

Interim Regret Minimization*

Wei He[†] Jiangtao Li[‡] Kexin Wang[§]

March 25, 2024

Abstract

We consider a robust version of monopoly pricing when the seller only knows the bound on valuations and the mean of the distribution of the buyer's value. The seller seeks to minimize interim regret, the forgone expected revenue due to not knowing the distribution of the buyer's value. The optimal pricing policy randomizes over a range of prices; the support of the pricing policy is bounded away from zero.

KEYWORDS: Robust mechanism design, distributional uncertainty, interim regret, regret minimization.

*We thank Takashi Kunimoto and Jingyi Xue for helpful comments.

[†]Department of Economics, The Chinese University of Hong Kong, hewei@cuhk.edu.hk

[‡]School of Economics, Singapore Management University, jtli@smu.edu.sg

[§]School of Economics, Singapore Management University, kexinwang.2020@phdecons.smu.edu.sg

1 Introduction

We consider a robust version of monopoly pricing. The seller knows the bound on valuations and the mean of the distribution of the buyer’s value, but has non-Bayesian uncertainty about other aspects of the distribution. The seller sets an ambitious benchmark—the highest expected revenue if she knew the true distribution—and minimizes the interim regret, defined to be the difference between the highest expected revenue if she knew the true distribution and the actual expected revenue from her pricing policy. In other words, the interim regret is the foregone revenue due to not knowing the distribution of the buyer’s value.

Clearly, if the seller knew the distribution of the buyer’s value, as is modeled in the Bayesian environment, then selecting the optimal pricing policy poses no interim regret, as it yields the highest expected revenue based on the known distribution. However, if the seller does not know the distribution, as is modeled in the current paper, then any pricing policy will leave the seller facing various kinds of exposures of interim regret. The seller in our model chooses a pricing policy that minimizes the worst-case interim regret, where the worst case is taken over all distributions that are perceived to be plausible.

A typical approach to this kind of worst-case analysis is to solve for a saddle point of the corresponding zero-sum game between the seller and the adversarial Nature. However, the zero-sum game associated with the interim regret minimization problem does not admit a saddle point.¹ Intuitively, if there were a saddle point, then the seller’s pricing policy in the saddle point necessarily achieves the lowest interim regret with respect to Nature’s choice in the saddle point. It follows that this pricing policy generates the highest expected revenue with respect to Nature’s choice and the interim regret of this pricing policy with respect to Nature’s choice is 0. This further implies that the seller’s worst-case interim regret from this pricing policy is 0, which clearly cannot be true.

To solve the interim regret minimization problem, we first consider an auxiliary problem where the seller minimizes what we call modified interim regret. The key difference of the interim regret and the modified interim regret is that, while the interim regret is convex in the distribution, the modified interim regret is linear in the

¹On a technical level, this is due to the fact that the benchmark—the highest expected revenue if the seller knew the distribution—is convex in the distribution.

distribution. This feature enables us to construct a saddle point of the auxiliary game. For any pricing policy, while the modified interim regret and the interim regret with respect to specific distributions may be different, we show that the worst-case modified interim regret and the worst-case interim regret are the same. This facilitates solving the original interim regret minimization problem via the saddle point of the auxiliary game. We show that the optimal pricing policy (that minimizes the worst-case interim regret) randomizes over a range of prices; the support of the pricing policy is bounded away from zero.

Our paper is not the first to consider the concept of interim regret minimization. To the best of our knowledge, [Bergemann and Schlag \(2005\)](#) is the pioneering study that addresses the interim regret minimization problem within a design context. They consider the pricing problem of a seller who only knows that the true distribution is in the neighborhood of a given model distribution and solve both the ex post regret minimization problem and the interim regret minimization problem. The robust mechanism design literature has by and large neglected the notion of interim regret minimization (and focused on ex post regret). We believe that this concept deserves attention, because it captures precisely the forgone expected revenue due to not knowing the distribution of the buyer's valuation. On the flip side, one can also interpret the worst-case interim regret as the highest possible gain that can be achieved via expending resources to learn the distribution of the buyer's value. Thus, the seller could use the worst-case interim regret as a measure to determine whether it is worthwhile to expend resources to acquire information about the distribution. Perhaps more importantly, the notion of interim regret permits an alternative interpretation with a continuum of buyers; if the monopoly has to use the same pricing policy against a continuum of buyers, even if we consider the complete information scenario where the seller knows the values of the buyers, the appropriate benchmark is the revenue of the best pricing policy with respect to the distribution and the notion of the interim regret seems more plausible than the alternative notion of ex post regret.

In terms of the solution concept, [Arieli, Babichenko, Talgam-Cohen, and Zabarnyi \(2023\)](#) is closely related to our paper, albeit in a rather different setting. They study information aggregation with a decision maker aggregating recommendations from several agents. The decision maker knows the prior distribution over states and the marginal distribution of each agent's recommendation but does not know the correlation. They study interim regret, which they define to be the difference between

the benchmark—the performance of the optimal aggregation rule in the Bayesian model—and the performance of her aggregation rule.

The interim regret minimization approach complements the ex post regret minimization approach in the existing literature. The most closely related papers that study ex post regret include the seminal work of [Bergemann and Schlag \(2008\)](#) and [Bergemann and Schlag \(2011\)](#). [Bergemann and Schlag \(2008\)](#) consider the problem of pricing a single object when the seller only knows the support of the possible valuations.² [Bergemann and Schlag \(2011\)](#) consider a seller who only knows that the true distribution is in the neighborhood of a given model distribution. In other settings, [Guo and Shmaya \(2023a\)](#) analyze the regret-minimizing project choice problem and [Guo and Shmaya \(2023b\)](#) study how to regulate a monopolistic firm.

[Carrasco, Luz, Kos, Messner, Monteiro, and Moreira \(2018\)](#) offer a systematic analysis of maxmin mechanism design when the seller only knows some moment conditions. In particular, Proposition 5 in their paper solves for the maxmin optimal mechanism that achieves the highest revenue guarantee, when the seller only knows the bound on valuations and the distribution’s mean. The ex post regret minimization problem, where the distribution is subject to the mean restriction, can be shown to reduce to this maxmin problem. The current paper takes the complementary approach of interim regret minimization.

2 Preliminaries

A seller (she) seeks to sell a single indivisible object to a buyer (he). The buyer has a value v , his willingness to pay for the object, belonging to $V = [0, 1]$. The buyer’s net utility of purchasing the good at price p is given by

$$u(v, p) = v - p.$$

Clearly, the buyer purchases the good at price p if $v \geq p$ and does not purchase if $v < p$. The seller’s profit from a deterministic price p if the buyer’s value is v is then

$$\pi(p, v) = p \mathbf{1}_{\{v \geq p\}},$$

²[Zhang \(2022\)](#) extends the analysis to the case of multiple goods and multiple bidders.

where $\mathbf{1}_{\{v \geq p\}}$ is the indicator function specifying

$$\mathbf{1}_{\{v \geq p\}} = \begin{cases} 1, & \text{if } v \geq p, \\ 0, & \text{if } v < p. \end{cases}$$

The buyer's value v is private information to the buyer and unknown to the seller. The seller knows that $v \in V$, has an estimate of the mean $m \in (0, 1)$ of the distribution of the buyer's value, but has non-Bayesian uncertainty about other aspects of the distribution. In other words, any distribution F is perceived to be plausible as long as $\text{supp}(F) \in V$ and F is consistent with m . Let

$$\mathcal{F} := \left\{ F \in \Delta V \mid \int_V v \, dF(v) = m \right\}$$

denote the collection of such distributions, where ΔV is the set of all probability distributions on V .

Consider a Bayesian seller who knows the distribution F of the buyer's value and maximizes the expected revenue with respect to F . The expected revenue from using a deterministic price p against F is

$$\pi(p, F) = \int_V \pi(p, v) \, dF(v) = p \cdot (1 - F(p-)),$$

where $F(p-)$ denotes the measure on the interval $[0, p)$. If the seller chooses a mixed pricing policy $\Phi \in \Delta \mathbb{R}_+$, then the expected revenue of the seller against F is

$$\pi(\Phi, F) = \int_0^\infty \pi(p, F) \, d\Phi(p).$$

Let $\pi(F)$ denote the Bayesian seller's highest expected revenue against F . It follows from standard arguments in mechanism design that the optimal selling mechanism is a posted-price mechanism. Let $p^*(F) \in \mathbb{R}_+$ denote the optimal posted price against F . Formally,

$$\pi(F) = \max_p p \cdot (1 - F(p-)) = \pi(p^*(F), F).$$

The seller in our model has non-Bayesian uncertainty about the distribution of the buyer's value. She sets an ambitious benchmark—the highest expected revenue if she knew the true distribution—and minimizes the interim regret, which we define to

be the difference between the highest expected revenue if she knew the true distribution and the actual expected revenue from her pricing policy. In other words, the interim regret is the foregone revenue due to not knowing the distribution of the buyer's value. Formally, the interim regret of a deterministic price $p \in \mathbb{R}_+$ if the true distribution is F is defined as follows:

$$r(p, F) = \pi(F) - \pi(p, F).$$

The interim regret of a mixed pricing strategy $\Phi \in \Delta\mathbb{R}_+$ if the true distribution is F is defined as follows:

$$r(\Phi, F) = \pi(F) - \pi(\Phi, F).$$

Since the seller does not know the distribution of the buyer's value, her objective is to choose a pricing policy with low regret for all distributions that are perceived to be plausible. Let

$$r(\Phi) = \sup_{F \in \mathcal{F}} r(\Phi, F)$$

denote the interim regret of a pricing policy Φ in the worst-case scenario. The seller chooses an optimal pricing policy to minimize the worst-case interim regret. Formally, the seller solves the following interim regret minimization problem:

$$\inf_{\Phi \in \Delta\mathbb{R}_+} r(\Phi) = \inf_{\Phi \in \Delta\mathbb{R}_+} \sup_{F \in \mathcal{F}} r(\Phi, F).$$

Remark. *This minimization problem can be viewed as a two-player zero-sum game between the seller and Nature, where the seller chooses a pricing policy and Nature chooses the distribution of the buyer's value. Let $\Delta\mathcal{F}$ denote the set of probability measures on \mathcal{F} . For any $\tau \in \Delta\mathcal{F}$, let F_τ denote the distribution induced by τ :*

$$F_\tau(x) = \int_{\mathcal{F}} F(x) \, d\tau(F).$$

Clearly, $F_\tau \in \mathcal{F}$. Let $\pi(\tau)$ denote the Bayesian seller's highest expected revenue against τ , and $r(\Phi, \tau)$ the interim regret of the pricing policy Φ against τ . As the seller does not observe the outcome of the randomization, we have $\pi(\tau) = \pi(F_\tau)$. Consequently,

$$r(\Phi, \tau) = \pi(\tau) - \pi(\Phi, \tau).$$

Note that for any pricing policy Φ and any $\tau \in \Delta\mathcal{F}$, we have

$$r(\Phi, \tau) = \pi(\tau) - \pi(\Phi, \tau) = \pi(F_\tau) - \pi(\Phi, F_\tau) = r(\Phi, F_\tau).$$

In words, for the interim regret minimization problem, it is inconsequential whether Nature uses a mixed strategy.

Denote by \mathcal{F}^e the set of extreme points of \mathcal{F} . For any $v \in V$, let δ_v denote the Dirac measure on v . For any x, y such that $0 \leq x < m < y \leq 1$, let $\delta_{x,y}$ denote the binary-support distribution that puts probability $\frac{y-m}{y-x}$ on x and probability $\frac{m-x}{y-x}$ on y (so that $\delta_{x,y} \in \mathcal{F}$). [Winkler \(1988\)](#) shows that $F \in \mathcal{F}^e$ if and only if $F = \delta_m$ or $F = \delta_{x,y}$ for some x, y where $0 \leq x < m < y \leq 1$. Proposition 1 below shows that, for any pricing policy (deterministic or mixed), to solve for its worst-case interim regret, it is without loss of generality to consider \mathcal{F}^e .

Proposition 1. For any $\Phi \in \Delta\mathbb{R}_+$,

$$r(\Phi) = \sup_{F \in \mathcal{F}} r(\Phi, F) = \sup_{F \in \mathcal{F}^e} r(\Phi, F).$$

Proof. Since $r(\Phi, F) = \pi(F) - \pi(\Phi, F)$, $\pi(F)$ is convex in F , and $\pi(\Phi, F)$ is linear in F , we know that $r(\Phi, F)$ is a convex function of F . Clearly, \mathcal{F} is compact and convex. Thus, we can conclude that

$$\sup_{F \in \mathcal{F}} r(\Phi, F) = \sup_{F \in \mathcal{F}^e} r(\Phi, F).$$

□

3 The optimal pricing policy

In this section, we consider the interim regret minimization problem of the seller. As is well-known, this minimization problem can be viewed as a two-player zero-sum game between the seller and Nature, where the seller chooses a pricing policy Φ and Nature chooses the distribution F of the buyer's value. For any (Φ, F) , the seller's payoff is $-r(\Phi, F)$, and Nature's payoff is $r(\Phi, F)$.

If this zero-sum game admits a Nash equilibrium, then the optimal pricing policy can be determined via the equilibrium. However, it is not hard to see that this game

does not admit a Nash equilibrium. To wit, suppose that there exists a Nash equilibrium (Φ^*, F^*) . It must be that

$$r(\Phi^*, F^*) \leq r(\Phi, F^*), \forall \Phi \in \Delta\mathbb{R}_+.$$

Since $r(\Phi, F^*) = \pi(F^*) - \pi(\Phi, F^*)$ for any $\Phi \in \Delta\mathbb{R}_+$, we have

$$\pi(\Phi^*, F^*) \geq \pi(\Phi, F^*), \forall \Phi \in \Delta\mathbb{R}_+.$$

Thus, Φ^* generates the highest expected revenue against the distribution F^* , and

$$r(\Phi^*, F^*) = \pi(F^*) - \pi(\Phi^*, F^*) = \pi(\Phi^*, F^*) - \pi(\Phi^*, F^*) = 0.$$

Since F^* is Nature's best response to Φ^* , that is,

$$r(\Phi^*, F^*) \geq r(\Phi^*, F), \forall F \in \mathcal{F},$$

we then have $r(\Phi^*, F) = 0$ for all $F \in \mathcal{F}$. This is a contradiction.

To solve the interim regret minimization problem, we proceed as follows. In Section 3.1, we first construct an auxiliary game where we modify the payoffs of the seller and Nature in the original zero-sum game; we refer to Nature's payoff in the auxiliary game as the modified interim regret (to be defined shortly). The key difference of the interim regret and the modified interim regret is that, while the interim regret is convex in F , the modified interim regret is linear in F . This enables us to construct a saddle point for the auxiliary game. Section 3.2 presents the construction of the saddle point. Finally, we show in Section 3.3 that for any pricing strategy, the interim regret in the worst case and the modified interim regret in the worst case are the same. This facilitates solving the original interim regret minimization problem via the saddle point of the auxiliary game.

3.1 An auxiliary game

Define $m\pi(\tau) = \int_{\mathcal{F}} \pi(F) d\tau(F)$ and let

$$\begin{aligned} mr(\Phi, \tau) &= m\pi(\tau) - \pi(\Phi, \tau) \\ &= \int_{\mathcal{F}} \pi(F) d\tau(F) - \int_{\mathcal{F}} \pi(\Phi, F) d\tau(F) \end{aligned}$$

$$= \int_{\mathcal{F}} r(\Phi, F) \, d\tau(F).$$

We interpret $mr(\Phi, \tau)$ as the modified interim regret where we modify the benchmark to be the highest expected revenue the seller could achieve if she knew not only Nature's strategy but also the outcome of Nature's randomization. We emphasize that this is not our preferred way of thinking about interim regret, and the only role of the modified interim regret is to facilitate solving our original interim regret minimization problem. Since $\pi(F)$ is a convex function of F , we have, for any Φ and any τ ,

$$mr(\Phi, \tau) = m\pi(\tau) - \pi(\Phi, \tau) = \int_{\mathcal{F}} \pi(F) \, d\tau(F) - \pi(\Phi, \tau) \geq \pi(F_\tau) - \pi(\Phi, F_\tau) = r(\Phi, \tau).$$

This should not come as a surprise, since in calculating the modified benchmark $m\pi(\tau)$ we assume that the seller could choose the optimal price to depend on the outcome of Nature's randomization. Clearly, if τ is degenerate on some F , then $m\pi(\tau) = \pi(\tau)$ and $mr(\Phi, \tau) = r(\Phi, \tau)$ for any Φ .

We consider the following auxiliary game between the seller and Nature, where the seller chooses a pricing strategy $\Phi \in \Delta\mathbb{R}_+$ and Nature chooses a probability measure $\tau \in \Delta\mathcal{F}$. For any (Φ, τ) , the seller's payoff is $-mr(\Phi, \tau)$, and Nature's payoff is $mr(\Phi, \tau)$. Crucially, in this auxiliary game, the players' payoffs are linear in Nature's strategy, that is, for any Φ and any τ ,

$$mr(\Phi, \tau) = \int_{\mathcal{F}} mr(\Phi, F) \, d\tau(F).$$

A pair of strategies (Φ^*, τ^*) is a Nash equilibrium of the auxiliary game if it forms a saddle point, that is,

$$mr(\Phi^*, \tau) \leq mr(\Phi^*, \tau^*) \leq mr(\Phi, \tau^*), \quad \forall \Phi \in \Delta\mathbb{R}_+, \quad \forall \tau \in \Delta(\mathcal{F}).$$

Clearly, if such a saddle point (Φ^*, τ^*) exists, then

$$\begin{aligned} \Phi^* &\in \arg \min_{\Phi \in \Delta\mathbb{R}_+} \max_{\tau \in \Delta(\mathcal{F})} mr(\Phi, \tau), \text{ and} \\ mr(\Phi^*, \tau^*) &= \min_{\Phi \in \Delta\mathbb{R}_+} \max_{\tau \in \Delta(\mathcal{F})} mr(\Phi, \tau). \end{aligned}$$

In this case, we say that the pricing strategy Φ^* attains the minimax modified interim regret.

3.2 The saddle point of the auxiliary problem

In this subsection, we explicitly construct a saddle point for the auxiliary problem. We first simplify the analysis via some preliminary observations. Proposition 2 below shows that Nature necessarily uses a mixed strategy in the Nash equilibrium of the auxiliary game; τ^* must be randomization over the set \mathcal{F} (rather than a Dirac measure on some $F \in \mathcal{F}$).

Proposition 2. *There does not exist a saddle point (Φ^*, τ^*) of the auxiliary game where τ^* is degenerate on some $F^* \in \mathcal{F}$.*

The logic is similar to that when we argue there does not exist a saddle point for the zero-sum game corresponding to the original interim regret minimization problem. Suppose to the contrary, there exists a saddle point (Φ^*, F^*) of the auxiliary game. That is, $mr(\Phi^*, F^*) \leq mr(\Phi, F^*)$ for all Φ . It follows that $\pi(\Phi^*, F^*) \geq \pi(\Phi, F^*)$ for all Φ . This further implies that $mr(\Phi^*, F^*) = 0$ and $mr(\Phi^*, \tau) = 0$ for all τ . We arrive at a contradiction.

Proposition 3 shows that, to identify the Nash equilibrium, it suffices to consider Nature's strategies that randomize over the extreme points of the set \mathcal{F} .

Proposition 3. *If the auxiliary game admits a saddle point (Φ^*, τ^*) , then the auxiliary problem admits a saddle point $(\Phi^*, \tilde{\tau})$ where $\tilde{\tau} \in \Delta(\mathcal{F}^e)$.*

Proof. Suppose that the auxiliary game admits a saddle point (Φ^*, τ^*) . For each $G \in \text{supp}(\tau^*)$, there exists a probability distribution τ_G over the extreme points of \mathcal{F} such that

$$G(x) = \int_{\mathcal{F}^e} F(x) d\tau_G(F).$$

Consider the probability distribution $\tilde{\tau}$ constructed via replacing each $G \in \text{supp}(\tau^*)$ by the corresponding probability distribution τ_G over the extreme points of \mathcal{F} . By construction, $\tilde{\tau} \in \Delta(\mathcal{F}^e)$ and $F_{\tilde{\tau}} = F_{\tau^*}$. In what follows, we show that $(\Phi^*, \tilde{\tau})$ is also a saddle point of the auxiliary game.

We first show that Φ^* is the seller's best response to $\tilde{\tau}$. Since Φ^* is the seller's best response to τ^* , it must be that Φ^* generates the highest expected revenue with respect to τ^* . Since the expected revenue of any pricing policy is the same with respect to $\tilde{\tau}$ and τ^* , it follows that Φ^* generates the highest expected revenue with respect to $\tilde{\tau}$. This further implies that Φ^* is the seller's best response to $\tilde{\tau}$. Formally, since

(Φ^*, τ^*) is a saddle point, we have

$$mr(\Phi^*, \tau^*) = m\pi(\tau^*) - \pi(\Phi^*, \tau^*) \leq mr(\Phi, \tau^*) = m\pi(\tau^*) - \pi(\Phi, \tau^*)$$

for any $\Phi \in \Delta\mathbb{R}_+$. Since $F_{\tilde{\tau}} = F_{\tau^*}$, we have $\pi(\Phi, \tilde{\tau}) = \pi(\Phi, \tau^*)$ for any $\Phi \in \Delta\mathbb{R}_+$. Thus,

$$mr(\Phi^*, \tilde{\tau}) = m\pi(\tilde{\tau}) - \pi(\Phi^*, \tilde{\tau}) \leq mr(\Phi, \tilde{\tau}) = m\pi(\tilde{\tau}) - \pi(\Phi, \tilde{\tau})$$

for any $\Phi \in \Delta\mathbb{R}_+$.

Next, we show that $\tilde{\tau}$ is Nature's best response to Φ^* , that is, $mr(\Phi^*, \tilde{\tau}) \geq mr(\Phi^*, \tau)$ for all τ . Since τ^* is Nature's best response to Φ^* and deviating from τ^* to $\tilde{\tau}$ can only increase the modified interim regret, we can conclude that $\tilde{\tau}$ is also a best response to Φ^* . Formally,

$$\begin{aligned} mr(\Phi^*, \tilde{\tau}) &= \int_{\mathcal{F}} \pi(F) d\tilde{\tau}(F) - \pi(\Phi^*, \tilde{\tau}) \\ &= \int_{\mathcal{F}} \int_{\mathcal{F}^e} \pi(F) d\tau_G(F) d\tau^*(G) - \pi(\Phi^*, \tilde{\tau}) \\ &= \int_{\mathcal{F}} m\pi(\tau_G) d\tau^*(G) - \pi(\Phi^*, \tilde{\tau}) \\ &\geq \int_{\mathcal{F}} \pi(G) d\tau^*(G) - \pi(\Phi^*, \tau^*) \\ &= mr(\Phi^*, \tau^*) \\ &\geq mr(\Phi^*, \tau) \end{aligned}$$

for all τ , where the second line follows from the construction of $\tilde{\tau}$, the first inequality is because $m\pi(\tau_G) \geq \pi(G)$ for each $G \in \text{supp}(\tau^*)$ and $F_{\tilde{\tau}} = F_{\tau^*}$, and the second inequality is because (Φ^*, τ^*) is a saddle point. \square

Proposition 3 greatly simplifies our search for the saddle point, as we can without loss of generality focus on the extreme points of the set \mathcal{F} .

Towards F_{τ^*} Since a deterministic pricing policy will always leave the seller facing the dilemma of different kinds of exposures (upward exposure where the value of the buyer is very high and the regret arises from having offered a price too low relative to the valuation, and downward exposure where the value of the buyer is very low and the regret arises from having offered a price too high relative to the valuation),

she may “hedge” against regret and attempt to attain the minimal revenue loss due to not knowing the true distribution via a well-calibrated mixed pricing policy. The seller is likely to be indifferent in her pricing policy against F_{τ^*} , which implies that the marginal revenue must be zero over the range of prices which the seller offers. In the language of optimal monopoly pricing, the virtual utility of different prices with respect to F_{τ^*} has to be constant and equal to zero:

$$p - \frac{1 - F_{\tau^*}(p)}{f_{\tau^*}(p)} = 0.$$

Thus, we derive the following conjecture for the distribution F_{τ^*} :

$$F_{\tau^*}(v) = \begin{cases} 1 - \frac{a}{c} & \text{if } v \in [0, c), \\ 1 - \frac{a}{v} & \text{if } v \in [c, 1), \\ 1 & \text{if } v = 1, \end{cases}$$

where a and c are constants that satisfy $0 < a < c < m$, which will be determined in later analysis. For F_{τ^*} to be consistent with the mean restriction m , we must have

$$a(1 - \ln c) = m.$$

Clearly, the expected revenue of price p with respect to F_{τ^*} is

$$\pi(p, F_{\tau^*}) = p \cdot (1 - F_{\tau^*}(p-)) \begin{cases} < a & \text{if } p \in [0, c), \\ = a & \text{if } p \in [c, 1]. \end{cases}$$

In words, a is the highest expected revenue that the seller could achieve against F_{τ^*} , and the optimal pricing policy against F_{τ^*} has measure 0 on the interval $[0, c)$.

Towards τ^* We now explore the structure of $\tau^* \in \Delta(\mathcal{F}^e)$ that induces the distribution F_{τ^*} we identify above. The two possible mass points at 0 and 1 imply that Nature chooses a strategy which randomizes over the following special classes of demands $F \in \mathcal{F}^e$:

$$\delta_{0,y}, \forall y \in (m, 1),$$

$$\delta_{x,1}, \forall x \in [c, m),$$

$$\delta_{0,1},$$

δ_m .

In other words, each $y \in (m, 1)$ is matched with 0, each $x \in [c, m)$ is matched with 1, and any residual measure on $\{0\}$ and residual measure on $\{1\}$ will be matched with each other.

For any distribution $\delta_{0,y}$ where $y \in (m, 1)$, the probability measures on $\{0\}$ and $\{y\}$ are respectively $\frac{y-m}{y}$ and $\frac{m}{y}$. Since the density of F_{τ^*} at the point $y \in (m, 1)$ is $\frac{a}{y^2}$, the infinitesimal measure on $\{0\}$ needed to match y for some $y \in (m, 1)$ is

$$\frac{a}{y^2} \cdot \frac{\frac{y-m}{y}}{\frac{m}{y}} = \frac{a}{y^2} \cdot \frac{y-m}{m}.$$

It follows that the measure on $\{0\}$ needed to match y for all $y \in (m, 1)$ is

$$\int_m^1 \frac{a}{y^2} \cdot \frac{y-m}{m} dy = \frac{a}{m} \cdot (-\ln m + m - 1).$$

Since $F_{\tau^*}(0) = 1 - \frac{a}{c}$, the residual measure on $\{0\}$ is

$$r_0 := 1 - \frac{a}{c} - \frac{a}{m} \cdot (-\ln m + m - 1).$$

For any distribution $\delta_{x,1}$ where $x \in [c, m)$, the probability measures on $\{x\}$ and $\{1\}$ are respectively $\frac{1-m}{1-x}$ and $\frac{m-x}{1-x}$. Since the density of F_{τ^*} at the point $x \in [c, m)$ is $\frac{a}{x^2}$, the infinitesimal measure on $\{1\}$ needed to match x for some $x \in [c, m)$ is

$$\frac{a}{x^2} \cdot \frac{\frac{m-x}{1-x}}{\frac{1-m}{1-x}} = \frac{a}{x^2} \cdot \frac{m-x}{1-m}.$$

It follows that the measure on $\{1\}$ needed to match x for all $x \in [c, m)$ is

$$\int_c^m \frac{a}{x^2} \cdot \frac{m-x}{1-m} dx = \frac{a}{1-m} \cdot \left(-\ln\left(\frac{m}{c}\right) + \frac{m}{c} - 1 \right).$$

Since the measure F_{τ^*} on $\{1\}$ is a , the residual measure on $\{1\}$ is

$$r_1 := a - \frac{a}{1-m} \cdot \left(-\ln\left(\frac{m}{c}\right) + \frac{m}{c} - 1 \right).$$

For the distribution $\delta_{0,1}$, the probability measures on $\{0\}$ and $\{1\}$ are respectively $1 - m$ and m . We can readily verify (using $m = a - a \ln c$) that the residual measures

on $\{0\}$ and $\{1\}$ are compatible, that is,

$$\frac{1 - \frac{a}{c} - \frac{a}{m} \cdot (-\ln m + m - 1)}{1 - m} = \frac{a - \frac{a}{1-m} \cdot \left(-\ln\left(\frac{m}{c}\right) + \frac{m}{c} - 1\right)}{m}.$$

The values of a and c (to be determined in later analysis) guarantee that $r_0 \geq 0$ (and hence $r_1 \geq 0$). Thus, we have identified a probability measure τ^* over some particular demands $F \in \mathcal{F}^e$ such that the distribution F_{τ^*} is induced by τ^* .

Towards Φ^* We proceed to identify the pricing policy Φ^* such that Nature is indifferent among the demands in the support of τ^* . We first consider the modified interim regret with respect to the class of demands $\delta_{0,y}$ where $y \in (m, 1]$. For any pricing policy Φ , we have

$$\begin{aligned} & mr(\Phi, \delta_{0,y}) \\ &= m\pi(\delta_{0,y}) - \pi(\Phi, \delta_{0,y}) \\ &= m - \int_0^y p \cdot \frac{m}{y} d\Phi(p) \\ &= m - \frac{m}{y} \int_0^m p d\Phi(p) - \frac{m}{y} \int_m^y p d\Phi(p). \end{aligned}$$

As $mr(\Phi^*, \delta_{0,y})$ is the same for any $y \in (m, 1]$, we have the following candidate pricing policy defined via its density function:³

$$\phi^*(p) = \begin{cases} 0 & \text{if } p \in [0, c), \\ \frac{A+b}{p} + B & \text{if } p \in [c, m), \\ \frac{b}{p} & \text{if } p \in [m, 1], \end{cases}$$

where A , B , and b are parameters such that

$$\int_0^m p d\Phi^*(p) = (A + b)(m - c) + \frac{B}{2}(m^2 - c^2) = mb. \quad (1)$$

This ensures that $mr(\Phi^*, \delta_{0,y})$ is the same for any $y \in (m, 1]$. Indeed, by construction, we have

$$mr(\Phi^*, \delta_{0,y})$$

³It is reasonable to conjecture that Φ^* is atomless and the support of Φ^* is $[c, 1]$, the same as the support of F_{τ^*} minus the point $\{0\}$.

$$\begin{aligned}
&= m - \frac{m}{y} \int_0^m p \, d\Phi^*(p) - \frac{m}{y} \int_m^y p \, d\Phi^*(p) \\
&= m - \frac{m}{y} mb - \frac{m}{y} (y - m)b \\
&= m - mb
\end{aligned}$$

for all $y \in (m, 1]$.

Next, we consider the modified interim regret with respect to the class of demands $\delta_{x,1}$ where $x \in [c, m)$. Clearly, the highest expected revenue with respect to any such demand is $\max\{x, \frac{m-x}{1-x}\}$. We shall assume that $x \geq \frac{m-x}{1-x}$ for any $x \in [c, m)$. This is without loss, as our choice of c (to be determined in later analysis) ensures that this inequality holds. We have

$$\begin{aligned}
&mr(\Phi^*, \delta_{x,1}) \\
&= m\pi(\delta_{x,1}) - \pi(\Phi^*, \delta_{x,1}) \\
&= x - \int_0^x p \, d\Phi^*(p) - \int_x^1 p \cdot \frac{m-x}{1-x} \, d\Phi^*(p) \\
&= x - \int_0^m p \, d\Phi^*(p) + \frac{1-m}{1-x} \int_x^m p \, d\Phi^*(p) - \frac{m-x}{1-x} \int_m^1 p \, d\Phi^*(p) \\
&= x - mb + \frac{1-m}{1-x} \left((A+b)(m-x) + \frac{B}{2}(m^2 - x^2) \right) - \frac{m-x}{1-x} (1-m)b \\
&= x - mb + \frac{1-m}{1-x} \left(A(m-x) + \frac{B}{2}(m^2 - x^2) \right).
\end{aligned}$$

For $mr(\Phi^*, \delta_{x,1}) = m - mb$ for all $x \in [c, m)$, we must have

$$A = \frac{1+m}{1-m} \text{ and } B = -\frac{2}{1-m}.$$

It then follows from (1) that

$$b = \frac{1}{c} \left(A(m-c) + \frac{B}{2}(m^2 - c^2) \right) = \frac{(m-c)(1-c)}{(1-m)c}.$$

We obtain the pricing policy Φ^* as follows:

$$\Phi^*(p) = \begin{cases} 0 & \text{if } p \in [0, c), \\ \frac{m+c^2}{(1-m)c} \cdot \ln\left(\frac{p}{c}\right) - \frac{2}{1-m}(p-c) & \text{if } p \in [c, m], \\ \frac{(m-c)(1-c)}{(1-m)c} \cdot \ln\left(\frac{p}{m}\right) + \Phi^*(m) & \text{if } p \in (m, 1]. \end{cases} \quad (2)$$

Parameters We are now ready to determine the values of a and c . For Φ^* defined above to be a feasible pricing policy, let $c \in (0, m)$ be such that

$$\begin{aligned}
& \Phi^*(1) \\
&= \frac{(m-c)(1-c)}{(1-m)c} \cdot (-\ln m) + \Phi^*(m) \\
&= \frac{(m-c)(1-c)}{(1-m)c} \cdot (-\ln m) + \frac{m+c^2}{(1-m)c} \cdot \ln\left(\frac{m}{c}\right) - \frac{2}{1-m}(m-c) \\
&= -\frac{m+c^2}{(1-m)c} \cdot \ln c + \frac{1+m}{1-m} \cdot \ln m - \frac{2}{1-m}(m-c) \\
&= 1.
\end{aligned}$$

For F_{τ^*} to be consistent with the mean restriction m , let a be such that

$$a(1 - \ln c) = m.$$

In Appendix A, we verify that such values of a and c are unique, $0 < a < c < m$, F_{τ^*} and Φ^* are well-defined distributions, τ^* is a well-defined probability measure over \mathcal{F}^e , and $x \geq \frac{m-x}{1-x}$ for any $x \in [c, m)$.

Theorem 1. *The pricing policy Φ^* defined in (2) attains the minimax modified interim regret. The value of the minimax modified interim regret is $m - mb$.*

Proof. It suffices to show that (Φ^*, τ^*) is a saddle point of the auxiliary problem.

Φ^* is the seller's best response to τ^* . Clearly, the expected revenue of price p with respect to F_{τ^*} is

$$\pi(p, F_{\tau^*}) = p \cdot (1 - F_{\tau^*}(p-)) \begin{cases} < a & \text{if } p \in [0, c), \\ = a & \text{if } p \in [c, 1]. \end{cases}$$

Since the support of the pricing policy Φ^* is $[c, 1]$, we have $\pi(\Phi^*, F_{\tau^*}) \geq \pi(\Phi, F_{\tau^*})$ for any Φ . This further implies that

$$mr(\Phi, \tau^*) = m\pi(\tau^*) - \pi(\Phi, F_{\tau^*}) \geq m\pi(\tau^*) - \pi(\Phi^*, F_{\tau^*}) = mr(\Phi^*, \tau^*).$$

τ^* is Nature's best response to Φ^* . Given the pricing policy Φ^* , the modified interim regret with respect to any distribution $F \in \mathcal{F}^e$ is

$$mr(\Phi^*, F) = m\pi(F) - \pi(\Phi^*, F) = \pi(F) - \pi(\Phi^*, F).$$

We have already argued above (when constructing the pricing policy Φ^*) that $mr(\Phi^*, F)$ must equal to $m - mb$ when $F = \delta_{0,y}$ with $y \in (m, 1]$ or $F = \delta_{x,1}$ with $x \in [c, m)$. It is straightforward to verify that $mr(\Phi^*, \delta_m) = m - \int_0^m p \, d\Phi^*(p) = m - mb$. Thus, the modified interim regret of the pricing policy Φ^* with respect to any distribution F in the support of τ^* is $mr(\Phi^*, F) = m - mb$. It follows that

$$mr(\Phi^*, \tau^*) = \int_{\mathcal{F}} mr(\Phi^*, F) \, d\tau^*(F) = m - mb.$$

In what follows, we show that for the pricing policy Φ^* any $F \in \mathcal{F}^e$ cannot generate a modified interim regret that is strictly larger than $m - mb$. For any $\delta_{x,y}$, we have

$$mr(\Phi^*, \delta_{x,y}) = \max \left\{ x, y \frac{m-x}{y-x} \right\} - \int_0^x p \, d\Phi^*(p) - \int_x^y p \frac{m-x}{y-x} \, d\Phi^*(p).$$

There are three cases to consider. First, if $\delta_{x,y}$ is such that $x \leq y \frac{m-x}{y-x}$, we have

$$\begin{aligned} & mr(\Phi^*, \delta_{x,y}) \\ &= y \frac{m-x}{y-x} - \int_0^x p \, d\Phi^*(p) - \int_x^y p \frac{m-x}{y-x} \, d\Phi^*(p) \\ &= y \frac{m-x}{y-x} - \frac{y-m}{y-x} \int_0^x p \, d\Phi^*(p) - \frac{m-x}{y-x} \int_0^y p \, d\Phi^*(p) \\ &= y \frac{m-x}{y-x} - \frac{y-m}{y-x} \int_0^x p \, d\Phi^*(p) - \frac{m-x}{y-x} \cdot yb \\ &\leq y \frac{m-x}{y-x} (1-b) \\ &\leq m(1-b), \end{aligned}$$

where the fourth line follows from $\int_0^m p \, d\Phi^*(p) = mb$ and $\int_m^y p \, d\Phi^*(p) = (y-m)b$, and the last line holds since $y \frac{m-x}{y-x} \leq m$.

Second, if $\delta_{x,y}$ is such that $x > y \frac{m-x}{y-x}$ and $x \geq c$, then

$$mr(\Phi^*, \delta_{x,y})$$

$$\begin{aligned}
&= x - \int_0^x p \, d\Phi^*(p) - \int_x^y p \frac{m-x}{y-x} \, d\Phi^*(p) \\
&= x - \int_0^m p \, d\Phi^*(p) + \frac{y-m}{y-x} \int_x^m p \, d\Phi^*(p) - \frac{m-x}{y-x} \int_m^y p \, d\Phi^*(p) \\
&= x - mb + \frac{y-m}{y-x} \int_x^m \left(\frac{1+m}{1-m} + b - \frac{2}{1-m} p \right) \, dp - \frac{m-x}{y-x} \cdot (y-m)b \\
&= x - mb + \frac{y-m}{y-x} \cdot \left(\frac{1+m}{1-m} + b \right) \cdot (m-x) - \frac{y-m}{y-x} \cdot \frac{m^2-x^2}{1-m} - \frac{m-x}{y-x} \cdot (y-m)b \\
&= x - mb + \frac{y-m}{y-x} \cdot \frac{1-x}{1-m} \cdot (m-x) \\
&\leq m - mb,
\end{aligned}$$

where the inequality follows from the observation that $\frac{y-m}{y-x} \cdot \frac{1-x}{1-m} \leq 1$.

Finally, if $\delta_{x,y}$ is such that $x > y \frac{m-x}{y-x}$ and $x < c$, then

$$\begin{aligned}
&mr(\Phi^*, \delta_{x,y}) \\
&= x - \int_x^y p \frac{m-x}{y-x} \, d\Phi^*(p) \\
&= x - \frac{m-x}{y-x} \cdot \int_0^y p \, d\Phi^*(p) \\
&= x - \frac{m-x}{y-x} \cdot yb \\
&\leq x - \frac{m-x}{1-x} \cdot b \\
&\leq c - \frac{m-c}{1-c} \cdot b \\
&= m - mb,
\end{aligned}$$

where the fifth line follows from $\frac{m-x}{y-x} \cdot y \geq \frac{m-x}{1-x}$ and the last line follows from the definition of b .

Therefore, we can conclude that

$$mr(\Phi^*, F) \leq mr(\Phi^*, \tau^*) = m - mb, \forall F \in \mathcal{F}^e.$$

This further implies

$$mr(\Phi^*, \tau) = \int_{\mathcal{F}} r(\Phi^*, F) \, d\tau(F) \leq \sup_{F \in \mathcal{F}} r(\Phi^*, F) = \sup_{F \in \mathcal{F}^e} r(\Phi^*, F) \leq mr(\Phi^*, \tau^*), \forall \tau,$$

where the second equality follows from Proposition 1. □

3.3 The interim regret minimization problem

We are now ready to wrap up our analysis. Proposition 4 below shows that for any pricing policy $\Phi \in \Delta\mathbb{R}_+$, the worst-case modified interim regret and the worst-case interim regret are the same. It then follows immediately that Φ^* , which has the worst-case (modified) interim regret of $m - mb$, solves the seller's interim regret minimization problem. We present these findings below.

Proposition 4. *For any pricing policy $\Phi \in \Delta\mathbb{R}_+$, we have*

$$r(\Phi) = \sup_{\tau \in \Delta\mathcal{F}} mr(\Phi, \tau).$$

Proof. Fix any pricing policy $\Phi \in \Delta\mathbb{R}_+$. Since $mr(\Phi, \tau) \geq r(\Phi, \tau)$ for all τ , we have

$$\sup_{\tau \in \Delta\mathcal{F}} mr(\Phi, \tau) \geq \sup_{\tau \in \Delta\mathcal{F}} r(\Phi, \tau) \geq \sup_{F \in \mathcal{F}} r(\Phi, F) = r(\Phi).$$

Furthermore,

$$mr(\Phi, \tau) = \int_{\mathcal{F}} r(\Phi, F) d\tau(F) \leq \sup_{F \in \mathcal{F}} r(\Phi, F) = r(\Phi), \forall \tau.$$

It then follows that

$$r(\Phi) = \sup_{\tau \in \Delta\mathcal{F}} mr(\Phi, \tau).$$

□

Theorem 2. *The pricing policy Φ^* defined in (2) solves the seller's interim regret minimization problem. Formally,*

$$r(\Phi^*) = \min_{\Phi \in \Delta\mathbb{R}_+} r(\Phi) = m - mb.$$

Theorem 2 follows immediately from Theorem 1 and Proposition 4.

A Values of a and c

Uniqueness of a and c . We first show that there is a unique $c \in (0, m)$ such that

$$-\frac{m+c^2}{(1-m)c} \cdot \ln c + \frac{1+m}{1-m} \cdot \ln m - \frac{2}{1-m}(m-c) = 1. \quad (3)$$

Let

$$f(x) = -\frac{m+x^2}{(1-m)x} \cdot \ln x + \frac{1+m}{1-m} \cdot \ln m - \frac{2}{1-m}(m-x).$$

Since

$$\begin{aligned} f'(x) &= -\frac{x^2-m}{(1-m)x^2} \cdot \ln x - \frac{x^2+m}{(1-m)x^2} + \frac{2}{1-m} \\ &= \frac{1}{1-m} \cdot \left(1 - \frac{m}{x^2}\right) \cdot (1 - \ln x) \\ &< 0 \end{aligned}$$

for all $x \in (0, m]$, $f(m) = 0$, and $\lim_{x \rightarrow 0} f(x) = \infty$, we conclude that there is a unique $c \in (0, m)$ such that $f(c) = 1$. This further implies that $a = \frac{m}{1-\ln c}$ is unique and

$$b = \frac{(m-c)(1-c)}{(1-m)c} > 0.$$

Verifying $0 < a < c < m$. We already know that $c \in (0, m)$. Furthermore, it is obvious that $a = \frac{m}{1-\ln c} > 0$. In what follows, we show that $a < c$. Since $m = a(1 - \ln c)$, it suffices to show that $c(1 - \ln c) > m$. Consider the following function:

$$g(x) = x(2-x), \quad x \in [0, 1].$$

Since $g(x)$ is strictly increasing on $[0, m]$, $g(0) = 0$, and $g(m) = 2m - m^2 > m$, there exists a unique $\hat{c} \in (0, m)$ such that $g(\hat{c}) = m$. If we can show that $c > \hat{c}$, then we can conclude that

$$c(1 - \ln c) > \hat{c}(1 - \ln \hat{c}) > \hat{c}(2 - \hat{c}) = m,$$

where the first inequality holds because $x(1 - \ln x)$ is strictly increasing on $(0, 1)$, and the second inequality is because $\hat{c} - \ln \hat{c} > 1 - \ln 1 = 1$.⁴

We are left to show that $c > \hat{c}$. Rearranging (3), we have

$$2c - \frac{m + c^2}{c} \cdot \ln c = (1 + m)(1 - \ln m). \quad (4)$$

We shall show that

$$2\hat{c} - \frac{m + \hat{c}^2}{\hat{c}} \cdot \ln \hat{c} - (1 + m)(1 - \ln m) > 0.$$

For this purpose, we consider the following function:

$$\begin{aligned} h(x) &= 2x - 2 \ln x - \left(1 + (2x - x^2)\right) \cdot \left(1 - \ln(2x - x^2)\right) \\ &= x^2 + 2 \ln(2 - x) - (x - 1)^2 \cdot \ln(2x - x^2) - 1. \end{aligned}$$

Clearly, $h(1) = 0$ and

$$\begin{aligned} h(\hat{c}) &= 2\hat{c} - 2 \ln \hat{c} - \left(1 + (2\hat{c} - \hat{c}^2)\right) \cdot \left(1 - \ln(2\hat{c} - \hat{c}^2)\right) \\ &= 2\hat{c} - \frac{m + \hat{c}^2}{\hat{c}} \cdot \ln \hat{c} - (1 + m)(1 - \ln m), \end{aligned}$$

where the second line follows from $2\hat{c} - \hat{c}^2 = m$. Note that

$$\begin{aligned} h'(x) &= 2x - \frac{2}{2-x} - 2(x-1) \cdot \ln(2x-x^2) - (x-1)^2 \cdot \frac{2-2x}{2x-x^2} \\ &= -\frac{2(x-1)^2}{2-x} - 2(x-1) \cdot \ln(2x-x^2) - (x-1)^2 \cdot \frac{2-2x}{x(2-x)} \\ &< 0 \end{aligned}$$

for all $x \in (0, 1)$, where the last line uses the following inequality $\ln(2x - x^2) \leq \ln 1 = 0$ for all $x \in (0, 1]$. Thus, we have $h(\hat{c}) > h(1) = 0$. Since $2x - \frac{m+x^2}{x} \cdot \ln x$ is strictly decreasing on $(0, m)$, it follows from (4) and $h(\hat{c}) > 0$ that $c > \hat{c}$.

F_{τ^*} is well-defined distribution and $F_{\tau^*} \in \mathcal{F}$. Since $0 < a < c < m$ and $a(1 - \ln c) = m$, F_{τ^*} is a well-defined distribution and is consistent with the mean restriction m . Note that F_{τ^*} has two atoms at $\{0\}$ and $\{1\}$ respectively.

⁴The function $x - \ln x$ is strictly decreasing on $(0, 1)$.

Φ^* is a well-defined pricing policy. We have established that $\Phi^*(1) = 1$. We show that Φ^* is nondecreasing on $[0, 1]$. Recall that

$$\Phi^*(p) = \begin{cases} 0 & \text{if } p \in [0, c), \\ \frac{m+c^2}{(1-m)c} \cdot \ln\left(\frac{p}{c}\right) - \frac{2}{1-m}(p-c) & \text{if } p \in [c, m], \\ \frac{(m-c)(1-c)}{(1-m)c} \cdot \ln\left(\frac{p}{m}\right) + \Phi^*(m) & \text{if } p \in (m, 1]. \end{cases}$$

For any $p \in [m, 1]$, $\phi^*(p) = \frac{(m-c)(1-c)}{(1-m)c} \cdot \frac{1}{p} > 0$. For any $p \in [c, m]$, we have

$$\begin{aligned} \phi^*(p) &= \frac{m+c^2}{(1-m)c} \cdot \frac{1}{p} - \frac{2}{1-m} \\ &> \frac{m+c^2}{(1-m)c} \cdot \frac{1}{m} - \frac{2}{1-m} \\ &= \frac{m-2mc+c^2}{(1-m)mc} \\ &> \frac{(m-c)^2}{(1-m)mc} \\ &> 0. \end{aligned}$$

Note that Φ^* is an atomless pricing policy.

τ^* is a well-defined probability measure over \mathcal{F}^e . We are left to show that

$$r_0 = 1 - \frac{a}{c} - \frac{a}{m} \cdot (-\ln m + m - 1) \geq 0.$$

Since $a = \frac{m}{1-\ln c}$, this is equivalent to showing that

$$2 - \frac{m}{c} - \ln\left(\frac{c}{m}\right) \geq m. \quad (5)$$

Consider the function

$$i(x) = 2 - \frac{1}{x} - \ln x.$$

Since

$$i'(x) = \frac{1}{x^2} - \frac{1}{x} > 0$$

for any $x \in (0, 1)$, $\lim_{x \rightarrow 0} i(x) = -\infty$, and $i(1) = 1$, there is a unique $\tilde{c} \in (0, 1)$ such that $i(\tilde{c}) = m$. Since $i(x)$ is strictly increasing on $(0, 1)$, to show (5), it suffices to show that $\frac{c}{m} \geq \tilde{c}$, or equivalently, $c \geq m\tilde{c}$.

Recall that

$$2c - \frac{m + c^2}{c} \cdot \ln c = (1 + m)(1 - \ln m).$$

In what follows, we shall show that

$$\begin{aligned} & 2m\tilde{c} - \frac{m + m^2\tilde{c}^2}{m\tilde{c}} \cdot \ln(m\tilde{c}) \\ &= 2m\tilde{c} - \left(\frac{1}{\tilde{c}} + m\tilde{c}\right) \cdot (\ln m + \ln \tilde{c}) \\ &> (1 + m)(1 - \ln m). \end{aligned} \tag{6}$$

Since the function $2x - \frac{m+x^2}{x} \cdot \ln x$ is strictly decreasing on $(0, m)$, $c \in (0, m)$, and $0 < m\tilde{c} < m$, we conclude that $c > m\tilde{c}$.

We consider the following function:

$$j(x) = 2xi(x) - \left(\frac{1}{x} + xi(x)\right) \cdot (\ln x + \ln i(x)) - (1 + i(x)) \cdot (1 - \ln i(x)).$$

Since $i(1) = 1$, we have $j(1) = 0$. Furthermore, we have

$$j(\tilde{c}) = 2m\tilde{c} - \left(\frac{1}{\tilde{c}} + m\tilde{c}\right) \cdot (\ln m + \ln \tilde{c}) - (1 + m)(1 - \ln m).$$

Therefore, to show (6), or equivalently, $j(\tilde{c}) > j(1)$, it suffices to show that $j'(x) < 0$ for all $x \in [\tilde{c}, 1)$. This step is purely algebraic, and presented below:

$$\begin{aligned} j'(x) &= 2(xi(x))' \\ &- \left(-\frac{1}{x^2} + (xi(x))'\right) \cdot (\ln x + \ln i(x)) - \left(\frac{1}{x} + xi(x)\right) \cdot \left(\frac{1}{x} + \frac{i'(x)}{i(x)}\right) \\ &- i'(x) \cdot (1 - \ln i(x)) + (1 + i(x)) \cdot \frac{i'(x)}{i(x)}, \\ &= \frac{1}{x} - \frac{1}{x^2} - \ln x \\ &+ \left(\frac{1}{x^2} - 1 + \ln x\right) \cdot \ln x \\ &+ \left(\frac{1}{x^2} - 1 + \ln x + i'(x)\right) \cdot \ln i(x) \\ &+ (1 - x) \cdot \left(i(x) - \frac{1}{x}\right) \cdot \frac{i'(x)}{i(x)} \\ &- i'(x) \\ &< 0 \end{aligned}$$

for any $x \in [\tilde{c}, 1)$, where we substitute $(xi(x))' = 1 - \ln x$ to derive the second equality and the inequality follows from the following observations:

- (1) the function $\frac{1}{x} - \frac{1}{x^2} - \ln x$ is strictly increasing on $(0, 1)$ and equals 0 when $x = 1$,
- (2) the function $\frac{1}{x^2} - 1 + \ln x$ is strictly decreasing on $(0, 1)$ and equals 0 when $x = 1$,
- (3) $0 < m = i(\tilde{c}) \leq i(x) < i(1) = 1 < \frac{1}{x}$ for any $x \in [\tilde{c}, 1)$.
- (4) $i'(x) > 0$ for any $x \in (0, 1)$.

$x \geq \frac{m-x}{1-x}$ **for any** $x \in [c, m)$. Since $g(x) = x(2-x)$ is strictly increasing on $[0, m]$,

$$x(2-x) \geq c(2-c) > \hat{c}(2-\hat{c}) = m$$

for any $x \in [c, m)$. It follows that $x > \frac{m-x}{1-x}$ for any $x \in [c, m)$.

References

- ARIELI, I., Y. BABICHENKO, I. TALGAM-COHEN, AND K. ZABARNYI (2023): “A random dictator is all you need,” *arXiv preprint arXiv:2302.03667*.
- BERGEMANN, D., AND K. H. SCHLAG (2005): “Robust monopoly pricing: The case of regret,” Working paper.
- (2008): “Pricing without priors,” *Journal of the European Economic Association*, 6(2-3), 560–569.
- (2011): “Robust monopoly pricing,” *Journal of Economic Theory*, 146(6), 2527–2543.
- CARRASCO, V., V. F. LUZ, N. KOS, M. MESSNER, P. MONTEIRO, AND H. MOREIRA (2018): “Optimal selling mechanisms under moment conditions,” *Journal of Economic Theory*, 177, 245–279.
- GUO, Y., AND E. SHMAYA (2023a): “Regret-minimizing project choice,” *Econometrica*, 91(5), 1567–1593.
- (2023b): “Robust monopoly regulation,” Working paper.
- WINKLER, G. (1988): “Extreme points of moment sets,” *Mathematics of Operations Research*, 13(4), 581–587.
- ZHANG, W. (2022): “Auctioning multiple goods without priors,” Working paper.