Robust Monopoly Regulation

Yingni Guo and Eran Shmaya

Northwestern University
Stony Brook University

February 21, 2025

Introduction

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

Environment

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

$$\int_0^q P(z)dz - C(q).$$

- The maximal total surplus is given by

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

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

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

Regulatory Policies

- The regulator observes the firm's choice (q, p).
- A regulatory policy is given by an upper-semicontinuous function

$$\rho : [0, 1] \times [0, \bar{v}] \to \mathbb{R} \cup \{-\infty\}.$$

- If the firm chooses (q, p), it receives the revenue $\rho(q, p)$.
- This revenue is the sum of the market revenue qp and any tax or subsidy, $\rho(q, p) - qp$, imposed by the regulator.
- We assume that $\rho(0, 0) \geq 0$, so the firm can stay out of business without suffering a negative profit.

Regulatory Policies

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

Regulatory Policies

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

Regulatory Policies

- By Taxation Principle, it is without loss of generality to work directly with the revenue function $\rho(q, p)$.
- Fix a policy ρ and a pair of (P, C) functions. If the firm sells q units at price p, then consumers' surplus and the firm's profit are given by

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

- We say that (q, p) is a firm's best response to (P, C) under policy ρ if it maximizes the firm's profit over all feasible quantity-price pairs.
- The firm may have multiple best responses. Its participation constraint implies that $FP(\rho, P, C, q, p) \geq 0$ for every best response (q, p).
- The regulator's payoff is a weighted sum, $CS + \alpha FP$, of consumers' surplus and the firm's profit.
- The parameter $\alpha \in [0, 1]$ is the welfare weight the regulator puts on the firm's profit.

Complete Information

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

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

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

Claim (1)

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

- Under complete information, the regulator only compensates cost and achieves $FP = 0$.

Regret

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

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

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

Claim (2)

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

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

- To see it,

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

- Efficiency loss and redistribution concern.

The Regulator's Problem

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

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

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

Lower Bound on Worst-Case Regret

- We begin with an example in Figure 1 to illustrate the trade-off between (i) protecting consumers' surplus and (ii) mitigating underproduction.
- Suppose that the regulator constrains how much consumers' surplus the firm can extract by imposing a price cap at k.
- This price cap has opposing implications for the two market scenarios in Figure 1.
- In the left panel, every consumer has the highest value $\bar{v}$, and the cost is zero.
- The firm will price at k and serve all consumers. There is no distortion since all consumers are served, as they should be. The firm's profit is k, so regret is $(1-\alpha)k$.
- The lower the price cap k, the lower this regret from the firm's profit.
- In the right panel, every consumer still has the highest value $\bar{v}$, but now the firm has a fixed cost of k.
- It is a firm's best response not to produce. The firm's profit is zero, but distortion is $\bar{v}-k$, which is the maximal total surplus that could have been generated. Thus, regret is $\bar{v}-k$.
- The lower the price cap k, the higher this regret from distortion. Combining these two market scenarios, we conclude that the worst-case regret given the price cap k is at least $\max\{(1-\alpha)k, \bar{v}-k\}$.

Lower Bound on Worst-Case Regret

- We now provide a lower bound on the worst-case regret based on this trade-off.

Claim (3)

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

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

Lower Bound on Worst-Case Regret

- Building on Claim 3, our next theorem incorporates overproduction and the trade-off between mitigating underproduction and limiting potential overproduction.
- Theorem 1 refines the lower bound on the worst-case regret in Claim 3.

Theorem (1 Lower Bound on Worst-Case Regret)

Let

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

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

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

Lower Bound on Worst-Case Regret

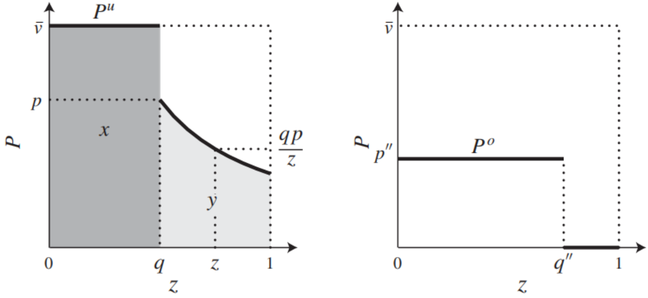

FIGURE 2. INVERSE-DEMAND FUNCTIONS P^u AND P^o IN THE PROOF OF THEOREM 1

Lower Bound on Worst-Case Regret

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

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

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

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

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

Lower Bound on Worst-Case Regret

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

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

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

Lower Bound on Worst-Case Regret

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

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

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

Optimal Policy

- To illustrate our optimal policy and the intuition behind it, we begin with two polar cases of the weight α : $\alpha = 0$ and $\alpha = 1$. We then extend the intuition from these two cases to any weight between zero and one.
- We first consider the case that $\alpha = 0$. One unit of the firm's profit translates into one unit of regret. Compared to the cases in which $\alpha > 0$, the regulator with $\alpha = 0$ has the strongest redistribution objective and thus the strongest incentive to constrain the firm's revenue.
- We next show that it is optimal to cap the firm's average revenue at $f_0 \bar{v} = \bar{v}/2$.

Claim (4 Optimal Policy for $\alpha = 0$)

Assume that $\alpha = 0$. The policy

$$\rho(q, p) = \min \{q f_0 \bar{v}, qp\} = \min\{q\bar{v}/2, qp\}$$

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

Optimal Policy

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

Claim (5 Optimal Policy for $\alpha = 1$)

Assume that $\alpha = 1$. The policy

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

achieves the worst-case regret r_1.

Optimal Policy

Theorem (2 Optimal Policy for $0 \le \alpha \le 1$)

Let

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

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

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

achieves the worst-case regret r_α.

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

Optimal Policy

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

Theorem (3 Properties of Any Optimal Policy)

Let ρ be an optimal policy. Then,

- *(Cap on Average Revenue): $\rho(q,p) \leq q f_\alpha \bar{v}$ for every $q \leq q_\alpha$.*
- *(Subsidy): If $s_\alpha > 0$, then there exists some (q,p) such that $\rho(q,p) \geq qp + s_\alpha$.*
- *(Cap on Total Subsidy): $\rho(q,p) \leq qp + r_\alpha$ for every (q,p).*

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

When Does the Firm Clear the Market?

- In this subsection, we discuss whether and when the firm wants to clear the market.
- First, the optimal policy (7) is weakly increasing in price p, so if (q, p) is a best response of the firm under this policy, then $(q, P(q))$ is also a best response. This shows that under policy (7), the firm always has a best response that clears the market.
- Second, we have shown that the cap on average revenue can be implemented by imposing a price cap. The price cap implementation (8) is no longer weakly increasing in p, so under this policy the firm may not have a best response that clears the market. For instance, suppose that $\alpha = 0$ and that the regulator imposes a price cap at $\bar{v}/2$. Suppose also that the firm's inverse-demand function is $P(q) = \bar{v}$ and its cost function is $C(q) = \bar{v}q^2/2$. The firm's best response is $(q, p) = (1/2, \bar{v}/2)$, which cannot clear the market.
- Third, a common assumption in the monopoly-regulation literature is that the firm has decreasing average cost. We next show that, if we assume decreasing average cost, then even under the price cap implementation (8), the firm has a best response that clears the market.

Proposition (1)

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

Incorporating Additional Knowledge

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

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

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

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

Incorporating Additional Knowledge

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

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

- On the other hand, the regulator knows that $P \leq \bar{P}$, so the total consumer value that the firm has created is at most $\bar{V}(q)$. We illustrate the value of $\Theta(q, p)$ in the left panel of Figure 4 and that of $\bar{V}(q)$ in the right panel. In our baseline model, $\Theta(q, p)$ equals the market revenue qp since $\underline{P} \equiv 0$, and $\bar{V}(q)$ equals $q\bar{v}$ since $\bar{P} \equiv \bar{v}$.
- We now adjust policy (7) in Theorem 2 to incorporate the bounds $(\underline{P}, \bar{P})$ on P. We replace $q\bar{v}$ and qp with the more general terms of $\bar{V}(q)$ and $\Theta(q, p)$, respectively.

Theorem (4 Optimal Policy Given Bounds $(\underline{P}, \bar{P})$)

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

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

achieves the minimal worst-case regret.

Conclusion

- How to regulate a monopolistic firm when information asymmetries are pronounced and multidimensional is an important and thus far underexplored topic.
- We study such a setting in which the regulator has very little information about the firm's demand and cost scenarios.
- We show that the optimal policy features three instruments:
  - (i) a cap on the firm's average revenue to protect consumers' surplus
  - (ii) a piece rate subsidy up to a target level to mitigate underproduction
  - (iii) a cap on the total subsidy to limit potential overproduction. When the consumers are sufficiently homogeneous or the regulator's redistributive objective is sufficiently strong, a cap on the firm's average revenue is sufficient, which can be implemented by a price cap policy.
- Our result thus provides another explanation for why price cap policies are popular.

Thanks!