

Robust Monopoly Regulation

Yingni Guo* Eran Shmaya†

October 7, 2019

Abstract

We study the regulation of a monopolistic firm using a robust-design approach. We solve for the policy that minimizes the regulator’s worst-case regret, where the regret is the difference between his complete-information payoff and his realized payoff. When the regulator’s payoff is consumers’ surplus, it is optimal to impose a price cap. The optimal cap balances the benefit from more surplus for consumers and the loss from underproduction. When his payoff is consumers’ surplus plus the firm’s profit, he offers a piece-rate subsidy in order to mitigate underproduction, but caps the total subsidy so as not to incentivize severe overproduction.

JEL: D81, D82, D86

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

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)

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

†MEDS, Kellogg School of Management, Northwestern University, e-shmaya@kellogg.northwestern.edu.

show that prices at monopoly hospitals are 12% higher than those in markets with four or five rivals. In order to protect consumer well-being, 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 may end up not producing. The regulator must balance the need to protect consumer well-being and the need to not distort the 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 socially optimal level and to set 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 the consumer demand or the technological capacity of the firm. How shall the regulatory policy be designed when the regulator knows considerably less about the industry than the firm does? If the regulator demands robustness and wants a policy that works "fairly well" in all circumstances, what shall this policy look like?

We address this classic problem of monopoly regulation (e.g., Baron and Myerson (1982)) with 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 charge a tax from it. Given a policy, the firm chooses its price and quantity to maximize its profit. The *regret* to the regulator is the difference between what he could have gotten if he had complete information about the industry and what he actually gets. The regulator evaluates a policy by its worst-case regret, i.e., the maximal regret he can incur across all possible demand and cost scenarios. The optimal policy minimizes the 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 where 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 robust 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 his payoff is consumers' surplus, then it is optimal to impose a price cap. A price cap bounds how much consumers' surplus that 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 his payoff is the total surplus of 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, subsidy might incentivize production above the efficient level. The optimal design of subsidy must balance the loss from underproduction and that from overproduction.

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 doesn't exceed the subsidy cap. This piece-rate subsidy up to the target price effectively lifts the firm's selling price, motivating the firm to serve more than just those consumers 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 the potential overproduction.

For intermediate payoffs, the regulator puts some weight on the firm's profit, but this weight is lower 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 price cap and the subsidy rule described above, leading to a regulatory policy with three distinctive features. First, the regulator will impose a price cap so the firm can't get more than the price cap per unit. Second, the firm gets a piece-rate subsidy instead of a lump-sum one. Third, the firm is subsidized up to the price cap subject to a cap on the total subsidy it will get. As the regulator puts more weight on the firm's profit, the level of this price cap increases.

Our contribution is threefold. First, we solve for an optimal regulatory policy. Second, our result explains why price cap regulation and piece-rate subsidy are common in practice. Third, we introduce the worst-case regret approach to the regulation problem. We advocate for this approach over the Bayesian approach for the two shortcomings of the Bayesian approach that are also emphasized in Armstrong and Sappington (2007). First, since the relevant information asymmetries can be difficult to characterize precisely, it is not clear how to formulate a prior. Second, since multi-dimensional screening problems are difficult to solve, the form of optimal regulatory policies is generally not known.

Related literature. This paper contributes to the literature on monopoly regulation. Caillaud et al. (1988) and Braeutigam (1989) provide an overview of earlier contributions in this field. Armstrong and Sappington (2007) discuss the recent developments. Our paper is closely related to Baron and Myerson (1982). The most significant difference is our approach to uncertainty. Baron and Myerson (1982) take a Bayesian approach to uncertainty, and assume that there is a one-dimensional cost parameter that is unknown to the regulator.

We take a non-Bayesian, worst-case regret approach, and assume that the regulator lacks information about both the demand and the cost functions.

Our paper contributes to the growing literature of mechanism design with worst-case objectives. Carroll (2019) provides a survey of recent theory in this field. Most of this literature assumes that the designer aims to maximize his worst-case payoff. We assume that the designer aims to minimize his worst-case regret. From this aspect, we are closely related to Hurwicz and Shapiro (1978), Bergemann and Schlag (2008, 2011), Renou and Schlag (2011), and Beviá and Corchón (2019). Hurwicz and Shapiro (1978) show that a 50-50 split is an optimal sharecropping contract when the optimality criterion involves the ratio of the designer's payoff to the first-best total surplus. Bergemann and Schlag (2008, 2011) examine robust monopoly pricing and argue that minimizing the worst-case regret is more relevant than maximizing the worst-case payoff, since the latter criterion suggests pricing to the lowest-value buyer. 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 and find within this class the contest for which the designer's worst-case regret is minimized.

Minimizing the worst-case regret is more relevant a criterion in our setting as well for two reasons. First, the 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, instead, the regulator's goal should be to protect surplus in situations where there is some surplus to protect. The notion of regret catches this idea.

The worst-case regret approach goes back at least to Savage (1954). Under this approach, when a decision maker has to choose some action facing uncertainty, he chooses the action

that minimizes the worst-case regret across all possible realizations of the uncertainty. The 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 has uncertainty about the demand and cost functions 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, 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 judgements. Under the policy that minimizes the worst-case regret no member of the group faces a large regret, so no member will feel that the suggestion is a serious mistake. Seminal game theory papers in which players try to 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).

Our work also contributes to the delegation literature (e.g., Holmström (1977); Holmström (1984)). When the regulator cares only about consumers' surplus, it is optimal to simply impose a price cap. To our knowledge, we are the first to show that a delegation contract — a contract that doesn't use money — is optimal in a contracting environment where both parties can transfer money to each other.

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 $(q, p) \in [0, 1] \times [0, \bar{v}]$ is *feasible* if and only if it is below the inverse-demand function, i.e., $p \leq V(q)$. The 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 *social optimum* is given by:

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

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

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

To simplify notation, we omitted the dependence of OPT on V, C and the dependence of DSTR on V, C and q . We will do the same for some other terms in this section when no confusion arises.

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 $\rho(q, p)$. The firm's revenue from the policy, $\rho(q, p)$, includes the revenue qp from the marketplace, and any tax or subsidy, $\rho(q, p) - qp$, imposed by the regulator. We give three examples of regulatory policies:

1. A regulator who decides not to intervene will choose $\rho(q, p) = qp$, so the firm's revenue $\rho(q, p)$ equals its revenue from the marketplace.
2. The regulator can give the firm a lump-sum subsidy s 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 $q \geq \tilde{q}$.
3. The regulator can require that the firm get no more than k per unit. The policy is $\rho(q, p) = \min(qp, qk)$. If the firm prices above k , it pays a tax of $q(p - k)$ to the regulator.

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:

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

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 from the firm goes to the consumers. 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.

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 $\text{FP} \geq 0$ 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 for some fixed parameter $\alpha \in [0, 1]$.

Complete information

Fix an inverse-demand function V and a cost function C . We let CIP denote the regulator's complete-information payoff, which is what the regulator would achieve if he could tailor the policy for these inverse-demand and cost functions. Formally,

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

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

The following claim shows that the regulator's complete-information payoff is the social optimum. He would ask the firm to produce the socially optimal quantity and give the firm a revenue equal to its cost. As a result, the maximum surplus is generated, all of which goes

to consumers.

Claim 1. *Fix an inverse-demand function V and a cost function C . Then $\text{CIP} = \text{OPT}$.*

Proof. First, the regulator's complete-information payoff is at most OPT . Indeed,

$$\text{CS} + \alpha\text{FP} \leq \text{CS} + \text{FP} \leq \text{OPT},$$

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

Second, let q^* denote a quantity that achieves the social optimum. The regulator can achieve OPT 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 which gives $\text{CS} = \text{OPT}$ and $\text{FP} = 0$, so $\text{CS} + \alpha\text{FP} = \text{OPT}$. \square

Regret

When the regulator does not know (V, C) , the policy will usually not give the regulator his complete-information payoff. The regulator's *regret* is the difference between what he could have gotten under complete information and what he actually gets. The following step allows

us to express regret in terms of distortion and the firm's profit:

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

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

Regret has a natural interpretation in our setting. DSTR represents the loss in efficiency, since the regulator wishes the firm to produce as efficiently as possible. $(1 - \alpha)\text{FP}$ represents the loss in his redistribution objective, since the regulator wishes that more surplus goes to consumers rather than to the firm.

The regulator's problem

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

$$\underset{\rho}{\text{minimize}} \quad \max_{V, C, q, p} \text{RGRT}$$

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

Formulating the regulator's problem as a minimax problem is our only 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). We instead consider environments where the regulator knows only the range of consumers' values, which is much easier to figure out than formulating a prior.

Remark 1. In the definition of CIP 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 CIP in Claim 1 or the solution to the regulator’s problem in Theorems 3.1 to 3.3.¹

3 Main result

We first provide 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 upper-bound discussions center on the tradeoff between giving more surplus to consumers, mitigating underproduction, and mitigating overproduction.

Suppose that the regulator imposes a price cap k . A price cap advances the regulator’s redistribution objective by bounding how much consumers’ surplus the firm can extract, but it may worsen the problem of underproduction. There is a price cap level that balances these two forces. Explicitly, consider a market in which every consumer has the highest value \bar{v} . If the cost is zero, the firm will price at k and serve all consumers. There is no distortion since all consumers are served, as it should be, but the firm’s profit is k . The regret is $(1 - \alpha)k$. The lower k is, the lower the regret is. On the other hand, if 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 the distortion is $\bar{v} - k$, which is the surplus that could have been made. The regret equals this distortion. The lower k is, the higher the regret is. We let k_α be the price cap such that these two levels

¹If in the definition of CIP we assumed that the firm breaks ties against the regulator, we would define

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

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

of regret are equalized, so $k_\alpha = \bar{v}/(2 - \alpha)$ as depicted in the left panel of Figure 1.

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

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

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

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, $(1 - \alpha)qk_\alpha - qp \log q$, is the possible regret from underproduction if the revenue to the firm is too low. The second term, $q(k_\alpha - p)$, is the possible regret from overproduction if the revenue is too high. No matter how the policy is designed, the regulator has to suffer from one of these two. Since the worst-case regret is at least the minimum of these two for every (q, p) , we can take the maximum over $q \in [0, 1]$ 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:

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

The middle and right panels of Figure 1 depict the values of r_α and q_α .

Theorem 3.2 (Optimal policy). *Let*

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

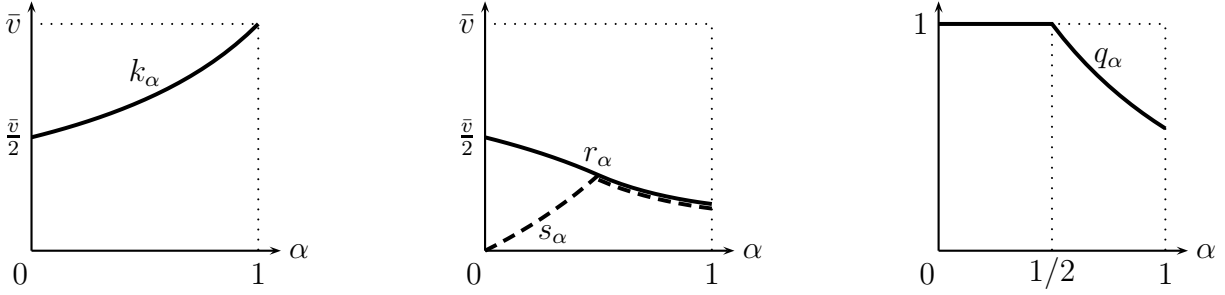


Figure 1: Values of k_α , r_α , q_α , and s_α

The policy

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

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

We first provide some intuition as to how a policy of the form (6) addresses the three goals of giving more surplus to consumers, mitigating underproduction, and mitigating overproduction simultaneously. First, the firm can't get more than k_α for each unit it sells. This caps how much consumers' surplus that 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 subsidy is also under control.

Depending on how much the regulator cares about the firm's profit, he 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 1.

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

Note that $s_\alpha = 0$ when $\alpha = 0$. Hence, the policy $\rho(q, p) = q \min(\bar{v}/2, p)$ is optimal when $\alpha = 0$, and it simply requires that the firm get less than $\bar{v}/2$ per unit.

The optimal policy in Theorem 3.2 features three properties. First, the fact 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. To see the price cap more explicitly, consider the policy $\tilde{\rho}$ given by:

$$\tilde{\rho}(q, p) = \begin{cases} \rho(q, p) & p \leq k_\alpha \\ -\infty & p > k_\alpha. \end{cases}$$

The policy $\tilde{\rho}$ induces similar behavior to that of ρ in the sense that (q, p) is a best response to ρ if and only if $(q, \min(p, k_\alpha))$ is a best response to $\tilde{\rho}$ and these responses give the same consumers' surplus. Therefore, by Theorem 3.2, $\tilde{\rho}$ is also optimal. 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 . 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 the cost functions except for monotonicity, semicontinuity, and the range of consumers' values (between 0 and \bar{v}). 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. For instance, the regulator may know that the firm has a constant marginal cost together with a fixed cost, but doesn't know these cost levels. This is the type of cost functions 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., $C(q)$ is constant for any $q > 0$. (See remark 2 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. 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.

Of course, the regulator may do better than r_α if he has significant knowledge about the industry. We believe that incorporating the regulator's additional knowledge is an exciting direction for future research, which will demonstrate the adaptability of the worst-case regret approach. Our analysis and policy serve as the very first step toward understanding the optimal policy for any particular industry.

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 prefer 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 property, by subsection 4.1 Theorem 3.1 remains correct under this decreasing average assumption on the cost function, 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

5.1 Preliminaries

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

$$V_{q,p}(z) = \begin{cases} \bar{v}, & \text{if } z \leq q \\ \frac{qp}{z}, & \text{if } q < z \leq 1, \end{cases} \quad \text{and } W_{q,p}(z) = \begin{cases} p, & \text{if } z \leq q \\ 0, & \text{if } q < z \leq 1, \end{cases}$$

as shown in Figure 2. The inverse demand $W_{q,p}$ 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.

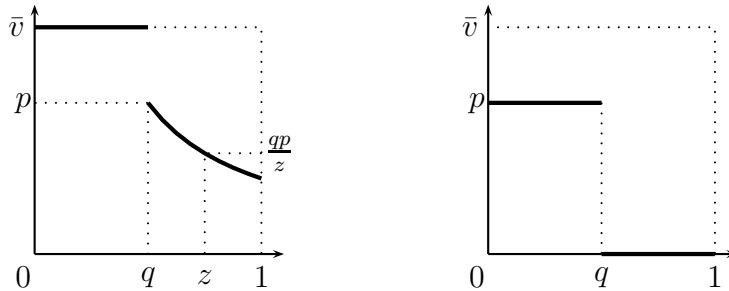


Figure 2: $V_{q,p}$ and $W_{q,p}$ demand

To understand the role of $V_{q,p}$ 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 $V_{q,p}$ and the cost is zero, then selling q units at price \bar{v} is a firm's best response. This response causes distortion of $\int_q^1 V_{q,p}(z) dz = -qp \log q$ due to underproduction. The following lemma shows that this is the worst distortion that can happen when the firm is unregulated.

Lemma 5.1. *Assume that an unregulated firm sells \bar{q} units at a price \bar{p} such that $\bar{p} \geq \sup_{z > \bar{q}} V(z)$. Let*

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

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

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

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

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

where

$$\varphi(q) = \begin{cases} \frac{1}{e}, & \text{if } q < \frac{1}{e} \\ -q \log q, & \text{if } q \geq \frac{1}{e}, \end{cases}$$

is the least decreasing majorant of $q \mapsto -q \log q$.

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

Proof of Lemma 5.1. We can assume that $\text{FP}_{\bar{q}} = 0$, otherwise replace C with \tilde{C} such that $\tilde{C}(z) = C(z)$ if $z \leq \bar{q}$ and $\tilde{C}(z) = C(z) + \text{FP}_{\bar{q}}$ if $z > \bar{q}$.

Let $q' \in \arg \max_q \int_{\bar{q}}^q V(z) dz - (C(q) - C(\bar{q}))$. Let $c' = C(q') - C(\bar{q})$.

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

$$V(z) \leq \frac{\bar{q}\bar{p} + C(z) - C(\bar{q})}{z} \leq \frac{\bar{q}\bar{p} + c'}{z},$$

for $\bar{q} < z \leq q'$.

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 $q'' = \bar{q} + \frac{c'}{\bar{p}}$, so $q'' \leq q'$. Then

$$\begin{aligned} \int_{\bar{q}}^{q'} V(z) dz - c' &\leq (q'' - \bar{q})\bar{p} + \int_{q''}^{q'} \frac{\bar{q}\bar{p} + c'}{z} dz - c' \leq (q'' - \bar{q})\bar{p} + \int_{q''}^1 \frac{\bar{q}\bar{p} + c'}{z} dz - c' \\ &= \int_{q''}^1 \frac{\bar{q}\bar{p} + c'}{z} dz = -q''\bar{p} \log q'' \leq \varphi(\bar{q})\bar{p}, \end{aligned}$$

where the first step uses the fact that $V(z) \leq \bar{p}$ for $\bar{q} < z \leq q''$, the second to last step follows from $\bar{q}\bar{p} + c' = q''\bar{p}$, and the last step follows from $\bar{q} \leq q''$ and the definition of φ . \square

For $\bar{q} = 0$ and $\bar{p} = \bar{v}$ the lemma 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 social optimum. We are unaware of previous statements of this corollary, but similar arguments to those in the proof of Lemma 5.1 with zero cost appeared in Roesler and Szentes (2017).

Corollary 5.2. *For an unregulated firm which best responds to (V, C) , we have*

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

5.2 Lower bound 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) and all the firm's best responses (q, p) to (V, C) under ρ .

For a policy ρ let $\bar{\rho}(q) = \max_{q' \leq q, p' \leq \bar{v}} \rho(q', p')$ be the maximal revenue the firm can get under ρ from selling q units or less, 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 3, $\bar{\rho}(q)$ 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.

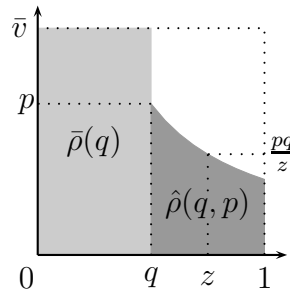


Figure 3: Definitions of $\bar{\rho}(q)$ and $\hat{\rho}(q, p)$

We first show that the worst-case regret under a policy is at least the maximum subsidy that this policy offers.

Claim 2. *Fix a policy ρ . Then $\text{WCR}(\rho) \geq \rho(q, p) - qp$ for every q, p .*

Proof. If $\rho(q, p) \leq qp$ the assertion follows from the fact that regret is nonnegative. Assume that $\rho(q, p) \geq qp$ and consider the inverse-demand function $W_{q,p}$ and a fixed cost $\rho(q, p)$. Then the firm will produce q units at price p , with $\text{FP} = 0$ and

$$\text{RGRT} = \text{DSTR} = \rho(q, p) - qp,$$

because of overproduction. □

We then show that, if the firm doesn't receive sufficiently more revenue from producing more, there is sizable regret due to underproduction.

Claim 3. *Fix a policy ρ . Let $q \leq \underline{q} \in [0, 1]$ and let $p \in [0, k_\alpha]$. 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 $V_{q,p}$ and a cost function such that producing \underline{q} units or less 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 $\text{FP} = \bar{\rho}(\underline{q})$ 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 $V_{q,p}$ and zero cost. The firm will produce at most q units, with $\text{FP} = \bar{\rho}(q) \geq \bar{\rho}(\underline{q}) + (q - \underline{q})k_\alpha$ and $\text{DSTR} \geq -qp \log q$ 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 2 and 3, we show that the regulator suffers sizable regret either from underproduction or from overproduction.

Claim 4. Fix a policy ρ . Let $\underline{q} \leq q \in [0, 1]$ and let $p \in [0, k_\alpha]$. 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 $q'p' \leq qp$, $q' \geq q$ and $\rho(q', p') = \hat{\rho}(q, p)$. By Claim 2

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

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 3. □

5.3 Proof of Theorem 3.1

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 . This follows from Claim 4 with $\underline{q} = 0$.

Remark 2. The proof of Claim 4 for the case of $\underline{q} = 0$ relies only on fixed cost functions. So does the proof of Theorem 3.1.

5.4 Upper bound on worst-case regret

We consider a policy of the form

$$\rho(q, p) = \min(qk, qp + s). \tag{7}$$

We bound the regret from (7) separately for the case of overproduction and the case of underproduction.

Claim 5. *The regret from overproduction under (7) is at most*

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

Proof. Let q^* be a socially optimal 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

$$\text{DSTR} = \bar{c} - \int_{q^*}^q V(z) \, dz \leq \bar{c} - (q - q^*)p, \quad (8)$$

and

$$\bar{c} \leq \rho(q, p) - \rho(q^*, p^*) \quad (9)$$

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 (3), the fact that $\bar{c} \leq C(q)$ and (8); the second inequality follows from (9); the third inequality follows from the fact that $p \leq p^*$ and the fact that $p \mapsto \rho(q, p)$ is monotone increasing; in the fourth inequality, $\rho(q, p) - \rho(q^*, p) \leq (q - q^*)\rho(q, p)/q$ because $q^* \leq q$ and $q \mapsto \rho(q, p)/q$ is decreasing; the fifth inequality follows from $\rho(q^*, p^*) \leq (1 - \alpha)q^*k$, $\rho(q, p) \leq qp + s$, and $q \leq 1$.

□

Claim 6. *The regret from underproduction under (7) is at most*

$$\max_q q \max((1 - \alpha)k, \bar{v} - k) - q \left(k - \frac{s}{q} \right) \log q.$$

Proof. Let q^* be a socially optimal 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 was 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 5.1 there exists some $z^* \in [\bar{q}, q^*]$ such that

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

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

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

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

$$\int_{\bar{q}}^{q^*} V(z) dz \leq C(q^*) - C(q) + \rho(q, p) - \bar{q}k + \varphi(\bar{q})\bar{p}. \quad (10)$$

Therefore

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

where the first inequality follows from (10) and $V(z) \leq \bar{v}$, and the second from $\rho(q, p) \leq qk$.

It follows that

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

for some $\bar{q} \leq q' < 1$ such that $\varphi(\bar{q}) = -q' \log q'$. Here the last inequality follows from the fact that $\bar{p} \leq k - s/\bar{q} \leq k - s/q'$. \square

5.5 Proof of Theorem 3.2

Note first that from (5) with $q = 1$ and $p = 0$ we get $r_\alpha \geq (1 - \alpha)k_\alpha = \bar{v} - k_\alpha$.

Consider the policy (7) with $k = k_\alpha$ and $s_\alpha \leq s \leq r_\alpha$.

Since $s \leq r_\alpha$ and $\bar{v} - k_\alpha = (1 - \alpha)k_\alpha \leq r_\alpha$ it follows from Claim 5 that the regret from overproduction is at most r_α .

By Claim 6 to prove that the regret from underproduction is at most r_α it is sufficient to prove that $(1 - \alpha)qk_\alpha - q(k_\alpha - s/q) \log q \leq r_\alpha$ for every $q \in [0, 1]$. For $q = 1$ this follows from the fact that $(1 - \alpha)k_\alpha \leq r_\alpha$. Let $q < 1$ and let $p = k_\alpha - s/q$. Let $q' > q$. Then $q'(k_\alpha - p) > s$

and by the assumption on s this implies that $(1 - \alpha)q'k_\alpha - q'p \log q' \leq r_\alpha$. Since this is true for every $q' > q$ it follows by continuity that $(1 - \alpha)qk_\alpha - qp \log q \leq r_\alpha$, as desired.

5.6 Proof of Theorem 3.3

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 4 with $q = q_\alpha$ and $p = p_\alpha$,

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

2. Suppose that $\rho(q, p) < qp + s_\alpha$ for every q, p . Then there exists some $q \in [0, 1], p \in [0, k_\alpha]$ such that $(1 - \alpha)qk_\alpha - qp \log q > r_\alpha$ and $q(k_\alpha - p) > \max_{q', p'}(\rho(q', p') - p'q') \geq \hat{\rho}(q, p) - qp$, implying $\hat{\rho}(q, p) < qk_\alpha$. By Claim 3 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 2

References

- Armstrong, Mark, and David EM Sappington.** 2007. "Recent developments in the theory of regulation." *Handbook of industrial organization*, 3: 1557–1700.
- 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.
- Carroll, Gabriel.** 2019. “Robustness in Mechanism Design and Contracting.” *Annual Review of Economics*, 11(1): 139–166.
- 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.
- 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. M. Boyer and R. Kihlstrom, 115–141. Amsterdam:North-Holland.
- Hurwicz, Leonid, and Leonard Shapiro.** 1978. “Incentive Structures Maximizing Residual Gain under Incomplete Information.” *The Bell Journal of Economics*, 9(1): 180–191.
- Linhart, P.B., and R. Radner.** 1989. “Minimax-regret strategies for bargaining over several variables.” *Journal of Economic Theory*, 48(1): 152–178.
- Renou, Ludovic, and Karl H. Schlag.** 2011. “Implementation in minimax regret equilibrium.” *Games and Economic Behavior*, 71(2): 527 – 533.
- 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.
- Weitzman, Martin L.** 1974. “Prices vs. Quantities.” *The Review of Economic Studies*, October, 41(4): 477–491.