Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Selling to the mean

Nenad Kos and Matthias Messner

Working Paper n. 551

This Version: June 19, 2015

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

Selling to the mean*

Nenad Kos [†]
Bocconi University
IGIER

Matthias Messner[‡]
Bocconi University,
IGIER, CESifo

June 19, 2015

Abstract

We study optimal selling strategies of a seller who is poorly informed about the buyer's value for the object. When the maxmin seller only knows that the mean of the distribution of the buyer's valuations belongs to some interval then nature can keep him to payoff zero no matter how much information the seller has about the mean. However, when the seller has information about the mean and the variance, or the mean and the upper bound of the support, the seller optimally commits to a randomization over prices and obtains a strictly positive payoff. In such a case additional information about the mean and/or the variance affects his payoff.

JEL Code: C72, D44, D82.
Keywords: Optimal mechanism design, Robustness, Incentive compatibility, Individual rationality, Ambiguity aversion.

---

*For helpful discussions and comments we thank Dirk Bergemann, Alfredo Di Tillio, Johannes Hoerner, Larry Samuelson and Alexander Wolitzky. We would also like to thank seminar audiences at Boston College, Pittsburgh University and RUD 2015 in Milan.

[†]E-mail: nenad.kos@unibocconi.it
[‡]E-mail: matthias.messner@unibocconi.it

1 Introduction

The prevalent approach to modeling trading rules and other economic institutions through mechanism design, and game theory in general, assumes a great deal of structure on the form of player' beliefs. The seller knows the distribution of buyers' valuations, the government the distribution of agents' productivities, and the regulators the distribution of the firms' cost, to name a few. These beliefs come from a black box of the Bayesian approach, the assumption that the beliefs coincide with a true distribution is then appended almost without exception; see Morris (1994) for a departure. The drawbacks of the described approach have been prominently put forward in Wilson (1987), who argued that too much information is commonly known among the players. The common knowledge assumption was relaxed in the mechanism design framework, and a new wave of literature started, by Bergemann and Morris (2005). They studied the implementability on all type spaces. Majority of the subsequent literature followed in same vein, that is, characterizing incentive compatible mechanisms under various assumptions on the agents' beliefs about each other. Rather little attention, however, has been dedicated to the designer and the design of optimal mechanism from his perspective.[1]

We study the robust of optimal trading rule in an environment where a profit maximizing seller is selling a good to a single buyer. The seller has very limited information about the buyer's valuation. He is armed merely with the information about the mean of the distribution of the buyer's valuations.[2] The information about the mean is represented by the interval to which the mean belongs. The seller then evaluates any trading mechanism by taking the expectation with respect to the type distribution that yields the smallest profit among the distributions with a mean in the interval. In other words, the seller is a maxmin expected utility maximizer who believes that nature can choose any distribution over non-negative valuations with a mean in a fixed interval. Wolitzky (2014) provides an interesting interpretation for this information structure, where a model of a seller who knows only the mean of the buyer's distribution arises from model of information acquisition.

Our findings can be summarized as follows. First we show a somewhat daunting result that the seller who only knows that nature draws a non-negative valuation from a distribution with a mean in some interval, $[\underline{\mu}, \bar{\mu}]$, expects zero payoff. This is so irrespective of the interval to which the mean belongs and the mechanism he chooses. Additional information about the mean has no affect on such a seller's payoff. The picture drastically changes once one assumes

[1]Some notable exceptions are Bergemann and Schlag (2008) and Bergemann and Schlag (2011) who study optimal mechanism for sale of an object by a seller under max-min regret and max-min preferences. We provide a more substantial overview of the literature in the subsection Related Literature below.

[2]One of the first thing a student is thought in an econometrics class is how to estimate a mean. Yet very little is known about how agents faced with such limited information should behave. As a thoroughly explored framework, profit maximization in bilateral trade provides a natural testing ground.

that there is an upper bound, $\bar{\theta}$, on the support of nature's distributions. We show that the uniquely optimal selling strategy has the seller committing to a randomization over prices in interval $[\underline{\tau}, \bar{\theta}]$ for some $\underline{\tau} > 0$ with a cdf of the order $\ln t$ and no mass points. Interestingly, the seller randomizes over prices in such a way that when valuations are distributed on the same interval as prices, the seller's payoff depends only on the mean of the distribution of valuations. By making his payoff independent of the information about the parameters of the distribution that he does not know, the seller, in a sense, insures himself against the missing information about the distribution of valuations.

We obtain several comparative statics results. Given the interval of means that the seller entertains as possible, his expected payoff decreases when $\bar{\theta}$ increases. When nature cannot assign high values to the agents, somewhat curiously, the seller's payoff increases. This is due to the fact that by lowering $\bar{\theta}$ one restricts the set of possible distributions for nature. On the other hand, the seller's payoff is increasing in the lower bound of possible means, $\underline{\mu}$, of the distribution over valuations. Namely, for a given distribution over prices the seller's payoff cannot decrease if $\underline{\mu}$ increases.

Bounding the support of the distributions over valuations from above implicitly restricts the variance of those distributions. We therefore also explore the problem in which seller knows that the mean of valuations belongs to an interval $[\underline{\mu}, \bar{\mu}]$, and the variance to $[\underline{\sigma}^2, \bar{\sigma}^2]$; but there is no upper bound on the support. We show that the unique optimal strategy for the seller is to randomize over some interval $[\underline{\tau}, \bar{\tau}]$ with density of the form $h(x) = a/x + b$. In this case, he expects the valuations to be drawn from a distribution that has the support contained in the same interval over which he randomizes, and more importantly, his payoff then depends only on the mean and the variance of the distribution of valuations. Thus again, he insures himself against the information about the distribution of valuations that he does not have. Of course, he believes that nature chooses a distribution with the lowest possible mean, $\underline{\mu}$, and the highest possible variance, $\bar{\sigma}^2$. The seller's optimal payoff is shown to be increasing in $\underline{\mu}$ and decreasing in $\bar{\sigma}^2$.

For a given amount of knowledge about the mean, knowing the upper bound of the support of the distributions or knowing the upper bound of the variance is payoff equivalent for the seller. More precisely, given $[\underline{\mu}, \bar{\mu}]$, for every $\bar{\theta}$ there exists a $\bar{\sigma}^2$ ($\underline{\sigma}^2$ is irrelevant), and vice-versa, such that the seller's payoff is the same regardless of which of the two pieces of information he possesses. That notwithstanding, the seller uses different pricing schemes in the two cases. The price distribution in the case when he knows the upper bound on the support first order stochastically dominates the price distribution he uses when he knows the payoff equivalent upper bound on the variance of the nature's distribution. Whether the seller has information about the variance or the upper bound of the support of the nature's distribution is thus potentially identifiable from the data.

The above described results imply that when the seller only has information about the mean of the distribution over valuations, additional information about the mean has no effect on his payoff. However, when he has information about the mean and the variance, or the mean and the upper bound, then the additional information about the two parameters can be beneficial to the seller.

1.1 Related Literature

Wilson's critique and Bergemann and Morris (2005) have initiated a large body of literature on robust mechanism design. For an in-depth review see Bergemann and Morris (2013). Our paper is closely related to the work of Bergemann and Schlag (2011). They consider the problem of a seller selling a single good to a buyer. The seller is a max-min expected utility maximizer with imperfect information about the distribution over the valuations: he knows that the valuations are distributed in an epsilon neighborhood of some distribution. In their environment nature has a dominant strategy, therefore a deterministic take-it-or-leave-it price is optimal.[3] Auster (2013) analyses a model with common values in which the seller is privately informed, and ambiguity is of the similar form as in Bergemann and Schlag (2011). Garrett (2014) studies a model of cost-based procurement in which the principal is uncertain about the agent's effort cost function.

The environment that is closest to the one analyzed in this paper is studied in Wolitzky (2014). He considers efficiency in a bilateral trade model in which the buyer and the seller know only the mean of each other's valuations. He shows that under some parameters the efficient trade is possible and characterizes when exactly that is the case. Our paper, on the other hand, is concerned with profit maximization and the value of information for the seller. Also, unlike in Wolitzky (2014) we allow for the information about the variance. Carroll (2012) studies the problem of providing robust incentives for information acquisition. In his model the decision making max-min expected utility maximizing principal is incentivizing an expert to acquire costly information.

Lopomo, Rigotti, and Shannon (2009) explore robustness of mechanisms under incomplete preferences, as in Bewley (1986). Castro and Yannelis (2012) approach the problem from a different perspective and show that every efficient allocation rule is incentive compatible if and only if the agents have max-min preferences.

Robustness in the context of moral hazard has been explored in Lopomo, Rigotti, and

[3] López-Cunat (2000), Bergemann and Schlag (2008) and Bergemann and Schlag (2011) explore the seller's problem when he is minimizing his regret.

Shannon (2011), Chassang (2013) and Carroll (2015), to name a few.

Our paper is also related to the growing literature on mechanism design under ambiguity aversion. Though in that literature, unlike in the present paper, the buyers are the ones who are ambiguity averse. See for example Bose, Ozdenoren, and Pape (2006), Bose and Daripa (2009) and Bodoh-Creed (2012). More recently Bose and Renou (2014) and Di Tillio, Kos, and Messner (2014) have shown that in such environments the seller might benefit from using non-standard mechanism.

2 Environment and terminology

Preferences and information: There is a seller who wants to sell a single unit of a good to a buyer. We denote the probability with which the good is transferred to the agent by x and write τ for the transfer to be be paid by the agent. Throughout the paper we will slightly abuse terminology and refer to x also as an 'allocation'.

The buyer is a risk neutral expected utility maximizing agent whose valuation for the good is denoted by θ. Thus, if he receives the good with probability x and pays the transfer τ in exchange, his payoff is

$$x\theta - \tau.$$

If, instead, the buyer decides not to participate in the mechanism his payoff is 0.

The seller is uncertain about the agent's valuation of the good, θ. He only has partial information about the distribution of the buyer's valuations. More specifically, we assume that the seller knows that the buyer's valuations are non-negative and drawn from a distribution which has a mean that belongs to the interval $[\underline{\mu}, \bar{\mu}]$. We denote the set of all such distributions by $\mathcal{F}(\underline{\mu}, \bar{\mu})$. Later we will explore environments in which the seller knows additional features of the distributions.

The seller does not assign any value to the good, he only cares about the revenue he can obtain. In addition, we assume that the seller evaluates mechanisms with respect to their worst case expected value; i.e. it is this worst case expected revenue that he seeks to maximize. Therefore, if the transfer function in an incentive compatible mechanism is t and the seller believes that the buyer's valuation is drawn from a distribution in the set $\mathcal{F}(\underline{\mu}, \bar{\mu})$, then his payoff is

$$\inf_{F(\underline{\mu}, \bar{\mu}) \in \mathcal{F}} E_F[t].$$

An axiomatic representation of maxmin preferences is provided in Gilboa and Schmeidler

(1989).

3 Unbounded Valuations

First we study the case where the seller only knows that the buyers valuations are distributed over $\mathbb{R}_+$ and the bounds on the mean of the distribution of values $[\underline{\mu}, \bar{\mu}]$. The seller chooses an incentive compatible and individually rational mechanisms. With one agent it is without loss of generality to consider randomizations over posted prices.[4]

The set of distributions over prices is denoted by $\mathcal{H}$ and its generic element by H. We require that $\int_0^\infty \theta dH$ exists and is finite.

Proposition 1. *Suppose that the set of possible types is $\mathbb{R}_+$. Then*

$$\inf_{F \in \mathcal{F}(\underline{\mu}, \bar{\mu})} U(H, F) = 0,$$

for every distribution over prices $H \in \mathcal{H}$. Consequently,

$$\sup_{H \in \mathcal{H}} \inf_{F \in \mathcal{F}(\underline{\mu}, \bar{\mu})} U(H, F) = 0.$$

Proof. All missing proofs can be found in Appendix A. □

The highest payoff the seller can expect from a given randomization over the prices is the expected price. This would be the case if he were to sell with probability one. However, nature can choose a binary distribution over valuation 0 and some high valuation θ with the mean $\underline{\mu}$. [5] The probability that nature needs to assign to θ is then decreasing in θ. Since the seller in that case sells only when the buyer's valuation is θ, and his expected payoff is bounded above by

[4]Clearly a randomization over prices is an incentive compatible and individually rational mechanism. On the other hand, starting from any incentive compatible and individually rational mechanism one can obtain a randomization over prices. That is, in an incentive compatible mechanism the allocation rule is monotonic, therefore the cumulative distribution over prices can be defined to be equal to the allocation rule. If no type gets the object with probability one, this can be replicated by assigning probability to high prices. Since both the allocation rule from the original mechanism, as well as randomization over prices assign the object with the same probability to each type the two mechanisms have the same transfers up to a constant. Since any randomization over (non-negative) prices leaves type 0 with zero payoff, the randomization over prices achieves at least as high a payoff as the original mechanism. See Skreta (2006) and Kos and Messner (2013) for how to deal with the cases where the distribution are not continuous. The seller could potentially also try to obtain the information about the distribution of the valuations of the buyer, rather than just ask for his type. Since such information is payoff irrelevant for the buyer, the seller will not be able to elicit it. Randomization over prices is also studied in Bergemann and Schlag (2011).

[5]The seller's problem can be thought of as a zero sum game between the seller and nature. We elaborate on this later in the paper.

the expected price, nature can keep the seller to an arbitrarily low positive payoff by making sure that the seller sells with a very small probability (i.e. by choosing θ high enough).

The daunting consequence of the above result is that knowledge of the mean itself – or the interval it resides in – has no effect on the seller's payoff. Regardless of the extent of knowledge the seller has about the mean, nature can hold him to payoff zero.

The above result hinges on the assumption that the seller is willing to consider arbitrarily high valuations. Next we explore the case of a seller who believes that only valuations up to some upper bound $\bar{\theta}$ are possible.

4 Bound on the valuations

In this section we study the environment in which there is an upper bound on the buyer's valuations. In particular, the seller believes that buyer's valuations are distributed on the interval $[0, \bar{\theta}]$, for some $\bar{\theta} > 0$, with a distribution that has a mean in $[\underline{\mu}, \bar{\mu}]$. The set of possible distributions of valuations is denoted by $\mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})$.

4.1 Optimal posted prices

We start our analysis by exploring the seller's optimal posted price. The appeal of posted prices stems from their simplicity, their empirical relevance and not least from the fact that they are optimal for the seller in a standard setting with ambiguity neutral sellers (see Myerson (1981) and Riley and Zeckhauser (1983)). While the optimality property of posted prices does not carry over to our environment they still constitute a natural benchmark.

Let $u(\tau, F)$ be the seller's expected payoff when he offers the price τ and the valuations are drawn from the distribution F. Given the price τ, type θ of the buyer will acquire the good if and only if $\theta > \tau$.[6] Therefore

$$u(\tau, F) = \tau[1 - F(\tau)].$$

The seller's optimization problem is then

$$\sup_{\tau \geq 0} \inf_{F \in \mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})} u(\tau, F). \tag{1}$$

[6]Throughout the analysis we maintain the assumption that the agent does not buy the good if he is indifferent between buying and not buying. As will become clear, this technical assumption serves the purpose of guaranteeing that a best response for nature exists. This assumption does not affect the solution and the value of the problem.

The buyer's behavior is simple: he buys the good if his value is above the offered price and does not otherwise. The seller's problem can therefore be thought of as a game against the malevolent nature. The seller chooses the price τ after which nature chooses a distribution $F \in \mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})$ with the objective to minimize the seller's payoff.

It is never optimal for the principal to set a price equal to zero or a price that exceeds the lower bound of the expectation of the agent's type, $\underline{\mu}$. A price of zero implies a payoff equal to zero. The same is true for a price above or at $\underline{\mu}$. In the latter case nature can hold the principal to payoff zero by choosing the distribution that puts the entire probability mass on $\underline{\mu}$. This distribution trivially satisfies the mean constraint while at the same time generating with probability one a type that is not willing to buy at the price set by the seller. On the other hand, any price τ' that lies strictly between 0 and $\underline{\mu}$ guarantees the principal a strictly positive payoff. Indeed, any type distribution that respects the mean constraint must put strictly positive weight on the interval $[\underline{\mu}, \bar{\theta}]$. Therefore, for a price τ' satisfying $0 < \tau' < \underline{\mu}$ there is always – irrespective of what type distribution nature chooses – a set of types with positive probability that is willing to buy the good. But then the seller's maxmin expected revenue must be strictly positive too.

To find the optimal price, we first determine nature's best response to each possible price $\tau < \underline{\mu}$. Nature's goal is to chose a distribution that minimizes the seller's expected payoff. Buyers with valuations smaller or equal to τ do not buy,[7] while buyers with valuations above τ buy the good. Nature's objective is therefore to place as much probability mass as possible at or below τ subject to the constraints that the expected value of the distribution must be in $[\underline{\mu}, \bar{\mu}]$ and that no set of types outside $[0, \bar{\theta}]$ can be given positive probability. This objective is best achieved by choosing the type distribution with mean $\underline{\mu}$ that concentrates all the weight on the points τ and $\bar{\theta}$.[8]

The probability mass, p, that this distribution assigns to the type τ is defined by

$$p\tau + (1 - p)\bar{\theta} = \underline{\mu},$$

therefore

$$p = \frac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - \tau}.$$

[7]Remember that the buyer does not buy the good when he is indifferent.

[8]For a fixed price τ, nature is maximizing $F(\tau)$ subject to the constraint $\underline{\mu} = F(\tau)E[\theta|\theta \leq \tau] + (1 - F(\tau))E[\theta|\theta > \tau]$. Given that $\tau < \underline{\mu}$, $F(\tau)$ is maximized when $E[\theta|\theta \leq \tau] = \tau$ and $E[\theta|\theta > \tau] = \bar{\theta}$.

Using this we can rewrite Problem 1 as

$$\max_{\tau \in [0, \underline{\mu}]} \left(1 - \frac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - \tau}\right) \tau.$$

It is straightforward to verify that the objective function of this problem is strictly concave in τ, and equal to zero on the two extremes of the interval $[0, \underline{\mu}]$. Solving for τ yields

$$\tau^* = \sqrt{\bar{\theta}} \left[\sqrt{\bar{\theta}} - \sqrt{\bar{\theta} - \underline{\mu}} \right].$$

Let F^* be the type distribution corresponding to the principal's most pessimistic belief associated with the optimal posted price (i.e. nature's best response to the optimal price). It is straightforward to verify that the weight that this distribution puts on $\bar{\theta}$ is $1 - p = \tau^*/\bar{\theta}$. Since this is the probability with which the principal expects the agent to buy the good, his maxmin expected revenue is $u(\tau^*, F^*) = \tau^{*2}/\bar{\theta}$. We summarize these findings in the following proposition.

Proposition 2. *The optimal posted price is*

$$\tau^* = \bar{\theta} - \sqrt{\bar{\theta}(\bar{\theta} - \underline{\mu})}.$$

Given this price the seller expects that the agent will buy with probability $\tau^*/\bar{\theta}$, therefore his expected revenue is

$$u(\tau^*, F^*) = [\tau^*]^2/\bar{\theta} = \left(\sqrt{\bar{\theta}} - \sqrt{\bar{\theta} - \underline{\mu}} \right)^2.$$

Comparative statics. The price that the seller charges and the payoff that he realizes are increasing in the lower bound of the expected type, $\underline{\mu}$. Indeed, for every price the seller's payoff (weakly) increases after an increase in $\underline{\mu}$. Notice also that the seller's payoff is independent of $\bar{\mu}$.

At first it might seem somewhat surprising, that preventing nature from assigning high valuations to a buyer might help the seller, but in our environment this is indeed the case. The seller's payoff is decreasing in the upper bound on the support of nature's distributions, $\bar{\theta}$: the larger the upper limit of the set of possible types, the weaker becomes the restriction that the mean constraint imposes on nature's ability to concentrate mass on low types. That is, increasing $\bar{\theta}$ increases the set of nature's strategies.

4.2 Random Pricing Policies

In this section we explore what the seller can achieve by committing to a randomization over posted prices. Clearly the seller cannot fare worse when we expand his set of possible mechanisms to randomizations over posted prices. We motivate the following analysis with an example.

Example 1. Assume that the agent's valuations are in the interval $[0, 1]$ (i.e. $\bar{\theta} = 1$), and that the seller knows that the expected value of the agent's valuation is at least $3/4$ (i.e. $\underline{\mu} = 3/4$). The result from the preceding subsection implies that the seller's optimal posted price is $\tau^* = 1/2$ and $u(\tau^*, F^*) = 1/4$.

Now suppose that the seller instead of using a posted price commits to a fifty-fifty randomization over the prices $\tau_1 = 1/3$ and $\tau_2 = 1/2$.[9] Given the seller's randomization over prices, nature will not put a strictly positive probability mass below τ_1. Since the agent does not buy the good for any valuation $\theta \leq \tau_1$, any such mass could be pushed up to τ_1 to get a new distribution with an expected value strictly above $\underline{\mu}$. In the resulting distribution one could shift some of the mass from above τ_1 to τ_1 and still have the expected value larger or equal to $\underline{\mu}$. Under the new distribution the seller would, however, obtain a strictly smaller payoff. Expanding on the above reasoning one can see that nature will only ever randomize over τ_1, τ_2 and $\bar{\theta} = 1$.

It is easy to see that nature never chooses a distribution with a mean strictly above $\underline{\mu}$. Therefore, there are just three possible candidates for the nature's distribution. A distribution on $\{\tau_1, 1\}$ with mean $\underline{\mu}$, a distribution over $\{\tau_2, 1\}$ with mean $\underline{\mu}$, and distributions with the support $\{\tau_1, \tau_2, 1\}$ and mean $\underline{\mu}$. It is easy to see, that one only needs to compare nature's payoff from the first two distributions. Namely, if one of them yields a higher payoff then the other, a distribution with full support cannot be optimal.[10]

If nature randomizes over $\{1/2, 1\}$ it must do so with probability 0.5, due to $\underline{\mu} = 3/4$. In this case the seller's payoff is 7/24. That is, with probability half he charges $1/3$ and sells irrespective of the type. With probability half he charges $1/2$ and sells half of the time (when the valuation is 1). On the other hand, when nature chooses a distribution on $\{1/3, 1\}$ it assigns the probability $3/8$ to the value $1/3$. In that case the seller's payoff is 25/96. Nature, thus, fares better when randomizing over $\{1/3, 1\}$. But even in that case the seller's expected payoff is

[9] One can think of this mechanism as the seller committing to a randomization over prices. Then the randomization is executed, and finally the buyer decides whether he wants to buy the object at the price.

[10] Suppose the distribution over $\{\tau_1, 1\}$ is F_1 and the distribution over $\{\tau_2, 1\}$ is F_2. If it is the case that nature gets a higher payoff from F_2 then F_1 and a distribution F_3 with a support $\{\tau_1, \tau_2, 1\}$ is optimal, it has to be the case that nature gets a higher payoff from F_3 then from F_1. Since F_1's support is contained in F_3's, one can write F_3 as a convex combination of F_1 and some other distribution on $\{\tau_1, \tau_2, 1\}$, i.e., $F_3 = \alpha F_1 + (1 - \alpha) F_4$. Since both F_1 and F_3 have mean $\underline{\mu}$, so does F_4. But then F_4 is a viable distribution for nature that gives a strictly higher payoff than F_3, contradicting that the latter is optimal.

strictly larger than his maxmin payoff of 0.25 achieved under the optimal (deterministic) price.

The above example shows that the seller benefits from posting a random pricing policy. In the maxmin problem nature gets to choose its preferred strategy after the seller picks a mechanism. When the seller posts a price nature has the upper hand due to the informational advantage of knowing what action the seller chose. By randomizing over prices the seller makes it harder for nature to target its reply. Thus stochastic pricing levels the playing field by decreasing nature's second mover advantage.

In the rest of this section we characterize the optimal random pricing policy. We describe a random pricing policy by the cumulative distribution function over prices H that it induces. We write $\mathcal{H}(\bar{\theta})$ for the set of all distribution functions with support in $[0, \bar{\theta}]$. Let $U(H, F)$ be the principal's expected payoff when he chooses the price distribution H and the buyer's type is drawn from the distribution F, i.e.

$$U(H, F) = \int_0^{\bar{\theta}} u(\tau, F) dH(\tau),$$

where

$$u(\tau, F) = \tau(1 - F(\tau)).$$

The principal's problem can now be written as

$$\max_{H \in \mathcal{H}(\bar{\theta})} \inf_{F \in \mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})} U(H, F). \tag{2}$$

Instead of directly tackling this problem, we solve for a saddle point of the payoff functional $U(H, F)$. That is, we look for a pair of distributions (F^*, H^*) such that

$$U(H^*, F) \geq U(H^*, F^*) \geq U(H, F^*)$$

holds for all feasible pairs (H, F). A standard result states that if such a saddle point exists, then

$$H^* \in \arg \max_{H \in \mathcal{H}(\bar{\theta})} \left\{ \inf_{F \in \mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})} U(H, F) \right\}$$

and

$$U(H^*, F^*) = \sup_{H \in \mathcal{H}(\bar{\theta})} \inf_{F \in \mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})} U(H, F).$$

One can think of the seller's optimization problem also as the problem of finding a subgame perfect equilibrium of a sequential zero-sum game played between the seller and nature in which the seller moves first and nature's payoff is the negative of the seller's.[11] Instead of solving directly for such a subgame perfect equilibrium we solve for a Nash equilibrium (H^*, F^*) of the simultaneous move version of this zero-sum game, which corresponds to a saddle point of the payoff functional U. The properties of a saddle point imply that the seller's equilibrium strategy in the simultaneous move game, H^*, is also his maxmin strategy (i.e. his equilibrium strategy in the subgame perfect equilibrium of the sequential game). See Appendix B for the corresponding formal result.

In what follows we derive properties of a saddle point, should one exist. As we go along, we impose several additional assumptions to facilitate the derivation. These assumptions will, of course, be verified before our analysis concludes.

Characterizing nature's distribution F^*: Suppose there exists a saddle (H^*, F^*). If the stochastic posted price H^* solves

$$\max_{\mathcal{H}(\bar{\theta})} U(H, F^*),$$

then it must be the case that each element in its support yields the saddle value $U^* = U(H^*, F^*)$. That is,

$$u(\tau, F^*) = \tau[1 - F^*(\tau)]$$

must be constant in τ over the support of H^*.[12]

We conjecture that the support of H^* is some interval $[\underline{\tau}, \bar{\theta}]$ for $\underline{\tau}$ such that $0 < \underline{\tau} < \underline{\mu}$. If this is the case, then this interval must also be the support of nature's strategy F^*. To see this, observe first that $[\underline{\tau}, \bar{\theta}]$ must be contained in the support of F^* for otherwise there would be an open subset of $[\underline{\tau}, \bar{\theta}]$ over which nature would not randomize and therefore F^* would be

[11] In this game between, the seller has a possibility to commit. Therefore if the seller commits to a randomization over prices, nature has to choose a distribution before the outcome of the realization is revealed.

[12] In the saddle (equilibrium) of the auxiliary simultaneous move zero sum game between the seller and nature the seller needs to be indifferent over all the prices in the support of his distribution. Of course, from a mechanism design perspective (in the original problem) the seller has commitment and thus can offer mechanisms in which he randomizes over prices even if he is not indifferent between them.

constant. But then $\tau[1-F^*(\tau)]$ would not be constant over the same open interval, contradicting the assumption that $[\underline{\tau}, \bar{\theta}]$ is the support of H^*. It remains to be argued that it is suboptimal for nature to assign any probability mass to types strictly below $\underline{\tau}$. Notice that the expected revenue generated by any type below $\underline{\tau}$ is 0, which is the same expected revenue that is generated by type $\underline{\tau}$ itself. Therefore, shifting the probability mass from the interval $[0, \underline{\tau})$ to $\underline{\tau}$ does not alter the expected revenue of the seller. At the same time such a shift relaxes nature's mean constraint. This slack in the mean constraint can be exploited to shift probability mass that is placed above $\underline{\mu}$ down to $\underline{\tau}$. Doing so strictly reduces the seller's expected revenue.

The above arguments imply that there is some $U^* > 0$ such that $[1 - F^*(\tau)]\tau = U^*$ for all $\tau \in (\underline{\tau}, \bar{\theta})$, and $F^*(\tau) = 0$ for all $\tau < \underline{\tau}$. The first one of these two conditions is equivalent to

$$F^*(\tau) = 1 - U^*/\tau \tag{3}$$

for all $\tau \in (\underline{\tau}, \bar{\theta})$. Thus, once $\underline{\tau}$ and U^* are pinned down so is F^*. Consider first the choice of U^* for given $\underline{\tau}$. Notice that the seller can guarantee himself a payoff that is arbitrarily close to $\underline{\tau}$ by setting a price just below $\underline{\tau}$. This implies that $U^* \geq \underline{\tau}$. At the same time, in order for the condition $F^*(\tau) = 1 - U^*/\tau \geq 0$ to be satisfied for all $\tau \in (\underline{\tau}, \bar{\theta})$, we must have $U^* \leq \underline{\tau}$. Combining these two observations yields

$$U^* = \underline{\tau}. \tag{4}$$

Equations (3) and (4) have two important implications. First,

$$F^*(\underline{\tau}) = 0,$$

which means that there is no atom at $\underline{\tau}$. And second,

$$\lim_{\tau \nearrow \bar{\theta}} F^*(\tau) = 1 - \underline{\tau}/\bar{\theta},$$

meaning that F^* has an atom of mass $\underline{\tau}/\bar{\theta}$ at $\bar{\theta}$.

We are left to determine $\underline{\tau}$. The type distribution F^* derived in the preceding paragraphs for a given $\underline{\tau}$, is a feasible type distribution only if it satisfies the mean constraint

$$\int_0^{\bar{\theta}} \theta dF^* = \underline{\mu},$$

which after using the functional form of F^* becomes,

$$\underline{\tau}\left[1 + \ln(\bar{\theta}) - \ln(\underline{\tau})\right] = \underline{\mu}.$$

The expression $\underline{\tau}\left[1+\ln(\bar{\theta})-\ln(\underline{\tau})\right]-\underline{\mu}$ is strictly negative for values of $\underline{\tau}$ close to zero and strictly positive for values of $\underline{\tau}$ close to $\underline{\mu}$. Moreover, it is continuous and strictly increasing in $\underline{\tau}$. Thus, there can be only one value of $\underline{\tau}$ for which the above condition is satisfied. In the proof of Proposition 4 below we show that solving for this unique value delivers

$$\underline{\tau} = -\frac{\underline{\mu}}{W_{-1}(-\underline{\mu}/e\bar{\theta})},$$

where W_{-1} is the lower branch of the Lambert W function. For further details regarding the Lambert W function see Appendix C.

Characterizing the seller's randomization over prices H^*: Having determined the distribution of the types, F^*, that renders the seller indifferent, we move to the question of how the posted prices need to be distributed for F^* to be an optimal choice for nature.

When characterizing nature's strategy we required that the seller be indifferent between all the points in the support of his strategy. This approach can not be applied in the same way when characterizing the seller's optimal strategy. That is, it does not make sense to require that the nature be indifferent between all types, or degenerate distributions corresponding to them, in the support of its distribution. Such degenerate distributions might not represent viable strategies for nature due to the mean restriction. We show that if F^* is to be nature's best response, nature has to be indifferent over a particular set of two-point-distributions. Furthermore, nature's indifference on this set implies indifference over all the distributions with mean $\underline{\mu}$ and support in $[\underline{\tau}, \bar{\theta}]$. We have encountered such distributions in the analysis of optimal (deterministic) posted prices.

Definition 1. *For each $\tau \in [\underline{\tau}, \bar{\theta}]$, let F_τ be a distribution that concentrates the probability mass on at most two points and has an expected value $\underline{\mu}$. If $\tau \in (\underline{\mu}, \bar{\theta}]$, the support of F_τ are the two points $\underline{\tau}$ and τ; if instead $\tau \in [\underline{\tau}, \underline{\mu})$, then the support of F_τ is $\{\tau, \bar{\theta}\}$. Finally, if $\tau = \underline{\mu}$, then $\underline{\mu}$ is the only element of the support of F_τ. We refer to F_τ as the $\tau-$distribution.*

The distribution F_τ, for $\tau < \underline{\mu}$, assigns the largest possible weight to τ among all distributions with expected value $\underline{\mu}$. Indeed, the weight on $\underline{\tau} \leq \tau < \mu$ can be increased, while preserving the mean at $\underline{\mu}$ as long as it is possible to place sufficient mass somewhere above $\underline{\mu}$. Thus the limit is reached if the entire mass which is not placed on τ itself is concentrated on the most effective counterweight, namely the point $\bar{\theta}$. Similarly, for $\underline{\mu} < \tau \leq \bar{\theta}$ the best counterweight in $[\underline{\tau}, \bar{\theta}]$ is $\underline{\tau}$. In what follows we denote the weight that the τ-distribution assigns to τ by $p(\tau)$.

For later reference:

$$p(\tau) = \begin{cases} \dfrac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - \tau} & \text{if } \underline{\tau} \le \tau < \underline{\mu}, \\ \dfrac{\underline{\mu} - \underline{\tau}}{\tau - \underline{\tau}} & \text{if } \underline{\mu} < \tau \le \bar{\theta}. \end{cases} \tag{5}$$

The following lemma shows that indifference over all τ−distributions is a necessary condition for the optimality of F^*.

Lemma 1. *If the type distribution F^* is a best response to a price distribution $\hat{H}$, then F_τ is a best response to $\hat{H}$ for every $\tau \in [\underline{\tau}, \bar{\theta}]$.*

This result follows from the following observation. Suppose that there is a saddle point where nature uses a distribution F^* with support $[\underline{\tau}, \bar{\theta}]$. If the distribution F_τ gives a strictly higher payoff to nature, then F_τ is a profitable deviation. If, on the other hand, it gives a strictly smaller payoff than F^*, then so do distributions $F_{\tau'}$ for every τ' in a small enough neighborhood of τ. But then nature would be better of by reducing the probability F^* assigns to all such $F_{\tau'}$.[13]

Next we show that indifference over the set of all τ-distributions is not only necessary but also sufficient for F^* to be a best reply to a price distribution $\hat{H}$. More specifically, the following lemma tells us that indifference over all τ−distributions implies indifference over all distributions with mean $\underline{\mu}$ and support in $[\underline{\tau}, \bar{\theta}]$. It is easy to see that nature is worse off by choosing a distribution that assigns mass below $\underline{\tau}$ or has a mean above $\underline{\mu}$.

Lemma 2. *Suppose that H is a price distribution such that every τ-distribution yields the same payoff $\bar{U}$, i.e. $U(H, F_\tau) = \bar{U}$, for all $\tau \in [\underline{\tau}, \bar{\theta}]$. Then*

$$U(H, F) = \bar{U},$$

for all F with mean $\underline{\mu}$ and support in $[\underline{\tau}, \bar{\theta}]$.

Lemmata 1 and 2 imply that type distribution F^* is a best reply to a price distribution H if and only if the latter yields the same payoff against all τ-distributions. We have yet to determine, however, the price distribution that yields the same payoff for all F_τ.

[13]More precisely, F^* can be written as a convex combination of a distribution G, which is a randomization over $F_{\tau'}$ for τ' in a small neighborhood around τ, and some other distribution H. Notice that H also has the mean $\underline{\mu}$. Since G yields a strictly smaller payoff to nature than F, H must yield a higher payoff, and thus represent a profitable deviation. This is the continuous counterpart of what was discussed in footnote 10.

Consider first the case where the buyer's valuations are drawn from F_τ, for $\tau \leq \underline{\mu}$. Equation (5) implies

$$p(\tau) = \frac{\bar{\theta} - \mu}{\bar{\theta} - \tau}.$$

Moreover, assume that the distribution according to which the seller picks the price, H^*, is twice differentiable.[14] The seller's payoff is

$$p(\tau) \int_{\underline{\tau}}^{\tau} t dH^*(t) + (1 - p(\tau)) \int_{\underline{\tau}}^{\bar{\theta}} t dH^*(t).$$

With probability $p(\tau)$ nature draws the value τ in which case the seller sells at a price t if $t < \tau$. With probability $1 - p(\tau)$ nature draws valuation $\bar{\theta}$ and the seller sells for every price $t < \bar{\theta}$. If the payoff is to remain constant across all τ-distributions, it must be the case that the derivative of the payoff with respect to τ is equal to zero. That is,

$$p'(\tau) \int_{\underline{\tau}}^{\tau} t h^*(t) dt + p(\tau) \tau h^*(\tau) - p'(\tau) \int_{\underline{\tau}}^{\bar{\theta}} t h^*(t) dt = 0,$$

must hold for all τ in $(\underline{\tau}, \underline{\mu})$. Equality (5) implies $p(\tau)/p'(\tau) = \bar{\theta} - \tau$, which can be used to simplify the above equation to

$$(\bar{\theta} - \tau)\tau h(\tau) - \int_{\tau}^{\bar{\theta}} t h^*(t) dt = 0. \tag{6}$$

Condition (6), in turn, can hold for every τ only if the derivative of the left hand side is zero over the entire interval $(\underline{\tau}, \underline{\mu})$:

$$(\bar{\theta} - 2\tau) h^*(\tau) + (\bar{\theta} - \tau)\tau \frac{dh^*(\tau)}{d\tau} + \tau h^*(\tau) = 0, \tag{7}$$

or equivalently,

$$h^*(\tau) + \tau \frac{dh^*(\tau)}{d\tau} = 0. \tag{8}$$

This latter expression is a simple differential equation the solutions of which are of the form

$$h(\tau) = \frac{a}{\tau},$$

where a is a constant. Notice that (6) holds for any such solution. Thus, (7) is not only

[14]This assumption will be verified later, as will be the question of whether there exist equilibria in which H^* is not twice differentiable.

necessary but also sufficient for (6).

For $\tau \in (\mu, \bar{\theta})$, the seller's payoff when he adopts the price distribution H^* and nature replies with the distribution F_τ, $\tau \in (\mu, \bar{\theta})$, is

$$p(\tau) \int_{\underline{\tau}}^{\tau} t h^*(t) dt.$$

Nature draws the valuation τ with probability $p(\tau)$ and valuation $\underline{\tau}$ with the remaining probability. In the first case the seller sells at a price t if t is below τ. In the second case, the seller does not sell the good. Following the above procedure one obtains the same differential equation

$$h^*(\tau) + \tau \frac{dh^*(\tau)}{d\tau} = 0.$$

The preceding analysis implies that there is a twice differentiable price distribution against which F^* is a best response if and only if there is a constant a such that the function $h(\tau) = a/\tau$, is a density, i.e., h integrates to 1 over $[\underline{\tau}, \bar{\theta}]$.[15] This pins down the constant a:

$$a = \frac{1}{\ln(\bar{\theta}) - \ln(\underline{\tau})} > 0.$$

The following result characterizes the seller's optimal behavior.

Proposition 3. *Suppose the seller knows that valuations are drawn from a distribution with a mean in* $[\underline{\mu}, \bar{\mu}]$ *and with a support that is contained in* $[0, \bar{\theta}]$*, where* $\underline{\mu} \leq \bar{\mu} \leq \bar{\theta}$*. Then, it is optimal for the seller to commit to a randomization over prices described by the distribution*

$$H^*(\tau) = \begin{cases} 0, & \text{if } \tau < \underline{\tau} \\ \dfrac{\ln(\tau) - \ln(\underline{\tau})}{\ln(\bar{\theta}) - \ln(\underline{\tau})}, & \text{if } \underline{\tau} \leq \tau \leq \bar{\theta}, \end{cases}$$

where $\underline{\tau}$ *is the unique solution to* $\underline{\tau}\left[1 + \ln(\bar{\theta}) - \ln(\underline{\tau})\right] = \underline{\mu}$.

We have shown that (H^*, F^*) is a saddle point of the seller's payoff functional given in (2). Since in the saddle both the seller and nature randomize over the interval $[\underline{\tau}, \bar{\theta}]$ and the seller sells with probability one when he charges the price $\underline{\tau}$, his expected payoff is $\underline{\tau}$.

Interestingly, the randomization over prices H^* makes the seller's payoff dependent only on the mean of nature's distribution; bar a class of suboptimal distributions. More precisely, when

[15] Strictly speaking, the above argument establishes that the density must be a_1/τ on $[\underline{\tau}, \underline{\mu})$, and a_2/τ on $[\underline{\mu}, \bar{\theta}]$, for some two positive constants a_1 and a_2. It is then easy to verify that nature's indifference over all binary distributions F_τ implies $a_1 = a_2$.

the seller randomizes over prices with density $h(\tau) = a/\tau$ as specified above and the valuations are distributed with a distribution F, the seller's payoff is

$$\int_{\underline{\tau}}^{\bar{\theta}} x(1 - F(x))h(x)dx = a \int_{\underline{\tau}}^{\bar{\theta}} (1 - F(x))dx$$
$$= a\mu - a \int_{0}^{\underline{\tau}} (1 - F(x))dx,$$

where μ is the mean of distribution F, and where the last equality follows from $\int (1 - F(x))dx = \mu$. It is thus optimal for nature to set $F(\tau) = 0$ for all $\tau \leq \underline{\tau}$. In words, nature does not want to assign positive probability to valuations below $\underline{\tau}$. But then the seller's payoff is simply $a\mu$. The seller, in a sense, insures himself against the information about the parameters of the type distribution he does not have by making his payoff dependent only on the mean.

Uniqueness. We strengthen the above result by showing that (H^*, F^*) is the unique saddle of U and therefore H^* the unique optimal strategy for the seller. We first argue that F^* is the only type distribution that can be part of a saddle point. Suppose that $(\hat{H}, \hat{F})$ is another saddle point. Then, $(H^*, \hat{F})$ is a saddle point too; this is a standard result, an immediate consequence, for example, of Lemma 36.2 in Rockafellar (1997) (see also Appendix B). Since the support of H^* is the entire interval $[\underline{\tau}, \bar{\theta}]$, the seller must be indifferent between all prices in the interval, i.e.

$$(1 - \hat{F}(\tau))\tau = U^*$$

for all $\tau \in [\underline{\tau}, \bar{\theta}]$. We have seen earlier that F^* is the only distribution function which satisfies this condition. Therefore, it must be the case that $\hat{F} = F^*$.

Next we argue $\hat{H} = H^*$. First we show that if $(\hat{H}, \hat{F})$ is a saddle point, then $\hat{H}$ is strictly increasing and continuous on $[\underline{\tau}, \bar{\theta}]$.

Lemma 3. *If $(\hat{H}, \hat{F})$ is a saddle point then $\hat{H}$ has no mass points. Moreover, the support of $\hat{H}$ is $[\underline{\tau}, \bar{\theta}]$.*

As mentioned before, if there is a second solution to the seller's problem, $\hat{H}$, then $(\hat{H}, F^*)$ must also form a saddle. Moreover, the two saddles have the same value. Lemma 1 then implies that when the seller uses $\hat{H}$, all the τ-distributions must give the same payoff , i.e.,

$$p(\tau) \int_{\underline{\tau}}^{\tau} t dH(t) + (1 - p(\tau)) \int_{\underline{\tau}}^{\bar{\theta}} t dH(t) = U^*, \tag{9}$$

for all $\tau \in [\underline{\tau}, \underline{\mu}]$ and $H \in \{\hat{H}, H^*\}$ and

$$p(\tau) \int_{\underline{\tau}}^{\tau} t dH(t) = U^*, \tag{10}$$

for all $\tau \in (\underline{\mu}, \bar{\theta}]$ and $H \in \{\hat{H}, H^*\}$. Equality (10) implies $\int_{\underline{\tau}}^{\tau} t d\hat{H} = \int_{\underline{\tau}}^{\tau} t dH^*$ for all $\tau \in (\underline{\mu}, \bar{\theta}]$. Moreover, using $\int_{\underline{\tau}}^{\bar{\theta}} t d\hat{H} = \int_{\underline{\tau}}^{\bar{\theta}} t dH^*$ with (9) yields

$$\int_{\underline{\tau}}^{\tau} t d\hat{H}(t) = \int_{\underline{\tau}}^{\tau} t dH^*,$$

for all $\tau \in [\underline{\tau}, \underline{\mu}]$. Since $\hat{H}$ is continuous by Lemma 3, it follows that $\hat{H}(\tau) = H^*(\tau)$ for all $\tau \in [\underline{\tau}, \bar{\theta}]$.[16]

The above analysis implies that the saddle point (H^*, F^*) identified in Proposition 3 is the unique saddle point of the seller's problem given by (2). The result in Appendix B then implies that the seller's optimal strategy is unique.

Proposition 4. *The randomization over prices H^* identified in Proposition 3 is the unique optimal selling strategy for the seller when he believes that nature chooses a type distribution with a mean in $[\underline{\mu}, \bar{\mu}]$ and with a support in $[0, \bar{\theta}]$.*

4.2.1 Comparative statics

Throughout the analysis we fixed the lower bound on the mean of the distributions of valuations, $\underline{\mu}$, and the upper bound on the buyer's valuation $\bar{\theta}$. The analysis showed that the upper bound on the mean of valuations, $\bar{\mu}$, plays no role. The seller's pessimism leads him to expect nature to choose a distribution with the smallest possible mean $\underline{\mu}$. Whether $\bar{\mu}$ is just slightly above $\underline{\mu}$ or equal to $\bar{\theta}$ is then irrelevant.

The lower bound on the mean $\underline{\mu}$ allows for more interesting comparative statics. A higher $\underline{\mu}$ forces nature to assign more weight to higher types. Since those types are willing to pay higher prices the seller's payoff increases. In more detail: for any seller's randomization over prices his payoff cannot decrease if $\underline{\mu}$ increases, and his payoff will increase if the randomization over prices assigns positive probability to prices below $\underline{\mu}$. More can be said in the extreme cases. When $\underline{\mu}$ is close to zero there are virtually no restrictions left on the seller's pessimism.

[16]Suppose two continuous distributions functions $\hat{H}$ and H^* do not coincide. Then there must be points τ' and τ'' such that i) the two functions coincide on $[\underline{\tau}, \tau'] \cup \{\tau''\}$ and ii) one strictly exceeds the other for every $\tau \in (\tau', \tau'')$. But that means that if one conditions on prices below τ'', one conditional distribution strictly first order stochastically dominates the other. Consequently, the corresponding (conditional) expected prices must be strictly ranked as well.

Consequently, his maxmin payoff must go to zero. On the other hand, when $\underline{\mu}$ is close to $\bar{\theta}$ the buyer's valuation is with high probability close to $\bar{\theta}$, and so is therefore the seller's payoff.

As $\bar{\theta}$ increases, the buyer can have higher valuation. However, the seller's maxmin expected payoff decreases. An increase of $\bar{\theta}$, ceteris paribus, expands nature's set of strategies. In particular, the weight that any $\tau-$extreme distribution with $\tau < \underline{\mu}$ assigns to τ increases because the counterweight $\bar{\theta}$ is larger. But if nature is able to concentrate more probability mass on low types the seller's maxmin payoff must decrease. In fact, when $\bar{\theta}$ goes to infinity the seller's payoff collapses to zero.

Proposition 5. *Denote the maximal payoff that the seller can achieve when nature has to choose a price distribution from $\mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})$ by $U^*(\underline{\mu}, \bar{\mu}, \bar{\theta})$. U^* is strictly increasing in $\underline{\mu}$, constant in $\bar{\mu}$, and strictly decreasing in $\bar{\theta}$. Furthermore, $U^*(0, \bar{\mu}, \bar{\theta}) = 0$, $U^*(\bar{\theta}, \bar{\theta}, \bar{\theta}) = \bar{\theta}$ and $\lim_{\bar{\theta}\to\infty} U^*(\underline{\mu}, \bar{\mu}, \bar{\theta}) = 0$.*

The above result is illustrated in the following figure. The panel on the left hand of the figure shows the seller's optimal payoff U^* in dependence of $\underline{\mu}$ with $\bar{\theta}$ fixed at 1.

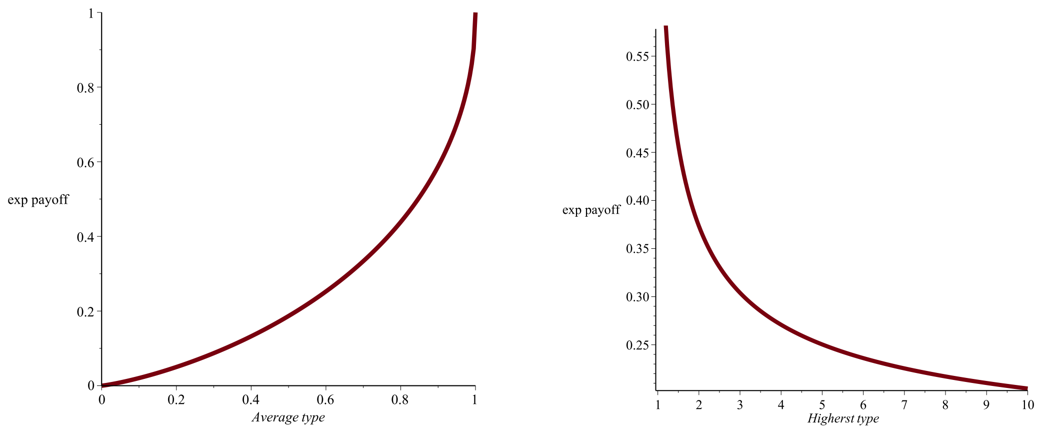

Figure 1: U^* as function of $\underline{\mu}$ (left) and $\bar{\theta}$ (right)

The result on the behavior of U^* has immediate implications for how the seller's optimal strategy, H^* varies with the parameters $\underline{\mu}$ and $\bar{\theta}$. This is so because U^* coincides with $\underline{\tau}$, and from Proposition 3 we know that $\underline{\mu}$ enters H^* only through $\underline{\tau}$. The following proposition provides the details. We write $H^*_{\underline{\mu},\bar{\theta}}$ for the optimal price distribution in the environment characterized by the parameters $\underline{\mu}$ and $\bar{\theta}$.[17]

[17] We drop $\bar{\mu}$ from the list of parameters since the optimal choices do not depend on it.

Proposition 6. *The following holds:*

i) If $\underline{\mu}' > \underline{\mu}$, then $H^*_{\underline{\mu}',\bar{\theta}}$ first order stochastically dominates its counterpart $H^*_{\underline{\mu},\bar{\theta}}$.

ii) Let $\bar{\theta}' > \bar{\theta}$. There exists a $\tilde{\tau}$ such that $H^*_{\underline{\mu},\bar{\theta}'}(\tau) > (<) H^*_{\underline{\mu},\bar{\theta}}(\tau)$ for all $\tau < (>)\tilde{\tau}$.

Part i) of the above proposition is illustrated in the following figure. The panel on the left hand of the figure shows the price distributions corresponding to the parameter pairs $(\underline{\mu}, \bar{\theta}) = (1/3, 1)$ (blue curve) and $(\underline{\mu}, \bar{\theta}) = (2/3, 1)$ (red curve). The right hand panel shows the type distributions for the same two parameter pairs (with the same coloring).

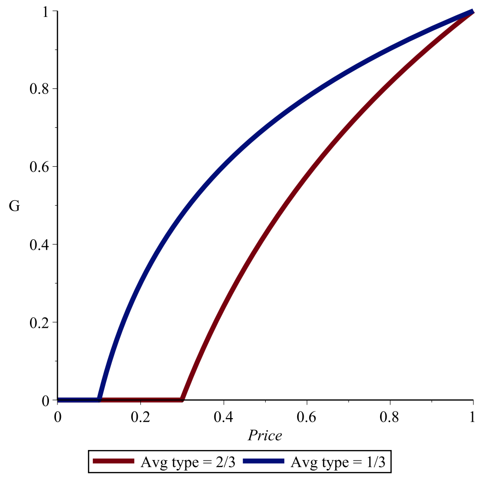

Figure 2: Price distributions $H^*_{\frac{1}{3},1}$ and $H^*_{\frac{2}{3},1}$.

4.3 Posted prices vs. randomized prices

Deterministic posted prices are desirable due to their simplicity. However, the above analysis shows that the seller achieves a higher profit by using a randomized pricing policy. In this section we study how large the loss is when the seller forgoes the possibility to randomize over prices and how this loss varies with the parameters of the model.

In order to simplify the interpretation of the results we express the loss as a fraction of the maximally achievable payoff (i.e. the payoff that the optimal randomized price generates), and denote it by $\rho(\underline{\mu}, \bar{\theta})$. It is straightforward to show that both the optimal payoff with randomizations, U^*, and its counterpart in the case of (deterministic) posted prices are homogeneous of degree one in $(\underline{\mu}, \bar{\theta})$. Therefore, we normalize $\bar{\theta} = 1$ and analyze how $\rho(\underline{\mu}, 1)$ varies with $\underline{\mu}$.

Proposition 7. *The relative loss* $\rho(\cdot, 1)$ *is strictly decreasing and satisfies* $\rho(0, 1) = 1$ *and* $\rho(1, 1) = 0$. *That is, the principal's relative loss is large (*100%*) when* $\underline{\mu}$ *converges to* 0*, and it vanishes only when his knowledge of the mean type implies knowledge of the exact type:* $\underline{\mu} = \bar{\theta} = 1$.

When $\underline{\mu}$ is close to 1 most of the probability over valuations must be close to 1 too, therefore there is little loss in using deterministic prices. One might then think that the same result would obtain when $\underline{\mu}$ is close to the lower bound of the support, but this is not the case. The relative loss is strictly decreasing and approaches 100% when $\underline{\mu}$ approaches 0. This is a consequence of the fact that when the seller offers a deterministic price τ, nature optimally uses a distribution on $\{\tau, 1\}$.

The above result is illustrated in the following figure.

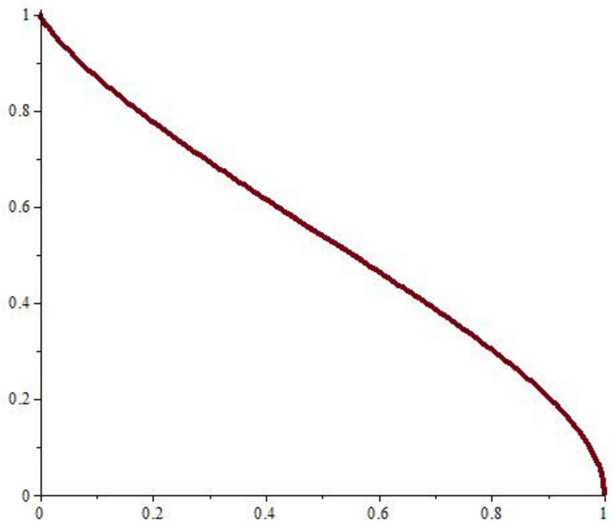

Figure 3: The relative loss of foregoing the option to randomize

5 Variance

The bound on the support of nature's distribution implicitly imposes a bound on the variance of the same distribution. Therefore it is natural to explore what the optimal selling mechanism are when in addition to bounding the mean of nature's distribution, one also bounds the variance.

Suppose that on top of knowing that the mean of the type distribution belongs to the interval $[\underline{\mu}, \bar{\mu}]$ the seller also has some information about the dispersion. In particular, assume that he knows that the variance is in the interval $[\underline{\sigma}^2, \bar{\sigma}^2]$.

From the analysis in the preceding sections we know that the relevant part of the information with respect to the mean is its lowest possible value, $\underline{\mu}$. In particular, the beliefs that determine the seller's behavior have mean $\underline{\mu}$. This insight carries over to the analysis of this section. Indeed, any type distribution with a mean $\tilde{\mu} > \underline{\mu}$ can be shifted downward until it satisfies the mean constraint with equality.[18] Such a shift can at most lead to a reduction of the variance (depending on whether or not the shift creates a mass point at 0) and could at most decrease the seller's payoff. In what follows we therefore only consider type distributions with mean $\underline{\mu}$ to start with. Given this restriction it is convenient to simplify notation by writing μ instead of $\underline{\mu}$.

To derive the seller's optimal mechanism we proceed again as in the case with a bound on the highest type. That is, we look for a saddle point (H^*, F^*) of the seller's payoff functional U. We conjecture that the saddle has certain properties, derive how the saddle should behave given the assumptions, and at the end verify that the derived object is indeed a saddle.

Suppose the seller randomizes over some interval $[\underline{\tau}, \bar{\tau}]$. If that is the case, then nature's distribution must be such that the seller is indifferent between all the prices in the same interval. Following the reasoning from the case where there was an upper bound on the support of nature's distributions, it can be shown that the indifference implies that nature's distribution F^* must be of the form

$$F^*(\tau) = 1 - \underline{\tau}/\tau, \tag{11}$$

for all $\underline{\tau} \leq \tau < \bar{\tau}$, with a mass point $\bar{\tau}/\underline{\tau}$ at $\bar{\tau}$. In addition, nature's distribution must satisfy the mean constraint

$$\int_{\underline{\tau}}^{\bar{\tau}} \tau dF^* = \mu,$$

which using the functional form for F^* yields

$$\underline{\tau}[1 + \ln(\bar{\tau}) - \ln(\underline{\tau})] = \mu, \tag{12}$$

and the variance constraint. We further conjecture that nature chooses a distribution with the highest possible variance. We have shown before that when unrestricted by variance nature can hold the seller to an arbitrarily low payoff by mixing over two points, for example 0 and $\hat{\tau}$ where the latter is chosen very large, with a distribution that has mean μ. If the variance of nature's optimal distribution was strictly smaller than $\bar{\sigma}^2$, then the nature could achieve a higher payoff by mixing this distribution with the before mentioned binary distribution and achieve a

[18] Of course, the probability mass distributed over $[0, \tilde{\mu} - \mu)$ has to be concentrated on the point 0.

higher payoff. The variance constraint reads

$$\int_{\underline{\tau}}^{\bar{\tau}} \tau^2 dF^* - \mu^2 = \bar{\sigma}^2,$$

which, using the functional form for F^*, simplifies to

$$\underline{\tau}[2\bar{\tau} - \underline{\tau}] - \mu^2 = \bar{\sigma}^2. \tag{13}$$

The following lemma establishes that (12) and (13) have for each pair of parameters $(\mu, \bar{\sigma}^2)$ a unique solution in terms of $\underline{\tau}$ and $\bar{\tau}$. The assumption that the seller randomizes over an interval, therefore, fully pins down the type distribution.

Lemma 4. *The system comprised of equations* (12) *and* (13) *has a unique solution for each pair* $(\mu, \bar{\sigma}^2)$.

Thus far we have assumed that the seller randomizes over an interval, and derived nature's distribution that makes the seller indifferent over all prices in an interval. Notice that given the restrictions on mean and variance, and assuming nature chooses a distribution with the highest variance, nature's distribution is uniquely pinned down. We are left to determine how the seller should randomize in order for the distribution F^* to be optimal for nature.

In the case with an upper bound on valuations, H^* was chosen so that nature was indifferent among all distributions with mean μ. To identify this price distribution we have exploited the fact that general indifference holds for nature if it holds on the set of τ-distributions. This class of distributions allowed for a one dimensional parametrization, which we used to characterize the density of the price distribution through a simple differential equation.

An analogous approach works when a variance constraint is imposed. However, binary distributions cannot be used here. A binary distribution corresponding to a τ too far above or below μ might have too high a variance. The relevant set of elementary distributions over which we need to guarantee indifference is a class of ternary distributions that have mean μ and variance $\bar{\sigma}^2$. In particular, for each type $\tau \in [\underline{\tau}, \bar{\tau}]$ this class contains a distribution whose support includes τ itself and two points from the set $\{\underline{\tau}, \mu, \bar{\tau}\}$ (ternary τ-distribution).

More precisely, the index set of (ternary) τ-distributions with mean μ and variance $\bar{\sigma}^2$, can be split into three segments, $[\bar{\tau}, \tau_1]$, $[\tau_1, \tau_2]$ and $[\tau_2, \bar{\tau}]$ depending on τ. For low values of τ, $\tau \in [\underline{\tau}, \tau_1]$, those are the distributions on $\{\tau, \mu, \bar{\tau}\}$. For the intermediate values, $\tau \in [\tau_1, \tau_2]$; the relevant distributions have the support $\{\underline{\tau}, \tau, \bar{\tau}\}$. Finally, for the high values, $\tau \in [\tau_2, \bar{\tau}]$, the distributions are on $\{\underline{\tau}, \mu, \tau\}$. The threshold τ_1 is the largest τ in $[\underline{\tau}, \bar{\tau}]$ for which there exists a distribution with support $\{\tau, \mu, \bar{\tau}\}$, mean μ, and variance $\bar{\sigma}^2$. On the other hand, τ_2 is the smallest $\tau \in [\underline{\tau}, \bar{\tau}]$ such that there exists a distribution with support $\{\underline{\tau}, \mu, \tau\}$, mean μ, and variance $\bar{\sigma}^2$.

The condition that the payoff across these distributions has to be constant can again be translated into a differential equation for the density of the price distribution. Solving this differential equation yields the result that the density is of the form

$$h(\tau) = a/\tau + b,$$

for $\tau \in [\underline{\tau}, \bar{\tau}]$, where a and b are some constants.

Two conditions are used to determine the constants a and b. The first is to incentivize nature not to assign probability outside of $[\underline{\tau}, \bar{\tau}]$. The next lemma shows that this will be the case precisely when $a/\bar{\tau} + b = 0$, for $a > 0$ and $b < 0$. The second condition is the requirement that h is indeed a density: $\int h(t)dt = 1$.

Lemma 5. *Suppose that the seller randomizes over prices on the interval $[\underline{\tau}, \bar{\tau}]$ with distribution that has density*

$$h(\tau) = \frac{a}{\tau} + b,$$

for $\tau \in [\underline{\tau}, \bar{\tau}]$, 0 otherwise, and $h(\bar{\tau}) = 0$. Then, for every nature's distribution F that satisfies the mean and the variance restriction, there exists a distribution $\hat{F}$ with the same mean and variance, such that the support of $\hat{F}$ is contained in $[\underline{\tau}, \bar{\tau}]$ and the distribution $\hat{F}$ yields a smaller payoff to the seller. Moreover, if F is nature's distribution with support in $[\underline{\tau}, \bar{\tau}]$, then the seller's payoff is

$$a\mu + 0.5b\mu^2 + 0.5b\sigma^2,$$

where μ is the mean and σ^2 the variance of distribution F.

The above lemma shows that when the seller randomizes over prices with density $h(\tau) = a/\tau + b$ on $[\underline{\tau}, \bar{\tau}]$ that satisfies $h(\bar{\tau}) = 0$, nature optimally chooses a distribution whose support is contained in $[\underline{\tau}, \bar{\tau}]$. More importantly, the seller's payoff in the latter case depends only on the mean and the variance of nature's distribution. With other words, the seller can by randomizing over prices with a density h as specified above, in a sense, insure himself against the information about the parameters of the distribution he does not have.

The requirements $h(\bar{\tau}) = \frac{a}{\bar{\tau}} + b = 0$ and $\int_{\underline{t}}^{\bar{t}} h(t)dt = 1$ fully pin down the constants a and b:

$$a = \frac{\bar{\tau}}{\bar{\tau}[\ln(\bar{\tau}) - \ln(\underline{\tau})] - [\bar{\tau} - \underline{\tau}]},$$
$$b = -\frac{1}{\bar{\tau}[\ln(\bar{\tau}) - \ln(\underline{\tau})] - [\bar{\tau} - \underline{\tau}]}.$$

It is easy to see that $a > 0$ and $b < 0$ using the inequality $\ln(x) \leq x - 1$, where the inequality is strict for $x > 0$.

The upcoming result provides a characterization of the seller's optimal pricing policy when he has information only about the mean and the variance of the distribution over valuations.

Let H^* and F^* be two distributions with support $[\underline{\tau}, \bar{\tau}]$ where the pair $\{\underline{\tau}, \bar{\tau}\}$ is the unique solution of the system

$$\underline{\tau}[1 + \ln(\bar{\tau}) - \ln(\underline{\tau})] = \mu,$$
$$\underline{\tau}[2\bar{\tau} - \underline{\tau}] - \mu^2 = \bar{\sigma}^2.$$

Let

$$H^*(\tau) = \frac{\bar{\tau}[\ln(\tau) - \ln(\underline{\tau})] - (\tau - \underline{\tau})}{\bar{\tau}[\ln(\bar{\tau}) - \ln(\underline{\tau})] - (\bar{\tau} - \underline{\tau})},$$

for $\tau \in [\underline{\tau}, \bar{\tau}]$, with $H^*(\tau) = 0$ for $\tau < \underline{\tau}$, and $H^*(\tau) = 1$ for $\tau > \bar{\tau}$, and

$$F^*(\tau) = 1 - \frac{\underline{\tau}}{\tau},$$

for $\tau \in [\underline{\tau}, \bar{\tau})$, and $F^*(\tau) = 1$ for $\tau \geq \bar{\tau}$.

Proposition 8. *If the seller believes that the buyer's valuations are drawn from a distribution with mean μ and variance in the interval $[\underline{\sigma}^2, \bar{\sigma}^2]$ it is optimal for him to randomize over prices with distribution H^*. The seller's max-min expected payoff is $\underline{\tau}$.*

Proof. We started by conjecturing that the seller randomizes over an interval. Assuming further that nature chooses a distribution with the highest possible variance enabled us to derive the unique nature's distribution, F^*, that makes the seller indifferent over some interval. As a byproduct we obtained the interval over which the seller randomizes, i.e. $\underline{\tau}$ an $\bar{\tau}$.

Lemma 5 shows that when the seller randomizes with H^* it is optimal for nature to choose a distribution with a support in $[\underline{\tau}, \bar{\tau}]$. Moreover, once nature chooses such a distribution the seller's payoff depends only on the mean and the variance of the distribution. In particular it is increasing in the mean and decreasing in the variance. The distribution F^* with the mean $\bar{\sigma}^2$

therefore minimizes the seller's payoff, given his randomization over the prices, and is therefore optimal for nature. The pair (H^*, F^*) is thus indeed a saddle. □

Uniqueness. Here we argue that the above established saddle, and therefore the optimal selling strategy for the seller, are unique. Suppose that there was another saddle $(\hat{H}, \hat{F})$. Then $(H^*, \hat{F})$ is a saddle too. The reasoning behind why $\hat{F} = F^*$ is the same as in the case of upper bound on the support. The idea is, the support of H^* is $[\underline{\tau}, \bar{\tau}]$, therefore this must also be the support of $\hat{F}$. But then the seller must be indifferent over all the prices in the interval. This pins down $\hat{F}$. In particular, $\hat{F} = F^*$.

On the other hand, if $(\hat{H}, \hat{F})$ is a saddle, so is $(\hat{H}, F^*)$. In a result paralleling Lemma 1 we show that if the distribution of valuations F^* is a best response to a price distribution $\hat{H}$, then nature is indifferent between ternary distributions F_τ for every $\tau \in [\underline{\tau}, \bar{\tau}]$ (when the seller chooses $\hat{H}$).

Lemma 6. *If the type distribution F^* is a best response to a price distribution $\hat{H}$, then also F_τ is a best response to $\hat{H}$ for every $\tau \in [\underline{\tau}, \bar{\tau}]$.*

The above Lemma implies

$$p_\tau(\tau) \int_{\underline{\tau}}^{\tau} t d\hat{H} + p_\tau(\mu) \int_{\underline{\tau}}^{\mu} t d\hat{H} + p_\tau(\bar{\tau}) \int_{\underline{\tau}}^{\bar{\tau}} t d\hat{H} = U^*,$$

for all $\tau \in [\underline{\tau}, \tau_1]$, where $p_\tau(x)$, for $x \in \{\tau, \mu, \bar{\tau}\}$, is the probability the ternary distribution F_τ assigns to point x; similar expressions are obtained for $\tau \in [\tau_1, \tau_2]$ and $\tau \in [\tau_2, \bar{\tau}]$. The left hand side of the above equation can, therefore, be thought of as a constant function and is thus differentiable in τ. Since $p_\tau(\cdot)$ is differentiable in τ for $\tau \in (\underline{\tau}, \tau_1)$ one can show that H must be differentiable. Even more, we argue that the fact that the left hand side of the above equation is identical to U^* implies that H is three times differentiable.

Lemma 7. *If $(F^*, \hat{H})$ is a saddle, then $\hat{H}$ must be three times differentiable.*

Differentiating (19) three times one obtains a differential equation, the solutions of which are of the form $h(\tau) = a/\tau + b$, for $\tau \in [\underline{\tau}, \bar{\tau}]$ and 0 elsewhere.

Lemma 8. *If $(\hat{H}, F^*)$ is a saddle, then $\hat{H}$ must have a density of the form $h(\tau) = a/\tau + b$ on $[\underline{\tau}, \bar{\tau}]$ and 0 otherwise.*

Thus far we have shown that in any saddle $(\hat{H}, \hat{F})$, $\hat{F} = F^*$ and $\hat{H}$ must have density of the form $h(\tau) = a/\tau + b$. Proposition 8 showed that for

$$a = \frac{\bar{\tau}}{\bar{\tau}[\ln(\bar{\tau}) - \ln(\underline{\tau})] - [\bar{\tau} - \underline{\tau}]},$$
$$b = -\frac{1}{\bar{\tau}[\ln(\bar{\tau}) - \ln(\underline{\tau})] - [\bar{\tau} - \underline{\tau}]},$$

one obtains an optimal distribution. Here we show that this is the only combination of a and b that yields an optimal distribution over prices.

Lemma 9. *Let $(\hat{H}, \hat{F})$ be a saddle, where $\hat{H}$ has a density of the form $h(\tau) = a/\tau + b$ on $[\underline{\tau}, \bar{\tau}]$ and 0 otherwise. Then*

$$a = \frac{\bar{\tau}}{\bar{\tau}[\ln(\bar{\tau}) - \ln(\underline{\tau})] - [\bar{\tau} - \underline{\tau}]},$$
$$b = -\frac{1}{\bar{\tau}[\ln(\bar{\tau}) - \ln(\underline{\tau})] - [\bar{\tau} - \underline{\tau}]}.$$

The above sequence of lemmata delivers the following result.

Proposition 9. *The distribution over prices H^* is the unique optimal seller's distribution.*

Proof. Suppose there was a saddle $(\hat{H}, \hat{F}) \neq (H^*, F^*)$. We have argued in the text that $\hat{F} = F^*$. Moreover, $(\hat{H}, F^*)$ must also be a saddle. Therefore $\hat{H}$ is three times differentiable by Lemma 7. In turn, $\hat{H}$ has a density of the form $h(\tau) = a/\tau + b$, for $\tau \in [\underline{\tau}, \bar{\tau}]$ and 0 otherwise by Lemma 8. Constants a and b are then pinned down by Lemma 9. Therefore $\hat{H} = H^*$, which contradicts $(\hat{H}, \hat{F}) \neq (H^*, F^*)$. □

Comparative Statics. The seller's optimal strategy is independent of $\underline{\sigma}^2$. What matters is the upper bound on the variance of the type distribution $\bar{\sigma}^2$. The seller fears nature choosing a distribution that places most of the weight on very low types. To preserve the mean μ, it must pick with some probability large types. Thus, what the seller fears are strongly dispersed distributions which are characterized by a large variance. The following result formalizes the above idea.

Proposition 10. *The seller's payoff is strictly decreasing in $\bar{\sigma}^2$.*

Proof. Since an increase of $\bar{\sigma}^2$ means that the variance constraint on nature is relaxed it follows that nature must be able to do weakly better against any strategy of the seller. In particular, this is true also vis-a-vis the equilibrium pricing strategy. Thus, the seller's payoff cannot possibly increase as $\bar{\sigma}^2$ increases.

In order to show that it must actually strictly decrease remember that the seller's payoff is $\underline{\tau}$ and notice that the condition

$$\underline{\tau}\left[1 + \ln(\bar{\tau}) - \ln(\underline{\tau})\right] = \mu$$

implies

$$\frac{d\underline{\tau}}{d\bar{\sigma}^2}\left[1 + \ln(\bar{\tau}) - \ln(\underline{\tau})\right] + \underline{\tau}\left[\frac{d\bar{\tau}}{d\bar{\sigma}^2}\frac{1}{\bar{\tau}} - \frac{d\underline{\tau}}{d\bar{\sigma}^2}\frac{1}{\underline{\tau}}\right] = 0,$$

which can be rewritten in the following more compact format (using $1 + \ln(\bar{\tau}) - \ln(\underline{\tau}) = \mu/\underline{\tau}$)

$$\frac{d\underline{\tau}}{d\bar{\sigma}^2}\frac{\mu - \underline{\tau}}{\underline{\tau}} + \frac{d\bar{\tau}}{d\bar{\sigma}^2}\frac{\underline{\tau}}{\bar{\tau}} = 0$$

Therefore, whenever $\underline{\tau}$ is constant in $\bar{\sigma}^2$ so is $\bar{\tau}$. But if neither of the two limits of the support of F^* change, then F^* itself remains unchanged and so also its variance stays the same. A constant variance of F^* is in contradiction with the fact that the variance constraint must be binding. □

While the seller attaches no value to information about the mean only, as long as he contemplates as possible all non-negative valuations, additional information about the mean or the variance can have an effect on the seller's payoff when he has some information about the mean and the variance to start with.

An interesting implication of Proposition 8 is that whenever the variance constraint is binding a high enough upper bound on the set of possible types can be imposed without altering the problem's solution or value. Thus, in a sense one can replace the bounds on the set of feasible types with a bound on the variance of the types and vice versa. More precisely, in the problem where the seller knows the mean belongs to the interval $[\underline{\mu}, \bar{\mu}]$ and the upper bound on the valuations is $\bar{\theta}$, the seller and nature randomize over the interval $[\underline{\tau}, \bar{\theta}]$ where $\underline{\tau}$ is defined by

$$\underline{\tau}\left[1 + \ln(\bar{\theta}) - \ln(\underline{\tau})\right] = \underline{\mu}.$$

If the seller knew instead that the mean belongs to the interval $[\underline{\mu}, \bar{\mu}]$ and the variance to $[\underline{\sigma}^2, \bar{\sigma}^2]$

he would randomize over the interval $[\underline{\tau}, \bar{\tau}]$ where $\underline{\tau}$ and $\bar{\tau}$ are defined by

$$\underline{\tau}[1 + \ln(\bar{\tau}) - \ln(\underline{\tau})] = \underline{\mu},$$
$$\underline{\tau}[2\bar{\tau} - \underline{\tau}] - \mu^2 = \bar{\sigma}^2.$$

The two intervals of randomization coincide when $\bar{\tau} = \bar{\theta}$ and therefore $\bar{\sigma}^2 = \underline{\tau}[2\bar{\theta} - \underline{\tau}] - \underline{\mu}^2$. Furthermore, in both cases the seller's payoff is precisely $\underline{\tau}$. Therefore the seller knowing $\mu \in [\underline{\mu}, \bar{\mu}]$ and that the possible valuations are in $[0, \bar{\theta}]$ is payoff equivalent for the seller to knowing that $\mu \in [\underline{\mu}, \bar{\mu}]$ and $\sigma^2 \in [\underline{\sigma}^2, \bar{\sigma}^2]$ for the $\bar{\sigma}^2$ as derived above (and any $\underline{\sigma}^2 \leq \bar{\sigma}^2$).

This notwithstanding, the two types of constraints are not equivalent. In fact, there is no pair $\bar{\theta}$ and $\bar{\sigma}^2$ such that the $\bar{\sigma}^2$-constrained and the $\bar{\theta}$ constrained problem yield the same solution to the seller's problem. The following proposition elaborates on this point. It states that whenever the two types of constraint lead to the same support for the price and type distributions, the price distribution that corresponds to the case with a bound on the type set first order stochastically dominates the price distribution that corresponds to the case with a constraint on the variance. Therefore, one could discern from the data whether the seller's information is about the mean and the variance or about the mean and the upper bound on the valuations.

Proposition 11. *Let μ, $\bar{\theta}$ and $\bar{\sigma}^2$ be such that the pairs $(\mu, \bar{\theta})$ and $(\mu, \bar{\sigma}^2)$ induce the same supports for the price distribution $[\underline{\tau}, \bar{\theta}]$. If the two optimal price distributions are denoted by $H^*_{\bar{\theta}}$ and $H^*_{\bar{\sigma}^2}$, respectively, then for all $\underline{\tau} < \tau < \bar{\theta}$,*

$$H^*_{\bar{\sigma}^2}(\tau) > H^*_{\bar{\theta}}(\tau).$$

The above result is illustrated in the following figure which shows the price distributions for the pairs $(\mu, \bar{\theta}) = (1/2, 1)$ and $(\mu, \sigma^2) = (1/2, 0.088)$.

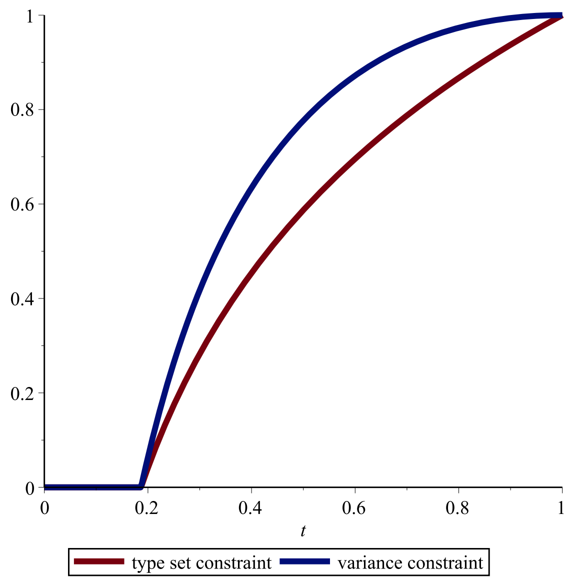

Figure 4: Optimal price distributions: bounded variance (blue), bounded type set (red)

6 Conclusion

We showed that the seller who only has information about the mean expects payoff zero regardless of which mechanism he uses. However, the choice of the seller's mechanism does become important when he has information both about the mean and the variance. In that case more information benefits the seller. However, even when he knows the mean and the variance he entertains as possible a rather "large" set of distributions. To connect the model with the Bayesian setting it would be of interest to explore the seller's optimal mechanisms when he also has the information about higher order moments. This is left for future research.

7 Appendix A

Proof of Proposition 1. Let H be a distribution over prices for which $\int_0^\infty \tau dH$ exists and is finite. Let $\tilde{F}$ be the binary distribution that assigns probability p to some $\theta > 0$, probability $1 - p$ to 0, and has mean $\underline{\mu}$; thus, $p = \frac{\underline{\mu}}{\theta}$. The seller's payoff from H can be bounded as follows

$$\begin{aligned}\inf_{F\in\mathcal{F}(\underline{\mu},\bar{\mu})} U(H,F) &\leq U(H,\tilde{F}) \\ &= p\int_0^\theta tdH \\ &= \frac{\underline{\mu}}{\theta}\int_0^\theta tdH \\ &\leq \frac{\underline{\mu}}{\theta}\int_0^\infty tdH,\end{aligned}$$

where the second line is due to nature randomizing over values 0 and θ, and the fact that the seller receives the price t when the said price is below θ.

Since θ can be chosen arbitrarily large and $\int_0^\infty tdH$ is finite, the upper bound on the seller's payoff can be made arbitrarily small. That is, $\inf_{F\in\mathcal{F}(\underline{\mu},\bar{\mu})} U(H,F) = 0$. □

Proof of Lemma 1. If for some $\tau \in [\underline{\tau}, \bar{\theta}]$, F_τ gives nature a strictly higher payoff against $\hat{H}$ than F^*, then F^* is not a best response to $\hat{H}$. Thus if F^* is to form a saddle with $\hat{H}$, then it must be the case that $U(\hat{H}, F_\tau) \geq U(\hat{H}, F^*)$ for every $\tau \in [\underline{\tau}, \bar{\theta}]$.

Next we argue that no F_τ can give nature a strictly smaller payoff than F^*. Suppose, to the contrary, that $U(\hat{H}, F_{\hat{\tau}}) > U(\hat{H}, F^*)$ for some $\hat{\tau} \in [\underline{\tau}, \underline{\mu})$ (the case $\hat{\tau} \in [\mu, \bar{\theta}]$ is very similar and is thus omitted). We claim that this implies that there is a (small) $\varepsilon > 0$, such that $U(\hat{H}, F_{\tau'}) > U(\hat{H}, F^*)$ holds for every $\tau' \in (\hat{\tau}, \hat{\tau} + \epsilon)$. Since

$$U(\hat{H}, F_\tau) = p(\tau)\int_{\underline{\tau}}^{\tau} xd\hat{H}(x) + (1 - p(\tau))\int_{\underline{\tau}}^{\bar{\theta}} xd\hat{H}(x),$$

it follows that for every $\tau' > \hat{\tau}$,

$$U(\hat{H}, F_{\hat{\tau}}) - U(\hat{H}, F_{\tau'}) = (p(\tau') - p(\hat{\tau}))\int_{\hat{\tau}}^{\bar{\theta}} xd\hat{H}(x) - p(\tau')\int_{\hat{\tau}}^{\tau'} xd\hat{H}(x).$$

When τ' converges to $\hat{\tau}$ from above the first term of the above expression vanishes, while the second may not if $\hat{H}$ has atoms at $\hat{\tau}$ or around it. Therefore, either $U(\hat{H}, F_{\tau'})$ converges to $U(\hat{H}, F_\tau)$ (when $\hat{H}$ is atomfree around $\hat{\tau}$) or, $U(\hat{H}, \tau) < U(\hat{H}, \tau')$ for every $\tau' > \hat{\tau}$ that is sufficiently close to $\hat{\tau}$; notice that in this latter case, $F_{\tau'}$ gives nature a strictly lower payoff against $\hat{H}$ than F^*. Thus, in either of the two cases there exists an $\varepsilon > 0$ such that nature obtains a smaller payoff from $F_{\tau'}$ for $\tau' \in [\hat{\tau}, \hat{\tau} + \epsilon)$ than it does from F^*. Let $\tau_l = \hat{\tau}$ and $\tau_h = \hat{\tau} + \varepsilon/2$.

The next step of our argument is to show that F^* can be written as a convex combination of two distribution functions, F_α and G, where the latter is defined as a mixture of $F_{\tau'}$-distributions with $\tau' \in [\tau_l, \tau_h]$. Thus, by construction G is a type distribution that yields to nature a strictly lower payoff when played against $\hat{H}$ than F^* does. But then F_α must be a distribution that does better against $\hat{H}$ than F^*, meaning that $(\hat{H}, F^*)$ cannot be a saddle of U.

Let

$$G(t) = \frac{1}{\tau_h - \tau_l} \int_{\tau_l}^{\tau_h} F_\tau(t) d\tau,$$

for $t \in [0, \bar{\theta}]$, be a uniform mixture over F_τ for $\tau \in [\tau_l, \tau_h]$. It is easy to verify that G is a distribution on $[0, \bar{\theta}]$ with mean $\underline{\mu}$. Using the definitions of F_τ one can rewrite G as follows

$$\begin{aligned} G(t) &= \frac{1}{\tau_h - \tau_l} \int_{\tau_l}^{t} p_\tau d\tau \\ &= \frac{1}{\tau_h - \tau_l} \int_{\tau_l}^{t} \frac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - \tau} d\tau, \end{aligned}$$

for $t \in [\tau_l, \tau_h]$ and $G(t) = \frac{1}{\tau_h - \tau_l} \int_{\tau_l}^{\tau_h} \frac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - \tau} d\tau$, for $t \in (\tau_h, \bar{\theta})$ (also, $G(t) = 0$ for $t < \tau_l$). For later reference, the derivative

$$G'(t) = \frac{1}{\tau_h - \tau_l} \frac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - t}$$

attains its maximum on $[\tau_l, \tau_h]$ at τ_h where $G'(\tau_h) = (\bar{\theta} - \underline{\mu})/[(\bar{\theta} - \tau_h)(\tau_h - \tau_l)]$.

Define

$$F_\alpha = \frac{1}{1 - \alpha} [F^* - \alpha G].$$

We will argue that for sufficiently small α, F_α is a distribution. Clearly $F_\alpha(0) = 0$ and $F_\alpha(\bar{\theta}) = 1$. Since both F^* and G are right continuous so must be F_α. The main point is to show that F_α is nondecreasing for small enough (but positive) α. Towards that, notice that

$$F'_\alpha(t) = \frac{1}{1 - \alpha} [F'(t) - \alpha G'(t)],$$

at all the points of differentiability of F and G. The critical region of t is $[\tau_l, \tau_h]$ and the point

$t = \bar{\theta}$ where F_α exhibits a jump that it inherits from F^* and G^*. Below this interval $G = 0$ and above it G is again constant except for the jump at $\bar{\theta}$.

For $t \in [\tau_l, \tau_h]$

$$
\begin{aligned}
F'_\alpha(t) &= \frac{1}{1-\alpha}[F'(t) - \alpha G'(t)] \\
&= \frac{1}{1-\alpha}\left[\frac{\underline{\tau}}{t} - \alpha \frac{1}{\tau_h - \tau_l}\frac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - t}\right] \\
&\geq \frac{1}{1-\alpha}\left[\frac{\underline{\tau}}{\tau_h} - \alpha \frac{1}{\tau_h - \tau_l}\frac{\bar{\theta} - \underline{\mu}}{\bar{\theta} - \tau_h}\right],
\end{aligned}
$$

where the second equality uses the above derived expressions for F' and G', and the inequality is due to the fact that the derivative $F'(t)$ is minimized and the derivative G' maximized at τ_h. The expression in the last bracket is clearly positive for α small enough. Notice also that for α small enough the jump that F_α exhibits at $\bar{\theta}$ must be an upward jump.

Now fix some (sufficiently small but strictly positive) α such that F_α is a distribution function and rewrite the defining equation of F_α as follows

$$F^* = (1-\alpha)F_\alpha + \alpha G.$$

Since both G and F^* have mean $\underline{\mu}$ this equation implies that also F_α must have mean $\underline{\mu}$. Therefore, F_α represents a viable strategy for nature. Since playing G against $\hat{H}$ yields for nature a strictly smaller payoff than playing F^*, it must be the case that F_α yields a strictly higher payoff against $\hat{H}$ than F^* does. But then F^* cannot be the best reply to $\hat{H}$ which contradicts our starting assumption. □

Proof of Lemma 2. Let H be as in the statement of the result and take any type distribution $F \in \mathcal{F}(\underline{\mu}, \bar{\mu}, \bar{\theta})$ with the support in $[\underline{\tau}, \bar{\theta}]$. The seller's expected payoff from (H, F) is

$$U(H,F) = \int_{\underline{\tau}}^{\bar{\theta}} \int_{\underline{\tau}}^{\tau} t dH(t) dF(\tau). \tag{14}$$

By assumption for every τ the payoff under the corresponding τ-extreme distribution is equal to some constant c. That is

$$
\begin{aligned}
p(\tau)\int_{\underline{\tau}}^{\tau} t dH(t) + (1 - p(\tau))\int_{\underline{\tau}}^{\bar{\theta}} t dH(t) &= c \qquad \text{for all } \underline{\tau} \leq \tau < \underline{\mu}, \\
p(\tau)\int_{\underline{\tau}}^{\tau} t dH(t) &= c \qquad \text{for all } \underline{\mu} < \tau \leq \bar{\theta}.
\end{aligned}
$$

Observe that the second line implies that the term $\int_{\underline{\tau}}^{\bar{\theta}} t dH(t)$ in the first line can be replaced by $c/p(\bar{\theta})$. Thus, the integrand of the outer integral in (14) is equal to

$$\int_{\underline{\tau}}^{\tau} t dH(t) = \begin{cases} \frac{c}{p(\tau)} - \frac{1-p(\tau)}{p(\tau)} \frac{c}{p(\bar{\theta})} & \text{if } \underline{\tau} \leq \tau < \underline{\mu} \\ \frac{c}{p(\tau)} & \text{if } \underline{\mu} < \tau \leq \bar{\theta}. \end{cases} \tag{15}$$

Using (5) it is straightforward to show that both expressions on the right hand side simplify to

$$\int_{\underline{\tau}}^{\tau} t dH(t) = c \frac{\tau - \underline{\tau}}{\underline{\mu} - \underline{\tau}}.$$

Substituting this into (14) we obtain

$$U(H, F) = \int_{\underline{\tau}}^{\bar{\theta}} \int_{\underline{\tau}}^{\tau} t dH(t) dF(\tau) = \frac{c}{\underline{\mu} - \underline{\tau}} \int_{\underline{\tau}}^{\bar{\theta}} (\tau - \underline{\tau}) dF(\tau) = c, \tag{16}$$

where the third equality follows from the fact that the expectation of τ under F is equal to $\underline{\mu}$. □

Proof of Proposition 3. In the text preceding Proposition 3 we established that (H^*, F^*) is a saddle. The fact that H^* is then optimal for the seller follows from the result in Appendix B. □

Proof of Lemma 3. *Step 1:* No mass point.

If $(\hat{H}, \hat{F})$ is a saddle point, so is $(\hat{H}, F^*)$, which in turn implies $U(H^*, F^*) = U(\hat{H}, F^*)$. Suppose $\hat{H}$ has a mass point, at $\hat{\tau}$, and let $\Delta = \hat{H}(\hat{\tau}) - H(\hat{\tau}-)$ be the size of the mass point at $\hat{\tau}$. The mass point at $\hat{\tau}$ implies that nature's payoff from playing $F_{\hat{\tau}}$ is larger than its payoff from playing F_τ for $\tau \in (\hat{\tau}, \hat{\tau} + \varepsilon)$ for $\varepsilon > 0$ small enough. To see this, assume first that $\hat{\tau} < \underline{\mu}$. In this case the seller's payoff when he randomizes according to $\hat{H}$ and nature replies with F_τ, is

$$p(\tau) \int_{\underline{\tau}}^{\tau} x d\hat{H}(x) + (1 - p(\tau)) \int_{\underline{\tau}}^{\bar{\theta}} x d\hat{H}(x),$$

where $p(\tau) = (\bar{\theta} - \mu)/(\bar{\theta} - \tau)$ is the probability that nature assigns to τ.

Therefore, for $\tau \in (\hat{\tau}, \hat{\tau} + \varepsilon)$

$$U(\hat{H}, F_{\hat{\tau}}) - U(\hat{H}, F_\tau) = (p(\hat{\tau}) - p(\tau)) \int_{\hat{\tau}}^{\bar{\theta}} x d\hat{H}(x) - p(\tau) \int_{\hat{\tau}}^{\tau} x d\hat{H}(x).$$

The first term can be made arbitrarily small by choosing ε small enough. The integral in the second term, on the other hand, is bounded below by $\hat{\tau}\Delta$ because of the atom at $\hat{\tau}$. Therefore, if ε is small enough then $U(\hat{H}, F_{\hat{\tau}}) - U(\hat{H}, F_{\tau}) < 0$. Since nature's objective is to minimize the seller's payoff, this inequality implies that nature is not not indifferent among all τ-distributions when facing a seller who adopts the price distribution $\hat{H}$. But then $(\hat{H}, F^*)$ cannot be a saddle due to Lemma 1.

A perfectly analogous argument applies also in the case $\bar{\theta} > \hat{\tau} \geq \underline{\mu}$. In principle this leaves the possibility of a jump at $\bar{\theta}$. But the price $p = \bar{\theta}$ generates a zero revenue and is thus never optimal. We can therefore conclude that $\hat{H}$ cannot exhibit any mass points.

Step 2: The support of $\hat{H}$ is $[\underline{\tau}, \bar{\theta}]$.

The seller has no incentive to charge a price larger than $\bar{\theta}$ with positive probability (no type buys at such a price), or a price lower than $\underline{\tau}$. In the latter case, the seller would be strictly better off by increasing the price to $\underline{\tau}$, in which case he would still sell the good with probability one. The support of $\hat{H}$ is therefore contained in $[\underline{\tau}, \bar{\theta}]$.

Now suppose that there is some hole in the support. We argue that F^* then cannot be an optimal response to $\hat{H}$ since nature can improve upon it by reallocating the mass that F^* assigns to the hole. Suppose that nature first shifts the mass that F^* assigns to the hole to the upper limit of the hole. Doing so yields a type distribution F' which when played against $\hat{H}$ generates the same payoff as F^*. But then nature can manipulate F' further by shifting a strictly positive share of the mass that F' concentrates on the interval $(\underline{\mu}, \bar{\theta}]$ to the lowest type $\underline{\tau}$ to obtain distribution F'' with the mean $\underline{\mu}$. Clearly, this second shift strictly reduces the seller's payoff. □

Proof of Proposition 4. The result was proven in the text preceding the statement of the result. Here we show how $\underline{\tau}$ can be expressed as a function of Lambert function. The condition that defines the lower limit of the support of the distributions H^* and F^* can be solved to yield the expression

$$\underline{\tau} = -\frac{\mu}{W_{-1}(-\mu/e\bar{\theta})}.$$

To show this, remember that the defining condition of $\underline{\tau}$ is

$$\underline{\tau}(1 + \ln(\bar{\theta}) - \ln(\underline{\tau})) = \mu.$$

Dividing both sides of the equation by $\underline{\tau}$ and evaluating them with the exponential function we can rewrite this condition as follows

$$\underline{\tau}^{-1}e\bar{\theta} = e^{\mu\underline{\tau}^{-1}} \quad \Leftrightarrow \quad -e\bar{\theta}(-\mu\underline{\tau}^{-1}) = \mu e^{\mu\underline{\tau}^{-1}} \quad \Leftrightarrow \quad -\mu\underline{\tau}^{-1}e^{-\mu\underline{\tau}^{-1}} = -\frac{\mu}{e\bar{\theta}}.$$

Remember that the lower branch of the Lambert W function, W_{-1}, is defined on the interval $(-1/e, 0)$ and describes the corresponding part of the inverse of the function$z$, defined by $z(W) = We^W$. Using this function we can rewrite the above equation as follows

$$z\left(-\frac{\mu}{\underline{\tau}}\right) = -\frac{\mu}{e\bar{\theta}} \qquad \Leftrightarrow \qquad W_{-1}\left(-\frac{\mu}{e\bar{\theta}}\right) = -\frac{\mu}{\underline{\tau}}$$

which is easily solved for $\underline{\tau}$

$$\underline{\tau} = -\frac{\mu}{W_{-1}\left(-\frac{\mu}{e\bar{\theta}}\right)}.$$

□

Proof of Proposition 5. U^* coincides with $\underline{\tau}$. Taking the partial derivative of $\underline{\tau}$ with respect to μ yields

$$\frac{\partial\underline{\tau}(\mu, \bar{\theta})}{\partial\mu} = \left[-W_{-1}(-\mu/e\bar{\theta}) - \frac{\mu}{e\bar{\theta}}W'_{-1}(-\mu/e\bar{\theta})\right] / \left[W_{-1}(-\mu/e\bar{\theta})\right]^2.$$

From the implicit definition of the function W_1 through the equation

$$z = W_{-1}(z)e^{W_{-1}(z)},$$

for $z \in [-1/e, 0)$, it immediately follows that

$$W'_{-1}(z) = \frac{W_{-1}(z)}{z[1 + W_{-1}(z)]}.$$

Since W_1 is strictly smaller than -1 on its entire domain we can conclude that its derivative, W'_{-1}, must be strictly negative everywhere. This in turn means that $\partial\underline{\tau}/\partial\mu$ is the sum of two strictly positive expressions und thus must be strictly positive.

The partial derivative of $\underline{\tau}$ with respect to $\bar{\theta}$ is

$$\frac{\partial\underline{\tau}(\mu, \bar{\theta})}{\partial\bar{\theta}} = \left[\frac{\mu^2}{e\bar{\theta}^2}W'_{-1}(-\mu/e\bar{\theta})\right] / \left[W_{-1}(-\mu/e\bar{\theta})\right]^2.$$

Thus this partial derivative shares the sign of W'_{-1}, which means that it is negative. □

Proof of Proposition 6. Part i) is a rather immediate implication of Proposition 5. Both distribution functions are everywhere decreasing in $\underline{\tau}$ (except at the upper limit $\tau = \bar{\theta}$). Since $\underline{\tau}$ is an increasing function of $\underline{\mu}$ the result follows.

As for part ii) notice that

$$\frac{\partial H^*(\tau; \underline{\mu}, \bar{\theta})}{\partial \bar{\theta}} = -\frac{\underline{\tau}_{\bar{\theta}}}{\underline{\tau}}\left[\ln(\bar{\theta}) - \ln(\tau)\right] - \frac{1}{\bar{\theta}}\left[\ln(\tau) - \ln(\underline{\tau})\right],$$

where $\underline{\tau}_{\bar{\theta}}$ indicates the partial derivative of $\underline{\tau}$ with respect to $\bar{\theta}$. Both terms of the difference are continuous in τ. The first one is strictly decreasing in τ and the second one is strictly increasing in τ. Finally, for $\tau = \underline{\tau}$ the first term is strictly positive while the second one is equal to zero. The converse is true for $\tau = \bar{\theta}$. Thus, there must exist exactly one point, $\tilde{\tau}$ were the derivative is equal to zero. For smaller values of τ the derivative must be strictly positive and for larger values it is strictly negative.

□

Proof of Proposition 7. The relative loss ρ is given by

$$\rho(\underline{\mu}, 1) = \frac{\underline{\tau} - \left(1 - \sqrt{1 - \underline{\mu}}\right)^2}{\underline{\tau}}.$$

Using $\underline{\tau} = -\underline{\mu}/W_{-1}(-\underline{\mu}/e\bar{\theta})$ we can rewrite ρ as follows

$$\rho(\underline{\mu}, 1) = 1 + \frac{W_{-1}(-\underline{\mu}/e)\left(1 - \sqrt{1 - \underline{\mu}}\right)^2}{\underline{\mu}}.$$

Since $W_{-1}(-1/e) = -1$ it immediately follows that $\rho(1, 1) = 0$.

As for the behavior of $\rho(\underline{\mu}, 1)$ when $\underline{\mu}$ goes to 0 notice that we can write

$$\rho(\underline{\mu}, 1) = 1 + \frac{\alpha(\underline{\mu})}{\beta(\underline{\mu})}$$

where

$$\alpha(\underline{\mu}) = \frac{\left(1 - \sqrt{1 - \underline{\mu}}\right)^2}{\underline{\mu}} \qquad \text{and} \qquad \beta(\underline{\mu}) = \frac{1}{W_{-1}(-\underline{\mu}/e)}.$$

Using L'Hopital's Rule it can be shown that $\alpha(\underline{\mu}) \to 0$ as $\underline{\mu}$ converges to 0. Moreover, since $\lim_{x \uparrow 0} W_{-1}(x) = -\infty$ it follows that also $\beta(\underline{\mu})$ tends to zero when $\underline{\mu}$ goes to zero. To compute the limit of the ratio of α and β we therefore apply L'Hopital's Rule once more. Doing so yields

$$\lim_{\underline{\mu}\to 0} \frac{\alpha(\underline{\mu})}{\beta(\underline{\mu})} = \lim_{\underline{\mu}\to 0} \frac{\alpha'(\underline{\mu})}{\beta'(\underline{\mu})} = \frac{1/4}{-\infty} = 0.$$

In order to see that we are entitled to apply l'Hopital's rule notice that both α and β are differentiable at all points $\underline{\mu} > 0$. That β is differentiable at all $\underline{\mu}$ follows from the fact that the (lower branch of the) Lambert W function is differentiable on $(-1/e, 0)$. Moreover, it is also straightforward to verify that $\beta'(\underline{\mu}) \neq 0$ for all $\underline{\mu} > 0$.

It remains to be shown that $\rho(\cdot, 1)$ is decreasing . Computing the derivative of $\rho(\cdot, 1)$ yields

$$\frac{\partial \rho(\underline{\mu}, 1)}{\partial \underline{\mu}} = \frac{W_{-1}(-\underline{\mu}/e)\left[1 - \sqrt{1-\underline{\mu}}\right]\left[W_{-1}(-\underline{\mu}/e)\left(1 - \sqrt{1-\underline{\mu}}\right) + \underline{\mu}\right]}{\left[1 + W_{-1}(-\underline{\mu}/e)\right]\underline{\mu}^2\sqrt{1-\underline{\mu}}}.$$

Since $W_{-1} \leq -1$ this expression is negative if and only if

$$-W_{-1}(-\underline{\mu}/e) > \frac{\underline{\mu}}{\left(1 - \sqrt{1-\underline{\mu}}\right)}. \tag{17}$$

Both sides of this inequality are decreasing functions of $\underline{\mu}$ and both go to 1 as $\underline{\mu}$ goes to 1. Moreover, while the rhs goes to 2 as $\underline{\mu}$ approaches 0, the lhs diverges to ∞. Thus, the above inequality holds at least in a neighborhood of 0. Suppose that contrary to our hypothesis there is some $\underline{\mu}' \in (0, 1)$ at which the two sides of the inequality cross (are equal), i.e.

$$-W_{-1}(-\underline{\mu}'/e) = \frac{\underline{\mu}'}{\left(1 - \sqrt{1-\underline{\mu}'}\right)}.$$

At $\underline{\mu}'$ the slope of the lhs is

$$\frac{d(-W_{-1}(-\underline{\mu}'/e))}{d\underline{\mu}} = -\frac{W_{-1}(-\underline{\mu}'/e)}{\underline{\mu}'\left(1 + W_{-1}(-\underline{\mu}'/e)\right)} = -\frac{1}{\sqrt{1-\underline{\mu}'}\left(1 - \sqrt{1-\underline{\mu}'}\right)}.$$

The slope of the rhs is instead

$$\frac{-2 + \underline{\mu}' + 2\sqrt{1-\underline{\mu}'}}{2\left(1 - \sqrt{1-\underline{\mu}'}\right)^2\sqrt{1-\underline{\mu}'}}$$

Dividing this expression by the slope of the lhs we obtain

$$1 - \frac{\underline{\mu}'}{2\left(1 - \sqrt{1 - \underline{\mu}'}\right)} < 1.$$

Thus, at every point of intersection of the two sides of (17) the lhs is decreasing faster than the rhs. Consequently, there can be at most one point of intersection. Since the two curves meet at $\underline{\mu} = 1$ it follows that there can be no intersection point in $(0, 1)$. □

Proof of Lemma 4. Solving (13) for $\bar{\tau}$ yields

$$\bar{\tau} = \frac{\mu^2 + \bar{\sigma}^2}{2\underline{\tau}} - \frac{\underline{\tau}}{2}.$$

Plugging this expression back into (12) delivers

$$\underline{\tau}\left[1 + \ln\left(\frac{\mu^2 + \bar{\sigma}^2}{2\underline{\tau}^2} + \frac{1}{2}\right)\right] = \mu.$$

Write $f(\underline{\tau})$ for the expression on the lhs of the above equation. Since for $\underline{\tau} = \mu$ the expression under the logarithm is larger than 1 it follows that $f(\mu) > \mu$. On the other hand, if $\underline{\tau}$ tends to zero then $f(\underline{\tau})$ converges to the limit of $\underline{\tau}\ln(1/\underline{\tau}^2)$. Using L'Hopital's rule it can be shown that the latter expression goes to 0 as $\underline{\tau}$ converges to 0. Since f is a continuous function of $\underline{\tau}$ we can conclude that there must be some admissible value of $\underline{\tau}$ such that $f(\underline{\tau}) = \mu$.

As for uniqueness of the solution it is more convenient to first solve (12) for $\underline{\tau}$ and then subsitute it into (13). We know from the preceding section that (12) has a unique solution in $\underline{\tau}$ and that this solution is given by $\underline{\tau}(\bar{\tau}, \mu) = -\mu/W_{-1}(-\mu/\bar{\tau}e)$.

Substituting this solution into (13) we obtain

$$\underline{\tau}(\bar{\tau}, \mu)[\bar{\tau} - \underline{\tau}(\bar{\tau}, \mu)] = \bar{\sigma}^2 + \mu^2. \tag{18}$$

If the lhs of this equation is strictly increasing in $\bar{\tau}$ then the equation can have at most one solution. The derivative of the lhs of (18) with respect to $\bar{\tau}$ is positive if and only if

$$\frac{\underline{\tau}'\bar{\tau}}{\underline{\tau}} - \underline{\tau}' \geq -1.$$

where $\underline{\tau}'$ denotes the partial derivative of $\underline{\tau}(\bar{\tau}, \mu)$ with respect to $\bar{\tau}$. Using the definition of $\underline{\tau}(\bar{\tau}, \mu)$

this inequality can be rewritten as

$$-2\frac{\bar{\tau}}{\mu}W_{-1}(-\mu/\bar{\tau}e) \leq \frac{\bar{\tau}}{\mu}[W_{-1}(-\mu/\bar{\tau}e)]^2 + 1.$$

Since $W_{-1}(-1/e) = -1$, it follows that this inequality is satisfied with equality at the lowest possible value for $\bar{\tau}$, $\bar{\tau} = \mu$. For $\bar{\tau} > \mu$, $W_{-1}(-\mu/\bar{\tau}e) < -1$. Thus, both sides of the inequality increase with $\bar{\tau}$. But due to the quadratic expression on the rhs, that side grows faster than the lhs, therefore we can conclude that for every μ, $\underline{\tau}(\cdot, \mu)$ is strictly increasing. □

Proof of Lemma 5. Suppose nature chooses a distribution F. The seller's payoff is then

$$\begin{aligned}\int_{\underline{\tau}}^{\bar{\tau}} x(1-F(x))h(x)dx &= \int_{\underline{\tau}}^{\bar{\tau}} (a+bx)(1-F(x))dx \\ &= \int_0^{\infty} (a+bx)(1-F(x))dx - \int_0^{\underline{\tau}} (a+bx)(1-F(x))dx - \int_{\bar{\tau}}^{\infty} (a+bx)(1-F(x))dx \\ &= a\mu + 0.5b\sigma^2 + 0.5b\mu^2 - \int_0^{\underline{\tau}} (a+bx)(1-F(x))dx - \int_{\bar{\tau}}^{\infty} (a+bx)(1-F(x))dx,\end{aligned}$$

where μ is the mean and σ^2 the variance of distribution F, and the last equality follows from

$$\int_0^{\infty} x^k(1-F(x))dx = \frac{1}{k+1}\int_0^{\infty} x^{k+1}dF,$$

which implies $\int_0^{\infty}(1-F(x))dx = \mu$ and $\int_0^{\infty} x(1-F(x))dx = (\sigma^2+\mu^2)/2$.

In the next couple of lines we focus on the last two terms of the seller's payoff; integrals $\int_0^{\underline{\tau}}(a+bx)(1-F(x))dx$ and $\int_{\bar{\tau}}^{\infty}(a+bx)(1-F(x))dx$. Assumption $a/\bar{\tau} + b = 0$ together with the requirement that h is a density (thus non-negative on $[\underline{\tau}, \bar{\tau}]$) implies $a > 0$, and therefore, $a + b\tau < 0$ for $\tau > \bar{\tau}$, and $a + b\tau > 0$ for $\tau \leq \underline{\tau}$. Nature, who wishes to minimize the seller's payoff, therefore optimally sets $F(\tau) = 0$ for $\tau \leq \underline{\tau}$ and $F(\tau) = 1$ for $\tau \leq \bar{\tau}$. The above calculation implies that for every distribution F that satisfies the mean and the variance restriction, nature prefers a distribution $\hat{F}$ with the same mean and variance and support within the interval $[\underline{\tau}, \bar{\tau}]$. It is easy to see that such an $\hat{F}$ always exist.

Next we focus on distributions whose support is within $[\underline{\tau}, \bar{\tau}]$. Let F be some such distribution. The seller's payoff then simplifies to

$$a\mu + 0.5b\mu^2 + 0.5b\sigma^2.$$

Since a and b are such that $a + b\tau > 0$ for $\tau \in [\underline{\tau}, \bar{\tau})$, the seller's payoff is increasing in the mean on $[\underline{\tau}, \bar{\tau})$, implying that nature wants to decrease the mean as much as possible. On the other hand, $b < 0$ implies that nature wants to increase the variance until it hits the upper bound

$\bar{\sigma}^2$. □

Proof of Lemma 6. If for some $\tau \in [\underline{\tau}, \bar{\tau}]$, F_τ gives nature a strictly higher payoff against $\hat{H}$ than F^*, then F^* is not a best response to $\hat{H}$. Thus if F^* is to form a saddle with $\hat{H}$, then it must be the case that $U(\hat{H}, F_\tau) \geq U(\hat{H}, F^*)$ for every $\tau \in [\underline{\tau}, \bar{\tau}]$.

Next we argue that no F_τ can give nature a strictly smaller payoff than F^*. Suppose, to the contrary, that $U(\hat{H}, F_{\hat{\tau}}) > U(\hat{H}, F^*)$ for some $\hat{\tau} \in [\underline{\tau}, \tau_1)$ (the two cases with $\hat{\tau} \geq \tau_1$ are very similar and thus omitted). We claim that this implies that there is a (small) $\bar{\varepsilon} > 0$, such that for every $\varepsilon < \bar{\varepsilon}$ the ternary distribution with support $\{\hat{\tau} + \varepsilon, \mu + \varepsilon, \bar{\tau}\}$, mean μ and variance $\bar{\sigma}^2$, generates a strictly smaller payoff against $\hat{H}$ than F^*. With a slight abuse of notation we write F_ε for this type of ternary distribution. Notice that

$$U(\hat{H}, F_\varepsilon) = p_\varepsilon(\hat{\tau} + \varepsilon) \int_{\underline{\tau}}^{\hat{\tau}+\varepsilon} x d\hat{H}(x) + p_\varepsilon(\mu + \varepsilon) \int_{\underline{\tau}}^{\mu+\varepsilon} x d\hat{H}(x) + p_\varepsilon(\bar{\tau}) \int_{\underline{\tau}}^{\bar{\tau}} x d\hat{H}(x).$$

Thus it follows that for every $0 < \varepsilon < \bar{\varepsilon}$,

$$\begin{aligned} U(\hat{H}, F_0) - U(\hat{H}, F_\varepsilon) = {} & (p_0(\hat{\tau}) - p_\varepsilon(\hat{\tau} + \varepsilon)) \int_{\underline{\tau}}^{\hat{\tau}} x d\hat{H}(x) - p_\varepsilon(\hat{\tau} + \varepsilon) \int_{\hat{\tau}}^{\hat{\tau}+\varepsilon} x d\hat{H}(x) \\ & + (p_0(\mu) - p_\varepsilon(\mu + \varepsilon)) \int_{\underline{\tau}}^{\mu} x d\hat{H}(x) - p_\varepsilon(\mu + \varepsilon) \int_{\mu}^{\mu+\varepsilon} x d\hat{H}(x) \\ & + (p_0(\bar{\tau}) - p_\varepsilon(\bar{\tau})) \int_{\underline{\tau}}^{\bar{\tau}} x d\hat{H}(x). \end{aligned}$$

When ε converges to 0 from above the first terms in all three lines of the above expression vanish, while the other two do not if $\hat{H}$ has atoms at $\hat{\tau}$ and μ or around these points. Therefore, either $U(\hat{H}, F_\varepsilon)$ converges to $U(\hat{H}, F_0)$ (when $\hat{H}$ is atom free around $\hat{\tau}$ and μ) or, $U(\hat{H}, F_0) < U(\hat{H}, F_\varepsilon)$ for every $\varepsilon > 0$ that is sufficiently close to 0; notice that in this latter case, F_ε gives nature a strictly lower payoff against $\hat{H}$ than F^*. Thus, in either of the two cases there exists an $\bar{\varepsilon} > 0$ such that nature obtains a smaller payoff from F_ε for $\varepsilon \in [0, \bar{\varepsilon})$ than it does from F^*.

The next step of our argument shows that F^* can be written as a convex combination of two distribution functions, F_α and G, where the latter is defined as a mixture of F_ε-distributions with $\varepsilon \in [0, \bar{\varepsilon}]$. Thus, by construction G is a type distribution that yields to nature a strictly lower payoff when played against $\hat{H}$ than F^* does. But then F_α must be a distribution that does better against $\hat{H}$ than F^*, meaning that $(\hat{H}, F^*)$ cannot be a saddle of U.

For each $t \in [0, \bar{\tau}]$ define $G(t)$ as uniform mixture over all F_ε distributions. That is, set

$$G(t) = \frac{1}{\bar{\varepsilon}} \int_0^{\bar{\varepsilon}} F_\varepsilon(t) d\varepsilon,$$

for all $t \in [0, \bar{\tau}]$. It is easy to verify that G is a distribution on $[0, \bar{\tau}]$ with mean μ and variance $\bar{\sigma}^2$. Clearly, G is constant over the intervals $[0, \hat{\tau})$, $[\hat{\tau} + \bar{\varepsilon}, \mu)$ and $[\mu + \bar{\varepsilon}, \bar{\tau})$. Over the intervals $[\hat{\tau}, \hat{\tau} + \bar{\varepsilon})$ and $[\mu, \mu + \bar{\varepsilon})$, G is strictly increasing. Notice that in the first case we can write

$$G(t) = \frac{1}{\bar{\varepsilon}} \int_{\hat{\tau}}^{t} p_{\tau - \hat{\tau}}(\tau) d\tau,$$

while in the second case we have

$$G(t) = G(\hat{\tau} + \bar{\varepsilon}) + \frac{1}{\bar{\varepsilon}} \int_{\mu}^{t} p_{\tau - \mu}(\tau) d\tau.$$

In either of the two cases G is differentiable in the interiors of the respective intervals and the slope of G is proportional either to $p_{t-\hat{\tau}}(t)$ or $p_{t-\mu}(t)$. Since both these expressions are probabilities this means that $G'(t)$ is bounded above. Notice that the same conclusion can be drawn regarding the left and right hand derivatives of G at the endpoints of the two intervals (where G is not differentiable).

Next define

$$F_\alpha = \frac{1}{1-\alpha}[F^* - \alpha G].$$

We argue that for sufficiently small α, F_α is a distribution. Clearly, $F_\alpha(0) = 0$ and $F_\alpha(\bar{\theta}) = 1$. Since both F^* and G are right continuous so must be F_α. It thus only remains to show that F_α is nondecreasing for sufficiently small (but positive) α. Towards that, notice that

$$F'_\alpha(t) = \frac{1}{1-\alpha}[F^{*\prime}(t) - \alpha G'(t)],$$

at all the points of differentiability of G and F^*. Furthermore, $F^{*\prime}(t) = \underline{\tau} t^{-1}$ is bounded away from 0. Since $G'(t)$ is bounded above it follows that for sufficiently small α, $F'_\alpha(t)$ is positive at all points where it is defined. The only points where F'_α is not well defined are the end points of the intervals over which G is increasing and the point $\bar{\tau}$ where G exhibits a jump. Since F^* also exhibits a jump at $\bar{\tau}$ it is easy to see that $F^* - \alpha G$ must jump upward at $\bar{\tau}$ if α is small enough. As for the endpoints of the intervals where G is increasing, we can apply the same kind of argument as in the case of the points where G is differentiable, using left and right hand derivatives.

Now fix some (sufficiently small but strictly positive) α such that F_α is a distribution function and rewrite the defining equation of F_α as follows

$$F^* = (1-\alpha)F_\alpha + \alpha G.$$

Since both G and F^* have mean μ and variance $\bar{\sigma}^2$ this equation implies that F_α must also have mean μ and variance $\bar{\sigma}^2$. Therefore, F_α represents a viable strategy for nature. Since playing G against $\hat{H}$ yields for nature a strictly smaller payoff than playing F^*, it must be the case that F_α yields a strictly higher payoff against $\hat{H}$ than F^* does. But then F^* cannot be the best reply to $\hat{H}$ which contradicts our starting assumption. □

Proof of Lemma 7. We start with a more detailed description of the relevant ternary distributions. The ternary τ-distribution, F_τ, satisfies the following conditions: i) the support of F_τ is composed of τ itself and two elements from the set $\{\underline{\tau}, \mu, \bar{\tau}\}$, ii) the mean of F_τ is μ and the variance of F_τ is $\bar{\sigma}^2$.

More specifically, the support of F_τ is $\{\tau, \mu, \bar{\tau}\}$ if $\tau \leq \tau_1 = \mu - \sigma^2/(\bar{\tau} - \mu)$ and $\{\tau, \mu, \underline{\tau}\}$ if $\tau \geq \tau_2 = \mu + \sigma^2/(\mu - \underline{\tau})$; for all remaining $\tau \in [\underline{\tau}, \bar{\tau}]$ the support is $\{\underline{\tau}, \tau, \bar{\tau}\}$. The probability weights of the elements of the support are

$$p_\tau(\tau) = \begin{cases} \sigma^2/[(\bar{\tau}-\tau)(\mu-\tau)] & \text{if } \tau \in [\underline{\tau}, \tau_1] \\ \left[\sigma^2 + (\mu-\underline{\tau})(\bar{\tau}-\mu)\right] / \left[(\bar{\tau}-\tau)(\tau-\underline{\tau})\right] & \text{if } \tau \in (\tau_1, \tau_2) \\ \sigma^2/[(\tau-\underline{\tau})(\tau-\mu)] & \text{if } \tau \in [\tau_2, \bar{\tau}], \end{cases}$$

$$p_\tau(\underline{\tau}) = \begin{cases} 0 & \text{if } \tau \in (\underline{\tau}, \tau_1] \\ \left[\sigma^2 - (\bar{\tau}-\mu)(\mu-\tau)\right] / \left[(\bar{\tau}-\underline{\tau})(\tau-\underline{\tau})\right] & \text{if } \tau \in (\tau_1, \tau_2) \\ \sigma^2/\left[(\mu-\underline{\tau})(\tau-\underline{\tau})\right] & \text{if } \tau \in [\tau_2, \bar{\tau}], \end{cases}$$

$$p_\tau(\mu) = \begin{cases} 1 - \sigma^2/[(\bar{\tau}-\mu)(\mu-\tau)] & \text{if } \tau \in [\underline{\tau}, \tau_1] \\ 0 & \text{if } \tau \in (\tau_1, \tau_2), \tau \neq \mu \\ 1 - \sigma^2/[(\mu-\underline{\tau})(\tau-\mu)] & \text{if } \tau \in [\tau_2, \bar{\tau}], \end{cases}$$

$$p_\tau(\bar{\tau}) = \begin{cases} \sigma^2/[(\bar{\tau}-\mu)(\bar{\tau}-\tau)] & \text{if } \tau \in [\underline{\tau}, \tau_1] \\ \left[\sigma^2 + (\mu-\underline{\tau})(\mu-\tau)\right] / \left[(\bar{\tau}-\tau)(\bar{\tau}-\underline{\tau})\right] & \text{if } \tau \in (\tau_1, \tau_2) \\ 0 & \text{if } \tau \in [\tau_2, \bar{\tau}). \end{cases}$$

In the following figure we represent the probability weights for the parameter combination $\underline{\tau} = 0.1, \mu = 1, \bar{\tau} = 2, \sigma^2 = 0.25$.

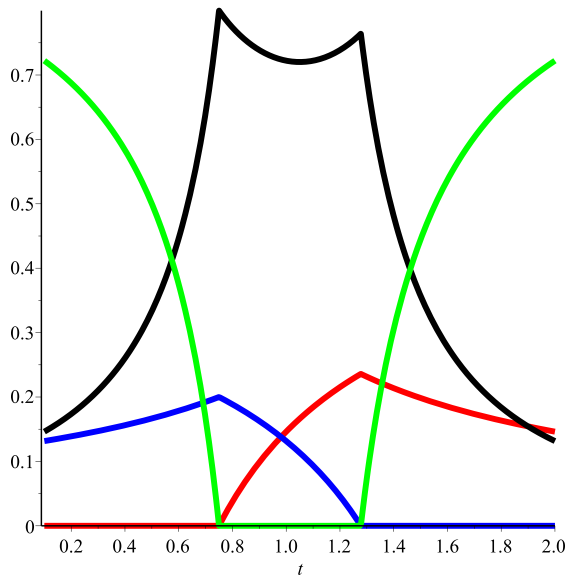

Figure 5: $p_\tau(\tau)$ (black), $p_\tau(\underline{\tau})$ (red), $p_\tau(\mu)$ (green) and $p_\tau(\bar{\tau})$ (blue)

Notice that the probabilities are differentiable at all points except τ_1 and τ_2; at those two points they exhibit kinks but are still continuous.

By Lemma 6 we know that each F_τ produces the saddle value when played against an optimal price distribution $\hat{H}$. That is, $U(\hat{H}, F_\tau)$ is equal to U^* for all τ, or more explicitly,

$$p_\tau(\tau) \int_{\underline{\tau}}^{\tau} td\hat{H} + p_\tau(\mu) \int_{\underline{\tau}}^{\mu} td\hat{H} + p_\tau(\bar{\tau}) \int_{\underline{\tau}}^{\bar{\tau}} td\hat{H} = U^*, \tag{19}$$

for $\tau \in [\underline{\tau}, \tau_1]$; the other two cases are similar. Being constant in τ the sum on the lhs of (19) is twice differentiable with respect to τ. Since the second and the third term of this sum are themselves differentiable with respect to τ (except at the points τ_1 and τ_2), so must be the first term. Since the first term is the product of two functions of τ, where the first one, $p_\tau(\tau)$, is again differentiable (except at the two points τ_1 and τ_2), this means that also the second one, $\int_{\underline{\tau}}^{\tau} tdH$, must be differentiable with respect to τ.

Since the optimal price distribution $\hat{H}$ is continuous,[19] we can use integration by parts to write

$$\int_{\underline{\tau}}^{\tau} td\hat{H} = \tau\hat{H}(\tau) - \underline{\tau}\hat{H}(\underline{\tau}) - \int_{\underline{\tau}}^{\tau} \hat{H}(t)dt. \tag{20}$$

This expression as a whole is differentiable with respect to τ, and since the second and the third

[19]Suppose it was not, and suppose that there is a discontinuity at some $\tilde{\tau}$ in $[\underline{\tau}, \tau_1)$ (the other cases are handled similarly). Then a ternary distribution $F_{\tau'}$ for a τ' close enough but above $\tilde{\tau}$ would yield strictly smaller payoff to nature than $\tilde{\tau}$. That would contradict Lemma 6.

term are each individually differentiable it follows that so must be the first one. The derivative of (20) is $\tau H'(t)$, or in our notation, $\tau h(t)$.

Using (19) we can write

$$\int_{\underline{\tau}}^{\tau} t d\hat{H} = g(\tau),$$

where

$$g(\tau) = \frac{1}{p_\tau(\tau)} U^* - \frac{p_\tau(\mu)}{p_\tau(\tau)} \int_{\underline{\tau}}^{\mu} t d\hat{H} - \frac{p_\tau(\bar{\tau})}{p_\tau(\tau)} \int_{\underline{\tau}}^{\bar{\tau}} t d\hat{H}.$$

Inspection of the function g reveals that it depends on τ only through the probabilities. Since those are twice differentiable wrt τ (with exception of the points τ_1 and τ_2) and since by (20) we know that the first derivative of g is equal to $\tau h(\tau)$, it follows that $h(\tau)$ must be differentiable except possibly at $\tau = \tau_1$ and at $\tau = \tau_2$.

A similar argument can be used to show that h' is differentiable. □

Proof of Lemma 8. Lemma 6 implies that if $(\hat{H}, F^*)$ is a saddle, then nature is indifferent between all the ternary F_τ distributions. In particular, for $\tau \geq \tau_2$, the seller's payoff from any such ternary distribution is

$$p_\mu(\tau) \int_{\underline{\tau}}^{\mu} t d\hat{H} + p_\tau(\tau) \int_{\underline{\tau}}^{\tau} t d\hat{H} = U^*, \tag{21}$$

the other two cases are handled in a similar fashion. By Lemma 7, $\hat{H}$ is (at least) three times differentiable. Differentiating (21) with respect to τ yields

$$p'_\mu(\tau) \int_{\underline{\tau}}^{\mu} t d\hat{H} + p'_\tau(\tau) \int_{\underline{\tau}}^{\tau} t d\hat{H} + p_\tau(\tau)\tau h(\tau) = 0. \tag{22}$$

for $\tau \in (\tau_2, \bar{\tau})$. Dividing through by $p'_\tau(\tau)$ and taking once more the derivative with respect to τ we obtain an expression of the type

$$A(\tau) \int_{\underline{\tau}}^{\mu} t d\hat{H} + B(\tau)\tau h(\tau) + C(\tau) \left[h(\tau) + \tau h'(\tau)\right] = 0,$$

for all $\tau \in (\tau_2, \bar{\tau})$. Cumbersome but straightforward calculations show that $B(\tau)/A(\tau) = -(\mu - \underline{\tau})$ and $C(\tau)/A(\tau) = (2\tau - \mu - \underline{\tau})(\mu - \underline{\tau})$. Thus if we divide the above equation by $A(\tau)$ and differentiate again we obtain

$$-(\mu-\underline{\tau})[h(\tau)+\tau h'(\tau)]+(\mu-\underline{\tau})[h(\tau)+\tau h'(\tau)]+[2\tau-\mu-\underline{\tau}](\mu-\underline{\tau})[2h'(\tau)+\tau h''(\tau)]=0,$$

which simplifies to

$$2h'(\tau)+\tau h''(\tau)=0,$$

for all $\tau \in (\tau_2, \bar{\tau})$. Solving this differential equation yields

$$h(\tau)=\frac{a}{\tau}+b,$$

for all $\tau \in (\tau_2, \bar{\tau})$. □

Proof of Lemma 9. We argue that $h(\bar{\tau})$ must be equal to 0. This, together with the requirement that h integrates to 1, pins down a and b.

Suppose that $(\hat{H}, F^*)$ was a saddle with $h(\bar{\tau}) = a/\bar{\tau} + b > 0$. Then for a small enough $\epsilon > 0$, $a/\tau + b > 0$ for all $\tau \in (\bar{\tau}, \bar{\tau} + \epsilon)$. In the proof of Lemma 5 we show that the seller's payoff when using a distribution with a density of the the form $h(\tau) = a/\tau + b$ on $[\underline{\tau}, \bar{\tau}]$, and the valuations are distributed with distribution F, can be written as

$$a\mu + 0.5b\sigma^2 + 0.5b\mu^2 - \int_0^{\underline{\tau}} (a+bx)(1-F(x))dx - \int_{\bar{\tau}}^{\infty} (a+bx)(1-F(x))dx.$$

Since $a + bx > 0$ for $x \in [\bar{\tau}, \bar{\tau} + \epsilon]$, nature would then strictly benefit from deviating to a distribution $\tilde{F}$ that has the same mean and variance as F^* but also assigns some probability mass to the interval $[\bar{\tau}, \bar{\tau} + \epsilon]$. Such a deviation can be found as long as F^* does not have variance 0, which cannot happen in a saddle anyways. Therefore $h(\bar{\tau}) > 0$ cannot be the case in a saddle. □

Proof of Proposition 11. Since the optimal randomization over prices has the same support in both cases, $\bar{\tau} = \bar{\theta}$. The optimal distribution for the seller when he knows the mean is μ and the variance bounded above by $\bar{\sigma}^2$ can be written as

$$H^*_{\bar{\sigma}^2}(\tau) = \frac{\bar{\theta}[\ln \tau - \ln \underline{\tau}] - (\tau - \underline{\tau})}{\bar{\theta}[\ln \bar{\theta} - \ln \underline{\tau}] - (\bar{\theta} - \underline{\tau})},$$

and the optimal distribution in the case when he knows that the valuations are distributed with a distribution with mean μ and support in $[0, \bar{\theta}]$

$$H^*_{\bar{\theta}}(\tau) = \frac{\ln \tau - \ln \underline{\tau}}{\ln \bar{\theta} - \ln \underline{\tau}}.$$

The difference $H^*_{\bar{\sigma}^2}(\tau) - H^*_{\bar{\theta}}(\tau)$ has the same sign as

$$(\ln \tau - \ln \underline{\tau})(\bar{\theta} - \underline{\tau}) - (\ln \bar{\theta} - \ln \underline{\tau})(\tau - \underline{\tau}),$$

which in turn has the same sign as

$$\frac{\ln \tau - \ln \underline{\tau}}{\tau - \underline{\tau}} - \frac{\ln \bar{\theta} - \ln \underline{\tau}}{\bar{\theta} - \underline{\tau}}.$$

The last expression is positive if $(\ln \tau - \ln \underline{\tau})/(\tau - \underline{\tau})$ is decreasing in τ. Differentiating the expression yields

$$\begin{aligned} \frac{d}{d\tau}\left(\frac{\ln \tau - \ln \underline{\tau}}{\tau - \underline{\tau}}\right) &= \frac{1 - \frac{\underline{\tau}}{\tau} + \ln \frac{\underline{\tau}}{\tau}}{(\tau - \underline{\tau})^2} \\ &\leq 0, \end{aligned}$$

for $\tau \geq \underline{\tau}$ and the inequality is strict for $\tau > \underline{\tau}$. The inequality follows from $\ln x \leq x - 1$, with strict inequality for $x \neq 1$. □

8 Appendix B

In this appendix we report the basic saddle point result on which parts of our analysis rely. For a more detailed discussion see Rockafellar (1970).

We start by introducing some notation. Let X and Y be two arbitrary sets and let K be a function that maps the Cartesian product $X \times Y$ into the extended real numbers, i.e.

$$K : X \times Y \to [-\infty, +\infty].$$

Define

$$f(x) = \sup_{y \in Y} K(x, y) \qquad \text{and} \qquad g(y) = \inf_{x \in X} K(x, y).$$

Proposition 12. *A pair* $(\bar{x}, \bar{y})$ *satisfies the saddle point condition*

$$K(x, \bar{y}) \geq K(\bar{x}, \bar{y}) \geq K(\bar{x}, y) \qquad \text{for all } x \in X, y \in Y,$$

if and only if $\bar{x}$ *solves*

$$\bar{x} \in \arg\min_{x \in X} f(x), \qquad \text{and} \qquad \bar{y} \in \arg\max_{y \in Y} g(y),$$

and the saddle-value of K exists, i.e. one has

$$\inf_{x \in X} f(x) = \sup_{y \in Y} g(y).$$

Proof. Suppose first that $\bar{x} \in \arg\min_{x \in X} f(x)$, $\bar{y} \in \arg\max_{y \in Y} g(y)$ and $f(\bar{x}) = g(\bar{y})$. By the definition of f and g it follows that

$$f(x) \geq K(x, y) \geq g(y) \qquad \text{for all } \ x \in X, \ y \in Y.$$

In particular, the inequalities must hold for $\bar{x}$ and $\bar{y}$, i.e.

$$f(\bar{x}) \geq K(\bar{x}, \bar{y}) \geq g(\bar{y}).$$

But since $f(\bar{x}) = g(\bar{x})$ we can conclude that

$$\sup_{y \in Y} K(\bar{x}, y) = K(\bar{x}, \bar{y}) = \inf_{x \in X} K(x, \bar{y}),$$

which implies the saddle point condition.

Conversely, suppose that the saddle point condition holds, i.e.

$$K(x, \bar{y}) \geq K(\bar{x}, \bar{y}) \geq K(\bar{x}, y) \qquad \text{for all } x \in X, y \in Y.$$

This immediately implies that

$$g(\bar{y}) = \inf_{x \in X} K(x, \bar{y}) = K(\bar{x}, \bar{y}) = \sup_{y \in Y} K(\bar{x}, y) = f(\bar{x}).$$

But since for all $x \in X$ and $y \in Y$ we have $f(x) \geq g(y)$ we also know that $f(\bar{x})$ is an upper bound of $g(y)$ and since this upper bound is reached for $y = \bar{y}$ we can conclude that $g(\bar{y}) \geq g(y)$ for all $y \in Y$. Finally, a perfectly analogous argument delivers $f(\bar{x}) \leq f(x)$ for all $x \in X$ and thus

$$\inf_{x \in X} f(x) = \sup_{y \in Y} \ g(y).$$

9 Appendix C

In this Appendix we report some basic facts regarding the Lambert W function.

The domain of the lower branch of the Lambert W function, W_{-1}, is the interval $[-1/e, 0)$. The defining equation for W_{-1} is

$$W_{-1}(x)e^{W_{-1}(x)} = x.$$

Implicit differentiation yields for all $x \in (-1/e, 0)$

$$x(1 + W_{-1}(x))\frac{dW_{-1}(x)}{dx} = W_{-1}(x)$$

or equivalently

$$\frac{dW_{-1}(x)}{dx} = \frac{W_{-1}(x)}{x(1 + W_{-1}(x))}.$$

References

Auster, S. (2013), "Bilateral trade under ambiguity." Working Paper.

Bergemann, D. and S. Morris (2005), "Robust mechanism design." *Econometrica*, 73, 1521–1534.

Bergemann, D. and S. Morris (2013), "An introduction to robust mechanism design." *Foundations and Trends in Microeconomics*, 8, 169 – 230.

Bergemann, D. and K. Schlag (2008), "Pricing without priors." *Journal of the European Economic Association*, 6, 560–569.

Bergemann, D. and K. Schlag (2011), "Robust monopoly pricing." *Journal of Economic Theory*, 146, 2527–2543.

Bewley, T. (1986), "Knightian decision theory: Part i." Mimeo.

Bodoh-Creed, Aaron (2012), "Ambiguous beliefs and mechanism design." *Games and Economic Behavior*, 75, 518–537.

Bose, Subir and Arup Daripa (2009), "A dynamic mechanism and surplus extraction under ambiguity." *Journal of Economic Theory*, 144(5), 2084–2114.

Bose, Subir, Emre Ozdenoren, and Andreas Pape (2006), "Optimal auction with ambiguity." *Theoretical Economics*, 1, 411–438.

Bose, Subir and Ludovic Renou (2014), "Mechanism design with ambiguous communication devices." *Econometrica*, 82, 1853–1872.

Carroll, G. (2012), "Robust incentives for information acquisition." Working Paper.

Carroll, Gabriel (2015), "Robustness and linear contracts." *The American Economic Review*, 105, 536–563.

Castro, L. De and N.C. Yannelis (2012), "Uncertainty, efficiency and incentive compatibility." Unpublished manuscript.

Chassang, Sylvain (2013), "Calibrated incentive contracts." *Econometrica*, 81, 1935–1971.

Di Tillio, A., N. Kos, and M. Messner (2014), "The design of ambiguous mechanisms." IGIER Working Paper.

Garrett, Daniel F (2014), "Robustness of simple menus of contracts in cost-based procurement." *Games and Economic Behavior*, 87, 631–641.

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin expected utility with a non-unique prior." *Journal of Mathematical Economics*, 18, 141–153.

Kos, Nenad and Matthias Messner (2013), "Extremal incentive compatible transfers." *Journal of Economic Theory*, 148, 134–164.

López-Cunat, Javier M (2000), "Adverse selection under ignorance." *Economic Theory*, 16, 379–399.

Lopomo, G., L. Rigotti, and C. Shannon (2009), "Uncertainty in mechanism design." Working paper.

Lopomo, G., L. Rigotti, and C. Shannon (2011), "Knightian uncertainty and moral hazard." *Journal of Economic Theory*, 146, 1148–1172.

Morris, Stephen (1994), "Trade with heterogeneous prior beliefs and asymmetric information." *Econometrica: Journal of the Econometric Society*, 1327–1347.

Myerson, R.B. (1981), "Optimal auction design." *Mathematics of Operations Research*, 6, 58–73.

Riley, John and Richard Zeckhauser (1983), "Optimal selling strategies: When to haggle, when to hold firm." *The Quarterly Journal of Economics*, 267–289.

Rockafellar, R Tyrrell (1997), *Convex analysis*. 28, Princeton university press.

Skreta, V. (2006), "Mechanism design for arbitrary type spaces." *Economic Letters*, 91, 293–299.

Wilson, R. (1987), *Game Theoretic Analysis of Trading Processes*. Cambridge Univeristy Press: Cambridge.

Wolitzky, A. (2014), "Mechanism design with maxmin agents: Theory and an application to bilateral trade." Working Paper.