

ROBUST MONOPOLY PRICING

By

Dirk Bergemann and Karl Schlag

July 2005

Revised April 2007

COWLES FOUNDATION DISCUSSION PAPER NO. 1527R



**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[‡]

First Version: May 2003

This Version: April 2007

Abstract

We consider a robust version of the classic problem of optimal monopoly pricing with incomplete information. In the robust version of the problem the seller only knows that demand will be in a 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. 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. The resulting optimal pricing policy under either criterion is robust to the model uncertainty. Finally we derive distinct implications of how a monopolist responds to an increase in ambiguity under each criterion.

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

JEL CLASSIFICATION: C79, D82

*The first author gratefully acknowledges support by NSF Grants #SES-0095321, #CNS-0428422 and a DFG Mercator Research Professorship at the Center of Economic Studies at the University of Munich. We thank the Editor, Eddie Dekel and three anonymous referees for valuable comments. We thank 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" at Yale University for many helpful comments. We thank the editor, Eddie Dekel, and three anonymous referees for their valuable comments.

[†]Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06511, dirk.bergemann@yale.edu and CEPR.

[‡]Department of Economics, European University Institute, schlag@iue.it

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. Naturally, 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. The optimal pricing policy is the most elementary instance of a revenue maximizing problem.

We investigate the robustness of the optimal selling policy by enriching the standard model to account for model ambiguity. Instead of assuming a given demand distribution from which the buyer is drawn, the seller is only assumed to believe that the demand distribution will be in the neighborhood of a given model distribution. The enlargement of the set of possible distributions represents the model ambiguity.

The objective of this paper is to demonstrate that we can relax the rigid Bayesian model by considering robust decision making. We maintain a formal approach by building on axiomatic decision theory and obtain interesting new insights for monopoly pricing. The methodological insight is that robustness is generated by considering decision making under multiple priors. We then present rich comparative statics results in terms of the response of prices to an increase in ambiguity and uncover a novel role for menu pricing. Thus, the analysis of the robust pricing problem leads to testable hypotheses regarding the behavior of the seller.

Currently, 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 & 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 was first suggested by Savage (1951) and axiomatized by Milnor (1954). The minimax regret criterion was recently adapted to multiple priors by Hayashi (2006) and by Stoye (2006). Here the decision maker takes the maximum of the expected regret as the prior varies and chooses an action that minimizes the maximum expected regret.

In this paper, we shall analyze the optimal pricing policy under both criteria. We analyze the optimal policies when the ambiguity is represented by a neighborhood around a given model distribution. We define the notion of a neighborhood through the usual metric of

weak convergence, the Prohorov metric. In the Prohorov metric two distributions are close to each other if they permit with *large* probability *small* changes in the valuations and with *small* probability *large* changes in the valuations. The analysis of the policies under the two decision criteria will reveal that either criterion leads to a robust policy in the following sense. We say that a candidate policy 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 ambiguity 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 distribution to the objective of the seller. When the decision maker is maximizing the minimum expected utility among the class of priors, the least favorable demand is always given by the distribution which puts maximal weight on the lowest quantiles subject to the restriction that the selected distribution is in the neighborhood of the model distribution. As the objective of nature is to minimize the revenue of the seller, the least favorable demand is the one which minimizes the potential revenue at any possible price level. In particular as we increase the ambiguity represented by an increase in the size of the neighborhood, the least favorable demand increases the weight on the lower quantiles of the distribution. In consequence the best response of the seller always consists in lowering her price deterministically.

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 seller and nature. The regret of the seller is the difference between the actual valuation of a buyer for the object and the actual revenue 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 revenue. In the equilibrium of the zero-sum game, the optimal pricing policy for 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 the 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 a true quantity in the case of a divisible object or a probability of obtaining the indivisible object.

The intuition regarding the price policy with regret is easy to establish in comparison to the revenue maximizing policy for a given distribution. 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 & Raviv (1981) and Riley & 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 revenue 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 ambiguity may now lead to an increase in the expected price. In the special case of linear model distribution we find that expected price increases if the optimal price under the model is low and decreases if the optimal price under the model is high.

From an axiomatic perspective, the maximin and minimax criteria represent different departures from the standard model of Anscombe & Aumann (1963). The maximin decision criterion emerges by replacing the independence axiom with the weaker certainty independence axiom and adding a convexity axiom. Certainty independence requires that preferences between two given acts remain unchanged when mixing both with some constant act. The minimax regret criterion emerges by maintaining the independence axiom but relaxing the axiom of independence of irrelevant alternatives. It is postulated that the most preferred choice does not change when a new act is added as long as the additional act does not change the best outcome that can be achieved in each state. The weaker version holds vacuously in perfect information environment, i.e. when the state is known before

the choice. The convexity axiom and a variation of the betweenness axiom completes the characterization. Either approach allows to consider sets of distributions. Both maximin utility and minimax regret criteria do not contradict subjective expected utility theory, and we may interpret them as alternative axiomatic systems for selecting subjective priors.¹

It should be pointed out that 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. Neither the axioms refer to foregone opportunities nor is it important whether or not ex post additional information becomes available. As in the case of maximin criterion of Gilboa & Schmeidler (1989), the minimax regret criterion in Hayashi (2006) and Stoye (2006) is completely characterized by a set of axioms.²

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. The statistical problem is to distinguish on the basis of a sample between two known alternative distributions. The model misspecification and consequent concern of robustness comes from the fact that each one 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. In the economic context, a recent article by Prasad (2003) shows that the standard optimal pricing policy is not robust to small model misspecifications.

A recent paper by Bose, Ozdenoren & Pape (2006) determines the optimal auction in the presence of an ambiguity averse seller and ambiguity averse bidders. As we consider the optimal pricing problem the ambiguity aversion of the buyers is immaterial as there is no strategic interaction across buyers. Lopomo, Rigotti & Shannon (2006) consider a general mechanism design setting when the agents, but not the principal, have incomplete preferences due to Knightian uncertainty. The notion of regret was investigated in mechanism design by Linhart & Radner (1989) in the context of bilateral trade as well

¹Klibanoff, M. Marinacci & 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 $\mathbf{■}$ representing ambiguity aversion. The additional elements, belief μ and ambiguity index $\mathbf{■}$, 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.

²In 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 to the more behavioral approaches to regret offered by Bell (1982) and Loomes & Sugden (1982).

as by Engelbrecht-Wiggans (1989) and Selten (1989) in the context of auctions. Linhart & Radner (1989) analyze minimax regret strategies in a bilateral bargaining framework. In contrast to the incomplete information environment here, the bulk of the analysis in Linhart & Radner (1989) is concerned with bilateral trade under complete information. In Engelbrecht-Wiggans (1989) and Selten (1989) the first and second price sealed bid auctions are analyzed incorporate regret for the bidders. Recently, Engelbrecht-Wiggans & Katok (2007) and Filiz & Ozbay (2006) present experimental evidence 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 criterion. In Section 4 we characterize the pricing policy under the minimax 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

A seller offer an object for sale to an unknown demand. The demand is either generated by a single large buyer or by many small buyers. In the paper we focus on the case of a single large buyer and later show how the results generalize naturally to the case of many small buyers. Accordingly, the seller faces a single potential buyer with value v for a unit of the object. The value v of the object is private information of the buyer and unknown to the seller. The valuation v of the buyer is an element of the unit interval, $v \in [0, 1]$.³ The marginal cost of production is constant and normalized to zero. The buyer wishes to buy at most one unit of the object.

The seller sets a price p , the *profit* of selling the object at price p 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:

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

³More generally, we assume that the value of the buyer is (known by the seller) contained in some closed interval which we normalize without loss of generality to $[0, 1]$.

In the standard monopoly problem with incomplete information, the seller maximizes the expected profit for a given prior F over valuations. The expected profit given a distribution F is:

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

By extension, if the seller chooses a random pricing policy $\Phi \in \Delta\mathbb{R}_+$, then the expected profit is:

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

We denote the probabilistic price that maximizes the profit for given distribution F by $\Phi^*(F)$ so

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

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

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

2.2 Ambiguity

In contrast to the standard model of monopoly pricing in which the seller acts as if the valuation of the buyer is drawn from a (subjective) distribution F , we assume that the seller faces ambiguity in the sense of Ellsberg (1961). The ambiguity 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 ambiguity is thus quantified by the size of the neighborhood around the model distribution.⁴

Given the model distribution F_0 we denote by $p_0 = p^*(F_0)$ a profit maximizing price at F_0 . For the remainder of the paper we shall assume that (i) p_0 is the unique maximizer of the profit function for the model distribution, (ii) the profit function, $\pi(p, F_0)$ at the model distribution F_0 is strictly concave near p_0 and (iii) the density f_0 is continuously

⁴This model of ambiguity permits at least two different interpretations. First, the ε neighborhood around the model distribution F_0 can be understood as a model with multiple priors. Second, the ε neighborhood can be viewed as an ε perturbation of the original model distribution F_0 . By considering the larger set of possible distributions the decision maker is protecting herself against measurement error and/or additional information which may slightly change the original model. We adopt throughout the first perspective, but it related to second perspective, prominent in statistical decision theory.

differentiable near p_0 . These regularity assumptions enable us 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 describe ε neighborhoods of the model distribution $F_0(v)$ by the *Prohorov neighborhood*, denoted by $\mathcal{P}_\varepsilon(F_0)$, and associated metric:

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

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

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

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 (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.⁵

Maximin Profit The seller maximizes the minimum profit by solving

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

Accordingly, we say that Φ_m *attains maximin profit*. We refer to F_m as a *least favorable demand* given Φ if

$$F_m \in \arg \min_{F \in \mathcal{P}_\varepsilon(F_0)} \pi(\Phi, F)$$

so the least favorable demand F_m minimizes profit under the policy Φ .

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

$$r(p, v) \triangleq v - p \mathbb{I}_{\{v \geq p\}} = v - \pi(p, v), \quad (2)$$

⁵The Prohorov metric applies to discrete and continuous distributions. In contrast, the Kullback-Leibler distance only defines neighborhoods for continuous distributions. A related model is the contamination “neighborhood” $\mathcal{N}_\varepsilon(F_0)$: $\mathcal{N}_\varepsilon(F_0) = \{F \mid F = (1 - \varepsilon)F_0 + \varepsilon H \text{ for some } H \in \Delta \mathbb{R}_+\}$. Yet the contamination “neighborhood” is not a neighborhood in the sense of the weak topology.

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:

$$r(\Phi, F) \triangleq \int r(p, v) d\Phi(p) dF(v) = \int v dF(v) - \int \pi(p, F) d\Phi(p). \quad (3)$$

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 .⁶

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

We refer to F_r as a *least favorable demand* given the pricing policy Φ if F_r maximizes regret under the pricing policy Φ_r :

$$F_r \in \arg \max_{F \in \mathcal{P}_\varepsilon(F_0)} r(\Phi_r, F).$$

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.⁷

⁶The fact that buyer value is contained in some known bounded set provides an upper bound on regret. If the support of $F \in \mathcal{P}_\varepsilon(F_0)$ would not be uniformly bounded then regret would be unbounded on $\mathcal{P}_\varepsilon(F_0)$ even if the support of F_0 is contained in $[0, 1]$. The neighborhood of the model F_0 puts restrictions on the support. Imposing upper bounds on the willingness to pay are natural once one thinks about realistic applications.

⁷Alternatively we could restrict the seller to offer the same price to all buyers. The present analysis of the single buyer then generalizes after imposing only that the marginal distribution of each buyer belongs to $\mathcal{P}_\varepsilon(F_0)$. The least favorable demand will then involve all buyers realizing the same valuation.

2.3 Robust Pricing Policy

For a given model distribution F_0 we identify a price policy as a class of probabilistic prices $\{\Phi_\varepsilon\}$ dependent on the size of the neighborhood ε .

Definition 1 (Robust Pricing Policy)

A pricing policy $\{\Phi^\varepsilon\}$ 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). The robust policy is allowed to depend on the size ε of the neighborhood.⁸ In contrast to minimax regret where profits are compared to best choices ex-post, robustness involves comparing expected profits to those attainable ex-ante when the valuation is drawn from a known distribution.

In the context of optimal monopoly pricing Prasad (2003) shows that the optimal policy is not robust if F_0 is a Dirac distribution. For a given model distribution F_0 , there are potentially many robust 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 Profit

We consider the problem of the monopolist who wishes to maximize the minimum profit for all distribution in the neighborhood of the model distribution F_0 . Following Neumann & Morgenstern (1953), the maximin pricing rule and the least favorable demand can be viewed as the equilibrium strategies of a game between the seller and adversarial nature (provided such an equilibrium exists). 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:

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

⁸The recent literature on robust decision making in macroeconomics, see Hansen & Sargent (2004) for a survey, uses the same notion of robustness for maximizing the minimum utility in intertemporal decision-making.

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

The objective of adversarial nature is to lower the expected revenue of the seller. For a given price p offered by the seller, the least favorable demand is achieved by increasing the cumulative probability of valuations strictly below v as much as possible given the neighborhood. The least favorable demand then minimizes the probability of sale by the seller. Given the model distribution F_0 and the size ε of the neighborhood the resulting distribution is uniquely determined for every p . The equilibrium analysis is now simplified by the fact that the least favorable demand does not depend on the probabilistic price of the seller. The least favorable demand is thus achieved by shifting the probabilities as far down as possible.

The construction of a least favorable distribution 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:

$$F_m(v) \triangleq \min \{F_0(v + \varepsilon) + \varepsilon, 1\}. \quad (4)$$

The first shift represents the possibility that small changes in valuations may occur with large probability, whereas the second shift represents the idea of large changes with a small probability.

Given that the demand F_m that minimizes profits does not depend on the offered prices, the monopolist acts as if the demand given by F_m . In consequence, the seller maximizes profits at F_m by choosing a deterministic price p_m where $p_m = p^*(F_m)$.

Proposition 1 (Maximin Profit)

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

It is then natural to ask how the optimal price will change with an increase in ambiguity. The rate of the change in the price depends on the curvature of the profit function at the model distribution. 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 have the following comparative static result.

Proposition 2 (Pricing under Maximin Profit)

The price p_m responds to an increase in ambiguity at $\varepsilon = 0$ by:

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

Accordingly, the maximin price responds to an increase in ambiguity 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 sales of the monopolist.

Consider now the profits attained by the maximin price p_m when facing some distribution F within the neighborhood of the model F_0 . These profits will be at least as high as those obtained when facing the least favorable demand F_m as the least favorable demand involves maximally decreasing all values within the neighborhood of the model. As we show that optimal profits when facing a known distribution are continuous in this distribution this means that profits achieved by p_m when facing F are close to those achieved by $p^*(F)$ when facing F . The maximin pricing rule thus qualifies as robust pricing rule.

Proposition 3

The pricing policy $\{p_m^\varepsilon\}$ consisting of the maximin prices is a robust policy.

4 Minimax Regret

4.1 Probabilistic Pricing

Next we consider the minimax regret problem of the seller. Analogous to case of maximin above, the minimax regret strategy Φ_r and the least favorable demand F_r are the equilibrium policies of a zero-sum game (provided such an equilibrium exists). 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:

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

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 least favorable 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. The regret of the seller depends critically on the price offered by the seller. 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. It now becomes evident that a single price will always expose the seller to substantial regret. Conversely, the least favorable 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 least favorable 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 ambiguity. 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 to this mass point.

Proposition 5 (Minimax Regret)

1. If ε is sufficiently small and $f_0(0) > 0$, then 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 satisfy $0 < a < b < c < 1$ and $a < p_0 < c$.
3. The boundary points a, b and c respond to an increase in ambiguity at $\varepsilon = 0$ as follows:

- (a) $\lim_{\varepsilon \rightarrow 0} a'(\varepsilon) = -\infty$,
- (b) $\lim_{\varepsilon \rightarrow 0} b'(\varepsilon) \in [-1, \frac{1}{2}]$ and,
- (c) $\lim_{\varepsilon \rightarrow 0} c'(\varepsilon) = \infty$.

We construct the minimax regret probabilistic price 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 $p \in [a, c]$. To protect against nature either undercutting or moving mass to highest possible prices the interval over which the seller randomizes increases substantially as ambiguity increases. On the other hand, the mass point remains close as ambiguity increases.

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 1: MINIMAX PRICING AND WORST CASE DEMAND

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 $\varepsilon = 0.04$, the impact of the ambiguity 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 the expected price to a marginal increase in ambiguity can be explained by the first order effects. For a small level of ambiguity, 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 ambiguity, the marginal change in regret can then be computed by holding the probabilistic price of the seller at the optimal price p_0 without ambiguity. Suppose then for the moment that $p_0 \leq \frac{1}{2}$. If the ambiguity 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 ambiguity would be $p_0 > \frac{1}{2}$, then the only modification would affect 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 ambiguity is now to find a probabilistic price which minimizes the additional regret

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

coming from the increase in ambiguity. Of course, the cost 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 ambiguity, 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 ambiguity 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

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

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 $\mathbb{E}[p_r]$ responds to an increase in ambiguity at $\varepsilon = 0$ by:

$$\frac{\partial}{\partial \varepsilon} \mathbb{E}[p_r] \Big|_{\varepsilon=0} = \begin{cases} -1 - \frac{f_0(p_0)+1}{\partial \pi^2(p_0, F_0)/\partial p^2} \geq -1 & \text{if } p_0 \leq \frac{1}{2} \\ -1 - \frac{f_0(p_0)-1}{\partial \pi^2(p_0, F_0)/\partial p^2} \leq -\frac{1}{2} & \text{if } p_0 > \frac{1}{2} \end{cases}. \quad (5)$$

We observe that for $p_0 > \frac{1}{2}$, the response of the expected price $\mathbb{E}[p_r]$ to an increase in ambiguity 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 ambiguity. As an implication from Proposition 6, 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 ambiguity has a downward effect on prices for all model distributions.

4.2 Menu Pricing

So far, our analysis assumed that the seller can only offer an indivisible object at some price p . We now extend the instruments of the seller and allow her to offer a menu of items. The

equilibrium policies with menus rather than single prices can be directly derived from the random pricing policies studied earlier and thus little new analysis will be necessary. 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 also emphasizes the role of robustness concerns in the optimal selling policies as menus would never be used in the standard setting for a given demand distribution.

If the allocative decision regards an indivisible object, or $x \in \{0, 1\}$, then a specific item on the menu assigns a probability of receiving the object at a corresponding price. If on the other hand, the allocative decision regards a continuous variable, or $x \in [0, 1]$, then a menu offers a variety of quantities at different prices. We observe that with the multiplicative utility $v \cdot x$ used here, the notions of probability and quantity are mathematically interchangeable. In a direct mechanism, a menu is a pair $(x(v), p(v))$ which maps a reported type v into a quantity $x(v)$ and price $p(v)$. We transform an equilibrium probabilistic price into a menu policy by defining the quantity assigned in the direct mechanism through:

$$x_r(v) \triangleq \Phi_r^*(v), \quad (6)$$

and the corresponding nonlinear prices as:

$$p_r(v) \triangleq \int_0^v y d\Phi_r(y). \quad (7)$$

By standard arguments recorded in Lemma 2 in the appendix this assignment of quantities to values defines an incentive compatible mechanism.

Form the point of view of menus, the minimax regret menu offered by seller then has three important characteristics. These properties can be described with reference to the mass point b : (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 given by the quantity offered at $v = b$, and given by $x^*(b)$. 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 of the seller. 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.

A natural 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 can offer lotteries or if the good is divisible 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 problems with ambiguity then offers an interesting and novel reason for menus. Despite the prevalence of menus, the literature currently offers only 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 & McAfee (1996) or if the buyers are budget constrained as in Che & Gale (2000).

The minimax regret response of the seller to an increase in ambiguity is perhaps even more 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 choice set for the buyers changes with an increase in the ambiguity. We define the size of the menu simply as the range of quantities offered by the seller (and accepted by some buyers) in equilibrium.

Proposition 7 (Menus and Ambiguity)

For small ambiguity:

1. *The size of the menu is increasing in ε .*
2. *The price per unit $p_r^*(v)/x_r^*(v)$ is decreasing in ε .*

As the ambiguity increases, the seller seeks to minimize her exposure 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 ambiguity. 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 ambiguity.

4.3 Robustness

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

Proposition 8 (Robustness)

If Φ_r attains minimax regret at F_0 for all sufficiently small ε then $\{\Phi_r^\varepsilon\}$ is robust at F_0 .

5 Conclusion

In this paper we analyzed robust pricing policies by a monopolist. We introduced robustness by allowing for multiple priors in the neighborhood of a model distribution. 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 revenue 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 ambiguity. The maximin policy uniformly maintains a deterministic price policy and uniformly lowers the price by an increase in ambiguity. In contrast, the minimax policy balances the downside versus the upside when responding to the ambiguity. Here the trade-off is optimally resolved by a probabilistic price. Importantly, the expected price does not necessarily decrease with an increase in ambiguity. 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 quantities 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.

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 wide class of design problems, such as the discriminating monopolist (as in Mussa & Rosen (1978) and Maskin & 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 had to make a decision as to whether or not to purchase the object. With the complete information by the buyer, there was no need to look for a robust purchasing rules. A substantial task would consequently arise by considering multi-agent design problems with incomplete information such as auctions, where it becomes desirable to “robustify” both the decisions of the buyers and the seller. The recent result by Segal (2003) and Chung & Ely (2003) regarding the sufficient conditions for the existence of dominant strategies for the bidders in optimal auctions might offer a first step in this direction. The complete solution of these problems poses a rich field for future research.

6 Appendix

The appendix contains some auxiliary results and 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 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 given p_m . ■

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 concavity assumptions on F_0 imply that we can apply the implicit function theorem at $\varepsilon = 0$ and this yields the statement to be proven. ■

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 $p_m = p^*(F_m)$ the proof is complete once we show that $\pi(p^*(F), F)$ is a continuous function of F with respect to the Prohorov neighborhood. Consider F, \widehat{F}_v such that $\widehat{F}_v \in \mathcal{P}_\varepsilon(F)$. Using the fact that

$$\widehat{F}_v(p) \leq F(p + \varepsilon) + \varepsilon,$$

we obtain

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

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

$$\pi(p^*(F), F) + 2\varepsilon \geq \pi(p^*(\widehat{F}_v), \widehat{F}_v) \geq \pi(p^*(F), F) - 2\varepsilon,$$

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{F} 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 continuity of

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

in γ the claim is established. ■

In order to derive the equilibrium policies in the case of small ambiguity 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 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 cumulative distribution functions only. In order to stay within ε distance of a given distribution function G one may first alter any value 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 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:

$$G(x - \varepsilon) \leq F_1(x) \leq G(x + \varepsilon), \quad F_2, F_3 \leq \varepsilon,$$

and

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

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

$$F_0(a - \varepsilon) - \varepsilon = 1 - \frac{b^2 f_0(b + \varepsilon)}{a}, \quad (8)$$

$$F_0(b + \varepsilon) = 1 - \frac{b^2 f_0(b + \varepsilon)}{b}, \quad (9)$$

$$F_0(c - \varepsilon) = 1 - \frac{b^2 f_0(b + \varepsilon)}{c}. \quad (10)$$

With respect to the existence of b , note that $b = p_0$ solves (9) 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 solution $b = b(\varepsilon)$ to (9) (with $b > 0$) exists for ε in a neighborhood of 0. To prove existence of c , define

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

Then $h(b) = F_0(b + \varepsilon) - F_0(b - \varepsilon)$ with

$$h'(b) = f_0(b + \varepsilon) - f_0(b - \varepsilon),$$

and

$$h''(b) = -\frac{2f_0(b + \varepsilon)}{b} - (f_0)'(b - \varepsilon) \approx -\frac{2f_0(p_0) + p_0 (f_0)'(p_0)}{p_0} < 0.$$

We note that $h(b) > 0$ by our earlier concavity assumptions on F_0 . Looking at the Taylor approximation of h near $v = b$ for small ε we obtain that there exists $c > b$ such that $h(c) = 0$ with $c \rightarrow p_0$ as $\varepsilon \rightarrow 0$. As for the existence of a , analogous calculations for $h(v) + \varepsilon$ show that there exists $a < b$ such that $h(a) + \varepsilon = 0$ 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 $2f_0(p_0) + p_0 (f_0)'(p_0) > 0$ we know that b is differentiable and by implicitly differentiating (9) we obtain:

$$b'(0) = -\frac{f_0(p_0) + p_0 (f_0)'(p_0)}{2f_0(p_0) + p_0 (f_0)'(p_0)} = -1 + \frac{f_0(p_0)}{2f_0(p_0) + p_0 (f_0)'(p_0)} \quad (11)$$

where $-1 \leq b'(0) \leq -1/2$. Next we show that a is differentiable. Since

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

we find that $b^2 f_0(b + \varepsilon) > a^2 f_0(a - \varepsilon)$ near $\varepsilon = 0$. Hence we can implicitly differentiate (8) to obtain

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

so

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

In particular we obtain that

$$\lim_{\varepsilon \rightarrow 0} a'(\varepsilon) = -\infty. \quad (13)$$

Similarly for c , we find that:

$$c'(\varepsilon) = -c \frac{c f_0(c - \varepsilon) + b f_0(b + \varepsilon)}{b^2 f_0(b + \varepsilon) - c^2 f_0(c - \varepsilon)}, \quad (14)$$

and hence

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

and in particular,

$$\lim_{\varepsilon \rightarrow 0} c'(\varepsilon) = \infty. \quad (15)$$

It now follows from (13) and (15) 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

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

where 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. Consider F_1 defined by

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

Then F_1 is a probability measure with $F_0(v - \varepsilon) \leq F_1(v)$. 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 $F_0(v - \varepsilon) \leq F_1(v) \leq F_0(v + \varepsilon)$.

Consider F_2 defined by:

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

Then

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

as

$$\frac{d}{dv}(v^2 f_0(v + \varepsilon)) = v^2 (f_0)'(v + \varepsilon) + 2v f_0(v + \varepsilon) > 0,$$

holds for ε sufficiently small and hence F_2 is weakly increasing with $F_2(1) = \varepsilon$. Since F_2 is also right continuous we obtain that F_2 is an additive probability measure.

Let F_3 be defined by

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

so $F_3(v)$ is an additive probability 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. Next we show that (G_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 $p \in [a, c]$. Similar to the calculations following the definition of F_1 it is easily shown that there exists $\eta > 0$ such that $1 - \frac{b^2 f_0(b + \varepsilon)}{v} < F_r(v)$ holds for all $v \in [p_0 - \eta, p_0 + \eta] \setminus [a, c]$ and all sufficiently small ε . Thus, for sufficiently small ε we obtain $[a, c] = \arg \max_{p \in [p_0 - \eta, p_0 + \eta]} \pi(p, F_r)$ and together with the upperhemicontinuity of profits that $[a, c] \subseteq \arg \max_p \pi(p, F_r)$.

Consider now the incentives of nature. Note that

$$r(G_r, F_r) = r(G_r, F_1) + r(G_r, F_2) - r(G_r, F_3), \quad (16)$$

where we choose F_2 and F_3 such that $F_2(1) = F_3(1) = \varepsilon$. In the following we show that each term in (16) is maximized separately. If nature could put all mass on a single value v , by construction of F_r nature would be indifferent over $v \in [a, b]$ and over $v \in [b, c]$. Since $r(G_r, v)$ is monotone increasing on $[0, a]$ and $[c, 1]$ it follows that $\arg \max_v r(G_r, v) \subseteq [a, b] \cup \{1\}$. For sufficiently small ε , $r(G_r, a) \approx p_0$ while $r(G_r, 1) \approx 1 - p_0$ and thus given $p_0 > \frac{1}{2}$ we obtain $[a, b] = \arg \max_v r(G_r, v)$.

Concerning F_3 let $\bar{v} = \inf \{v : F_0(v - \varepsilon) \geq \varepsilon\}$. We have to show that $r(G_r, \tilde{v}) \leq r(G_r, \hat{v})$ for $\tilde{v} \leq \bar{v} \leq \hat{v}$. Given the above it is sufficient to consider only $\tilde{v} = \bar{v}$ and $\hat{v} = c$ where $r(G_r, c) = c - \mathbb{E}(p_r)$. Let $\gamma \triangleq 2 \sup_{v>0} \frac{v}{F_0(v)}$. For v sufficiently small, $\gamma \geq \frac{v}{F_0(v)}$ and hence $r(G_r, \tilde{v}) = \tilde{v} \leq \varepsilon + \gamma F_0(\tilde{v} - \varepsilon) = \varepsilon(1 + \gamma)$. On the other hand, we showed in Step 1 that $\frac{\partial}{\partial \varepsilon} c|_{\varepsilon=0} = \infty$ and in the proof of Proposition 6 based only on arguments in Step 1 that $\frac{\partial}{\partial \varepsilon} \mathbb{E}(p_r)|_{\varepsilon=0} < \infty$ so $\frac{\partial}{\partial \varepsilon} r(G_r, c)|_{\varepsilon=0} = \infty$. Hence, $r(G_r, \tilde{v}) < r(G_r, c)$ for ε sufficiently small.

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 $r(G_r, v)$ is constant on $[b, c]$.

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

$$\begin{aligned} F_0(a - \varepsilon) - \varepsilon &= 1 - \frac{b^2 f_0(b + \varepsilon)}{a}, \\ F_0(b + \varepsilon) &= 1 - \frac{b^2 f_0(b + \varepsilon)}{b} + \varepsilon, \\ F_0(c - \varepsilon) - \varepsilon &= 1 - \frac{b^2 f_0(b + \varepsilon)}{c}, \end{aligned} \tag{17}$$

where

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

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 $F_r = F_1 + F_2 - F_3$ where

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

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

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

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 $r(G_r, v)$ so that F_2 puts all mass at $v = 1$. For the case of $p_0 = \frac{1}{2}$ Proposition 6 can be used to show that $r(G_r, 1) = 1 - \mathbb{E}[p_r] > r(G_r, a) = a$. As in the case where $p_0 > \frac{1}{2}$, $F_1(v) \leq F_0(v + \varepsilon)$ with tangency only at $v = b$ so F_1 again maximizes weight on $[a, b]$. $[a, b]$ is now only a local maximum of $r(G_r, v)$ but nevertheless it still follows easily that F_1 maximizes regret (use the fact that $F_0(b + \varepsilon) < F_0(c - \varepsilon)$). ■

Proof of Proposition 6. We obtain that

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

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

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

Inserting the value for $a'(\varepsilon)$, $b'(\varepsilon)$ and $c'(\varepsilon)$ from (11), (12) and (14) respectively, we obtain for $p_0 > \frac{1}{2}$:

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

The same operations yield the result for $p_0 < \frac{1}{2}$. ■

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 menu is increasing in ε for ε sufficiently small which proves (1). Next we verify (2). Assume $a < v < b$. Then $x^*(v) = \ln \frac{v}{a}$ and $p^*(v) = \int_a^v y \frac{1}{y} dy = v - a$ so given $a' < 0$ for ε small we obtain $\frac{\partial}{\partial \varepsilon} x^*(v) > 0$, $\frac{\partial}{\partial \varepsilon} p^*(v) > 0$ and

$$\frac{\partial p^*(v)}{\partial \varepsilon x^*(v)} = \frac{(v-a) \frac{1}{a} - \ln \frac{v}{a} a'(\varepsilon)}{(\ln \frac{v}{a})^2} < 0$$

as $\frac{d}{dv} \left((v-a) \frac{1}{a} - \ln \frac{v}{a} \right) = \frac{1}{a} - \frac{1}{v} > 0$. Thus, $x^*(v)v - p^*(v)$ is strictly increasing in ε .

Assume $b < v < c$. Then $x^*(v) = 1 - \ln \frac{c}{v}$ and $p^*(v) = v - a + \left(1 - \ln \frac{c}{a}\right) b = \mathbb{E}[p_r] + v - c$ so $\frac{\partial}{\partial \varepsilon} x^*(v) < 0$, $\frac{\partial}{\partial \varepsilon} p^*(v) < 0$ and

$$\frac{\partial p^*(v)}{\partial \varepsilon x^*(v)} = \frac{\frac{\partial}{\partial \varepsilon} \mathbb{E}[p_r]}{1 - \ln \frac{c}{v}} + \frac{\frac{1}{c} (\mathbb{E}[p_r] + v - c) - \left(1 - \ln \frac{c}{v}\right) c'(\varepsilon)}{\left(1 - \ln \frac{c}{v}\right)^2} < 0$$

where we use the fact that $c'(\varepsilon)$ is large and $\frac{d}{dv} \left(\frac{1}{c} (\mathbb{E}[p_r] + v - c) - (1 - \ln \frac{c}{v}) \right) = \frac{1}{c} - \frac{1}{v} < 0$ for ε small.

We obtain

$$\frac{\partial}{\partial \varepsilon} u(v) = \left(v - \frac{p^*(v)}{x^*(v)} \right) \frac{\partial}{\partial \varepsilon} x^*(v) - x^*(v) \frac{\partial}{\partial \varepsilon} \frac{p^*(v)}{x^*(v)} = \frac{c-v}{c} c'(\varepsilon) - \frac{\partial}{\partial \varepsilon} \mathbb{E}[p_r].$$

Since incentive compatibility implies that $x^*(v)v - p^*(v)$ is continuous in v and since x^* has an upwards jump at $v = b$ we obtain

$$\frac{p^*(b)}{x^*(b)} > \lim_{v \rightarrow b^-} \frac{p^*(v)}{x^*(v)}.$$

Clearly, $\frac{p^*(v)}{x^*(v)} > \frac{p^*(b)}{x^*(b)}$ for $v > b$ holds from above using right continuity of x^* . ■

The following lemma shows how probabilistic prices can be transformed into menus and vice versa.

Lemma 2 (Equivalence)

1. For any mixed pricing policy $\Phi(v)$ the menu $(x(v), p(v))$ is incentive compatible.
2. If $(x(v), p(v))$ is incentive compatible, then there exists a mixed pricing policy Φ such that $\pi(\Phi, v) \geq p(v)$ for all $v \in [0, 1]$.

Proof. First we show that if $g : [0, 1] \rightarrow [0, 1]$ is non decreasing then

$$vg(v) - \int_0^v s dg(s) - \int_0^v g(s) ds \equiv 0.$$

Let h be the left hand side of this equation. Clearly, $h(0) = 0$. Since g is non decreasing and bounded, h is differentiable almost everywhere which implies that $h' = 0$ almost everywhere. Consider some $\bar{v} \in [0, 1]$. If g is continuous at \bar{v} then so is h . Assume that g is not continuous at \bar{v} . Then

$$\bar{v}g(\bar{v}) - \int_0^{\bar{v}} s dg(s) = \lim_{v \uparrow \bar{v}} vg(v) + \bar{v} \left(g(\bar{v}) - \lim_{v \uparrow \bar{v}} g(v) \right) - \lim_{v \uparrow \bar{v}} \int_0^v s dg(s) - \bar{v} \left(g(\bar{v}) - \lim_{v < \bar{v}} g(v) \right)$$

so h is continuous at \bar{v} and thus $h \equiv 0$.

For the rest of the proof we can use a standard result on incentive compatibility, see Proposition 23.D.2 in Mas-Collel, Whinston & Green (1995). Part (1) follows immediately from the fact that F_p is nondecreasing and that $v\Phi(v) - \pi(\Phi, v) = \int_0^v \Phi(s) ds$ given our calculations above.

For part (2), notice that $x(v) \in [0, 1]$ and that incentive compatibility implies that $x(v)$ is non decreasing and $vx(v) - p(v) = \int_0^v x(s) ds$. Moreover, we can limit attention to menus where x is right continuous as otherwise there exists a right continuous incentive compatible menu $(\hat{x}(v), \hat{p}(v))_{v \in [0,1]}$ such that $\hat{p}(v) \geq p(v)$ for all v . As we consider x that is right continuous, Φ such that $\Phi(v) \triangleq x(v)$ for all v is a well defined mixed pricing policy and we obtain $p(v) = vx(v) - \int_0^v x(s) ds$. Our calculations above then imply that $\pi(\Phi, v) = p(v)$. ■

Proof of Proposition 8. Assume that \hat{p} 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

$$\pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(\hat{p}(F_0, \varepsilon), F_\varepsilon) \geq \gamma. \quad (18)$$

Assume that $(\hat{p}(F_0, \varepsilon), G_\varepsilon)$ is a saddle point of the regret problem (SP_r) given $\varepsilon > 0$. Then

$$r(\hat{p}(F_0, \varepsilon), G_\varepsilon) = \sup_{F \in \mathcal{P}_\varepsilon(F_0)} r(\hat{p}(F_0, \varepsilon), F),$$

and hence

$$\hat{p}(F_0, \varepsilon) = p^*(G_\varepsilon).$$

We can rewrite the left hand side of (18) as follows:

$$\begin{aligned} & \pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(\hat{p}(F_0, \varepsilon), F_\varepsilon) \\ = & \pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(p^*(G_\varepsilon), G_\varepsilon) + \pi(p^*(G_\varepsilon), G_\varepsilon) - \pi(p^*(G_\varepsilon), F_\varepsilon). \end{aligned} \quad (19)$$

Using (SP_r) we also obtain

$$0 \leq r(p^*(G_\varepsilon), G_\varepsilon) - r(p^*(G_\varepsilon), F_\varepsilon) = \int v dG_\varepsilon(v) - \int v dF_\varepsilon(v) + \pi(p^*(G_\varepsilon), F_\varepsilon) - \pi(p^*(G_\varepsilon), G_\varepsilon)$$

so that:

$$\pi(p^*(G_\varepsilon), G_\varepsilon) - \pi(p^*(G_\varepsilon), F_\varepsilon) \leq \int v dG_\varepsilon(v) - \int v dF_\varepsilon(v).$$

Entering this into (19) we obtain from (18) that:

$$\pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(p^*(G_\varepsilon), G_\varepsilon) + \int v dG_\varepsilon(v) - \int v dF_\varepsilon(v) \geq \gamma. \quad (20)$$

Since $F_\varepsilon, G_\varepsilon \in \mathcal{P}_\varepsilon(F_0)$ and since $h(v) = v$ is a continuous function and the Prohorov norm metricizes the weak* topology we obtain that

$$\int v dG_\varepsilon(v) - \int v dF_\varepsilon(v) < \gamma/2, \quad (21)$$

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 Prohorov neighborhood. Hence

$$\pi(p^*(F_\varepsilon), F_\varepsilon) - \pi(p^*(G_\varepsilon), G_\varepsilon) < \gamma/2 \quad (22)$$

if ε is sufficiently small. Comparing (20) to (21) and (22) yields the desired contradiction. ■

References

- Anscombe, F.J. & R.J. Aumann. 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics* 34:199–205.
- Bell, D.E. 1982. "Regret in Decision Making under Uncertainty." *Operations Research* 30:961–981.
- Bose, S., E. Ozdenoren & A. Pape. 2006. "Optimal Auctions with Ambiguity." *Theoretical Economics* 1:411–438.
- Che, Y.K. & I. Gale. 2000. "The Optimal Mechanism for Selling to a Budget Constrained Buyer." *Journal of Economic Theory* 92:198–233.
- Chung, K.-S. & J. Ely. 2003. "Implementation with Near-Complete Information." *Econometrica* 71:857–871.
- Deneckere, R.J. & R.P. McAfee. 1996. "Damaged Goods." *Journal of Economics and Management Strategy* 5:149–174.
- 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. & E. Katok. 2007. "Regret in Auctions: Theory and Evidence." *Economic Theory* p. forthcoming.
- Filiz, E. & E.Y. Ozbay. 2006. Auctions with Anticipated Regret: Theory and Experiment. Technical report Columbia University and New York University.
- Gilboa, I. & D. Schmeidler. 1989. "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics* 18:141–153.
- Hannan, J. 1957. Approximation to Bayes Risk in Repeated Play. In *Contributions to the Theory of Games*, ed. M. Dresher, A.W. Tucker & P. Wolfe. Princeton: Princeton University Press pp. 97–139.
- Hansen, L. Peter & T. Sargent. 2004. Misspecification in Recursive Macroeconomic Theory. Technical report University of Chicago, New York University, and Hoover Institution.

- Harris, M. & A. Raviv. 1981. "A Theory of Monopoly Pricing Schemes with Demand Uncertainty." *American Economic Review* 71:347–365.
- Hayashi, T. 2006. Regret Aversion and Opportunity Dependence. Technical report University of Texas.
- Huber, P.J. 1964. "Robust Estimation of a Location Parameter." *Annals of Mathematical Statistics* 35:73–101.
- Huber, P.J. 1965. "A Robust Version of the Probability Ratio Test." *Annals of Mathematical Statistics* 36:1753–1758.
- Klibanoff, P., M. Marinacci & S. Mukerji. 2005. "A Smooth Model of Decision-Making under Ambiguity." *Econometrica* 73:1849–1892.
- Linhart, P.B & R. Radner. 1989. "Minimax - Regret Strategies for Bargaining over Several Variables." *Journal of Economic Theory* 48:152–178.
- Loomes, G. & R. Sugden. 1982. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty." *Economic Journal* 92:805–824.
- Lopomo, G., L. Rigotti & C. Shannon. 2006. Uncertainty in Mechanism Design. Technical report.
- Mas-Collel, A., M.D. Whinston & J.R. Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.
- Maskin, E. & J. Riley. 1984. "Monopoly with Incomplete Information." *Rand Journal of Economics* 15:171–196.
- Milnor, J. 1954. Games Against Nature. In *Decision Processes*, ed. R.M. Thrall, C.H. Coombs & R.L. Davis. New York: Wiley.
- Mussa, M. & S. Rosen. 1978. "Monopoly and Product Quality." *Journal of Economic Theory* 18:301–317.
- Neumann, J. Von & O. Morgenstern. 1953. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- 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. & R. Zeckhauser. 1983. "Optimal Selling Strategies: When to Haggle, When to Hold Firm." *Quarterly Journal of Economics* 98:267–290.
- Savage, L.J. 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association* 46:55–67.
- Segal, I. 2003. "Optimal Pricing Mechanism with Unknown Demand." *American Economic Review* 93:509–529.
- 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. G. Debreu, W. Neufeind & W. Trockel. Heidelberg: Springer Verlag pp. 333–344.
- Stoye, J. 2006. Statistical Decisions under Ambiguity: An Axiomatic Analysis. Technical report New York University.
- Strassen, V. 1965. "The Existence of Probability Measures with Given Marginals." *Annals of Mathematical Statistics* 36:423–439.

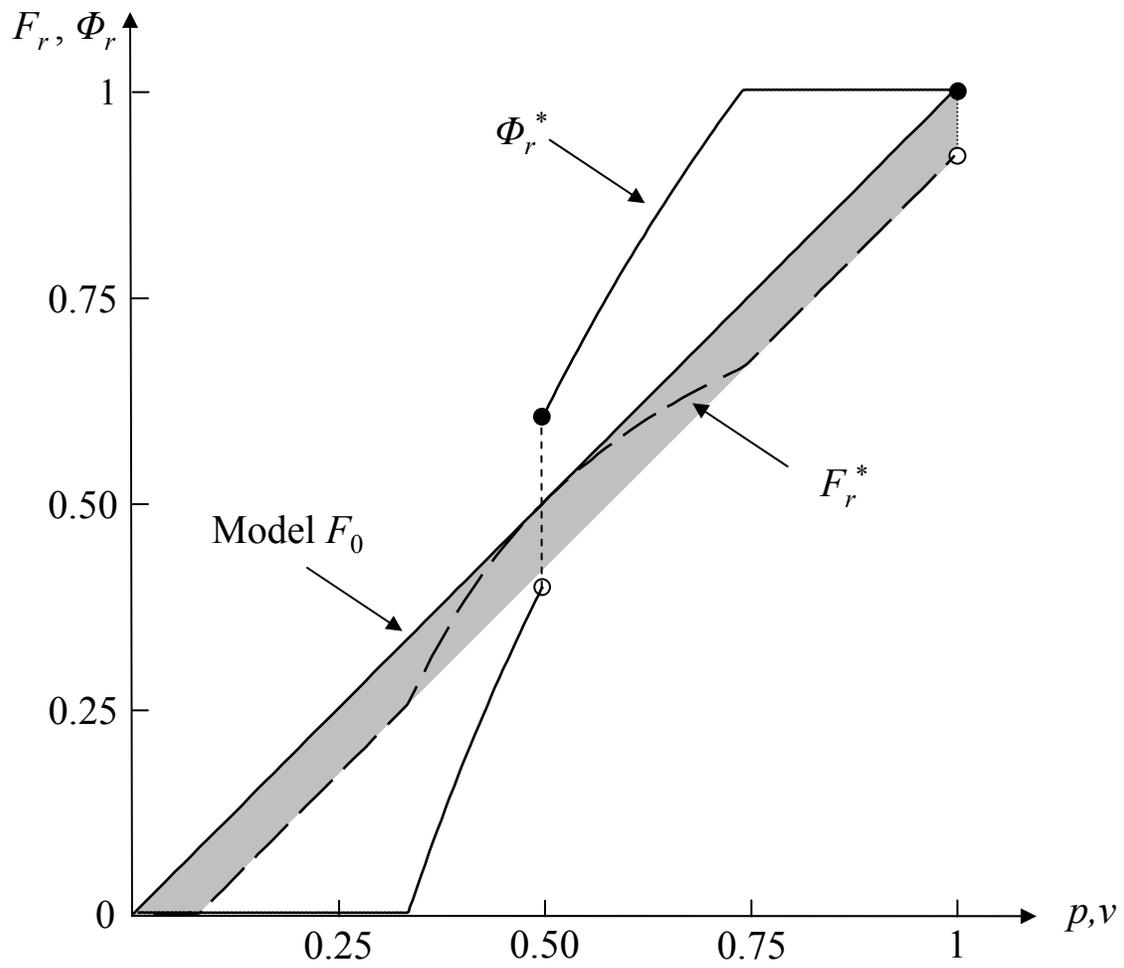


Figure 1. Optimal Pricing and Worst Case Demand with Uniform Model Density
 $(\varepsilon = 0.04)$