American Economic Review*, 78(5): 986–998.
- Lewis, Tracy R., and David E.M. Sappington.** 1988b. “Regulating a Monopolist with Unknown Demand and Cost Functions.” *The RAND Journal of Economics*, 19(3): 438–457.
- Manski, Charles F.** 2004. “Statistical Treatment Rules for Heterogeneous Populations.” *Econometrica*, 72(4): 1221–1246.
- Manski, Charles F.** 2006. “Search profiling with partial knowledge of deterrence.” *The Economic Journal*, 116(515): F385–F401.
- Manski, Charles F.** 2011. “Choosing Treatment Policies Under Ambiguity.” *Annual Review of Economics*, 3(1): 25–49.
- Renou, Ludovic, and Karl H. Schlag.** 2011. “Implementation in minimax regret equilibrium.” *Games and Economic Behavior*, 71(2): 527 – 533.
- 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.
- Savage, Leonard J.** 1954. *The Foundations of Statistics*. Wiley Publications in Statistics.

Stoye, Jörg. 2009. “Minimax regret treatment choice with finite samples.” *Journal of Econometrics*, 151(1): 70 – 81.

Weitzman, Martin L. 1974. “Prices vs. Quantities.” *The Review of Economic Studies*, October, 41(4): 477–491.