ELSEVIER

Notes

Robust monopoly pricing ☆

Dirk Bergemann [a,*], Karl Schlag [b]

[a] *Department of Economics, Yale University, 30, Hillhouse Avenue, New Haven, CT 06511, USA*
[b] *Department of Economics, University of Vienna, 1010 Vienna, Austria*

Abstract

We consider a robust version of the classic problem of optimal monopoly pricing with incomplete information. In the robust version, the seller faces model uncertainty and only knows that the true demand distribution is in the neighborhood of a given model distribution. We characterize the pricing policies under two distinct decision criteria with multiple priors: (i) maximin utility and (ii) minimax regret. The equilibrium price under either criterion is lower then in the absence of uncertainty. The concern for robustness leads the seller to concede a larger information rent to all buyers with values below the optimal price without uncertainty.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

In the past decade, the theory of mechanism design has found increasingly widespread applications in the real world, favored partly by the growth of the electronic marketplace and trading on the Internet. Many trading platforms, such as auctions and exchanges, implement key insights

of the theoretical literature. With an increase in the use of optimal mechanisms, the robustness of these mechanisms with respect to the model specification becomes an important issue. In this note, we investigate a robust version of the classic monopoly problem under incomplete information. The determination of the optimal monopoly price is the most elementary instance of a revenue maximization problem in mechanism design.

We analyze the robustness of the optimal selling policy by enriching the standard model to account for model uncertainty. In the classic model, the valuation of the buyer is drawn from a given prior distribution. In contrast, in the robust version, the seller only knows that the true distribution is in the neighborhood of a given model distribution. The size of the neighborhood represents the extent of the model uncertainty faced by the seller.

The optimal pricing policy of the seller in the presence of model uncertainty is an instance of decision-making with multiple priors. We therefore build on the axiomatic decision theory with multiple priors and obtain interesting new insights for monopoly pricing. The methodological insight is that robustness can be guaranteed by considering decision making under multiple priors. The strategic insight is that we are able predict how an increase in uncertainty effects the pricing policy by using exclusively the data of the model distribution.

There are two leading approaches to incorporate multiple priors into axiomatic decision making: maximin utility and minimax regret. In the maximin utility approach with multiple priors, due to Gilboa and Schmeidler [11], the decision maker evaluates each action by its minimum utility across all priors. The decision maker selects the action that maximizes the minimum utility. The minimax regret approach was axiomatized by Milnor [18] and recently adapted to multiple priors by Hayashi [13] and Stoye [24]. Here, the decision maker evaluates foregone opportunities using regret and chooses an action that minimizes the maximum expected regret among the set of priors.

The analysis of the optimal pricing under the two decision criteria reveals that either criterion leads to a robust policy in the sense of statistical decision theory. A family of policies, indexed by the size of the uncertainty, is said to be *robust*, if for any demand sufficiently close to the model distribution, the difference between the expected profit under the optimal policy for this demand and the expected profit under the candidate policy is arbitrarily small. While the optimal policies under *maximin utility* and *minimax regret* share the robustness property, the exact response to the uncertainty leads to distinct qualitative features under these two criteria.

The pricing policy of the seller is obtained as the equilibrium strategy of a zero-sum game between the seller and adversarial nature. In this construction, nature selects a *least favorable demand* given the objective of the seller. The choice by nature attempts to exploit the sensitivity of the objective function of the seller to the information regarding the demand. In consequence, the strategy of the seller is to minimize the sensitivity of his objective function with respect to the demand information. The sensitivity of the objective function to the private information shapes the equilibrium under either criterion. The central role of the information sensitivity is most immediate in the case of the maximin utility criterion, where the seller maximizes the minimal profit across a set of demand distributions. Consider for a moment the profit function of the seller at a candidate price p. The expected profit depends on the distribution of valuations *only* through the upper cumulative probability at price p, namely the probability that the valuation of the buyer is equal to or exceeds p. In particular, any variation of the distribution function which does not affect the upper cumulative probability at p, does not affect the value of the profit function. Given the sensitivity of the profit function, nature then seeks to minimize the upper cumulative probability. In turn, the seller minimizes the sensitivity to the information by choosing his optimal price as if nature would choose the lowest possible upper cumulative prob-

ability in the neighborhood of the model distribution. In consequence the equilibrium choice of the seller always consists in lowering her price relative to the optimal price in the absence of uncertainty.

The logic of the equilibrium is identical under the minimax regret criterion, the modifications that arise are due to the distinct informational sensitivity of the objective function. When we consider the minimax regret criterion, the notion of regret modifies the trade-off for the seller and for nature. The regret of the seller is the difference between the realized valuation of the buyer and the realized profit obtained by the seller. The regret of the seller can therefore be positive for two reasons: (i) a buyer has a low valuation relative to the price and hence fails to purchase the object, or (ii) he has a high valuation relative to the price and hence the seller could have realized a higher price. In turn, the expected regret is the difference between the expected valuation of the buyer and the expected profit, where the expected valuation represents the natural upper bound on the profits of the seller. Given this additive form of the objective function, and given a candidate price p, the expected regret therefore depends on the mean of the valuation and the upper cumulative probability at the candidate price p. The later element appears as in the maximin criterion, but the sensitivity to the mean of the demand distribution newly appears in the regret minimization problem. In equilibrium, the pricing policy of the seller has to minimize the exposure to these two different statistics of the demand distribution simultaneously. In particular, if the seller were to concern herself exclusively with the upper cumulative probability, and hence as in the profit maximization lower the price too much, then nature would take advantage by increasing the mean of the valuation and hence increase the regret from this new, second, source. The seller resolves the conflict between these two statistics by a random pricing policy which offers trades at a range of prices. The range of the prices, i.e. the support of the equilibrium price distribution is chosen so that the expected regret is equalized across all prices, and the frequency of the prices is chosen such that no other demand distribution can lead to a larger regret. The resulting randomized pricing policy still has the feature that, relative to the optimal price in the absence of uncertainty, the expected price paid by almost all buyers with valuations within the support of the mixed pricing policy decreases when uncertainty increases. Yet, the upper segment of the buyers see higher prices with positive probability.

This brief description of the equilibrium policies emphasizes the common determinants of the policies under maximin utility and minimax regret, and traces the divergent aspects to differences in the objective functions. We will return to these differences and their axiomatic foundations in the final section. The common concern for robustness leads to many shared features in the equilibrium policies. First, and most importantly, the equilibrium price is lower (at least with positive probability) then it would be in the absence of uncertainty. With maximin utility, the hedging concern is so strong that the lower price is quoted with probability one. With minimax regret, the hedging concern leads to a range of offers, below and above, the price in the absence of uncertainty. Second, in terms of the information rent, the concern for robustness leads the seller to concede a larger information rent to all buyers with value below the optimal price without uncertainty. In the conclusion we discuss the extent to which these arguments may carry over to more general mechanism design settings.

We conclude the introduction with a brief discussion of the directly related literature. A recent paper by Bose, Ozdenoren, and Pape [5] determines the optimal auction in the presence of an uncertainty averse seller and bidders. Lopomo, Rigotti, and Shannon [17] consider a general mechanism design setting when the agents, but not the principal, have incomplete preferences due to Knightian uncertainty. In related work, Bergemann and Schlag [3] consider the optimal monopoly problem under regret without *any* priors. There, the analysis is concerned with optimal

policies in the absence of information rather than robustness and responsiveness to uncertainty as in the current contribution. Linhart and Radner [15] analyzed bilateral trade under minimax regret. A related notion of regret was considered by Engelbrecht-Wiggans [8] in the context of auctions, and recently, Engelbrecht-Wiggans and Katok [9] and Filiz-Ozbay and Ozbay [10] present experimental evidence indicating concern for regret in first price auctions. In a complete information environment, Renou and Schlag [21] use minimax regret to analyze strategic uncertainty.

2. Model

Monopoly. The seller faces a single potential buyer with value $v \in [0, 1]$ for a unit of the object. The value v is private information to the buyer and unknown to the seller. The buyer wishes to buy at most one unit of the object. The marginal cost of production is constant and normalized to zero. The net utility of the buyer with value v of purchasing a unit of the object at price p is $v - p$. The *profit* of selling a unit of the object at a deterministic price $p \in \mathbb{R}_+$ if the valuation of the buyer is v is:

$$\pi(p, v) \triangleq p\mathbb{I}_{\{v \geqslant p\}},$$

where $\mathbb{I}_{\{v \geqslant p\}}$ is the indicator function specifying:

$$\mathbb{I}_{\{v \geqslant p\}} = \begin{cases} 0, & \text{if } v < p, \\ 1, & \text{if } v \geqslant p. \end{cases}$$

By extension, if the seller chooses a randomized pricing policy, represented by a probability distribution $\Phi \in \Delta\mathbb{R}_+$, then the expected profit when facing a buyer with value v equals:

$$\pi(\Phi, v) \triangleq \int \pi(p, v) \, d\Phi(p).$$

In the standard version of the monopoly with incomplete information, the seller maximizes the expected profit for a given prior F over valuations. For a given distribution F and deterministic price p the expected profit is:

$$\pi(p, F) \triangleq \int \pi(p, v) \, dF(v).$$

We note that the demand generated by the distribution F can either represent a single large buyer or many small buyers. Here we phrase the results in terms of a single large buyer, but the results generalize naturally to the case of many small buyers. With a random pricing policy Φ, the expected profit is given by:

$$\pi(\Phi, F) \triangleq \int \int \pi(p, v) \, dF(v) \, d\Phi(p).$$

A random pricing policy $\Phi^*(F)$ that maximizes the profit for a given distribution F solves:

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

A well-known result by [23] states that for every distribution F, there exists a deterministic price $p^*(F)$ that maximizes profits.

Uncertainty. In the robust version the seller faces uncertainty (or ambiguity) in the sense of [7]. The uncertainty is represented by a set of possible distributions. The set is described by a model

distribution F_0 and includes all distributions in a neighborhood of size ε of the model distribution F_0. The magnitude of the uncertainty is quantified by the size of the neighborhood around the model distribution. Given the model distribution F_0 we denote by p_0 a profit maximizing price at F_0: $p_0 \triangleq p^*(F_0)$. For the remainder of the paper we shall assume that at the model distribution F_0: (i) p_0 is the unique maximizer of the profit function $\pi(p, F_0)$ and (ii) the density f_0 is continuously differentiable near p_0. These regularity assumptions enable us to use the implicit function theorem for the local analysis.

We consider two different decision criteria that allow for multiple priors: maximin utility and minimax regret. In either approach, the unknown state of the world is identified with the value v of the buyer.

Neighborhoods. We consider the neighborhoods induced by the Prohorov metric, the standard metric in robust statistical decision theory (see [14]). Given the model distribution F_0, the ε neighborhood under the Prohorov metric, denoted by $\mathcal{P}_\varepsilon(F_0)$, is:

$$\mathcal{P}_\varepsilon(F_0) \triangleq \{F \mid F(A) \leqslant F_0(A^\varepsilon) + \varepsilon, \ \forall \text{ measurable } A \subseteq [0,1]\}, \tag{1}$$

where the set A^ε denotes the closed ε neighborhood of any measurable set A:

$$A^\varepsilon \triangleq \left\{x \in [0,1] \,\Big|\, \min_{y \in A} d(x,y) \leqslant \varepsilon\right\},$$

where $d(x,y) = |x - y|$ is the distance on the real line. We shall use the language of small neighborhood and ε-neighborhood in the following interchangeably.

The Prohorov metric has evidently two components. The additive term ε in (1) allows for a *small* probability of *large* changes in the valuations relative to the model distribution whereas the larger set A^ε permits *large* probabilities of *small* changes in the valuations. The Prohorov metric is a metric for weak convergence of probability measures. In the context of our demand model, the Prohorov metric gives a literal description of the two relevant sources of model uncertainty. With a large probability, the seller could misperceive the willingness to pay by a small margin, and with a small probability, the seller could be mistaken about the market parameters by a large margin.

Maximin utility. Under maximin utility, the seller maximizes the minimum utility, where the utility of the seller is simply the profit, by searching for

$$\Phi_m \in \arg\max_{\Phi \in \Delta\mathbb{R}_+} \ \min_{F \in \mathcal{P}_\varepsilon(F_0)} \pi(\Phi, F).$$

Accordingly, we say that Φ_m *attains maximin utility*. We refer to F_m as a *least favorable demand* (for maximin utility) if

$$F_m \in \arg\min_{F \in \mathcal{P}_\varepsilon(F_0)} \ \max_{\Phi \in \Delta\mathbb{R}_+} \pi(\Phi, F).$$

The least favorable demand F_m minimizes across all profit maximizing pricing policies. Occasionally, it is useful to explicitly state the dependence of the optimal policies Φ_m and F_m on the size ε of the neighborhood, in which case we write $\Phi_{m,\varepsilon}$ and $F_{m,\varepsilon}$.

Minimax regret. The *regret* of the monopolist at a given price p and valuation v is:

$$r(p,v) \triangleq v - p\mathbb{I}_{\{v \geqslant p\}} = v - \pi(p,v). \tag{2}$$

The regret of the monopolist charging price p facing a buyer with value v is the difference between the profit the monopolist could make if she were to know the value v of the buyer before setting her price and the profit she makes without this information. The regret is non-negative and can only vanish if $p = v$. The regret of the monopolist is strictly positive in either of two cases: (i) the value v exceeds the price p, the indicator function is $\mathbb{I}_{\{v \geqslant p\}} = 1$; or (ii) the value v is below the price p, the indicator function is $\mathbb{I}_{\{v \geqslant p\}} = 0$.

The expected regret of a random pricing policy Φ given a demand distribution F is:

$$r(\Phi, F) \triangleq \iint r(p, v) \, d\Phi(p) \, dF(v) = \int v \, dF(v) - \int \pi(p, F) \, d\Phi(p). \tag{3}$$

In the final expression of the expected regret in (3), we see that the expected regret is, as mentioned in the introduction, the difference between the expected valuation and the expected profit. It follows that the probabilistic pricing policy Φ is profit maximizing at F if and only if Φ minimizes (expected) regret when facing F. The pricing policy Φ_r attains *minimax regret* if it minimizes the maximum regret over all distributions F in the neighborhood of a model distribution F_0:

$$\Phi_r \in \arg\min_{\Phi \in \Delta \mathbb{R}_+} \max_{F \in \mathcal{P}_\varepsilon(F_0)} r(\Phi, F).$$

F_r is called a *least favorable demand* if

$$F_r \in \arg\min_{F \in \mathcal{P}_\varepsilon(F_0)} \max_{\Phi \in \Delta \mathbb{R}_+} r(\Phi, F) = \arg\max_{F \in \mathcal{P}_\varepsilon(F_0)} \left(\int v \, dF(v) - \max_{\Phi} \pi(\Phi, F) \right).$$

Thus, a least favorable demand maximizes the regret of a profit maximizing seller who knows the true demand. While the regret criterion seems to relate to foregone opportunities when the information is revealed ex post, this particular interpretation is solely an additional feature of the minimax regret model. In particular, the decision maker does not need the information to become available ex post to evaluate his expected regret.[1]

Robust policy. For a given model distribution F_0, we define a robust family of random pricing policies, $\{\Phi_\varepsilon\}_{\varepsilon > 0}$, which are indexed by the size of the neighborhood ε as follows.

Definition 1 *(Robust pricing policy).* A family of pricing policies $\{\Phi_\varepsilon\}_{\varepsilon > 0}$ is called *robust* if, for each $\gamma > 0$, there is $\varepsilon > 0$ such that $F \in \mathcal{P}_\varepsilon(F_0) \Rightarrow \pi(\Phi^*(F), F) - \pi(\Phi_\varepsilon, F) < \gamma$.

The above notion requires that for every, arbitrarily small, upper bound γ, on the difference in the profits between the optimal policy $\Phi^*(F)$ without uncertainty and an element of the robust family of policies $\{\Phi_\varepsilon\}$, we can find a sufficiently small neighborhood ε so that the robust policy Φ_ε meets the upper bound γ for all distributions in the neighborhood. An ideal candidate for a robust policy is the optimal policy $\Phi^*(F)$ itself. In other words, we would require that for each $\gamma > 0$, there is $\varepsilon > 0$ such that:

$$F \in \mathcal{P}_\varepsilon(F_0) \quad \Rightarrow \quad \pi(\Phi^*(F), F) - \pi(\Phi^*(F_0), F) < \gamma. \tag{4}$$

This notion of robustness, applied directly to the optimal policy $\Phi^*(F)$, constitutes the definition of α robustness in [20] where it is shown that the profit maximizing price in the optimal

---

[1] The axiomatic approach is distinct from the ex-post measure of regret due to Hannan [12] in the context of repeated games and from the behavioral approaches to regret due to Bell [2] and Loomes and Sugden [16].

monopoly problem is not robust to model misspecification.[2] One of the objectives here is to identify robust policies by considering decision making under multiple priors that do not suffer from such discontinuity in the profits.

3. Maximin utility

We consider a monopolist who maximize the minimum profit for all distributions in the neighborhood of the model distribution F_0. The pricing rule that attains maximin utility is the equilibrium strategy in a game between the seller and adversarial nature. The seller chooses a probabilistic pricing policy, a distribution $\Phi \in \Delta\mathbb{R}_+$, and nature chooses a demand distribution $F \in \mathcal{P}_\varepsilon(F_0)$. In this game, the payoff of the seller is the expected profit while the payoff of nature is the negative of the expected profit. A Nash equilibrium of this zero-sum game is a solution (Φ_m, F_m) to the saddle point problem:

$$\pi(\Phi, F_m) \leqslant \pi(\Phi_m, F_m) \leqslant \pi(\Phi_m, F), \quad \forall \Phi \in \Delta\mathbb{R}_+, \ \forall F \in \mathcal{P}_\varepsilon(F_0). \tag{SP_m}$$

The objective of adversarial nature is to lower the expected profit of the seller. For a given price p, the expected profit of the seller is

$$\pi(p, F) = \int \pi(p, v)dF(v) = p(1 - F(p)).$$

The profit minimizing demand, given p, is then achieved by decreasing the cumulative probability of valuations equal or larger than p by as much as possible within the neighborhood $\mathcal{P}_\varepsilon(F_0)$. The profit minimizing demand thus minimizes the probability of sale, the upper cumulative probability $1 - F(p)$. Given the model distribution F_0 and the size ε of the neighborhood, there is a unique distribution, which minimizes the probability, $1 - F(p)$, for all p in the unit interval *simultaneously*. We obtain this least favorable demand explicitly by shifting the probabilities as far down as possible, given the constraints imposed by the model distribution F_0 and the size ε of the neighborhood. We shift, for every v, the cumulative probability of the model distribution F_0 at the point $v + \varepsilon$ downwards to be the cumulative probability at the point v. In addition, we transfer the very highest valuations with probability ε to the lowest valuation, namely $v = 0$. This results in the distribution $F_{m,\varepsilon}$ given by:

$$F_{m,\varepsilon}(v) \triangleq \min\{F_0(v + \varepsilon) + \varepsilon, 1\}, \tag{5}$$

that is within the ε neighborhood of F_0. The first shift, generated by $v + \varepsilon$, represents small changes in valuations that occur with large probability. The second shift, generated by $F_0(\cdot) + \varepsilon$, represents large changes that occur with small probability. It is easily verified that $F_{m,\varepsilon}$ is a profit minimizing demand for *any* price p given the constraint imposed by the size of the neighborhood. In other words, the profit minimizing demand does not depend on the, possibly probabilistic, price p of the seller. Given that the profit minimizing demand $F_{m,\varepsilon}$ does not depend on the offered prices, the monopolist acts as if the demand is given by $F_{m,\varepsilon}$. In consequence, the seller maximizes profits at $F_{m,\varepsilon}$ by choosing a deterministic price $p_{m,\varepsilon}$ where $p_{m,\varepsilon} \triangleq p^*(F_{m,\varepsilon})$.

---

[2] The non-robustness is demonstrated in [20] by the following example: consider a Dirac distribution which puts probability one on valuation v. The optimal monopoly price p is equal to v. This policy is not robust to model misspecification, since if the true model puts probability one on a value arbitrarily close, but strictly below v, then the revenue is 0 rather than v.

Proposition 1 (*Maximin utility*). *For every $\varepsilon > 0$, there exists a pair $(p_{m,\varepsilon}, F_{m,\varepsilon})$, such that $p_{m,\varepsilon} \in [0, 1]$ attains maximin utility and $F_{m,\varepsilon}$ is a least favorable demand.*

An important aspect of the above result is that the construction of the profit minimizing demand does not require a local argument, and hence the above equilibrium characterization is valid for arbitrarily large neighborhoods around the model distribution. The construction of the least favorite demand, given by (5), also reveals that $F_{m,\varepsilon}$ is first-order stochastically dominated by any other distribution in the neighborhood $\mathcal{P}_\varepsilon(F_0)$. By this constructive argument, the result of Proposition 1 then extends to any notion of neighborhood (and/or generating metric) which forms a lattice (strictly speaking, we only need the semi-lattice property) with respect to first order stochastic dominance.[3] The optimal pricing policy of the seller can easily be recast as canonical mechanism design problem, using the language of virtual utility, as shown by Myerson [19] and Bulow and Roberts [6]. By using the incentive constraints to replace the transfers, the maximization problem of the seller can be represented as:

$$\int \left(\max_{x \in [0,1]} x(v) \left(v - \frac{1 - F(v)}{f(v)} \right) \right) f(v) \, dv. \tag{6}$$

For a given distribution F, the pointwise optimal solution $x^*(v) \in \{0, 1\}$ is to assign the object, $x^*(v) = 1$, if the virtual utility $v - (1 - F(v))/f(v)$ is positive, and to not assign the object, $x^*(v) = 0$, if it is negative. We can rewrite the above integral after disregarding the valuations which have negative virtual utility as they receive zero weight, $x^*(v) = 0$, in the optimal solution:

$$\int_{\{v \mid v - (1 - F(v))/f(v) \geqslant 0\}} \left(v f(v) - \left(1 - F(v) \right) \right) dv. \tag{7}$$

In this reformulation of the objective function of the seller, we see that the least favorable demand, as established in Proposition 1, depresses the mean valuation, conditional on allocating the good, and simultaneously depresses the information cost, or sensitivity to the private information, $1 - F(v)$. Thus, the least favorable demand generates the lowest feasible mean valuation, but the resulting allocation improves the *ex-post* efficiency as there are some intermediate types with willingness-to-pay v which will receive the object under $F_{m,\varepsilon}$, but would not receive it under any other distribution $F \in \mathcal{P}_\varepsilon(F_0)$.

How does the optimal price change with an increase in uncertainty? The rate of the change in the price depends on the curvature of the profit function at the model distribution F_0. By the assumption of concavity, we know that the curvature is negative. We can apply the implicit function theorem to the optimal price p_0 at the model distribution F_0 and obtain the following comparative static result.

Proposition 2 (*Pricing under maximin utility*). *The price $p_{m,\varepsilon}$ responds to an increase in uncertainty at $\varepsilon = 0$ by*:

$$\frac{dp_{m,\varepsilon}}{d\varepsilon} \Big|_{\varepsilon = 0} = -1 + \frac{1 - f_0(p_0)}{\partial \pi^2(p_0, F_0)/\partial p^2} < -\frac{1}{2}.$$

---

[3] We thank the editor for pointing out the relationship to the lattice property of the neighborhoods.

Accordingly, the price that attains maximin utility responds to an increase in uncertainty with a lower price. Marginally, this response is equal to -1 if the objective function is infinitely concave.

Consider now the profits realized by the price $p_{m,\varepsilon}$ – which attains maximin utility within the neighborhood $\mathcal{P}_\varepsilon(F_0)$ – at a given distribution $F \in \mathcal{P}_\varepsilon(F_0)$. By construction, these profits are at least as high as those obtained when facing the least favorable demand $F_{m,\varepsilon}$. We use the lower bound on the profits supported by F_m to show that the optimal profits are continuous in the demand distribution F. This implies that profits achieved by $p_{m,\varepsilon}$ when facing F are close to those achieved by $p^*(F)$ when facing F. The family of pricing rules that attain maximin utility thus qualify as robust.

Proposition 3 (*Robustness*). *The family of pricing policies $\{p_{m,\varepsilon}\}_{\varepsilon>0}$ is a robust family of pricing policies.*

4. Minimax regret

Next we consider the minimax regret problem of the seller, where a (probabilistic) pricing policy Φ_r and a least favorable demand F_r are the equilibrium policies of a zero-sum game. In this zero-sum game, the payoff of the seller is the negative of the regret while the payoff to nature is regret itself. A Nash equilibrium (Φ_r, F_r) is a solution to the saddle point problem:

$$r(\Phi_r, F) \leqslant r(\Phi_r, F_r) \leqslant r(\Phi, F_r), \quad \forall \Phi \in \Delta \mathbb{R}_+, \ \forall F \in \mathcal{P}_\varepsilon(F_0). \tag{SP$_r$}$$

The saddlepoint result permits us to link minimax regret behavior to payoff maximizing behavior under a prior as follows. If the minimax regret is derived from the equilibrium characterization in (SP$_r$) then any price chosen by a monopolist who minimizes maximal regret, is at the same time a price which maximizes expected profit against a particular demand, namely, the least favorable demand. In fact, the saddle point condition requires that Φ_r is a probabilistic price that maximizes profits given F_r and F_r is a regret maximizing demand given Φ_r.

In the equilibrium of the zero-sum game, the probabilistic price has to resolve the conflict between the regret which arises with low prices, against the regret associated with high prices. If she offers a low price, nature can cause regret with a distribution which puts substantial probability on high valuation buyers. On the other hand, if she offers a high price, nature can cause regret with a distribution which puts substantial probability at valuations just below the offered price. If we consider the formula of the expected regret at a deterministic price p, rather than a general random pricing policy Φ, as in (3),

$$r(p, F) = \int v \, dF(v) - \int \pi(p, v) \, dF(v) = \int v \, dF(v) - p\big(1 - F(p)\big) \tag{8}$$

we see this tension in terms of the expected valuation, the first term, and the expected profit, the second term. The first term is controlled by the mean valuation, whereas the second is controlled by the upper cumulative probability, $1 - F(p)$, the probability of a sale at price p. Now, the analysis of the maximin utility problem showed that a distribution which minimizes $1 - F(p)$ can be determined independently of p. The unique solution, given by the distribution $F_{m,\varepsilon}$ in (5) has the property that it is first order stochastically dominated by all other distributions in the neighborhood $\mathcal{P}_\varepsilon(F_0)$. An immediate consequence of the first order stochastic dominance is that the distribution $F_{m,\varepsilon}$ achieves the lowest mean valuation among all distributions of $F \in \mathcal{P}_\varepsilon(F_0)$. Now, as nature is seeking to maximize regret, the first term would suggest for nature to choose

the distribution with highest mean in $\mathcal{P}_\varepsilon(F_0)$, whereas the second term would suggest to choose the distribution with the lowest upper cumulative probability, namely $F_{m,\varepsilon}$, which by the above argument is also the distribution with the lowest mean. The relative importance of these two terms depends on the choice of the price p by the seller, and hence, in contrast to the case of maximin utility, the least favorable demand is the result of an equilibrium argument and cannot be constructed independently of the strategy of the seller (as it was the case with maximin utility). It also suggests that a deterministic pricing policy in the form of single price p exposes the seller to substantial regret, and that the seller can decrease her exposure by offering many prices in the form of a random pricing policy. We prove the existence of a solution to the saddlepoint problem (SP$_r$) using results from [22].

Proposition 4 (*Existence of minimax regret*). *A solution (Φ_r, F_r) to the saddlepoint condition* (SP$_r$) *exists.*

We should emphasize that the above existence result does not use local arguments, and establishes the existence of a Nash equilibrium for arbitrary neighborhoods, small or large. By contrast, the following explicit characterization of the equilibrium pricing strategy of the seller uses a local argument, namely the implicit function theorem, and hence is valid only for small neighborhoods.

The tension between the mean valuation and the upper cumulative probability changes the structure of the least favorable demand and the equilibrium pricing policy relative to the maximin utility analysis. Intuitively, nature seeks to accomplish two conflicting goals. First she tries to maximize the expected valuation, which represents the upper bound for the profit of the seller, and hence also the maximal value of regret, while, second, she attempts to minimize the expected profits. These goals are conflicting as an increase in the expected valuation ought to lead eventually to an increase in the surplus the seller can extract. Now for a given candidate price p, the expected profit only depends on the upper cumulative probability, $1 - F(p)$, at p. Now, to the extent that the upper cumulative probability is held constant at p (and so is p itself), nature would seek to maximize the expected valuation. But by the logic of the first order stochastic dominance, this means to maximize the upper cumulative probability everywhere but at p. This means, that in contrast to the least favorable demand, under maximin utility, the least favorable demand in the minimax regret will maximize rather than minimize $1 - F(p')$ for all $p' \neq p$. But this thought experiment suggests a discontinuity in the form of a downward jump of the upper cumulative probability to $1 - F(p)$ (or correspondingly an upward jump to $F(p)$) from the left, and a constant upper cumulative probability to the right of p, as long as permitted by the size of the neighborhood. But now observe, that if there were a flat segment in the probability distribution to the right of p, then the seller would a profitable deviation. By increasing the price from p to the largest price $p'' > p$ where the equality $F(p'') = F(p)$ would still prevail, the seller could raise his price without losing sales, clearly an improvement. It follows that in equilibrium, nature has to suspend the maximization of the upper cumulative probability precisely in the support of the prices offered by the seller, denoted by $[a, c]$ in the proposition below. In this interval, a real trade-off arises between the maximization of the expected valuation and the minimization of the profits. In particular, nature is attempting to maintain the prices offered sufficiently low and by the logic familiar from maximin utility this involves lowering the upper cumulative probability as much as feasible within the constraint imposed by the size of the neighborhood $\mathcal{P}_\varepsilon(F_0)$. The constraint on the choice set is going to be binding at some point $b \in [a, c]$, where nature cannot lower the upper probability, $1 - F(b)$, any further, and at this insensitive point b where

"undercutting" is infeasible, the random pricing strategy of the seller can place a single atom. Everywhere else, the random pricing strategy Φ_r keeps nature exactly indifferent with respect to local changes of the demand distribution.

Proposition 5 *(Minimax regret).*

1. *Given $\delta > 0$, for every ε sufficiently small, there exist a, b and c with $0 < a < b < c < 1$ and $p_0 - \delta < a < p_0 < c < p_0 + \delta$ such that a minimax regret probabilistic price Φ_r is given by:*

$$\Phi_r(p) = \begin{cases} 0 & \text{if } 0 \leqslant p < a, \\ \ln \frac{p}{a} & \text{if } a \leqslant p < b, \\ 1 - \ln \frac{c}{p} & \text{if } b \leqslant p \leqslant c, \\ 1 & \text{if } c < p \leqslant 1. \end{cases}$$

2. *The boundary points a, b and c respond to an increase in uncertainty at $\varepsilon = 0$:*
   (i) $\lim_{\varepsilon \to 0} a'(0) = -\infty$;
   (ii) $\lim_{\varepsilon \to 0} b'(0)$ *is finite; and*
   (iii) $\lim_{\varepsilon \to 0} c'(0) = \infty$.

The proof of Proposition 5 relies on a straightforward but lengthy application of the implicit function theorem and is provided in Proposition 5 of [4]. The least favorable demand makes the seller indifferent among all prices $p \in [a, c]$. As uncertainty increases, the interval over which the seller randomizes increases rapidly in order to protect against nature either undercutting or moving mass towards higher valuations. At the same time, the mass point b does not change drastically.

We now investigate the comparative static behavior of the random price $p_{r,\varepsilon}$ governed by the distribution $\Phi_{r,\varepsilon}$. The response of the expected price, $\mathbb{E}[p_{r,\varepsilon}]$, to a marginal increase in uncertainty can be explained by the first order effects. For a small level of uncertainty, we may represent the regret through a linear approximation $r^* = r_0 + \varepsilon \cdot \partial r^*/\partial \varepsilon$, where r_0 is the regret at the model distribution. The optimal response of the seller to an increase in uncertainty is now to find a probabilistic price which minimizes the additional regret $\varepsilon \cdot \partial r^*/\partial \varepsilon$ coming from the increase in uncertainty. Locally, the cost of moving the price away from the optimum is given by the second derivative of the objective function. With small uncertainty, the curvature of the regret is identical to the curvature of the profit function. The rate at which the minimax regret price responses to an increase in uncertainty is then simply the ratio of the response of the marginal regret to a change in price divided by the curvature of the profit function.

Corollary 1 *(Comparative statics with minimax regret).*

1. *The expected price $\mathbb{E}[p_{r,\varepsilon}]$ responds to an increase in uncertainty at $\varepsilon = 0$ by:*

$$\frac{d}{d\varepsilon}\mathbb{E}[p_{r,\varepsilon}]|_{\varepsilon=0} = \begin{cases} -1 - \frac{f_0(p_0)+1}{\partial \pi^2(p_0, F_0)/\partial p^2} > -1 & \text{if } p_0 \leqslant \frac{1}{2}, \\ -1 - \frac{f_0(p_0)-1}{\partial \pi^2(p_0, F_0)/\partial p^2} < -\frac{1}{2} & \text{if } p_0 > \frac{1}{2}. \end{cases} \tag{9}$$

2. *If ε is sufficiently small, then for any $v \in (a, c) \setminus b$,*

$$\frac{d}{d\varepsilon}\mathbb{E}[p_{r,\varepsilon} \mid p_{r,\varepsilon} \leqslant v] < 0.$$

By comparing Corollary 1 and Proposition 2, we find that the marginal response of the expected price $\mathbb{E}[p_{r,\varepsilon}]$ to an increase in uncertainty is identical under minimax regret and maximin profit for $p_0 > \frac{1}{2}$. In both cases, the (expected) equilibrium price is lower than under the model distribution F_0 without uncertainty. The difference between the two criteria arises at a low level of p_0 at which the seller is less aggressive in lowering her price under minimax regret. When the optimal price in the absence of uncertainty is low, $p_0 < \frac{1}{2}$, then the trade-off that nature faces in her two conflicting goals, namely to maximize the expected valuation, while, second, to minimize the expected profits, is resolved in favor of the former, namely to maximize the expected valuation. We discussed this trade-off above following Proposition 4. Now, if $p_0 < \frac{1}{2}$, and hence in the lower half of the support $[0, 1]$, an increase in the expected valuation is guaranteed to increase regret by more, namely $1 - p_0 > \frac{1}{2}$ than a lost sale which would increase regret only by $p_0 < \frac{1}{2}$. This explains the change in the derivative of the expected price $\mathbb{E}[p_{r,\varepsilon}]$ at the midpoint of the support $[0, 1]$. In fact, for the case of $p_0 < \frac{1}{2}$, it turns out that the expected price can be strictly increasing in ε. As nature finds it to her advantage to increase the mean value of the distribution, and hence also to increase the upper cumulative probability, the seller is responding to the increase in the demand distribution with a raise in the expected equilibrium price. But since the expected equilibrium price is close to $p_0 < \frac{1}{2}$, the resulting least favorable demand still leads to an increase in regret relative to the model distribution. For example, the increase in the price occurs if the model density is linear and strictly decreasing. This response of the equilibrium price to an increase in uncertainty stands in stark contrast to the maximin behavior where any increase in the size of the uncertainty has a downward effect on prices, regardless of the model distribution.

The change in the expected price, as given by Corollary 1(1) also represent the change in welfare to a buyer who purchases with probability one, or $v > c$. The impact of the uncertainty on a buyer whose purchase occurs with probability less than one, or $v < c$ is stated in Corollary 1(2). The derivative of the conditional price is defined everywhere but at b, where it has a discontinuity as the mass point at b changes its location with an increase in ε. It shows that the price *conditional* on a purchase decreases when the uncertainty increases. The decrease in the price *conditional on purchase* does not contradict the possible increase in the expected price. The increase in the unconditional price is driven by the changes in the support of Φ_r – in particular the increase of c – and increases (in the sense of first order stochastic dominance) of the unconditional distribution of prices. In [4], we established this result in terms of menu of prices, where we showed that for every type $v \in (a, c)$, the price paid per unit of the object is decreasing with an increase in uncertainty, see Proposition 7 in [4].[4]

We conclude by showing that the solution to the minimax regret problem also generates a robust family of policies in the sense of Definition 1.

Proposition 6 *(Robustness). If $\{\Phi_{r,\varepsilon}\}_{\varepsilon>0}$ attains minimax regret at F_0 for all sufficiently small ε, then $\{\Phi_{r,\varepsilon}\}_{\varepsilon>0}$ is a robust family of pricing policies.*

---

[4] In the menu representation, a buyer with willingness-to-pay v, pays a transfer $t_r(v)$ to receive the object with probability $q_r(v)$. The conditional expected price with random pricing here is equal to the per unit price in the menu representation there, or: $\mathbb{E}[p_{r,\varepsilon} \mid p_{r,\varepsilon} \leqslant v] = t_r(v)/q_r(v)$.

5. Discussion

We conclude by relating the pricing behavior of the seller in the incomplete information monopoly to the axiomatic foundations. Finally, we spell out how the insights from the specific model here might relate to more general models of mechanism design and uncertainty.

Axioms and behavioral implications. From an axiomatic perspective, the maximin utility and minimax regret criteria represent different departures from the standard model of [1] by allowing for multiple priors. The maximin utility criterion emerges when imposing the independence axiom only when mixing with constant actions. The minimax regret criterion allows the choice to be menu dependent by requiring independence of irrelevant alternatives only when the changes in the menu do not change the best outcome in any of the states (see [24]). Both criteria include an additional axiom to capture aversion to ambiguity by postulating that the decision maker will hedge against uncertainty by mixing whenever indifferent. This hedging leads the decision maker under either criterion to sometimes offer the object for sale at lower prices than he would at the model distribution F_0, but in the absence of uncertainty. Interestingly, the difference between maximin utility and minimax regret then arises in the strength of the hedging motive. In the absence of the axiom of independence of irrelevant alternatives, the choice under regret does depend on the opportunities, both in terms of missed sales and missed revenues, the respective downward and upward opportunities. In contrast, the maximin utility maintains the independence of irrelevant alternatives, and this leads the decision-maker to act as if the lowest distribution (in terms of first-order dominance) were the true distribution.

The choice of metric (and neighborhoods). We investigated the robust policies with respect to neighborhoods generated by the Prohorov metric. The question then arises as to how sensitive the results are to the choice of the metric. In particular, there are a number of other distances, such as the Levy metric or the bounded Lipschitz metric which also metrize the weak topology. Of course, these distances define different neighborhoods and hence different choice sets for nature. However, these distinct notions differ only in the support sets over which the distributions are evaluated. Therefore the comparative static results near $\varepsilon = 0$ are unaffected by the specific notion for the metric.

Beyond small neighborhoods. We analyzed the pricing policies when robustness is required with respect to small neighborhoods. But we required the assumption of small neighborhoods only in the use of the implicit function theorem for the comparative static results and in the explicit construction of the random pricing under minimax regret. In related work, Bergemann and Schlag [3], we considered the monopoly problem under minimax regret in the absence of *any* restrictions about the uncertainty, in other words, *very large* neighborhoods. The analysis there was notably easier as there were no constraints on the choice of strategy by nature. But interestingly, the distinct features of minimax regret strategy were preserved, namely the logarithmic distribution of the prices with a single mass point. This suggests that the intermediate case of large neighborhoods would support similar results for minimax regret. The associated analysis, however, is beyond the scope of this note, as it requires the use of general optimal control techniques to keep track of the multitude of constraints imposed by the neighborhood on the least favorable demand. But as we know that the minimax regret policy has to remain a random pricing policy, and as we have established the general form of the maximin policy, we already know that the distinct features of the minimax regret and maximin utility are preserved beyond small neighborhoods.

Beyond monopoly pricing. We analyzed robust policies in a simple mechanism design environment, namely the monopolistic sale of a single unit under incomplete information. The robust pricing policies displayed less sensitivity to private information and hedged against the uncertainty by offering sales at lower prices relative to the policy without uncertainty. We expect these insights to extend to mechanism design problems beyond the single good monopoly pricing as the logic of the equilibrium construction indicates. The common element of revenue maximizing mechanism, whether it pertains to single good, multi-unit goods (nonlinear pricing) or multi-person problems (auctions) is that the principal maximizes the virtual utility rather than the social utility. In all of the above problems, the virtual utility takes the form, $v - (1 - F(v))/f(v)$, where the later term represents the cost of private information to the seller. Now, we saw in the maximin utility that nature lowers the revenue by lowering the entire social surplus by minimizing the upper cumulative probability. But, as we saw then, this means that the virtual utility of each type v is increased, and hence that low types v will have positive virtual utility in the presence of uncertainty, whereas the would not have in the absence of uncertainty. It follows that agents with lower valuation v now receive the object or receive an assignment more generally. But as the participation constraint still binds for these types, it means that the prices will be lower, and overall the cost of the private information, represented by $(1 - F(v))/f(v)$ will be lower. It is this general aspect of the robust policy, namely lower prices through lower inverse hazard rates, that we expect to emerge in general allocation environments as well. In other words, the robust revenue maximizing policy is closer to the socially efficient allocation as the information cost, $(1 - F(v))/f(v)$, carries less weight, and the resulting prices are closer to the externality prices imposed by the efficient Groves mechanism. Given the prominent role of the notion of first-order stochastic dominance and the immediate link to the information cost $(1 - F(v))/f(v)$ in the argument presented here, it is conceivable that the present argument extends directly to the above mentioned, more general, allocation problems.

By contrast, in the minimax regret problem, we can expect the general downward trend of information cost and allocation policy to be attenuated relative to the maximin problem as adversarial nature attempts to maximize the potential for regret, namely the social value. But the exact determination of the robust policy appears to be much more difficult to establish in general environments, where we cannot exploit the specific structure of the allocation problem as in this note.

Appendix A

The appendix contains the proofs for the results in the main body of the text.

Proof of Proposition 1. As F_m is given by (5), we have that $\pi(p, F_m) \leqslant \pi(p, F)$ for all $F \in \mathcal{P}_\varepsilon(F_0)$. On the other hand, if $p_m = p^*(F_m)$, then $\pi(p_m, F_m) \geqslant \pi(p, F_m)$ holds for all p by definition of p_m. Jointly this implies that (p_m, F_m) is a saddle point of (SP$_m$) and p_m attains maximin payoff and F_m is a least favorable demand.   □

Proof of Proposition 2. For sufficiently small ε our assumptions on F_0 imply that F_m is differentiable near p_m. Since p_m is optimal given demand F_m, we find that p_m satisfies the associated first order conditions:

$$\frac{d}{dp}\big(p\big(1 - F_m(p)\big)\big)\big|_{p=p_m} = 0.$$

The earlier strict concavity assumption on $\pi(p, F_0)$ implies that we can apply the implicit function theorem at $\varepsilon = 0$ to the above equation to obtain

$$\left.\frac{dp_m}{d\varepsilon}\right|_{\varepsilon=0} = -1 + \frac{1 - f_0(p_0)}{-2 f_0(p_0) - p_0 f_0'(p_0)} = \frac{f_0(p_0) + p_0 f_0'(p_0) + 1}{-2 f_0(p_0) - p_0 f_0'(p_0)}.$$

Since $-2 f_0(p_0) - p_0 f_0'(p_0) < 0$, we observe that the l.h.s. of the above equation as a function of $f_0(p_0)$ is increasing in $f_0(p_0)$ and hence by taking the limit as $f_0(p_0)$ tends to infinity it follows that this expression is bounded above by $-1/2$. $\quad\square$

Proof of Proposition 3. We show that for any $\gamma > 0$, there exists $\varepsilon > 0$ such that $F \in \mathcal{P}_\varepsilon(F_0)$ implies $\pi(p^*(F), F) - \pi(p_m, F) < \gamma$. Note that $\pi(p_m, F) \geqslant \pi(p_m, F_m)$ and thus $\pi(p^*(F), F) - \pi(p_m, F) \leqslant \pi(p^*(F), F) - \pi(p_m, F_m)$. Since $\pi(p_m, F_m) = \pi(p^*(F_m), F_m)$ the proof is complete once we show that $\pi(p^*(F), F)$ is a continuous function of F with respect to the weak topology. Consider F, G such that $G \in \mathcal{P}_\varepsilon(F)$. Using the fact that $G(p) \leqslant F(p + \varepsilon) + \varepsilon$, we obtain

$$\pi\big(p^*(G), G\big) \geqslant \pi\big(p^*(F) - \varepsilon, G\big) = \big(p^*(F) - \varepsilon\big)\big(1 - G\big(p^*(F) - \varepsilon\big)\big)$$
$$\geqslant \big(p^*(F) - \varepsilon\big)\big(1 - F\big(p^*(F)\big) - \varepsilon\big) \geqslant \pi\big(p^*(F), F\big) - 2\varepsilon.$$

Since the Prohorov norm is symmetric and thus $F \in \mathcal{P}_\varepsilon(G)$, it follows that

$$\pi\big(p^*(F), F\big) + 2\varepsilon \geqslant \pi\big(p^*(G), G\big) \geqslant \pi\big(p^*(F), F\big) - 2\varepsilon,$$

and hence we have proven that $\pi(p^*(F), F)$ is continuous in F. $\quad\square$

Proof of Proposition 4. We apply Corollary 5.2 in [22] to show that a saddle point exists. For this we need to verify that the zero-sum game between the seller and nature is a compact Hausdorff game for which the mixed extension is both reciprocally upper semi continuous and payoff secure.

Clearly we have a compact Hausdorff game. Reciprocal upper semi continuity follows directly as we are investigating a zero-sum game. So all we have to ensure is *payoff security*. Payoff security for the monopolist means that we have to show for each (F_r, Φ_r) with $F_r \in \mathcal{P}_\varepsilon(F_0)$ and for every $\delta > 0$ that there exists $\gamma > 0$ and $\overline{\Phi}$ such that $F \in \mathcal{P}_\gamma(F_r)$ implies $r(\overline{\Phi}, F) \leqslant r(\Phi_r, F_r) + \delta$.

Let $\gamma \triangleq \delta/4$ and let $\overline{\Phi}$ be such that $\overline{\Phi}(p) \triangleq \Phi_r(p + \gamma)$. Then using the fact that $F(v) \geqslant F_r(v - \gamma) - \gamma$ we obtain

$$\int_0^1 v\, dF(v) \leqslant 2\gamma + \int_0^1 v\, dF_r(v).$$

Using the fact that $F(v) \leqslant F_r(v + \gamma) + \gamma$ we obtain

$$\pi(\overline{\Phi}, F) \geqslant \pi\big(\Phi_r(p + \gamma), \min\{F_r(v + \gamma) + \gamma, 1\}\big) \geqslant \pi(\Phi_r, F_r) - 2\gamma,$$

and hence $r(\overline{\Phi}, F) \leqslant r(\Phi_r, F_r) + \delta$. To show payoff security for nature we have to show for each (Φ_r, F_r) with $F_r \in \mathcal{P}_\varepsilon(F_0)$ and for every $\delta > 0$ that there exists $\gamma > 0$ and $\overline{F} \in \mathcal{P}_\varepsilon(F_0)$ such that $\Phi \in \mathcal{P}_\gamma(\Phi_r)$ implies $r(\Phi, \overline{F}) \geqslant r(\Phi_r, F_r) - \delta$.

Here we set $\overline{F} \triangleq F_r$. Given $\gamma > 0$ consider any $\Phi \in \mathcal{P}_\gamma(\Phi_r)$. All we have to show is that $\pi(\Phi, F_r) \leqslant \pi(\Phi_r, F_r) + \delta$ for sufficiently small γ. Note that $\Phi(p) \leqslant \Phi_r(p + \gamma) + \gamma$ implies

$$\pi(\Phi, F_r) \leqslant \gamma + \int (p+\gamma) \int_p^1 dF_r(v) d\Phi_r(p+\gamma) = \gamma + \int p \int_{p-\gamma}^1 dF_r(v) d\Phi_r(p)$$

$$= \gamma + \pi(\Phi_r, F_r) + \int p \int_{[p-\gamma,p)} dF_r(v) d\Phi_r(p)$$

$$\leqslant \gamma + \pi(\Phi_r, F_r) + \int \int_{[p-\gamma,p)} dF_r(v) d\Phi_r(p).$$

Given the continuity of the last integral term above in the boundary point γ, the claim is established. $\square$

Proof of Proposition 6. Assume that Φ_r attains minimax regret but is not robust. So there exists $\gamma > 0$, such that for all $\varepsilon > 0$, there exists F_ε such that $F_\varepsilon \in \mathcal{P}_\varepsilon(F_0)$ but

$$\pi\big(p^*(F_\varepsilon), F_\varepsilon\big) - \pi(\Phi_r, F_\varepsilon) \geqslant \gamma. \tag{10}$$

Assume that (Φ_r, F_r) is a saddle point of the regret problem (SP$_r$) given $\varepsilon > 0$. Then $\pi(\Phi_r, F_r) = \pi(p^*(F_r), F_r)$ and we can rewrite the left-hand side of (10) as follows:

$$\pi\big(p^*(F_\varepsilon), F_\varepsilon\big) - \pi(\Phi_r, F_\varepsilon) = \pi\big(p^*(F_\varepsilon), F_\varepsilon\big) - \pi\big(p^*(F_r), F_r\big)$$
$$+ \pi(\Phi_r, F_r) - \pi(\Phi_r, F_\varepsilon). \tag{11}$$

Using (SP$_r$) we also obtain

$$0 \leqslant r(\Phi_r, F_r) - r(\Phi_r, F_\varepsilon) = \int v \, dF_r(v) - \int v \, dF_\varepsilon(v) + \pi(\Phi_r, F_\varepsilon) - \pi(\Phi_r, F_r),$$

so that:

$$\pi(\Phi_r, F_r) - \pi(\Phi_r, F_\varepsilon) \leqslant \int v \, dF_r(v) - \int v \, dF_\varepsilon(v).$$

Entering this into (11) we obtain from (10) that:

$$\pi\big(p^*(F_\varepsilon), F_\varepsilon\big) - \pi\big(p^*(F_r), F_r\big) + \int v \, dF_r(v) - \int v \, dF_\varepsilon(v) \geqslant \gamma. \tag{12}$$

Since $F_\varepsilon, F_r \in \mathcal{P}_\varepsilon(F_0)$ and since $h(v) = v$ is a continuous function and the Prohorov norm metrizes the weak topology we obtain, if ε is sufficiently small, that

$$\int v \, dF_r(v) - \int v \, dF_\varepsilon(v) < \gamma/2. \tag{13}$$

In the proof of Proposition 3 we showed that $\pi(p^*(F), F)$ as a function of F is continuous with respect to the weak topology. Hence

$$\pi\big(p^*(F_\varepsilon), F_\varepsilon\big) - \pi\big(p^*(F_r), F_r\big) < \gamma/2, \tag{14}$$

if ε is sufficiently small. Comparing (12) to (13) and (14) yields the desired contradiction. $\square$

References

[1] F. Anscombe, R. Aumann, A definition of subjective probability, Ann. Math. Statist. 34 (1963) 199–205.
[2] D. Bell, Regret in decision making under uncertainty, Oper. Res. 30 (1982) 961–981.
[3] D. Bergemann, K. Schlag, Pricing without priors, J. Europ. Econ. Assoc. Papers Proc. 6 (2008) 560–569.
[4] D. Bergemann, K. Schlag, Robust monopoly pricing, Discussion paper 1527RR, Cowles Foundation for Research in Economics, Yale University, 2008.
[5] S. Bose, E. Ozdenoren, A. Pape, Optimal auctions with ambiguity, Theoret. Econ. 1 (2006) 411–438.
[6] J. Bulow, J. Roberts, The simple economics of optimal auctions, J. Polit. Economy 97 (1989) 1060–1090.
[7] D. Ellsberg, Risk, ambiguity and the savage axioms, Quart. J. Econ. 75 (1961) 643–669.
[8] R. Engelbrecht-Wiggans, The effect of regret on optimal bidding in auctions, Management Sci. 35 (1989) 685–692.
[9] R. Engelbrecht-Wiggans, E. Katok, Regret in auctions: Theory and evidence, Econ. Theory 33 (2007) 81–101.
[10] E. Filiz-Ozbay, E. Ozbay, Auctions with anticipated regret: Theory and experiment, Amer. Econ. Rev. 97 (2007) 1407–1418.
[11] I. Gilboa, D. Schmeidler, Maxmin expected utility with non-unique prior, J. Math. Econ. 18 (1989) 141–153.
[12] J. Hannan, Approximation to Bayes risk in repeated play, in: M. Dresher, A. Tucker, P. Wolfe (Eds.), Contributions to the Theory of Games, Princeton University Press, Princeton, 1957, pp. 97–139.
[13] T. Hayashi, Regret aversion and opportunity dependence, J. Econ. Theory 139 (2008) 242–268.
[14] P.J. Huber, Robust Statistics, John Wiley and Sons, New York, 1981.
[15] P. Linhart, R. Radner, Minimax – Regret strategies for bargaining over several variables, J. Econ. Theory 48 (1989) 152–178.
[16] G. Loomes, R. Sugden, Regret theory: An alternative theory of rational choice under uncertainty, Econ. J. 92 (1982) 805–824.
[17] G. Lopomo, L. Rigotti, C. Shannon, Uncertainty in mechanism design, Discussion paper, 2009.
[18] J. Milnor, Games against nature, in: R. Thrall, C. Coombs, R. Davis (Eds.), Decision Processes, Wiley, New York, 1954.
[19] R. Myerson, Optimal auction design, Math. Oper. Res. 6 (1981) 58–73.
[20] K. Prasad, Non-robustness of some economic models, Top. Theor. Econ. 3 (2003) 1–7.
[21] L. Renou, K. Schlag, Minimax regret and strategic uncertainty, J. Econ. Theory 145 (2010) 264–286.
[22] P.J. Reny, On the existence of pure and mixed strategy Nash equilibria in discontinuous games, Econometrica 67 (1999) 1029–1056.
[23] J. Riley, R. Zeckhauser, Optimal selling strategies: When to haggle, when to hold firm, Quart. J. Econ. 98 (1983) 267–290.
[24] J. Stoye, Axioms for minimax regret choice correspondences, Discussion paper, New York University, 2008.