

ROBUST MONOPOLY PRICING

By

Dirk Bergemann and Karl Schlag

July 2005

Revised April 2007

Revised September 2008

COWLES FOUNDATION DISCUSSION PAPER NO. 1527RR



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

Robust Monopoly Pricing*

Dirk Bergemann[†] Karl Schlag[‡]

September 2008

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 optimal pricing policy under two distinct, but related, decision criteria with multiple priors: *(i)* maximin expected utility and *(ii)* minimax expected regret. The resulting optimal pricing policy under either criterion yields a robust policy to the model uncertainty.

While the classic monopoly policy and the maximin criterion yield a single deterministic price, minimax regret always prescribes a random pricing policy, or equivalently, a multi-item menu policy. Distinct implications of how a monopolist responds to an increase in uncertainty emerge under the two criteria.

KEYWORDS: Monopoly, Optimal Pricing, Robustness, Multiple Priors, Regret.

JEL CLASSIFICATION: C79, D82

*The first author gratefully acknowledges support by NSF Grants #SES-0518929, #CNS-0428422 and a DFG Mercator Research Professorship at the Center of Economic Studies at the University of Munich. We thank Rahul Deb, Peter Klibanoff, Stephen Morris, David Pollard, Phil Reny, John Riley and Thomas Sargent for helpful suggestions. We are grateful to seminar participants at the California Institute of Technology, Columbia University, the University of California at Los Angeles, the University of Wisconsin and the Cowles Foundation Conference "Uncertainty in Economic Theory" for many comments.

[†]Department of Economics, Yale University, New Haven, CT 06511, dirk.bergemann@yale.edu.

[‡]Department of Economics, Universitat Pompeu Fabra, 08005 Barcelona, Spain, karl.schlag@upf.edu.

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 design models, the robustness of these mechanisms with respect to the model specification becomes an important issue. In this paper, we investigate a robust version of the classic monopoly problem of selling a product under incomplete information. Optimal monopoly pricing is the most elementary instance of a profit maximizing problem in mechanism design with incomplete information.

We investigate 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. We consider the neighborhoods induced by the Prohorov metric which is the standard metric in robust statistical decision theory (see the Huber (1981) and Hampel, Ronchetti, Rousseeuw, and Stahel (1986)). 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. The Prohorov metric incorporates exactly these two different types of deviations, allowing both for a large probability of small errors and a small probability of large errors.

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. The maximin utility approach with multiple priors is due to Gilboa and Schmeidler (1989). Here, the decision maker evaluates each action by its minimum expected utility across all priors. The decision maker selects the action that maximizes the minimum expected utility. The minimax regret approach, originating in Savage (1951), was axiomatized by Milnor (1954) and recently adapted to multiple priors by Hayashi (2008) and Stoye (2008). Here, the decision maker evaluates foregone opportunities using regret and chooses an action that minimizes the maximum expected regret among the set of priors.

From an axiomatic perspective, the maximin utility and minimax regret criteria represent different departures from the standard model of Anscombe and Aumann (1963) by allowing for multiple priors. The maximin utility criterion emerges by giving up the independence axiom and replacing it with the weaker certainty independence axiom and adding a convexity axiom. The minimax regret criterion emerges by maintaining the certainty independence axiom but relaxing the axiom of independence of irrelevant alternatives, to allow the choice to be menu dependent. A convexity axiom and a version of the betweenness axiom complete the characterization. Both the maximin utility and the minimax regret criteria can be interpreted as refinements of subjective expected utility theory.

The analysis of the optimal pricing under the two decision criteria reveals that either criterion leads to a family of robust policies in the following sense. We say that a candidate family of policies, indexed by the size of the uncertainty, is *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 response to the uncertainty leads to distinct qualitative features.

The pricing policy of the seller is obtained as the equilibrium strategy of a zero-sum game between the seller and adversarial nature. The strategy by nature selects the *least favorable demand* given the objective of the seller. Under maximin utility the seller is worse off when valuations are lower. The least favorable demand thus maximizes the weight on the lowest

valuations subject to the restriction that the selected distribution is in the neighborhood of the model distribution. In particular, as we increase the uncertainty represented by an increase in the size of the neighborhood, the buyers valuations as determined by the least favorable demand are lower in the sense of first order stochastic dominance. In consequence the best response of the seller always consists in lowering her price.

When we analyze the behavior under regret minimization, the optimal pricing policy is still determined by a zero-sum game between the seller and nature. The notion of regret modifies the trade-off for the seller and for nature. The regret of the seller is the difference between the actual valuation of a buyer for the object and the actual 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 does not purchase the object, or *(ii)* he has a high valuation relative to the price and hence the seller could have obtained a higher profit. In the equilibrium of the zero-sum game, the optimal pricing policy of the seller has to resolve the conflict between the regret which arises with low prices against the regret associated with high prices. If the seller offers a low price, nature can cause regret with a distribution which puts substantial probability on high valuation buyers. On the other hand, if the seller offers a high price, nature can cause regret with a distribution which puts substantial probability at valuations just below the offered price. It then becomes evident that a single price will always expose the seller to substantial regret. Consequently, the seller can decrease her exposure by offering many prices. This can either be achieved by a probabilistic price or, alternatively, by a menu of prices. With a probabilistic price, the seller diminishes the likelihood that nature will be able to cause large regret. Equivalently, the seller can offer a menu of prices and quantities. The quantity element in the menu can either represent the quantity of a divisible object or the probability of obtaining an indivisible object.

We provide additional intuition by contrasting the pricing policy under regret to the standard profit maximizing policy. An optimal policy for a given distribution of valuations is always to offer the entire object at a fixed price (a classic result by Harris and Raviv (1981) and Riley and Zeckhauser (1983)). In contrast, here the policy will offer many prices (with varying quantities). With a single price, the risk of missing a trade at a valuation just

below the given price is substantial. On the other hand, if the seller were simply to lower the price, she would miss the chance of extracting profit from higher valuation customers. She resolves this conflict by offering smaller trades at lower prices to the low valuation customers. The size of the trade is simply the probability by which a trade is offered or the quantity offered at a given price. In the game against nature, the seller will have to be indifferent between offering small and large trades. In terms of the virtual utility, the key notion in optimal mechanisms, this requires that the seller will receive *zero* virtual utility over a range of valuations. The resulting conditions on the distribution of valuations determine the least favorable demand. Importantly, an increase in uncertainty may now lead to an increase in the expected price. In the special case of a linear model distribution, we find that the expected price increases if the optimal price for the model distribution is low and decreases if the optimal price for the model distribution is high.

We conclude the introduction with a brief discussion of the directly related literature. The basic ideas of robust decision making (see Definition 1) were first formalized in the context of statistical inference, in particular, with respect to the classic Neyman-Pearson hypothesis testing framework. The statistical problem is to distinguish between two known distributions on the basis of a sample. The model misspecification and consequent concern for robustness come from the fact that each of the two distributions might be misspecified. Huber (1964), (1965) first formalized robust estimation as the solution to a minimax problem and an associated zero-sum game. A recent contribution by Prasad (2003) employs this notion of robustness to the optimal policy without uncertainty, where it is referred to as α -robustness, and demonstrates the non-robustness of some economic models. In particular, he shows that the profit maximizing price in the optimal monopoly problem considered here is not robust to model misspecification. The non-robustness is demonstrated by a simple example. Suppose the model distribution is a Dirac distribution, which put probability one on a particular valuation v . Then the optimal monopoly price p is equal to v . This policy is not robust to model misspecification, because if the true model puts probability one on a value arbitrarily close, but strictly below v , then the resulting revenue is 0 rather than v . One of the objectives of this paper to identify robust policies, but not necessarily the optimal policy without uncertainty, that do not suffer from such discontinuity in the

profits.¹

A recent paper by Bose, Ozdenoren, and Pape (2006) determines the optimal auction in the presence of an uncertainty averse seller and uncertainty averse bidders. Lopomo, Rigotti, and Shannon (2006) 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 (2008) 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. The notion of regret was investigated in mechanism design by Linhart and Radner (1989) in the context of bilateral trade as well as by Engelbrecht-Wiggans (1989) and Selten (1989) in the context of auctions. Recently, Engelbrecht-Wiggans and Katok (2007) and Filiz-Ozbay and Ozbay (2007) present experimental evidence indicating concern for regret in first price auctions.

The remainder of the paper is organized as follows. In Section 2, we present the model, the notion of robustness and the neighborhoods. In Section 3, we characterize the pricing policy under the maximin utility criterion. In Section 4, we characterize the pricing policy under the minimax regret criterion. We show that the resulting policies are robust under either criterion. Section 5 concludes with a discussion of some open issues. The appendix collects auxiliary results and the proofs.

2 Model

2.1 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

¹There is also a rapidly growing literature on robust decision making in macroeconomics, see Hansen and Sargent (2007) for a comprehensive introduction, that uses related notions of robustness for maximizing the minimum utility in the context of intertemporal decision-making.

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:

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

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

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

By extension, if the valuation of the buyer is v , a random pricing policy $\Phi \in \Delta \mathbb{R}_+$ yields an expected profit:

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

Given the risk neutrality of the buyer and the seller, a random pricing policy Φ by the seller can alternatively be represented as a menu policy $(q, t(q))$ where q is the probability that the buyer receives the object and $t(q)$ is the tariff that the buyer pays for the probability q . Given the random pricing policy Φ , we can define for every $p \in \text{supp}\{\Phi\}$:

$$q \triangleq \Phi(p), \tag{1}$$

and the corresponding nonlinear price $t(q)$ as:

$$t(q) \triangleq \int_0^p y d\Phi(y). \tag{2}$$

In the menu interpretation, q is either the probability of receiving the object if the object is indivisible or the quantity if the object is divisible.

In the classic monopoly problem with incomplete information, the seller maximizes the expected profit for a given prior F over valuations. In the robust version, we assume that the seller faces uncertainty (or ambiguity) in the sense of Ellsberg (1961). The uncertainty is represented by a set of possible distributions. We first introduce the basic notation for the classic monopoly model and then define the model with uncertainty. For given a distribution F and given 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. In this paper, 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 $\Phi \in \Delta\mathbb{R}_+$, the expected profit is given by:

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

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

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

A well-known result by Riley and Zeckhauser (1983) states that for every distribution F , there exists a deterministic price $p^*(F)$ that maximizes profits, so:

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

2.2 Uncertainty

We assume that the seller faces uncertainty (or ambiguity) in the sense of Ellsberg (1961). The uncertainty is represented by a set of possible distributions, where 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 thus 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 Given the model distribution F_0 , we define the ε neighborhoods, denoted by $\mathcal{P}_\varepsilon(F_0)$, through the *Prohorov metric*:

$$\mathcal{P}_\varepsilon(F_0) = \{F \mid F(A) \leq F_0(A^\varepsilon) + \varepsilon, \forall \text{ measurable } A \subseteq [0, 1]\}, \quad (3)$$

where the set A^ε denotes the closed ε neighborhood of any measurable set A .² Formally, the set A^ε is given by:

$$A^\varepsilon = \left\{ x \in [0, 1] \mid \min_{y \in A} d(x, y) \leq \varepsilon \right\},$$

where $d(x, y) = |x - y|$ is the distance on the real line. The Prohorov metric has evidently two components. The additive term ε in (3) 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.

Maximin Utility Under maximin utility, the seller maximizes the minimum utility, where the utility of the seller is simply the profit, by solving:

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

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 profits across all profit maximizing pricing policies.³

²See Dudley (2002) for the definition of the Prohorov metric and the link to weak convergence and Huber (1981) and Hampel, Ronchetti, Rousseeuw, and Stahel (1986) for its application in robust statistics.

³Klibanoff, Marinacci, and Mukerji (2005) propose a related and smooth model of ambiguity aversion by enriching the multiple prior model with a belief μ over distributions and with an increasing transformation φ representing ambiguity aversion. The additional elements, belief μ and ambiguity index φ , render the analysis of multiple priors richer but also substantially more complex. In addition, the one-dimensional representation of ambiguity in terms of the size of the neighborhood is not available anymore.

Minimax Regret The *regret* of the monopolist at a given price p and valuation v of a buyer is defined as:

$$r(p, v) \triangleq v - p\mathbb{I}_{\{v \geq p\}} = v - \pi(p, v), \quad (4)$$

The regret of the monopolist charging price p facing a buyer with value v is the difference between (i) the profit the monopolist could make if she were to know the value v of the buyer before setting her price and (ii) 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 then $\mathbb{I}_{\{v \geq p\}} = 1$; or (ii) the value v is below the price p , the indicator function is then $\mathbb{I}_{\{v \geq p\}} = 0$.

The expected regret with a random pricing policy Φ when facing a distribution F is given by:

$$r(\Phi, F) \triangleq \int r(p, v) d\Phi(p) dF(v) = \int v dF(v) - \int \pi(p, F) d\Phi(p). \quad (5)$$

Thus, the probabilistic price Φ is profit maximizing at F if and only if Φ minimizes (expected) regret when facing F . The pricing policy $\Phi_r \in \Delta\mathbb{R}_+$ *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

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

Thus, a least favorable demand maximizes the regret of a profit maximizing seller who knows the true demand. It should be pointed out that while this 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 additional information to become available ex post. As in the case of the maximin utility criterion of Gilboa and Schmeidler (1989), the minimax

regret criterion in Hayashi (2008) and Stoye (2008) is completely characterized by a set of axioms.⁴

The notion of regret naturally extends to the case of many buyers as follows. The regret of the seller facing n buyers is equal to the sum of the regret accrued over n buyers and n , possibly distinct, prices. While the seller is thus allowed to offer a different price to each buyer, the additivity of the regret implies that we can confine attention to price (distributions) which are identical across buyers.

2.3 Robust Pricing

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)

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 presents a formal criterion of robust decision making in the spirit of the statistical decision literature pioneered by Huber (1964). It 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 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. Each member Φ_ε in the robust family $\{\Phi_\varepsilon\}_{\varepsilon>0}$ is allowed to depend on the size ε of the neighborhood. A natural and 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:

$$F \in \mathcal{P}_\varepsilon(F_0) \Rightarrow \pi(\Phi^*(F), F) - \pi(\Phi^*(F_0), F) < \gamma. \quad (6)$$

⁴In particular, the axiomatic approach to minimax regret is distinct from the ex-post measure of regret due to Hannan (1957) in the context of repeated games or from the more behavioral approaches to regret offered by Bell (1982) and Loomes and Sugden (1982).

This notion of robustness, applied directly to the optimal policy $\Phi^*(F)$, constitutes the definition of α robustness in Prasad (2003) and his earlier mentioned example of the Dirac distribution shows that the optimal policy $\Phi^*(F)$ is in general not robust.⁵ For a given model distribution F_0 , there are potentially many robust families of pricing rules. Our objective is to select among these rules by considering decision making under multiple priors and then to show that the resulting pricing rules are robust in the above sense of statistical decision making.

3 Maximin Utility

We consider the problem of the monopolist who wishes to maximize the minimum profit for all distributions in the neighborhood of the model distribution F_0 . Following Von Neumann (1928), the pricing rule that attains maximin utility can be viewed as the equilibrium strategy in a game between the seller and adversarial nature. The seller chooses a probabilistic price Φ and nature chooses a demand distribution F from the set $\mathcal{P}_\varepsilon(F_0)$. In this game, the payoff of the seller is the expected profit while the payoff of nature is the negative of the expected profit. Formally, a Nash equilibrium of this zero-sum game can be characterized as a solution to the saddle point problem of finding (Φ_m, F_m) that satisfy:

$$\pi(\Phi, F_m) \leq \pi(\Phi_m, F_m) \leq \pi(\Phi_m, F), \quad \forall \Phi \in \Delta\mathbb{R}_+, \forall F \in \mathcal{P}_\varepsilon(F_0). \quad (\text{SP}_m)$$

In other words, at (Φ_m, F_m) the probabilistic price Φ_m is profit maximizing at F_m and F_m is a profit minimizing demand given Φ_m .

The objective of adversarial nature is to lower the expected profit of the seller. For a given price p offered by the seller, the profit minimizing demand given p is achieved by increasing the cumulative probability of valuations strictly below p as much as possible within the neighborhood. The profit minimizing demand then minimizes the probability of sale by the seller. Given the model distribution F_0 and the size ε of the neighborhood, the

⁵In Prasad (2003), the definition of α robustness evaluates the profits at the model distribution F_0 rather than at the elements F in the neighborhood $\mathcal{P}_\varepsilon(F_0)$ of the model distribution F_0 as in (6). This difference is irrelevant in the case of a failure of robustness, which is the focus in Prasad (2003), due to the symmetry property of the Prohorov distance.

resulting distribution is uniquely determined for every p (up to a set of measure 0). The equilibrium analysis is now simplified by the fact that the profit minimizing demand does not depend on the, possibly probabilistic, price of the seller. We obtain the least favorable demand 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.

The exact construction of the least favorable demand in the Prohorov metric is rather transparent. Given a model demand F_0 and a neighborhood size ε , 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 that is within the ε neighborhood of F_0 , with F_m given by:

$$F_m(v) \triangleq \min \{F_0(v + \varepsilon) + \varepsilon, 1\}. \quad (7)$$

The first shift represents the possibility that small changes in valuations may occur with large probability. The second shift represents the idea of large changes occurring with a small probability. It is easily verified F_m is a profit minimizing demand for any price given the constraint imposed by the size of the neighborhood. We illustrate the least favorable demand F_m and the price p_m that attain maximin utility below for a model distribution with uniform density on the unit interval and a neighborhood of size $\varepsilon = 0.05$. We visualize the uncertainty around the model demand F_0 by the grey shaded area, which represents the smallest set that contains all cumulative distributions that lie within the Prohorov neighborhood of the uniform distribution (see also Lemma 1 for a characterization of the distribution functions that lie within the Prohorov neighborhood.)

INSERT FIGURE 1: PRICING AND LEAST FAVORABLE DEMAND UNDER MAXIMIN UTILITY

Given that the profit minimizing demand F_m does not depend on the offered prices, the monopolist acts as if the demand is given by F_m . In consequence, the seller maximizes profits at F_m by choosing a deterministic price p_m where

$$p_m \triangleq p^*(F_m).$$

Proposition 1 (Maximin Utility)

For every $\varepsilon > 0$, there exists a pair (p_m, F_m) , such that $p_m \in [0, 1]$ attains maximin utility and F_m is a least favorable demand.

An important implication of the above result is that a deterministic pricing policy p_m can always attain maximin utility. In contrast, under minimax regret a random pricing policy will always be strictly preferred to a deterministic pricing policy.

We now ask how the optimal price will 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 earlier assumption of concavity, we know that the curvature is negative and given by:

$$\frac{\partial^2 \pi(p_0, F_0)}{\partial p^2} = -2f_0(p_0) - p_0 f_0'(p_0) < 0.$$

We can directly 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 responds to an increase in uncertainty at $\varepsilon = 0$ by:

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

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. As the profit function becomes less concave, the rate of the price change increases as the profit function of the seller becomes less sensitive to a (downward) change in price and a more aggressive response of the seller diminishes the impact that the least favorable demand has on the sales of the monopolist.

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 will be at least as high as those obtained when facing the least favorable demand F_m . We now 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 will imply 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 being robust.

Proposition 3 (Robustness)

The family of pricing policies $\{p_{m,\varepsilon}\}_{\varepsilon>0}$ is a robust family of pricing policies.

4 Minimax Regret

4.1 Random Pricing

Next we consider the minimax regret problem of the seller. In contrast to the case of maximin utility, we now find that the seller chooses to offer a random pricing policy. The minimax regret strategy Φ_r and the 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. That is, (Φ_r, F_r) can be characterized as a solution to the saddle point problem of finding (Φ_r, F_r) that satisfy:

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

The saddlepoint result permits us to link minimax regret behavior to payoff maximizing behavior under a prior as follows. When 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 .⁶

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. The regret of the seller depends critically on the price offered by the seller. If

⁶We emphasize that we consider a simultaneous move game between the seller and nature. In this static environment, the earlier discussed axiomatic foundations lead the decision-maker, here the seller, to be concerned with the expected regret of the mixed pricing rule. In contrast, in a multi-stage game, one might analyze the regret relative to a realized price to avoid time inconsistency by the decision-maker.

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. It now becomes evident that a single price will always expose the seller to substantial regret. Conversely, the regret maximizing demand will now typically depend on the price offered by the seller. In fact, the seller can decrease her exposure by offering many prices in form of a probabilistic price. In contrast to the maximin profit, the regret maximizing demand is the result of an equilibrium argument and cannot be constructed independently of the strategy of the seller. We shall prove the existence of a solution to the saddlepoint problem (SP_r) and thus existence of a probabilistic price attaining minimax regret using results from Reny (1999).

Proposition 4 (Existence of Minimax Regret)

A solution (Φ_r, F_r) to the saddlepoint condition (SP_r) exists.

The minimax regret probabilistic price of the seller has to respond to a set of possible distributions. With an adversarial nature, the minimax regret policy of the seller is to offer many prices. We might guess intuitively that even the lowest price offered by the seller is not very far away from p_0 , the optimal price for the model distribution. In consequence, the price might not be low enough to dissuade nature from “undercutting” by placing probability just below the lowest price offered by the seller. This in turn might suggest that an equilibrium of the minimax regret pricing game fails to exist, however contradicting Proposition 4 above. Equilibrium strategies will be established by using the constraints on the least favorable demand. Naturally, the seller will price close to the optimal price without uncertainty. A mass point in the pricing strategy of the seller will be placed precisely at the point where nature is constrained by the neighborhood to shift any additional probability from above to just below the mass point of the seller. The seller then places the remaining mass in a neighborhood $[a, c]$ of this mass point b to protect against an increase in regret through local increases in values near this mass point.

Proposition 5 (Minimax Regret)

1. Given $\delta > 0$, if ε is 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:

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

2. The boundary points a, b and c respond to an increase in uncertainty at $\varepsilon = 0$:

- (a) $\lim_{\varepsilon \rightarrow 0} a'(0) = -\infty$,
 (b) $\lim_{\varepsilon \rightarrow 0} b'(0)$ is finite,
 (c) $\lim_{\varepsilon \rightarrow 0} c'(0) = \infty$.

We construct a probabilistic price that attains minimax regret by means of the implicit function theorem, for which we need the differentiability of the density function near p_0 . 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 substantially in order to protect against nature either undercutting or moving mass to the highest possible prices. At the same time, the mass point b does not change drastically.

We now illustrate the equilibrium behavior with the uniform model distribution:

$$F_0(v) = v,$$

where the profit maximizing price p_0 under the model distribution is given by $p_0 = \frac{1}{2}$. We graphically represent the optimal behavior of the seller and nature for a small neighborhood.

INSERT FIGURE 2: PRICING AND LEAST FAVORABLE DEMAND UNDER MINIMAX REGRET

The interior curve in the above graph identifies the model distribution. Constraints induced by small changes in values cause the distribution function of F_r to be within an ε

bandwidth of the model distribution. The large changes of values, occurring with probability of at most ε , move the smallest valuation to the largest valuation, namely 1. The strategy of nature is then to place as little probability as necessary below the range of the prices offered by the seller and to shift values above the range as high as possible. Inside the range of prices offered by the seller, nature uses a density function which maintains the virtual utility of the seller at 0. In turn, the seller sets the density to make nature indifferent between all values above the mass point and all values below the mass point. Given the mass point set by the seller, nature shifts as much mass as possible below this point. We observe that even with the small neighborhood of $\varepsilon = 0.05$, the impact of the uncertainty on the probabilistic price is rather large and leads to a wide spread in the prices offered by the seller.

It remains to describe the comparative static of the probabilistic price and the regret of the seller as a function of the size of the neighborhood. The behavior of regret and of the expected price 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

$$r^* = r_0 + \varepsilon \frac{\partial r^*}{\partial \varepsilon},$$

where r_0 is the regret at the model distribution. For a small level of uncertainty, the marginal change in regret can then be computed by holding the probabilistic price of the seller at the optimal price p_0 without uncertainty. Suppose then for the moment that $p_0 \leq \frac{1}{2}$. If the uncertainty increases marginally, the constraints on the choice of a least favorable demand are relaxed. What precisely then can nature do, given the specification of neighborhood. First, nature can place the density $f_0(p_0)$ slightly below p_0 to marginally increase regret by $p_0 f_0(p_0)$, then nature can shift each value up by ε to marginally increase regret by 1 and finally shift mass from 0 to 1 to marginally increase regret by $1 - p_0$. The first two changes correspond to small changes in valuation with large probability, the third to large changes in the valuation with small probability. So the overall marginal effect on regret of an increase in ε near $\varepsilon = 0$ is $p_0 f_0(p_0) + 1 + (1 - p_0)$. If instead the optimal price without uncertainty were $p_0 > \frac{1}{2}$, then the robust modification would only pertain to the third element as nature would move mass from 0 to just below p_0 , so that the marginal increase

would be $p_0 f_0(p_0) + 1 + p_0$.

The optimal response of the seller to an increase in uncertainty is now to find a probabilistic price which minimizes the additional regret

$$\varepsilon \frac{\partial r^*}{\partial \varepsilon},$$

coming from the increase in uncertainty. Of course, the consequence of adjusting the price to minimize the marginal regret is that it changes the regret relative to the model distribution F_0 . 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 responds 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, or

$$\frac{\partial \mathbb{E}[\Phi_r]}{\partial \varepsilon} = \frac{\frac{\partial^2 r^*}{\partial \varepsilon \partial p}}{\frac{\partial^2 \pi(p_0, F_0)}{(\partial p)^2}}.$$

The next proposition shows that the above intuition can be made precise and shows its implication for the net utility of the buyer.

Proposition 6 (Comparative Statics with Minimax Regret)

The expected price $\mathbb{E}[\Phi_r]$ responds to an increase in uncertainty at $\varepsilon = 0$ by:

$$\frac{\partial}{\partial \varepsilon} \mathbb{E}[\Phi_r]_{\varepsilon=0} = \begin{cases} -1 - \frac{f_0(p_0)+1}{\partial \pi^2(p_0, F_0)/\partial p^2} > -1 & \text{if } p_0 \leq \frac{1}{2}, \\ -1 - \frac{f_0(p_0)-1}{\partial \pi^2(p_0, F_0)/\partial p^2} < -\frac{1}{2} & \text{if } p_0 > \frac{1}{2}. \end{cases} \quad (8)$$

We observe that for $p_0 > \frac{1}{2}$, the response of the expected price $\mathbb{E}[\Phi_r]$ to an increase in uncertainty is identical under regret minimization and profit maximization. The difference arises at a low level of p_0 at which the seller is less aggressive in lowering her price due to an increase in uncertainty. For the case of $p_0 \leq \frac{1}{2}$, it turns out that the expected price can be strictly increasing in ε . In fact, we find that in the class of linear densities the change in expected price as well as the change in the mass point is strictly positive if, and only if, the density is strictly decreasing. This has to be contrasted with the maximin behavior where any increase in size of the uncertainty has a downward effect on prices for all model distributions.

4.2 Menu Pricing

The equilibrium menu policy can be directly derived from the random pricing policy Φ_r . We identify the regret minimizing menu $(q, t_r(q))$ by determining the transfer price of every offered quantity q through the random pricing policy Φ_r . The resulting net utility for a buyer with value v is given by:

$$q \cdot v - t_r(q).$$

Specifically, the construction of the menu $(q, t_r(q))$ proceeds as follows. Every price $p \in \mathbb{R}_+$ of the random pricing policy Φ_r such that $p \in \text{supp}(\Phi_r)$, determines a probability q in the menu by:

$$q \triangleq \Phi(p), \tag{9}$$

and a corresponding nonlinear price $t_r(q)$ for the quantity q by:

$$t_r(q) \triangleq \int_0^p y d\Phi(y). \tag{10}$$

By the very construction of the transfer function $t_r(q)$, it follows that a buyer with value v will select the item q on the menu such that $q = \Phi(v)$. The self-selection condition for a buyer with value v is determined by choosing the quantity q , such that the net utility of the buyer is maximized, or

$$v - t'_r(q) = 0,$$

which occurs at $q = \Phi(v)$ as $t'_r(q) = v$ by (10). By the taxation principle in the theory of mechanism design, the menu $(q, t_r(q))$ can also be viewed as an incentive compatible allocation plan $(q_r(v), t_r(v))$ in the corresponding direct mechanism.

The equilibrium use of menus allows us to understand the selling policies from a different and perhaps more intuitive point of view. The optimality of menus emphasizes the concern for robustness as menus would never be used in the standard setting for a given demand distribution. The minimax regret menu offered by seller has three important characteristics. These properties can be described with reference to the mass point b in the random pricing policy Φ_r of Proposition 5: (i) low volume offers are made for buyers with low valuations, or $v < b$, (ii) a much higher offer is made for all buyers with valuation $v = b$, and (iii) even higher volume offers are made to buyers with large values $v > b$. We may think of

a standard offer as given by the quantity offered at $v = b$. In addition, the seller offers low volume downgrades and high volume upgrades. The expanded menu relative to the optimal single item menu for the model distribution, seeks to minimize the exposure to regret. Obviously, the seller loses profits on the high value buyers from making offers to the low value buyers by granting the high value buyers a larger information rent. The size of the information rent is kept small by offering menu items to the low value buyers only of substantially lower volume. This is the source of the gap in the quantities offered in the menu.

INSERT FIGURE 3: MENU PRICING UNDER MINIMAX REGRET

The response of the seller to an increase in uncertainty is informative when we consider menus. In a menu, the seller is offering many different choices to the buyers. An immediate question therefore is how the size of the menu and the associated prices change with an increase in the uncertainty. The size of the menu is simply the range of quantities offered by the seller (and accepted by some buyers) in equilibrium.

Proposition 7 (Menus and Uncertainty)

For small uncertainty ε :

1. *The size of the menu is increasing in ε .*
2. *The price per unit $t_r(v)/q_r(v)$ is decreasing in ε for every $v \in (a, c) \setminus b$.*

As the uncertainty increases, the seller seeks to minimize her exposure to regret by offering more choices to the buyers and hence increasing the probability of a sale, even if the sale is not “big” in terms of the sold quantity. For every given valuation v , the seller also increases the size of the deal offered. As larger deals are offered to buyers with lower valuations, it follows that the seller is willing to concede a larger information rent to buyers with higher valuations. In consequence, the average price per unit is decreasing as well. Jointly, these three properties imply that the seller is offering her products more aggressively and to a larger number of buyers with an increase in uncertainty. We observe that the monotonicity in the unit price holds even as the previous proposition showed that

the expected price may be increasing. The resolution of this apparent conflict comes from the fact that the seller is offering larger quantities in response to an increase in uncertainty.

An interesting comparison to a minimax regret decision maker is a risk averse decision maker. In particular, we could ask how the behavior of a risk averse seller would differ from the behavior of a minimax regret seller. Clearly, a risk averse seller would never find a probabilistic price optimal. However, if she were to be allowed to offer a menu, either of lotteries (in terms of probabilities of receiving the good) or different qualities of the good, then a risk averse seller might indeed offer a menu. The menu would consist of a set of possible quantity and price combinations. The difference with respect to the minimax regret seller would then be in the shape of the menu. In particular, if a risk averse seller were to face a continuous demand function (as expressed by F_0), then the optimal menu can be shown to be continuous. Yet, with a minimax regret seller, we saw that the optimal menu is discontinuous (at a single jump point) and essentially offers two (or three) classes of distinct service.

The minimax regret problem with uncertainty then offers an interesting and novel reason for menus to complement existing insights. The literature currently offers two leading explanations for menus in the standard monopoly setting: menus can be optimal if the marginal willingness to pay changes with the quantity offered as in Deneckere and McAfee (1996) or if the buyers are budget constrained as in Che and Gale (2000).

4.3 Robustness

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

Proposition 8 (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.

5 Conclusion

In this paper, we analyzed pricing policies of a monopolist which are robust to model uncertainty. The introduction of uncertainty about the true demand distribution formally lead to a decision theoretic model with multiple priors. The parsimonious representation of the uncertainty in terms of the neighborhood of a model distribution allowed us to deal with added complexity and maintain an intuitive understanding of how uncertainty affects optimal policies.

We analyzed the optimal pricing of a monopolist under two distinct, but related decision criteria with multiple priors: maximin profit and minimax regret. We showed that the solution under either criterion yields a robust solution in the statistical sense. The expected profit under either pricing rule is arbitrarily close to the optimal price for any distribution in a sufficiently small neighborhood of the model distribution. Despite the common robustness property, the prices respond differently to the uncertainty. The maximin policy uniformly maintains a deterministic price policy and uniformly lowers the price as a response to an increase in uncertainty. In contrast, the minimax policy balances the downside versus the upside when responding to the uncertainty. Here the trade-off is optimally resolved by a probabilistic price. Importantly, the expected price does not necessarily decrease with an increase in uncertainty. Interestingly, an equivalent policy to the probabilistic price is achieved by a menu. The menu offers a variety of quantities, ranging from small to large, to the buyer. By offering a menu, the seller can guarantee himself small deals on the downside and large deals on the upside. In consequence, the seller hedges to reduce maximal regret by offering multiple choices through a menu. A common feature of both models of decision making is that we can analyze how uncertainty influences pricing without adding degrees of freedom to the model. This renders our results parsimonious and falsifiable.

The problem of optimal monopoly pricing is in many respects the most elementary mechanism design problem. It would be of interest to extend the insights and apply the techniques developed here to a wider class of design problems, such as the discriminating monopolist (as in Mussa and Rosen (1978) and Maskin and Riley (1984)) and optimal auctions. The monopoly setting has the simplifying feature that the buyers have complete

information about their payoff environment. Given their known valuation and known price, each buyer simply has to make a decision as to whether or not to purchase the object. With the complete information of the buyer, there is no need to look for a robust purchasing rule. A substantial task would consequently arise by considering multi-agent design problems with incomplete information such as auctions, where it becomes desirable to simultaneously make the decisions of the buyers and the seller robust. The complete solution of these problems poses a rich field for future research.

6 Appendix

The appendix contains some auxiliary results as well as the proofs for the results in the main body of the text.

Proof of Proposition 1. As shown in the text, if F_m is such that

$$F_m(v) = \min \{F_0(v + \varepsilon) + \varepsilon, 1\},$$

then $\pi(p, F_m) \leq \pi(p, F)$ for all $F \in \mathcal{P}_\varepsilon(F_0)$. On the other hand, if $p_m = p^*(F_m)$, then $\pi(p_m, F_m) \geq \pi(p, F_m)$ holds for all p by the definition of p_m . Together this implies that (p_m, F_m) is a saddle point as described in (SP_m) and thus 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:

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

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

$$\frac{dp_m}{d\varepsilon}|_{\varepsilon=0} = -1 + \frac{1 - f_0(p_0)}{-2f_0(p_0) - p_0 f'_0(p_0)} = \frac{f_0(p_0) + p_0 f'_0(p_0) + 1}{-2f_0(p_0) - p_0 f'_0(p_0)}.$$

Since $-2f_0(p_0) - p_0 f'_0(p_0) < 0$, we observe that the lhs 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 $-1/2$. ■

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 $\pi(p_m, F) \geq \pi(p_m, F_m)$ and thus

$$\pi(p^*(F), F) - \pi(p_m, F) \leq \pi(p^*(F), F) - \pi(p_m, F_m).$$

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

$$G(p) \leq F(p + \varepsilon) + \varepsilon,$$

we obtain

$$\begin{aligned} \pi(p^*(G), G) &\geq \pi(p^*(F) - \varepsilon, G) = (p^*(F) - \varepsilon)(1 - G(p^*(F) - \varepsilon)) \\ &\geq (p^*(F) - \varepsilon)(1 - F(p^*(F)) - \varepsilon) \geq \pi(p^*(F), F) - 2\varepsilon. \end{aligned}$$

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

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

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

Proof of Proposition 4. We apply Corollary 5.2 in Reny (1999) 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 $\bar{\Phi}$ such that $F \in \mathcal{P}_\gamma(F_r)$ implies $r(\bar{\Phi}, F) \leq r(\Phi_r, F_r) + \delta$.

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

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

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

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

and hence

$$r(\bar{\Phi}, F) \leq r(\Phi_r, F_r) + \delta.$$

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 $\bar{F} \in \mathcal{P}_\varepsilon(F_0)$ such that $\Phi \in \mathcal{P}_\gamma(\Phi_r)$ implies $r(\Phi, \bar{F}) \geq r(\Phi_r, F_r) - \delta$.

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

$$\begin{aligned} \pi(\Phi, F_r) &\leq \gamma + \int (p + \gamma) \left(\int_p^1 dF_r(v) \right) d\Phi_r(p + \gamma) = \gamma + \int p \left(\int_{p-\gamma}^1 dF_r(v) \right) d\Phi_r(p) \\ &= \gamma + \pi(\Phi_r, F_r) + \int p \left(\int_{[p-\gamma, p)} dF_r(v) \right) d\Phi_r(p) \\ &\leq \gamma + \pi(\Phi_r, F_r) + \int \int_{[p-\gamma, p)} dF_r(v) d\Phi_r(p). \end{aligned}$$

Given the continuity of the integral term

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

in the boundary point γ , the claim is established. ■

In order to derive the equilibrium policies in the case of uncertainty we present a characterization of the Prohorov distance in Lemma 1 that builds on the following result of Strassen (1965).

Theorem (Strassen (1965)).

F and G have Prohorov distance less than or equal to ε if and only if there exist random variables X and Y such that X has distribution F , Y has distribution G and $\Pr(|Y - X| \leq \varepsilon) \geq 1 - \varepsilon$.

The two cumulative distributions F, G are close in terms of the Prohorov distance if and only if they are associated to two random variables that realize similar values with high probability. Our characterization describes the Prohorov distance in terms of monotone functions that are identified with positive additive measures and cumulative distribution functions respectively. In order to stay within ε distance of a given distribution function G one may first alter any value of $G(v)$ at v by at most ε , this creates a probability measure F_1 , and then move at most ε mass of the values. The new locations are described by a

measure F_2 while locations from where the mass has been taken is described by a measure F_3 .

Lemma 1 (Decomposition)

Consider $\varepsilon > 0$ and probability measures F and G . $F \in \mathcal{P}_\varepsilon(G)$ if and only if there exists a probability measure F_1 and positive additive measures F_2 and F_3 such that:

$$G(x - \varepsilon) \leq F_1(x) \leq G(x + \varepsilon) \quad \forall x,$$

and

$$F_2(1) = F_3(1) \leq \varepsilon,$$

and

$$F \triangleq F_1 + F_2 - F_3.$$

Proof. (\Leftarrow) Suppose F can be decomposed into F_1, F_2 and F_3 . We want to show that $F(A) \leq G(A^\varepsilon) + \varepsilon$. To this purpose, it is clearly sufficient to consider only closed sets A .

(a) We first prove the claim for $A = [x, y]$ with $0 \leq x \leq y \leq 1$. Given a probability measure H , let $H^-(\hat{v}) \triangleq \lim_{v \uparrow \hat{v}} H(v)$. Then

$$F_1([x, y]) = F_1(y) - F_1^-(x) \leq G(y + \varepsilon) - G^-(x - \varepsilon) = G([x, y]^\varepsilon).$$

Since $F_2([x, y]) \leq \varepsilon$ and $F_3([x, y]) \geq 0$ we obtain:

$$F([x, y]) = F_1([x, y]) + F_2([x, y]) - F_3([x, y]) \leq G([x, y]^\varepsilon) + \varepsilon.$$

(b) Next we consider $A = [x_1, y_1] \cup [x_2, y_2]$ with $y_1 + 2\varepsilon < x_2$ which implies that

$$[x_1, y_1]^\varepsilon \cap [x_2, y_2]^\varepsilon = \emptyset.$$

Using part (a) together with the fact that $A^\varepsilon = [x_1, y_1]^\varepsilon \cup [x_2, y_2]^\varepsilon$ holds for the $[\cdot]^\varepsilon$ operator, it follows that:

$$F_1(A) = F_1([x_1, y_1]) + F_1([x_2, y_2]) \leq G([x_1, y_1]^\varepsilon) + G([x_2, y_2]^\varepsilon) = G(A^\varepsilon).$$

Since $F_2(A) \leq \varepsilon$ and $F_3(A) \geq 0$, the claim is proven.

(c) The arguments in part (b) are easily generalized for any set A that can be decomposed into a finite union of disjoint closed intervals of distance greater than 2ε so $A = \cup_{k=1}^m [x_k, y_k]$ with $x_k \leq y_k < x_{k+1} - 2\varepsilon$ for $k \leq m - 1$.

(d) Finally, we show that we do not have to prove the statement for more general sets A . Notice that if $A_1^\varepsilon = A_2^\varepsilon$, $A_1 \subset A_2$ and $F(A_2) \leq G(A_2^\varepsilon) + \varepsilon$ then $F(A_1) \leq G(A_1^\varepsilon) + \varepsilon$. So we can restrict attention to proving the claim for closed sets A such that $A^\varepsilon = A_1^\varepsilon$ and $A \subseteq A_1$ implies $A = A_1$. Consider $x, y \in A$ such that $x < y \leq x + 2\varepsilon$. Then $\{A \cup [x, y]\}^\varepsilon = A^\varepsilon$ and hence $[x, y] \subseteq A$. It follows that A belongs to the class of sets investigated in part (c).

(\Rightarrow) Consider probability measures F and G with $\|F - G\| \leq \varepsilon$. We extend G to $[-\varepsilon, 1 + \varepsilon]$ such that $G(x) = 0$ for $-\varepsilon \leq x < 0$ and $G(x) = 1$ for $1 < x \leq 1 + \varepsilon$. Given the result of Strassen (1965), there exist random variables X and Y such that X has distribution F , Y has distribution G and $\Pr(|Y - X| \leq \varepsilon) \geq 1 - \varepsilon$.

Let Z_1 be the random variable with cdf F_1 such that $Z_1 \triangleq X$ if $|Y - X| \leq \varepsilon$ and $Z_1 \triangleq Y$ if $|Y - X| > \varepsilon$. Let $\varepsilon' \triangleq \Pr(|Y - X| > \varepsilon)$ so $\varepsilon' \leq \varepsilon$. Then $G(x - \varepsilon) \leq F_1(x) \leq G(x + \varepsilon)$. Let Z_2 be the random variable with cdf \widehat{F}_2 such that $Z_2 \triangleq 0$ if $|Y - X| \leq \varepsilon$ and $Z_2 \triangleq X$ if $|Y - X| > \varepsilon$. Let Z_3 be the random variable with cdf \widehat{F}_3 such that $Z_3 \triangleq 0$ if $|Y - X| \leq \varepsilon$ and $Z_3 \triangleq Y$ if $|Y - X| > \varepsilon$. Then $X = Z_1 + Z_2 - Z_3$ and $\widehat{F}_2(0), \widehat{F}_3(0) \geq 1 - \varepsilon'$. Let $F_i \triangleq \widehat{F}_i - (1 - \varepsilon')$ for $i = 2, 3$. Then F_2 and F_3 are positive additive measures with $F_2, F_3 \leq \varepsilon'$ and the proof is complete. ■

Proof of Proposition 5. We start by assuming $p_0 > \frac{1}{2}$. The proof proceeds in three steps. First we show the existence of the parameters a, b and c and use these to construct the least favorable demand F_r . Second, we decompose the least favorable demand by using Lemma 1 to show that it is close to F_0 . Third, we use this decomposition to verify that we have a saddle point.

Step 1. We start by showing that for sufficiently small ε , there exist parameters a, b, c

such that $a < b < c$ and $a < p_0 < c$ such that

$$F_0(a - \varepsilon) - \varepsilon = 1 - \frac{b^2 f_0(b + \varepsilon)}{a}, \quad (11)$$

$$F_0(b + \varepsilon) = 1 - \frac{b^2 f_0(b + \varepsilon)}{b}, \quad (12)$$

$$F_0(c - \varepsilon) = 1 - \frac{b^2 f_0(b + \varepsilon)}{c}. \quad (13)$$

Concerning the existence of b , note that $b = p_0$ solves (12) if $\varepsilon = 0$. As

$$\frac{d}{db} (1 - F_0(b + \varepsilon) - b f_0(b + \varepsilon))|_{\varepsilon=0} = -2f(p_0) - p_0 f'_0(p_0) < 0,$$

due to the strict concavity of profits at p_0 , the implicit function theorem implies that a continuous solution $b = b(\varepsilon)$ to (12) (with $b > 0$) exists for ε in a neighborhood of 0 with $b(0) = p_0$. To prove existence of c , define

$$h_\varepsilon(v) \triangleq 1 - \frac{b^2 f_0(b + \varepsilon)}{v} - F_0(v - \varepsilon) \quad \text{for } v > 0.$$

We need to show that there exists $c > b$ such that $h_\varepsilon(c) = 0$. Note that

$$h_\varepsilon(b) = F_0(b + \varepsilon) - F_0(b - \varepsilon),$$

and hence $h_\varepsilon(b) > 0$ for sufficiently small ε as $f(p_0) > 0$ by assumption in the statement of the proposition and as f is continuous in a neighborhood of p_0 . Moreover,

$$\begin{aligned} h'_\varepsilon(b) &= f_0(b + \varepsilon) - f_0(b - \varepsilon), \\ h''_\varepsilon(b) &= -\frac{2f_0(b + \varepsilon)}{b} - f'_0(b - \varepsilon) \approx -\frac{2f_0(p_0) + p_0 f'_0(p_0)}{p_0} < 0. \end{aligned}$$

In fact, one can easily show for v in a sufficiently small neighborhood around p_0 that $h'_\varepsilon(v)$ is small and that $h''_\varepsilon(v)$ is bounded below 0. Here one uses the fact that f is assumed to be continuously differentiable in the neighborhood of p_0 . Specifically one shows there exists $\tau > 0$ and a neighborhood of p_0 such that if $\eta > 0$ and ε is sufficiently small then $0 < h_\varepsilon(b) < \delta$ and $h'_\varepsilon(v) < \eta$ and $h''_\varepsilon(v) < -\tau$ for any v in this neighborhood. This means that $h_\varepsilon(v) < \eta + \eta(v - b) - \tau(v - b)^2$ for v in this neighborhood. Choosing η sufficiently small, one then shows that for sufficiently small ε there exists $c > b$ such that $h_\varepsilon(c) = 0$ where $c \rightarrow p_0$ as $\varepsilon \rightarrow 0$.

For the existence of a , similar calculations for $h_\varepsilon(v) + \varepsilon$ show that there exists $a < b$ such that $h_\varepsilon(a) + \varepsilon = 0$ with $a \rightarrow p_0$ as $\varepsilon \rightarrow 0$.

We can describe the local behavior of the parameters a, b and c by appealing to the implicit function theorem. Since $2f_0(p_0) + p_0f'_0(p_0) > 0$, we know that b is differentiable and by implicitly differentiating (12) we obtain:

$$b'(0) = -\frac{f_0(p_0) + p_0f'_0(p_0)}{2f_0(p_0) + p_0f'_0(p_0)} = -1 + \frac{f_0(p_0)}{2f_0(p_0) + p_0f'_0(p_0)}, \quad (14)$$

with $-1 < b'(0) \leq -1/2$. The second inequality follows from using the fact that $b'(0)$ as a function of $f_0(p_0)$ is increasing. Next, we show that a is differentiable. Since

$$\begin{aligned} \frac{b^2f_0(b+\varepsilon) - a^2f_0(a-\varepsilon)}{b-a} &= (b+a)f_0(b+\varepsilon) + a^2\frac{f_0(b+\varepsilon) - f_0(a-\varepsilon)}{b-a} \\ &\approx 2p_0f_0(p_0) + (p_0)^2f'_0(p_0), \end{aligned}$$

we find that $b^2f_0(b+\varepsilon) > a^2f_0(a-\varepsilon)$ near $\varepsilon = 0$. Hence we can implicitly differentiate (11) to obtain

$$a'(\varepsilon) = -a\frac{a + af_0(a-\varepsilon) + bf_0(b+\varepsilon)}{b^2f_0(b+\varepsilon) - a^2f_0(a-\varepsilon)}, \quad (15)$$

and so

$$\lim_{\varepsilon \rightarrow 0} \frac{b-a}{a} a'(\varepsilon) = -\frac{1 + 2f_0(p_0)}{2f_0(p_0) + p_0f'_0(p_0)}.$$

In particular, we obtain that

$$\lim_{\varepsilon \rightarrow 0} a'(\varepsilon) = -\infty. \quad (16)$$

Similarly for c , we find that:

$$c'(\varepsilon) = -c\frac{cf_0(c-\varepsilon) + bf_0(b+\varepsilon)}{b^2f_0(b+\varepsilon) - c^2f_0(c-\varepsilon)}, \quad (17)$$

and hence

$$\lim_{\varepsilon \rightarrow 0} \frac{c-b}{c} c'(\varepsilon) = \frac{2f_0(p_0)}{2f_0(p_0) + p_0f'_0(p_0)},$$

and in particular,

$$\lim_{\varepsilon \rightarrow 0} c'(\varepsilon) = \infty. \quad (18)$$

It now follows from (16) and (18) that $a < p_0 < c$.

Step 2. We now construct the least favorable demand on the basis of a, b and c .

Consider F_r given by:

$$F_r(v) \triangleq \begin{cases} 0 & \text{if } v \in [0, \varepsilon), \\ \max\{0, F_0(v - \varepsilon) - \varepsilon\}, & \text{if } v \in [\varepsilon, a], \\ 1 - \frac{b^2 f_0(b + \varepsilon)}{v}, & \text{if } v \in (a, c), \\ F_0(v - \varepsilon), & \text{if } v \in [c, 1), \\ 1, & \text{if } v = 1. \end{cases}$$

The definitions of a and c imply that F_r is continuous at a and c . It follows that F_r is a probability measure.

Next we show that $F_r \in \mathcal{P}_\varepsilon(F_0)$ by using Lemma 1. The proof is by construction and uses the decomposition of Lemma 1.

Define F_1 by:

$$F_1(v) \triangleq \begin{cases} 0, & \text{if } v \in [0, \varepsilon), \\ F_0(v - \varepsilon), & \text{if } v \in [\varepsilon, a], \\ \max\{F_r(v), F_0(v - \varepsilon)\}, & \text{if } v \in (a, b), \\ F_r(v), & \text{if } v \in [b, 1]. \end{cases}$$

Then F_1 is a probability measure with $F_0(v - \varepsilon) \leq F_1(v)$ for $v \in [\varepsilon, 1]$. By definition of b we obtain $F_r(b) = F_0(b + \varepsilon)$ and $F_r'(b) = \frac{d}{dv} F_0(v + \varepsilon)|_{v=b}$. Moreover, given

$$F_r''(v) = -\frac{2b^2 f_0(b + \varepsilon)}{v^3}$$

and

$$\frac{d^2}{dv^2} F_0(v + \varepsilon) = f_0'(v + \varepsilon),$$

strict concavity of profits near p_0 implies that $F_0''(v) < F_r''(v)$ for $v \in [a, c]$ and ε sufficiently small. Thus, for sufficiently small ε , as a and c are close to p_0 , we obtain $F_1(v) \leq F_0(v + \varepsilon)$ with equality if $v = b$. So $F_0(v - \varepsilon) \leq F_1(v) \leq F_0(v + \varepsilon)$.

Define F_2 by:

$$F_2(v) \triangleq \begin{cases} 0, & \text{if } v \in [0, a], \\ \varepsilon - \max\{F_0(v - \varepsilon) - F_r(v), 0\}, & \text{if } v \in (a, b), \\ \varepsilon, & \text{if } v \in (b, 1]. \end{cases}$$

Since F_2 is right continuous, we will obtain that F_2 is an additive measure once we show that

$$\frac{d}{dv} (F_r(v) - F_0(v - \varepsilon)) = \frac{b^2 f_0(b + \varepsilon)}{v^2} - f_0(v - \varepsilon) \geq 0 \quad (19)$$

holds if $v \in (a, b]$ and $F_r(v) < F_0(v - \varepsilon)$. Strict concavity of the profit function near $p = p_0$ can be used to show that $v^2 f_0(v - \varepsilon)$ is increasing in v for sufficiently small ε . We establish the inequality (19) for $v = x$ where $x \in (a, b)$ solves $F_r(x) = F_0(x - \varepsilon)$. Note that $x(1 - F_r(x)) = b^2 f_0(b + \varepsilon)$ and hence we aim to show that

$$x(1 - F_r(x)) - x^2 f_0(x - \varepsilon) = x(1 - F_0(x - \varepsilon) - x f_0(x - \varepsilon)) \geq 0,$$

which follows once we show that $x < p_0$. By invoking the implicit function theorem on $F_r(x) - F_0(x - \varepsilon) = 0$, we notice that the derivative with respect to x vanishes as $\varepsilon \rightarrow 0$ while the terms involving ε and $b(\varepsilon)$ combine to yield an expression that tends to $2f_0(p_0)$ as $\varepsilon \rightarrow 0$. Thus we find that the slope of x near $\varepsilon = 0$ is unbounded and as $x(\varepsilon) < b(\varepsilon)$ and $b'(0)$ is finite it follows that $x'(\varepsilon) \rightarrow -\infty$ as $\varepsilon \rightarrow 0$.

Define F_3 by:

$$F_3(v) \triangleq \begin{cases} 0, & \text{if } v \in [0, \varepsilon], \\ \min\{F_0(v - \varepsilon), \varepsilon\}, & \text{if } v \in (\varepsilon, 1], \end{cases}$$

and so $F_3(v)$ is an additive measure and $F_3(1) = \varepsilon$. Since $F_r = F_1 + F_2 - F_3$ we obtain from Lemma 1 that $F_r \in \mathcal{P}_\varepsilon(F_0)$.

Step 3. We now show that (Φ_r, F_r) is a saddle point. For the monopolist we verify easily that $\pi(p, F_r) = b^2 f_0(b + \varepsilon)$ for $p \in [a, c]$. Next we consider profits nearby but outside $[a, c]$. For $p < a$ we verify that $\pi(p, F_r) \leq b^2 f_0(b + \varepsilon)$ by showing that

$$\lim_{v \rightarrow a^-} \frac{d}{dv} F_r(v) < \lim_{v \rightarrow a^+} \frac{d}{dv} F_r(v)$$

which means that

$$f_0(a - \varepsilon) < \frac{b^2 f_0(b + \varepsilon)}{a^2}.$$

Similarly we need to show for the case where $p > c$ that

$$f_0(c - \varepsilon) > \frac{b^2 f_0(b + \varepsilon)}{c^2},$$

which is easily verified for sufficiently small ε by using the fact that $c'(0) = \infty$ while $b'(0)$ is finite. Thus we have shown that there exists $\eta > 0$ such that

$$[a, c] = \arg \max_{p \in [p_0 - \eta, p_0 + \eta]} \pi(p, F_r).$$

In particular, η can be chosen independently of ε once ε is sufficiently small. Upperhemi-continuity of the set of profit maximizing prices then implies that $[a, c] \subseteq \arg \max_p \pi(p, F_r)$.

Consider now the incentives of nature. The regret r is an additive function over the measures and so:

$$r(\Phi_r, F_r) = r(\Phi_r, F_1) + r(\Phi_r, F_2) - r(\Phi_r, F_3). \quad (20)$$

We show that nature maximizes each term in (20) and then sets $F_r = F_1 + F_2 - F_3$, subject to the constraints on F_1 , F_2 and F_3 specified in Lemma 1. Notice that we assume that $F_1(1) = F_2(1) = \varepsilon$. Consider first F_2 . By construction of Φ_r , regret $r(\Phi_r, \cdot)$ is constant over $v \in [a, b)$ and over $v \in [b, c]$. Since $r(\Phi_r, a) < r(\Phi_r, b)$ and since $r(\Phi_r, v)$ is monotone increasing on $[0, a]$ and $[c, 1]$ it follows that $\arg \max_v r(\Phi_r, v) \subseteq [a, b) \cup \{1\}$. For sufficiently small ε , $r(\Phi_r, a) \approx p_0$ while $r(\Phi_r, 1) \approx 1 - p_0$ and thus given $p_0 > \frac{1}{2}$ we obtain $[a, b) = \arg \max_v r(\Phi_r, v)$ and thus $\max_{F: F(1)=\varepsilon} r(\Phi_r, F) = r(\Phi_r, F_2)$.

Concerning F_3 let $v_1 = \inf \{v : F_0(v - \varepsilon) \geq \varepsilon\}$. We have to show that $r(\Phi_r, v_0) \leq r(\Phi_r, v_2)$ for $v_0 \leq v_1 \leq v_2$. Given the above it is sufficient to consider only $v_0 = v_1$ and $v_2 = c$ where $r(\Phi_r, c) = c - \mathbb{E}(\Phi_r)$. Let $\gamma \triangleq 2 \sup_{v > 0} (v/F_0(v))$, then

$$r(\Phi_r, v_1) = v_1 \leq \varepsilon + \gamma F_0(v_1 - \varepsilon) \leq \varepsilon(1 + \gamma).$$

On the other hand, we showed in Step 1 that $c'(0) = \infty$ and in the proof of Proposition 6, based only on arguments in Step 1, we find that $\frac{\partial}{\partial \varepsilon} \mathbb{E}(\Phi_r) |_{\varepsilon=0} < \infty$ so $\frac{\partial}{\partial \varepsilon} r(\Phi_r, c) |_{\varepsilon=0} = \infty$. Hence, $r(\Phi_r, v_1) < r(\Phi_r, c)$ for ε sufficiently small which shows that $\min_{F: F(1)=\varepsilon} r(\Phi_r, F) = r(\Phi_r, F_3)$.

Finally, consider F_1 . More mass cannot be allocated to regret maximizing values $[a, b)$ as $F_1(b) = F_0(b + \varepsilon)$, weight on values below a and above c are shifted up as far as possible as $F_v(v) = F_0(v - \varepsilon)$ for $v < a$ and $c < v < 1$ and allocation of F_1 for $F_1 \in (F_1(b), F_0(c - \varepsilon))$ will not influence regret as $r(\Phi_r, v)$ is constant on $[b, c]$. This completes the proof.

The case of $p_0 \leq \frac{1}{2}$ proceeds in an analogous manner. It is easily shown that there exist parameters a, b, c such that $a < b < c$ and $a < p_0 < c$ such that:

$$\begin{aligned} F_0(a - \varepsilon) - \varepsilon &= 1 - \frac{b^2 f_0(b + \varepsilon)}{a}, \\ F_0(b + \varepsilon) &= 1 - \frac{b^2 f_0(b + \varepsilon)}{b} + \varepsilon, \\ F_0(c - \varepsilon) - \varepsilon &= 1 - \frac{b^2 f_0(b + \varepsilon)}{c}, \end{aligned}$$

where

$$b'(0) = -\frac{f_0(p_0) + p_0 f_0'(p_0) - 1}{2f_0(p_0) + p_0 f_0'(p_0)} = -1 + \frac{f_0(p_0) + 1}{2f_0(p_0) + p_0 f_0'(p_0)}.$$

The least favorable demand F_r is now given by:

$$F_r(v) \triangleq \begin{cases} \max\{0, F_0(v - \varepsilon) - \varepsilon\}, & \text{if } v \in [0, a], \\ 1 - \frac{b^2 f_0(b + \varepsilon)}{v}, & \text{if } v \in (a, c), \\ \max\{0, F_0(v - \varepsilon) - \varepsilon\}, & \text{if } v \in [c, 1], \\ 1 & \text{if } v = 1, \end{cases}$$

decomposed as $F_r \triangleq F_1 + F_2 - F_3$ where we define the measures F_1, F_2, F_3 as follows:

$$F_1(v) \triangleq \begin{cases} F_0(v - \varepsilon), & \text{if } v \in [0, a], \\ 1 - \frac{b^2 f_0(b + \varepsilon)}{v} + \varepsilon, & \text{if } v \in (a, c), \\ F_0(v - \varepsilon), & \text{if } v \in [c, 1], \\ 1, & \text{if } v = 1. \end{cases}$$

$$F_2(v) \triangleq \begin{cases} 0, & \text{if } v \in [0, 1), \\ \varepsilon, & \text{if } v = 1, \end{cases}$$

and

$$F_3(v) \triangleq \min\{F_0(v - \varepsilon), \varepsilon\}, \text{ if } v \in [0, 1].$$

Lemma 1 can be applied to show that $F_r \in \mathcal{P}_\varepsilon(F_0)$. In contrast to the previous case of $p_0 > \frac{1}{2}$, now $v = 1$ maximizes $r(\Phi_r, v)$ so that F_2 puts all mass at $v = 1$. ■

Proof of Proposition 6. We obtain that

$$\mathbb{E}[\Phi_r] = \int_a^c \frac{1}{p} dp + b \left(1 - \int_a^c \frac{1}{p} dp \right) = c - a + b \left(1 - \ln \frac{c}{a} \right).$$

As a, b , and c are differentiable as shown in Step 1 of Proposition 5, we have:

$$\frac{\partial}{\partial \varepsilon} \mathbb{E}[\Phi_r] = \frac{b-a}{a} a'(\varepsilon) + \frac{c-b}{c} c'(\varepsilon) + \left(1 - \ln \frac{c}{a}\right) b'(\varepsilon).$$

Using (14), (15) and (17) we obtain for $p_0 > \frac{1}{2}$:

$$\frac{\partial}{\partial \varepsilon} \mathbb{E}[\Phi_r] |_{\varepsilon=0} = -1 + \frac{f_0(p_0) - 1}{2f_0(p_0) + p_0 f_0'(p_0)} < -\frac{1}{2}.$$

The same operations yield for $p_0 < \frac{1}{2}$ yield

$$\frac{\partial}{\partial \varepsilon} \mathbb{E}[\Phi_r] |_{\varepsilon=0} = -1 + \frac{f_0(p_0) + 1}{2f_0(p_0) + p_0 f_0'(p_0)} > -1,$$

which concludes the proof. ■

Proof of Proposition 7. Following Proposition 5, $\lim_{\varepsilon \rightarrow 0} a'(\varepsilon) = -\infty$ and $\lim_{\varepsilon \rightarrow 0} c'(\varepsilon) = \infty$ and therefore the size of the menu is increasing in ε for ε sufficiently small which proves (1). Next we verify (2). Assume $a < v < b$. Then $q_r(v) = \ln \frac{v}{a}$ and $t_r(v) = \int_a^v y \frac{1}{y} dy = v - a$ so given $a' < 0$ for ε small we obtain $\frac{\partial}{\partial \varepsilon} q_r(v) > 0$, $\frac{\partial}{\partial \varepsilon} t_r(v) > 0$ and

$$\frac{\partial t_r(v)}{\partial \varepsilon q_r(v)} = \frac{(v-a) \frac{1}{a} - \ln \frac{v}{a}}{\left(\ln \frac{v}{a}\right)^2} a'(\varepsilon) < 0,$$

as

$$\frac{d}{dv} \left((v-a) \frac{1}{a} - \ln \frac{v}{a} \right) = \frac{1}{a} - \frac{1}{v} > 0.$$

Assume $b < v < c$. Then $q_r(v) = 1 - \ln \frac{c}{v}$ and

$$t_r(v) = v - a + \left(1 - \ln \frac{c}{a}\right) b = \mathbb{E}[t_r] + v - c,$$

so $\frac{\partial}{\partial \varepsilon} q_r(v) < 0$, $\frac{\partial}{\partial \varepsilon} t_r(v) < 0$ and

$$\frac{\partial t_r(v)}{\partial \varepsilon q_r(v)} = \frac{\frac{\partial}{\partial \varepsilon} \mathbb{E}[t_r]}{1 - \ln \frac{c}{v}} + \frac{\frac{1}{c} (\mathbb{E}[t_r] + v - c) - \left(1 - \ln \frac{c}{v}\right)}{\left(1 - \ln \frac{c}{v}\right)^2} c'(\varepsilon) < 0,$$

where we use the fact that $c'(\varepsilon)$ is large and

$$\frac{d}{dv} \left(\frac{1}{c} (\mathbb{E}[t_r] + v - c) - \left(1 - \ln \frac{c}{v}\right) \right) = \frac{1}{c} - \frac{1}{v} < 0,$$

for sufficiently small ε which proves (2). ■

Proof of Proposition 8. 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

$$\pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(\Phi_r, F_\varepsilon) \geq \gamma. \quad (21)$$

Assume that (Φ_r, F_r) is a saddle point of the regret problem (SP_r) given $\varepsilon > 0$. Then

$$\pi(\Phi_r, F_r) = \pi(p^*(F_r), F_r).$$

We can rewrite the left hand side of (21) as follows:

$$\begin{aligned} & \pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(\Phi_r, F_\varepsilon) \\ = & \pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(p^*(F_r), F_r) + \pi(\Phi_r, F_r) - \pi(\Phi_r, F_\varepsilon). \end{aligned} \quad (22)$$

Using (SP_r) we also obtain

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

so that:

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

Entering this into (22) we obtain from (21) that:

$$\pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(p^*(F_r), F_r) + \int v dF_r(v) - \int v dF_\varepsilon(v) \geq \gamma. \quad (23)$$

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 that

$$\int v dF_r(v) - \int v dF_\varepsilon(v) < \gamma/2, \quad (24)$$

if ε is sufficiently small.

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

$$\pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(p^*(F_r), F_r) < \gamma/2, \quad (25)$$

if ε is sufficiently small. Comparing (23) to (24) and (25) yields the desired contradiction. ■

References

- ANSCOMBE, F., AND R. AUMANN (1963): “A Definition of Subjective Probability,” *Annals of Mathematical Statistics*, 34, 199–205.
- BELL, D. (1982): “Regret in Decision Making under Uncertainty,” *Operations Research*, 30, 961–981.
- BERGEMANN, D., AND K. SCHLAG (2008): “Pricing Without Priors,” *Journal of the European Economic Association Papers and Proceedings*, 6, 560–569.
- BOSE, S., E. OZDENOREN, AND A. PAPE (2006): “Optimal Auctions with Ambiguity,” *Theoretical Economics*, 1, 411–438.
- CHE, Y., AND I. GALE (2000): “The Optimal Mechanism for Selling to a Budget Constrained Buyer,” *Journal of Economic Theory*, 92, 198–233.
- DENECKERE, R., AND R. MCAFEE (1996): “Damaged Goods,” *Journal of Economics and Management Strategy*, 5, 149–174.
- DUDLEY, R. (2002): *Real Analysis and Probability*. Cambridge University Press, Cambridge, England.
- ELLSBERG, D. (1961): “Risk, Ambiguity and the Savage Axioms,” *Quarterly Journal of Economics*, 75, 643–669.
- ENGELBRECHT-WIGGANS, R. (1989): “The Effect of Regret on Optimal Bidding in Auctions,” *Management Science*, 35, 685–692.
- ENGELBRECHT-WIGGANS, R., AND E. KATOK (2007): “Regret in Auctions: Theory and Evidence,” *Economic Theory*, 33, 81–101.
- FILIZ-OZBAY, E., AND E. OZBAY (2007): “Auctions with Anticipated Regret: Theory and Experiment,” *American Economic Review*, 97, 1407–1418.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.

- HAMPEL, F. R., E. M. RONCHETTI, P. ROUSSEEUW, AND W. STAHEL (1986): *Robust Statistics - The Approach Based on Influence Functions*. John Wiley and Sons, New York.
- HANNAN, J. (1957): "Approximation to Bayes Risk in Repeated Play," in *Contributions to the Theory of Games*, ed. by M. Dresher, A. Tucker, and P. Wolfe, pp. 97–139. Princeton University Press, Princeton.
- HANSEN, L., AND T. SARGENT (2007): *Robustness*. Princeton University Press, Princeton.
- HARRIS, M., AND A. RAVIV (1981): "A Theory of Monopoly Pricing Schemes with Demand Uncertainty," *American Economic Review*, 71, 347–365.
- HAYASHI, T. (2008): "Regret Aversion and Opportunity Dependence," *Journal of Economic Theory*, 139, 242–268.
- HUBER, P. (1964): "Robust Estimation of a Location Parameter," *Annals of Mathematical Statistics*, 35, 73–101.
- (1965): "A Robust Version of the Probability Ratio Test," *Annals of Mathematical Statistics*, 36, 1753–1758.
- HUBER, P. J. (1981): *Robust Statistics*. John Wiley and Sons, New York.
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): "A Smooth Model of Decision-Making under Ambiguity," *Econometrica*, 73, 1849–1892.
- LINHART, P., AND R. RADNER (1989): "Minimax - Regret Strategies for Bargaining over Several Variables," *Journal of Economic Theory*, 48, 152–178.
- LOOMES, G., AND R. SUGDEN (1982): "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal*, 92, 805–824.
- LOPOMO, G., L. RIGOTTI, AND C. SHANNON (2006): "Uncertainty in Mechanism Design," Discussion paper.
- MASKIN, E., AND J. RILEY (1984): "Monopoly with Incomplete Information," *Rand Journal of Economics*, 15, 171–196.

- MILNOR, J. (1954): “Games Against Nature,” in *Decision Processes*, ed. by R. Thrall, C. Coombs, and R. Davis. Wiley, New York.
- MUSSA, M., AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18, 301–317.
- PRASAD, K. (2003): “Non-Robustness of some Economic Models,” *Topics in Theoretical Economics*, 3, 1–7.
- RENY, P. J. (1999): “On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games,” *Econometrica*, 67, 1029–1056.
- RILEY, J., AND R. ZECKHAUSER (1983): “Optimal Selling Strategies: When to Haggle, When to Hold Firm,” *Quarterly Journal of Economics*, 98, 267–290.
- SAVAGE, L. (1951): “The Theory of Statistical Decision,” *Journal of the American Statistical Association*, 46, 55–67.
- SELTEN, R. (1989): “Blame Avoidance as Motivating Force in the First Price Sealed Bid Private Value Auction,” in *Economic Essays in Honor of Werner Hildenbrand*, ed. by G. Debreu, W. Neuefeind, and W. Trockel, pp. 333–344. Springer Verlag, Heidelberg.
- STOYE, J. (2008): “Axioms for Minimax Regret Choice Correspondences,” Discussion paper, New York University.
- STRASSEN, V. (1965): “The Existence of Probability Measures with Given Marginals,” *Annals of Mathematical Statistics*, 36, 423–439.
- VON NEUMANN, J. (1928): “Zur Theorie der Gesellschaftsspiele,” *Mathematische Annalen*, 100, 295–300.

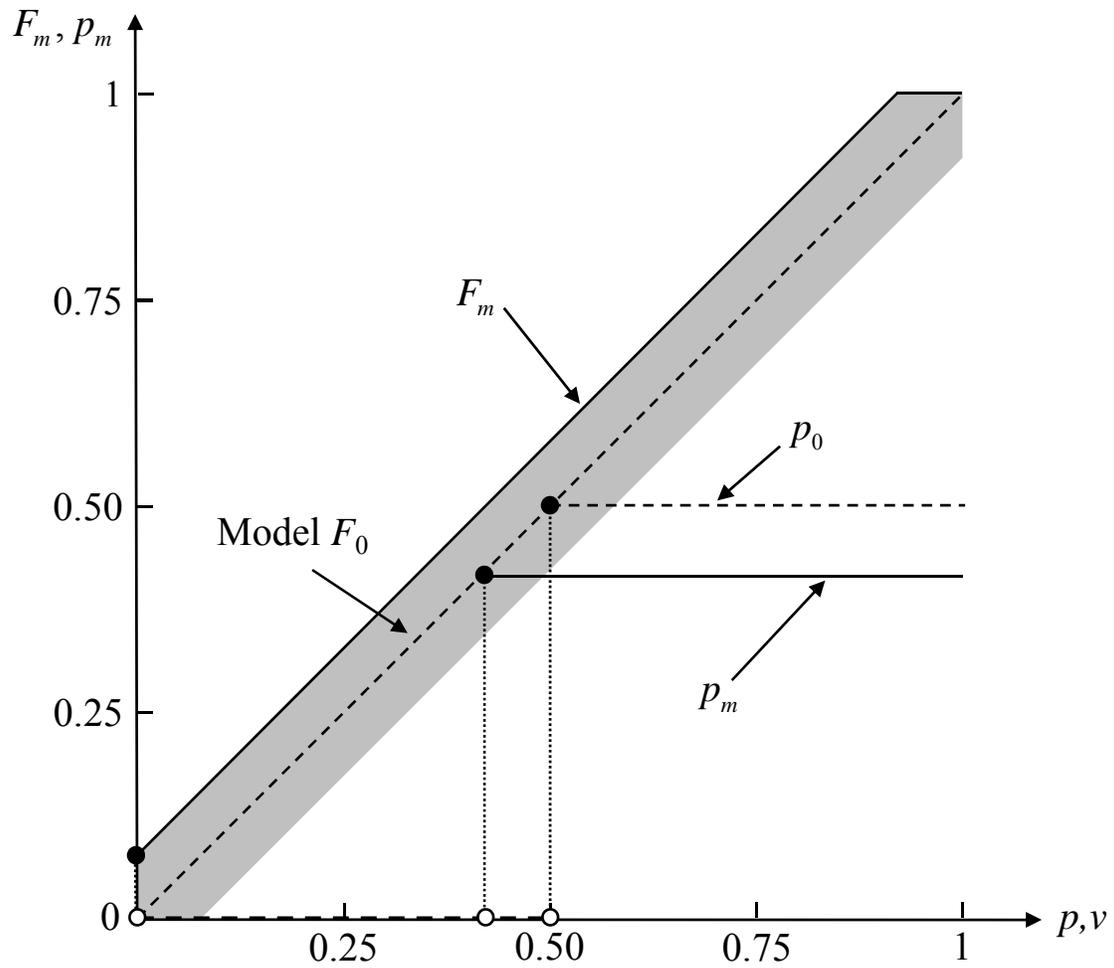


Figure 1. Maximin Pricing p_m and Least Favorable Demand F_m
 Model Distribution $F_0(v) = v$, Neighborhood $\varepsilon = 0.05$

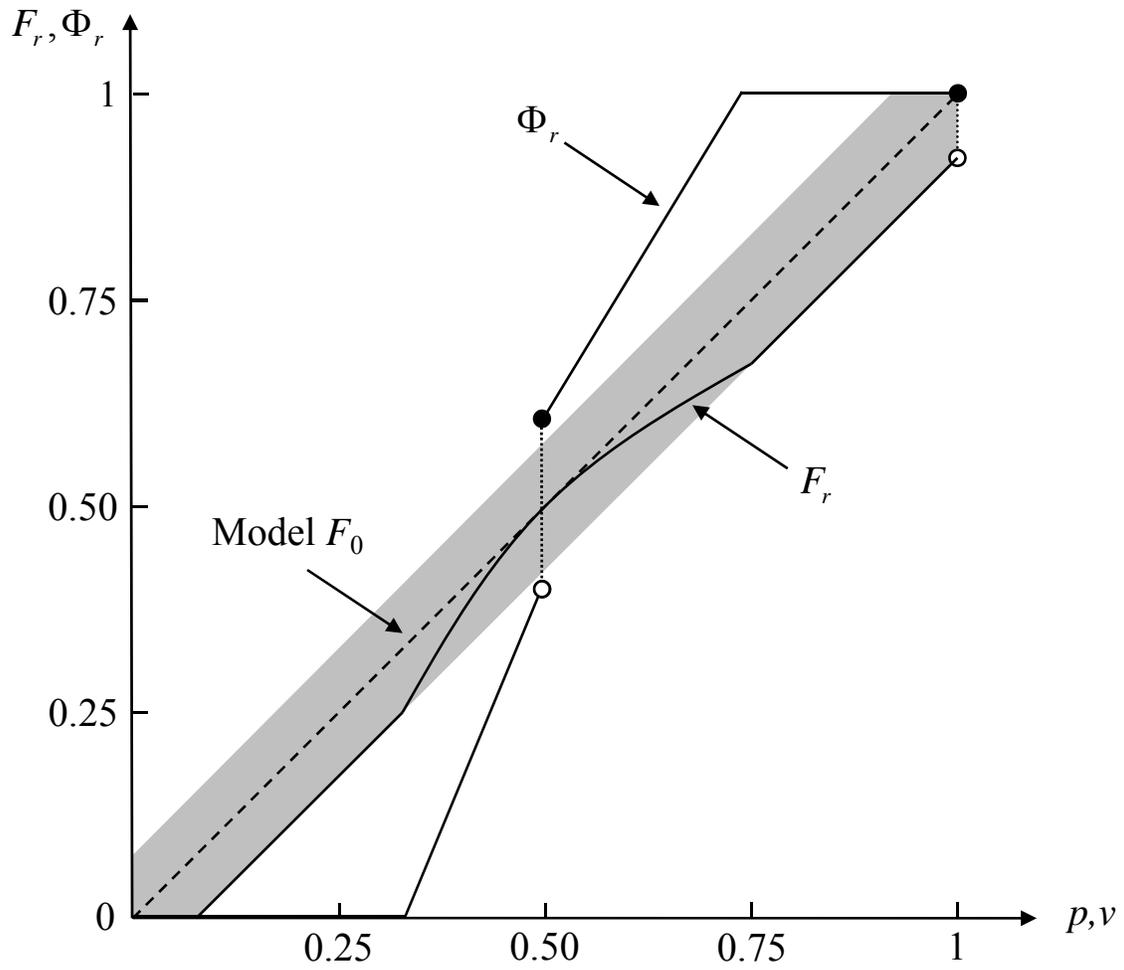


Figure 2. Minimax Pricing Φ_r and Least Favorable Demand F_r
 Model Distribution $F_0(v) = v$, Neighborhood $\varepsilon = 0.05$

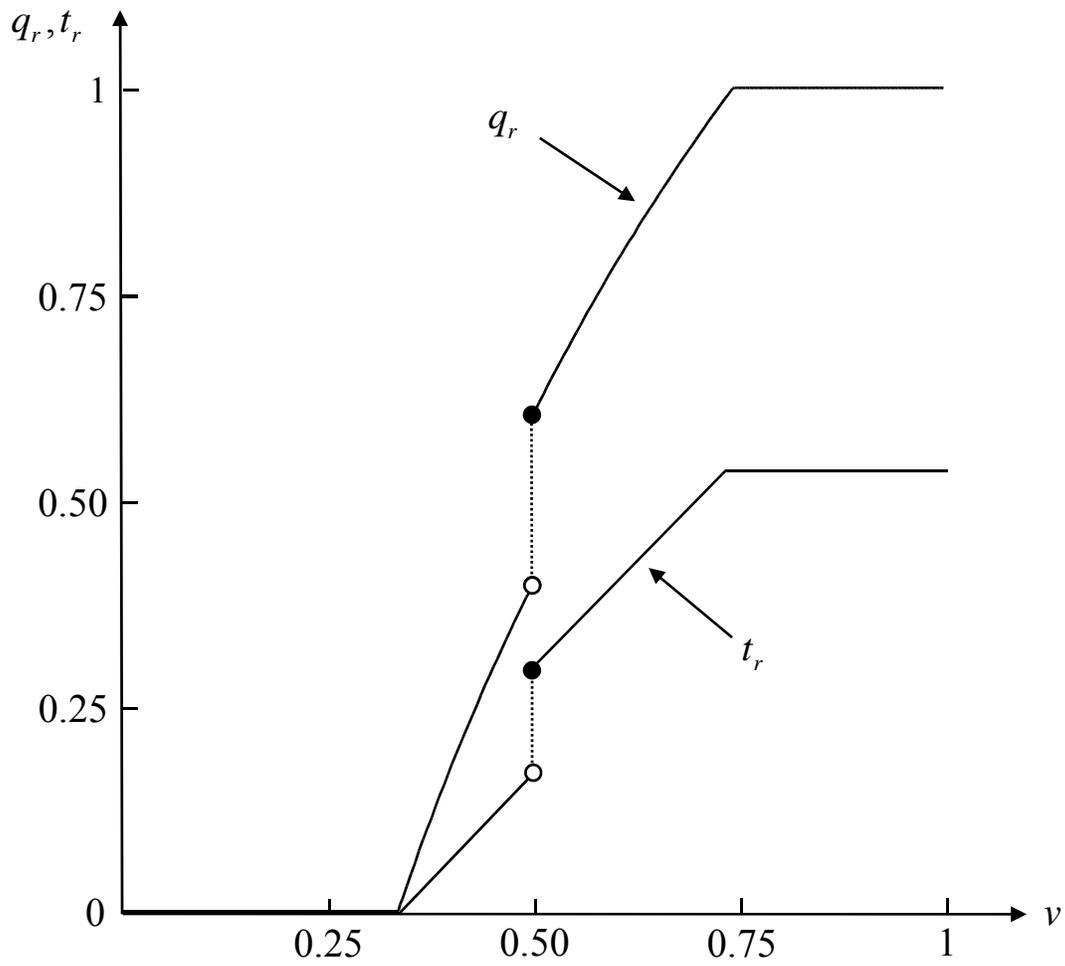


Figure 3. Minimax Menu Pricing (q_r, t_r)
 Model Distribution $F_0(v) = v$, Neighborhood $\varepsilon = 0.05$