A dynamic mechanism and surplus extraction under ambiguity

Journal Article

Citation:

Bose, S.; Daripa, A. (2009) A dynamic mechanism and surplus extraction under ambiguit – *Journal of Economic Theory* 144(5), pp. 2084-2114

Publisher version: http://dx.doi.org/10.1016/j.jet.2009.02.003

A Dynamic Mechanism and Surplus Extraction Under Ambiguity[1]

Subir Bose

University of Leicester, Leicester, UK

Arup Daripa

Birkbeck College, London University
Bloomsbury, London, UK

Final Version: January 2009
Published in *Journal of Economic Theory*, 2009

Abstract

We study the question of auction design in an IPV setting characterized by ambiguity. We assume that the preferences of agents exhibit ambiguity aversion; in particular, they are represented by the epsilon-contamination model. We show that a simple variation of a discrete Dutch auction can extract almost all surplus. This contrasts with optimal auctions under IPV without ambiguity as well as with optimal static auctions with ambiguity - in all of these, types other than the lowest participating type obtain a positive surplus. An important point of departure is that the modified Dutch mechanism is dynamic rather than static, establishing that under ambiguity aversion–even when the setting is IPV in all other respects–a dynamic mechanism can have additional bite over its static counterparts. A further general insight is that the standard revelation principle does not automatically extend to environments not characterized by subjective expected utility.

KEYWORDS: Ambiguity Aversion, Epsilon Contamination, Modified Dutch Auction, Dynamic Mechanism, Surplus Extraction, Revelation Principle

JEL CLASSIFICATION: D44

---

[1]We thank the associate editor and two anonymous referees for their insightful comments. We also thank Ludovic Renou, the participants of Birmingham Summer Theory Workshop 08, and of RUD 08 at Oxford for questions and comments that helped clarify our exposition. We are grateful to Peter Klibanoff for many valuable and detailed comments. Finally, we thank Sujoy Mukerji. Several discussions with him in the course of the work shaped our understanding of many of the issues involved.

1 Introduction

In the standard independent private values (IPV) setting bidders draw privately known valuations from a given distribution. Each bidder is assumed to maximize subjective expected utility, so that a bidder's beliefs about the values of any other bidder is represented by a unique prior (i.e. a unique distribution over the domain of values). In this setting Dutch auctions coincide with First Price Sealed Bid auctions, and optimal auctions leave all but the lowest participating type with a surplus. This is true whether bidders are risk neutral (Myerson (1981), Riley and Samuelson (1981)) or risk averse (Matthews (1983), Maskin and Riley (1984)).

As far as we are aware, Karni (1988) is the first to show that the equivalence between Dutch and First Price Sealed Bid auctions breaks down under non-expected utility preferences.

In this paper we relax the unique prior assumption and study the question of auction design in an IPV setting characterized by ambiguity: bidders have an imprecise knowledge of the distribution of values of others, and are faced with a set of priors. We also assume that their preferences exhibit ambiguity aversion; in particular we use the epsilon contamination representation, used extensively in the economics and statistics literature.

Several papers have studied auctions (or auction-like environments) when bidders have non-expected utility preferences (e.g. Karni and Safra 1986, 1989a,1989b; Karni 1988; Lo 1998; Nakajima 2004; Ozdenoren 2002; Volij 2002). The closest intellectual antecedents appear in the paper by Bose, Ozdenoren, and Pape (2006), who show that in the setting of ambiguity that we consider, the optimal *static* mechanism leaves buyer types with information rent, and the amount of rent varies continuously with the extent of the ambiguity. In contrast, our main result shows that in this setting of ambiguity averse buyers, the seller can use a simple variation of a discrete Dutch auction and extract almost all surplus. The important point of departure is that the modified Dutch mechanism we consider is *dynamic* rather than static, establishing that a dynamic mechanism can present the seller with additional surplus extraction opportunities under ambiguity aversion even in a setting that is captured by the IPV model in all other respects. This also shows that the introduction of ambiguity can pose difficulties for mechanism design, since the standard revelation principle may not extend to this setting.

In a seminal paper, Ellsberg (1961) showed that lack of knowledge about the distribution over states, often referred to as ambiguity, can affect the choice of a decision maker in a fundamental way that cannot be captured by a framework that assumes a unique prior. Several subsequent studies have underlined the importance of ambiguity aversion in understanding decision making behavior (e.g. Camerer and Weber (1992)). and models taking such aversion into account have provided important insights in a variety of economic applications including auctions.(1)

We model ambiguity aversion using the maxmin expected utility (MMEU) model of Gilboa and Schmeidler (1989). Here, the agents have a set of priors (instead of a single prior) on the underlying state space, and the payoff from any action is the minimum expected utility over the set of priors. In our setting, each buyer considers a set of distributions that contain the distribution from which the other buyer's valuation is drawn and each action (from the mechanism proposed by the seller) is evaluated based on the minimum expected utility over the set of distributions. The buyer then chooses the best action from the set of actions. To make a minimal departure from the standard model, we assume that the seller is ambiguity neutral(2) and both the buyers and the seller are risk neutral. In other words, apart from relaxing the unique prior assumption, our framework is as close to the standard IPV model as possible.(3),(4)

(1)For example, using such preferences Mukerji (1998) explains the incompleteness of contracts and Mukerji and Tallon (2004) explain the puzzling absence of wage indexation. An application to auction theory is developed by Lo (1998), who shows that if bidders are ambiguity averse, the revenue equivalence theorem (which holds in the standard IPV setting) is violated - sealed bid first price auctions raise more revenue than sealed bid second price auctions.

(2)Assuming the seller to be ambiguity neutral allows us to focus on revenue extraction as the seller's objective. This also allows us to compare our results directly with the standard results from mechanism design in a Bayesian setting. If the seller is also ambiguity sensitive, maximum surplus extraction is not necessarily the objective while designing the mechanism. Studying such issues are interesting, but beyond the scope of the work here.

(3)With multiple priors, the terms "independent" and "correlated" need to be used carefully. For the most part we avoid using these terms. The important point is that in the standard model, even with risk-neutrality, full surplus extraction is not possible when the beliefs do not depend on one's own valuation (i.e. in the independent case). Hence it is worth emphasizing that we consider the case where the sets of probability distributions are the same for every buyer and *do not depend* on a buyer's own valuations. As shown by Bose et al. (2006), the optimal *static* mechanism does not extract full surplus in this setting.

(4)As in other applied mechanism design papers, we start at the *interim* stage where agents know their own types (valuations, beliefs etc.). Of course, one could be interested in the ex ante stage also where

As noted before, we use a version of MMEU known as "epsilon-contamination." A recent paper by Kopylov (2008) provides an axiomatization for this formulation (see also Nishimura and Ozaki (2006)). The model we consider has a seller whose valuation of the object is (normalized to) zero. There are two potential buyers and the seller does not know either buyer's valuation but believes that the valuations are determined based on independent draws from the distribution $F(v)$ having support $[0,1]$. Each buyer knows his own valuation and ambiguity regarding the valuation of the other buyer is represented using the epsilon-contamination model. The specification is in widespread use for its intuitive qualities and analytical tractability. It is used extensively in the literature on robust statistics, starting with (as far as we are aware) Huber (1973). Examples from the economics literature include Chen and Epstein (2002), Chu and Liu (2002), Mukerji (1998), Nishimura and Ozaki (2004).

Formally, let $\mathcal{P}$ denote the set of all distributions on $[0,1]$. Fix any $\varepsilon \in (0,1]$. The set of priors is then given by

$$\mathcal{P}_B(\varepsilon) = \{G \in \mathcal{P} : G = (1-\varepsilon)F + \varepsilon L \text{ for some } L \in \mathcal{P}\}$$

Intuitively, for some $\varepsilon > 0$, the buyer puts a weight of $(1-\varepsilon)$ on the other buyer's valuation being drawn from the distribution F, but puts ε weight that the valuation could be drawn from some other distribution. As Kopylov (2008) shows, the weight $(1-\varepsilon)$ can be interpreted as the agent's confidence in the subjective belief F; alternatively, the weight ε can be thought of as an index of ambiguity aversion.

Let us now describe our Modified Dutch Mechanism (MDM). The seller declares a decreasing sequence of prices $\{p_1, .., p_n\}$ at the beginning. At stage k, provided the item has not been sold up to that point, the seller randomly (with equal probability of selecting any one buyer) approaches a buyer and offers the item at price p_k. This offer is secret in the sense that the other buyer is not made aware of this. If the approached

some draw by nature determines the agent's types. Note, however, that the ex ante stage requires more careful handling here compared to the standard case. For example, if the agents believe that nature draws valuations of both buyers from the same distribution, knowledge of own value provides some information about the set of distributions from which the other's valuation is drawn. To model the ex ante stage so that the problem remains similar in spirit to the standard IPV model, one would then need an assumption such as nature drawing values from distributions which are themselves chosen according to some independent but unknown (ambiguous) process.

It can be shown however, that even if nature draws both valuations from the same distribution, our results remain unaffected. However, since full surplus extraction is possible even in the unique prior case when beliefs vary with valuations, we do not focus on such situations.

buyer passes, the seller approaches the other buyer (also in secret) and offers the item at the same price p_k. If the second buyer refuses as well the game goes to stage $k+1$. If both buyers refuse at stage n, the seller keeps the item. We assume that the randomization is independent across periods.

As in any mechanism design exercise the seller commits to the mechanism (including the price offered in each period). Assuming that the buyers are approached randomly and secretly in every period only helps to keep the mechanism symmetric. However, this latter feature serves only an aesthetic purpose; the results do not depend on the mechanism being symmetric.

Our surplus extraction result states the following. Fix a preference parameter $\varepsilon > 0$. There is a $\delta^*(\varepsilon)$ such that for any given $\delta < \delta^*(\varepsilon)$ and any $\eta > 0$, the seller can construct an MDM (i.e. choose a price sequence) such that the mass of buyer types who do not buy is at most $[0, \eta]$ (i.e. the reserve type is at most η), and the types who buy do so at a price such that their ex post surplus is at most δ. Since both δ and η can be arbitrarily small, the seller can therefore extract almost full surplus.

The basic intuition for the result is that for any v and any price p where $p < v$, the buyer gets a *sure* payoff of $v - p$ from buying at p, whereas the payoff from waiting one more period is $v - p + \Delta_p$ times the probability that the current buyer obtains the item in the next period, where Δ_p is the difference between p and the next lower price. With epsilon contamination preferences, the buyer attaches at least probability ε that the item gets sold before he has the chance to obtain it next period. Thus the loss from waiting is at least $\varepsilon(v - p)$ whereas the gain from waiting is of the order Δ_p. For any given ε, since Δ_p can be made arbitrarily small independent of the value of ε, the gain from waiting can be made arbitrarily smaller than the loss from waiting. Put differently, even though purchasing at price p results in the ex post surplus being at most δ, the price sequence is *constructed* in such a way that this is still larger than the (expected) surplus from waiting to buy at a lower price.

Note that this cannot happen in the standard (i.e. the unique prior) model. There, for any type v, as long as the seller sells to types below v with positive probability, the surplus of type v cannot be made arbitrarily small. Roughly speaking, in the absence of ambiguity, given that F is smooth, the expected gain and expected loss from waiting shrink at the same rate as the price gap becomes smaller.

The above discussion also shows the importance of the special structure of the set of priors in our case. The ε-contamination model assigns an epsilon weight to even very unlikely events - and therefore a bidder always assigns an epsilon weight to the event that he loses if he waits for the price to drop by even a very small amount. We exploit precisely this feature to extract full surplus. Essentially, the full surplus extraction result requires that the set be such that even though expected gain from waiting can be made vanishingly small by making the price gaps small, the expected loss from waiting is bounded away from zero. We discuss this issue further in section 6.

As mentioned before, Bose et al. (2006) (specifically, see section 4) study optimal auction design for the same environment that we consider in this paper. They use the revelation principle to design the optimal mechanism. In the mechanism, types earn positive information rent, and further, the rents approach those found in the unique prior case as $\varepsilon \to 0$. Our dynamic mechanism, in contrast, extracts almost all rents for an arbitrarily small ε and is thus discontinuously different from the unique prior case.

Even though the *full* surplus extraction result is special to the specific properties of the ε-contamination, a general insight arising from our work is about the applicability of the standard revelation principle to settings with ambiguity. The revelation principle states that the outcome of any equilibrium from *any* mechanism can be replicated by a truthful equilibrium of a static direct mechanism. Now, in a standard expected utility setting (where, as usual, agents use Bayes Rule to update beliefs), preferences satisfy the property of dynamic consistency (over acts) and hence ex ante optimal decisions coincide with optimal decisions conditional on some event. Note that this is necessary for the revelation principle to hold: otherwise there would be no guarantee that a direct mechanism, which is by construction static, could replicate an equilibrium outcome of a dynamic mechanism which can provide extra information to agents before they make their decision. As we discuss with the help of a simple example in section (6.1.1), in settings with ambiguity, ex ante decisions can differ from conditional ones, and in our dynamic mechanism the agents must repeatedly make decisions conditioning on new information arising in course of the mechanism. We show that in the presence of ambiguity, dynamic mechanisms may have additional power over optimal static ones even in situations where the two types of mechanisms do not produce different results in ambiguity neutral settings. This also implies that the standard revelation principle can fail in environments not characterized by a unique prior.

While failure of dynamic consistency provides the additional power to the dynamic mechanism over optimal static ones, this also implies that our theory must provide a consistent way of reconciling preferences at different time periods. Recently, Siniscalchi (2006) has provided an axiomatic foundation of such sophisticated dynamic choice for maxmin utility and full Bayesian updating and we follow this approach as the foundation for dynamic behavior in our model. Agents are sophisticated in the sense that they correctly anticipate future preferences while making current plans. Further, (conditional) preferences are defined over trees, rather than Savage acts. As discussed in section 6.1.2, this gives us a coherent theory of dynamic choice, and allows us to abstract from the issue of dynamic consistency over acts. In constructing formal propositions, we repeatedly use the idea that agents are sophisticated and therefore make "correct" contingent plans, precluding the problem of future deviations from a current plan. To be specific, our equilibrium is characterized by "cut-off" types, where a cut-off type v_k is the lowest type that buys at price p_k. The set of types is then partitioned into intervals of the form $(v_{k-1}, v_k]$ such that types in $(v_{k-1}, v_k]$ plan to buy in period k at price p_k. We require equilibrium to be *perfect*, a particular implication of which is that when a type accepts the seller's offer as its (conditional) optimal action in that period, it *correctly* anticipates its own future behavior (and hence future payoffs) were it to reject the current offer.

While in a unique prior setting the optimal auction under IPV does not extract full surplus, a different strand of the literature considers environments with correlated types where it is possible to do so (see Crémer and McLean (1988), Crémer and McLean (1985),McAfee and Reny (1992)). Note however that unlike the mechanisms in this literature, ours do not involve any extraneous lotteries and satisfies limited liability.[5] It is also important to note that our mechanism meets better the criticism of the so-called Wilson doctrine– the idea that the mechanism should not require the designer to possess detailed knowledge of the environment–and is therefore an example of more robust mechanism design. The seller here does not need to know the exact distribution F, nor the specifics of the contaminating set of distributions (beyond the fact that it has certain properties - see section 6.3). We do assume in the formal model that the seller does know the value of ε; note however that the seller is not required to know the exact value of ε as long as it is known that there is a bound such that the true value of ε lies

[5]Robert (1991) shows that the Cremer-McLean result relies crucially on risk neutrality as well as limited liability. While we do not explicitly consider risk-aversion, it is easy to show that our basic result is unchanged if bidders are risk averse.

above this bound.

A final comment on the nature of the ambiguity in our model. Here each buyer faces ambiguity with respect to the other buyer's valuation distribution. There is neither any ambiguity about the strategy of the other buyer nor does the seller introduce any ambiguity through the mechanism. However, introducing any of these other sources of ambiguity would only make it easier for us to prove the central result. To see this, note that the central idea in the construction of the MDM is to give each buyer a choice at each stage between an ambiguous alternative and a sure payoff. The mechanism then exploits the ambiguity of one of the alternatives to make the sure payoff more attractive, which helps extract surplus. Adding extra sources of ambiguity does not change the sure payoff, but does affect the ambiguous alternative, which makes surplus extraction easier.

Earlier work in the area of robust Bayesian statistics have studied dynamic inference problems facing a decision maker with maxmin preferences. The literature shows that the juxtaposition of maxmin preferences with full Bayesian updating can give rise to surprising results (e.g. Augustin (2003), Grunwald and Halpern (2004), Seidenfeld (2004)). However, as far as we are aware, this paper is the first to study the question of dynamic mechanism design under such non-EU preferences.

The rest of the paper is organized as follows. The next section presents the model. Section 3 presents our mechanism, and characterizes equilibria in the induced game. The main result of the paper appears in section 4, and section 5 presents a numerical example. Section 6 discusses some aspects of the model, and section 7 concludes.

2 The Model

2.1 The Basic Set-up

There is a seller with one indivisible object for sale. The seller's valuation of the item is (normalized to) zero. There are two potential buyers with valuations of the object lying in the interval $[0, 1]$.[(6),(7)] Own valuation is private information of each buyer. Each buyer believes that the other's valuation is drawn from some distribution from a *set* of distributions on $[0, 1]$. The preferences of the buyers is represented by the maxmin expected utility (MMEU, henceforth) model of Gilboa and Schmeidler (1989). Briefly, if Ω is a set, $\mathcal{P}$ is a set of distributions on Ω, and $\mathcal{F}$ is a set of acts from Ω to the real line $\mathbb{R}$, then an act $f \in \mathcal{F}$ is evaluated according to the rule

$$\min_{p \in \mathcal{P}} \int u(f) dp$$

where u is some real valued function. In our context, we assume buyers are risk-neutral.

The seller is (risk and) ambiguity neutral[(8)] and has a prior over a buyer's valuation given by the distribution $F(v)$ with a continuous density $f(v) > 0$. As mentioned in the introduction, we model the set of priors representing buyer's ambiguous beliefs using the epsilon contamination model.

Let $\mathcal{P}$ denote the set of all distributions on $[0, 1]$. Fix any $\varepsilon \in (0, 1]$. The set of priors is then given by $\mathcal{P}_B(\varepsilon) = \{G \in \mathcal{P} : G = (1 - \varepsilon)F + \varepsilon L \text{ for some } L \in \mathcal{P}\}$.[(9)] As noted in the introduction, Kopylov (2008) has recently provided an axiomatic foundation for preferences to be represented by the epsilon contamination model. Note also that other than non-unique priors, the rest of the model conforms as closely as possible to the IPV model standard in auction theory.

(6) We could, for the sake of generality, represent the buyer's possible valuations to be the set $[\underline{v}, \overline{v}]$. However, we do allow the seller to have a non-trivial reserve price, and, as the result below shows, the normalization to the space $[0, 1]$ is harmless, and reduces algebraic clutter.

(7) Generalization to arbitrary $N > 2$ buyers is straightforward.

(8) See footnote (2).

(9) We use the same F to represent the seller's beliefs as well as to generate $\mathcal{P}_B$ to save on notation. However, F being focal is *inessential*; any other distribution in place of F to generate the set $\mathcal{P}_B$ would work just as well.

2.2 Extension to a dynamic setting

The Gilboa-Schmeidler model is atemporal. We need to extend this static choice model to suit the specific context of our dynamic mechanism. To this end, we assume that the buyer's have maxmin - in particular, epsilon contamination - preferences at every stage, and choose actions to maximize the minimum expected payoff from a set of updated distributions. In the setting with ambiguity, updating a distribution can be tricky. We discuss this issue in detail in section 6.2.

Specifically, based on the available information, an agent updates F and the distributions in the set $\mathcal{P}$, and then chooses the action that maximizes the minimum expected value of payoff where the minimizing set of distributions is obtained by taking a convex combination of updated F (with weight $(1-\varepsilon)$) and the updated distributions in $\mathcal{P}$ (with weight ε). This procedure is in keeping with the interpretation of ε that arises from the axioms of Kopylov (2008), where the weight $(1-\varepsilon)$ can be interpreted as the decision makers degree of confidence in her subjective belief F.(10)

The updating rule we use is the "full Bayesian" (also called "prior-by-prior") rule. This has been used in economics as well as in the extensive literature on robust statistics.(11) Even though this is probably the most well known rule, other rules have been proposed in the literature. We discuss this issue in more detail in section 6.2.

Finally, we assume that the dynamic behavior of the buyers is sophisticated. They form their decisions based on the entire game tree, correctly anticipate their own behavior at future dates, and form consistent plans. Recently, Siniscalchi (2006) has provided an axiomatic foundation of such sophisticated dynamic choice for maxmin utility and full Bayesian updating. We follow the same idea here and posit that the (conditional) preferences are defined over trees, rather than acts; we comment more on this in section 6.1 below.

(10) Alternatively, the weights ε and $(1-\varepsilon)$ can be used simply to generate the set of priors of the Gilboa-Schmeidler model. In this case the mixture is done *only once* (at the initial stage) to generate the set G, and in subsequent periods the distributions in G are updated. However, even though the algebraic expressions for the updated minimizing distributions would be different if one follows this alternative approach, this would not affect our main result.

(11) See, for example, Walley (1991), Rios and Ruggeri (2000), Epstein and Schneider (2003). For an axiomatization of this rule, see Jaffray (1994), Pires (2002). See Siniscalchi (2006) for an axiomatization for an approach that is closest to the one we take in this paper.

This completes the description of the general aspects of dynamic choice. Next, we specify the mechanism, which clarifies the precise nature of dynamic choice facing the bidders. The strategies of the bidders and the nature of equilibria induced by the mechanism are then discussed in section 3.2.

3 The Modified Dutch Mechanism

We now describe the Modified Dutch Mechanism (MDM). The mechanism works as follows. At the beginning, the seller declares a price sequence $\{p_1, p_2, ...p_n\}$ where p_t is the asking price in period t. In each period t, for $t = 1, 2, ..., n$, the seller randomly chooses a buyer to approach first and offers the object at price p_t. The randomizations are independent across periods and each buyer has equal chance of being asked first. If the buyer buys at that price the game is over; otherwise the seller approaches the other buyer and offers the same price. If the second buyer accepts, the game is over; otherwise we go to period $t + 1$ if $t < n$ and the game is over (and the seller keeps the item) if both buyers refuse even in period $t = n$. We assume that the buyers are asked in secret, so that in every period, a buyer, when asked by the seller, does not know whether he is being asked first or is being asked because the other buyer has refused the current offer.

The mechanism is a modification of a discrete price Dutch auction; in particular we assume–as is standard in dynamic auctions–that there is no discounting between periods. The seller's ex post payoff is the price at which the item is sold if it is sold and zero otherwise. The ex post payoff of a buyer of type v is $v - p$ if it obtains the item at price p and is zero otherwise.

We maintain the standard assumption of mechanism design literature that the seller, the mechanism designer in our context, can *commit* to the mechanism. In particular this means that the price sequence declared at the beginning of the game and the random procedure of approaching buyers every period is adhered to as the game progresses. Put differently, once a mechanism is chosen, only the two buyers - and not the seller - are the players in the game induced by the mechanism. We also make the standard assumption that all of the above is common knowledge.

3.1 The Price Sequence

As mentioned before, the MDM consists of a price sequence $\{p_0, p_1, ..., p_n\}$ where p_k is the asking price in period k. Our objective is to show that for any $\varepsilon > 0$, there is an MDM such that in the equilibrium of the game resulting from it, the seller can extract almost all surplus from almost all types. Let us start by describing how the required price sequence is constructed. For $\delta > 0$, let $\{p_0, p_1, ..., p_n\}$ be the price sequence where

$$\begin{aligned} p_0 &= 1 \text{ and} \\ p_k &= \frac{(1-\delta)^k}{(1-\delta+\varepsilon\delta/2)^{k-1}} \text{ for any } k > 0 \end{aligned} \tag{3.1}$$

We remind the reader that ε is a preference parameter; we explain the role of δ shortly. Let Δ_k denote the "price gap" $p_k - p_{k+1}$, where

$$\begin{aligned} p_0 - p_1 \equiv \Delta_0 &= \delta, \text{ and} \\ p_k - p_{k+1} \equiv \Delta_k &= \left(\frac{1-\delta}{1-\delta+\varepsilon\delta/2}\right)^k \frac{\varepsilon\delta}{2} \text{ for any } k > 0 \end{aligned} \tag{3.2}$$

Note therefore that both p_k, and the gap Δ_k, are decreasing in k. It also follows directly that

$$\lim_{n\to\infty} \sum_{k=0}^{n} \Delta_k = 1$$

Since in the limit the prices cover the entire unit interval, we have the following property, which is important for later results:

Property: *Given any $\eta \in (0, 1)$, there exists an integer T such that $\sum_{k=1}^{T} \Delta_k \geqslant 1 - \eta$.*

Given any $\eta \in (0, 1)$, let T_* be the smallest integer for which the above inequality is satisfied. We set $n = T_*$, which defines the last offered price p_n.

The basic idea of the surplus extraction result (derived in section 4) is then as follows. The prices are constructed so that they cover the entire unit interval in the limit. The parameter δ controls the price gaps Δ_k, $k \geqslant 0$. For any given positive ε - no matter how small - we can specify δ so that all price gaps are "small" compared to ε. This is the crucial feature that allows us to exploit the ambiguity sensitivity of bidders: for any ε, there is a δ^* such that for any choice of $\delta < \delta^*$ the gain from waiting can be made arbitrarily small compared to the loss from waiting, making it optimal for bidders to stop "within δ" of their values. This means all participating types get a surplus of at most δ.

The remaining question, then, is to see which types participate. By suitably choosing n, the price sequence can be designed so that in equilibrium, the lowest type that plans to purchase the good has valuation η. Therefore all types above η participate and obtain a surplus of at most δ. Since both δ and η can be made arbitrarily small, the result follows.

3.2 Strategies and equilibria

As explained above, the MDM results in a sequential (extensive form) game of incomplete information. A strategy of a type in this game is a plan to accept or reject the seller's offer at every information set (i.e. at every instance where the seller makes the offer) given the history of the game so far. An equilibrium is a pair of strategies, one for each buyer, satisfying the standard conditions: the pair is commonly known and each is a best response with respect to the other. Further restrictions on the structure of behavior of buyers are discussed below.

First, we make the standard assumption that the game itself is common knowledge. Each buyer faces ambiguity about the type of the other buyer, but, as in the standard models, knows how each type behaves in equilibrium. Previous research has studied static mechanisms in exactly this context and since our objective is to focus on the role played by dynamic mechanisms, we preserve the other aspects of the framework. Note that we are ruling out strategic ambiguity: players do not doubt each other's rationality. However, as noted in the introduction, other sources of ambiguity only makes it easier to show the surplus extraction result of this paper.

Second, as noted in section 2, buyers have maxmin preferences throughout the game, use the full Bayesian updating rule, and form consistent plans. As noted, Siniscalchi (2006) provides an axiomatic foundation of such sophisticated dynamic choice with maxmin utility and full Bayesian updating.

Third, the equilibrium strategy of a buyer is perfect in the sense that just like in the standard case, the same consistency requirement is imposed on the off-the-equilibrium-path information sets as well. A type's equilibrium decision in any period (i.e. to accept or to reject the seller's offer) is optimal not only with respect to the other buyer's strategy and the history of the game but also with respect to the knowledge of its own behavior at all future information sets, including those that will not occur if the type is to carry out its own equilibrium plan. We discuss this issue further in section 6.1 .

3.3 Characterizing Strategies

In this section, we discuss a particularly convenient way of representing strategies in the game induced by the MDM.

Recall that at each price p_k, $k \in \{1, \ldots, n\}$, a buyer, if asked by the seller, must choose one of two actions: accept or reject the seller's offer. A strategy of a type of buyer i is therefore a plan to accept or reject the seller's offer at each price given the history up to price p_{k-1}. We assume that a buyer type accepts when indifferent between accepting and rejecting and buys at the earlier period if indifferent between buying in two different periods.[12]

An important feature of the strategies is that the decisions to buy by different types must have a certain monotonicity property. Specifically, suppose that p_k is the highest price that a buyer of type v accepts. This means that the payoff $v - p_k$ is better than the best (maxmin) expected payoff from either not accepting the seller's offer at all or accepting some future price. Since all types start with the same set of priors and use the same rule to update the set, any type $v' > v$ must then also optimally accept the offer p_k rather than to continue. If p_k is the first price at which type v plans to accept, then the highest price that all higher types plan to accept must be at least as high as p_k. And similarly, the highest price that types below v accept is either p_k or a lower price.

For each price p_k there is a set of types (possibly empty) who buy at p_k. Note that

[12] Note that this need not be an entirely innocuous assumption. Since at every node, a buyer can either accept or reject, the strategies are pure. In a non-EU setting, an agent who is indifferent between two pure actions might nevertheless strictly prefer a randomization over them to either pure action (see Crawford (1990) for the seminal contribution). However, in our case allowing for randomization does not create any problem. This is because the action "accept" gives rise to a sure (and hence unambiguous) payoff. And an implication of the axiom certainty independence (see axiom A.2. in Gilboa and Schmeidler (1989)) is that (even though there may be gains from hedging ambiguous acts) there is no gain in hedging an ambiguous and an unambiguous act. A comment by an anonymous referee helped us simplify the exposition of this point significantly.

monotonicity implies that if p_k is the highest price accepted by types v and v', where $v > v'$, then the same is true of any type $v'' \in (v', v)$. Therefore such a strategy gives rise to a vector of n cut-offs $\{v_1, \ldots, v_n\}$ where $1 \geqslant v_1 \geqslant v_2 \geqslant \ldots \geqslant v_n \geqslant 0$, and where types in the interval $[v_1, 1]$ plan to buy at p_1, and types in the interval $[v_k, v_{k-1})$ plan to buy at p_k, $k \in \{2, \ldots, n\}$.

To continue, we see that any strategy satisfying monotonicity must give rise to a vector of n cut-offs as described above. Thus without loss of generality we can restrict attention to such strategies, and refer to these as "cut-off strategies." Note that any such cut-off strategy currently places no restriction on the parts of the strategies which specify actions at prices below the highest acceptable price. For a strategy to be part of a perfect equilibrium, further restrictions are required and we clarify these once we establish the next result.

Next, we define an "interior cut-off strategy."

Definition 1 ***Interior Cut-off Strategy:*** *A strategy of buyer i, $i \in \{1, 2\}$, is called an interior cut-off strategy if there exists a vector $\mathbf{v}^i = (v_1^i, \ldots, v_n^i)$, $0 \leqslant v_n^i < v_{n-1}^i < \ldots < v_1^i < 1$, such that for $k \geqslant 1$, the highest price accepted by the non-degenerate interval of types $[v_k^i, v_{k-1}^i)$ is p_k, where $v_0^i \equiv 1$.*

3.4 Characterizing Equilibria

In this section we discuss the properties of equilibria that results from the game induced by the MDM. We show that when the price sequence $\{p_1, ..., p_n\}$ is chosen appropriately, any equilibrium has the property that for every price, there are sets of types of positive measure for both buyers who plan to buy at that price. (For the rest of the paper, the phrase "positive measure" is used with respect to the distribution F.) We also define perfect cut-off strategy, i.e. cut-off strategies that are part of a perfect equilibrium, and show existence of a symmetric equilibrium where both buyers follow the same cut-off strategy.

For the rest of the section, we fix the preference parameter $\varepsilon > 0$.

The first result calculates the difference between the payoffs from buying at the current price and waiting for the next lower price. This calculation is useful later when we show that exactly such a calculation features in deriving equilibrium cut-off vectors.

Lemma 1 *Suppose the item has not been sold in periods* $1, \ldots, k-1$ *and in period* $k < n$ *the seller offers the item to buyer* i *at price* p_k *(given by equation* (3.1)*). Suppose* j *follows an interior cut-off strategy that gives rise to a vector of cut-offs* $\mathbf{v}^j = (v_1^j, \ldots, v_n^j)$*. For any type* v *of* i *the difference in payoff from buying immediately at price* p_k *versus waiting one period to buy at price* p_{k+1} *is*

$$G_k^i(v) = v - p_k - (1-\varepsilon)(v - p_{k+1})H_k^i \tag{3.3}$$

where

$$H_k^i \equiv \frac{F(v_k^j) + F(v_{k+1}^j)}{F(v_k^j) + F(v_{k-1}^j)} \tag{3.4}$$

where $v_0^j \equiv 1$.

The proof is given below. Derivation of the conditional probabilities H_k^i used in the proof is provided in section **A**.1 in the appendix.

Proof: If buyer i accepts the price p_k, the payoff is $v - p_k$. If the buyer waits to buy in period $k+1$ and manages to obtain the item then the ex post payoff is $v - p_{k+1}$.

It is shown in section **A**.1 in the appendix that H_k^i is the probability under the distribution F that i obtains the item at p_{k+1} given that he refuses the current offer of p_k. Under epsilon contamination preference, the buyer's expected payoff from waiting one period is given by $(1-\varepsilon)(v - p_{k+1})H_k^i$. Therefore $G_k^i(v)$ is as specified. $\|$

The next result shows that all equilibrium strategies are interior cut-off strategies whenever the price gaps are small. In other words, this shows that for each price there is a positive measure of types of each bidder who plan to buy at that price. The result plays a crucial role in characterizing all equilibria.

Proposition 1 *There exists* $\overline{\delta} > 0$ *such that for all* $\delta < \overline{\delta}$*, the equilibrium strategies of both buyers are interior cut-off strategies.*

The formal proof is relegated to the appendix (section **A**.2). Here we provide a brief sketch of the basic idea behind the result. Suppose the strategy followed by buyer j is not an interior cut-off strategy. In other words, the strategy has "gaps" in the sense that there are prices such that no type of buyer j plans to buy at those prices.

For example, suppose there are no types of j who would accept offers of $p_{k-\ell}$ through p_k, but there are types of j who buy at $p_{k-\ell-1}$ and also types who buy at p_{k+1}. Let

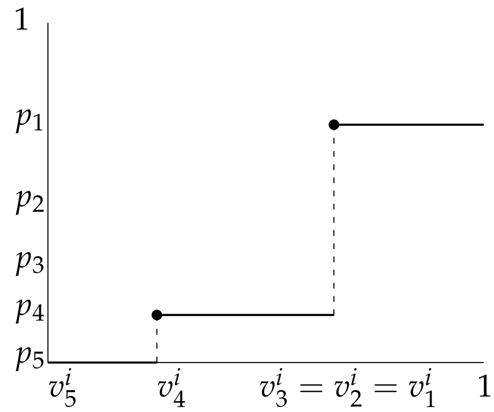

Figure 1: A cut-off strategy for buyer i under $n = 5$ with gaps at p_2 and p_3 - there are no types of bidder i who buys at p_2 or p_3. Our results rule out all gaps in equilibrium.

$v^j_{k-\ell-1}$ be the lowest type of j who buys at price $p_{k-\ell-1}$. Clearly, this type is indifferent between buying at $p_{k-\ell-1}$ and waiting till the price drops to p_{k+1}.

Now, according to the supposed equilibrium, all types in $(p_{k+1}, v^j_{k-\ell-1})$ refuse price offers $p_{k-\ell}$ through p_k. Note first that if ℓ is at least 1, and j does not plan to buy at prices $p_{k-\ell}$ through p_k, the best response of i should be not to buy at prices $p_{k-\ell}$ through p_{k-1}. It is possible that some type of i may want to buy at price p_k; however, the important point is that a gap from j will give rise to a corresponding gap from i.

Consider now $v^j_{k-\ell-1}$, the lowest type buying at $p_{k-\ell-1}$. The type is indifferent between buying at $p_{k-\ell-1}$ and waiting till p_{k+1}. Therefore a type just below (but arbitrarily close to) $v^j_{k-\ell-1}$ is approximately indifferent between those options. An important part of the argument is showing that as δ becomes small so that the gaps between prices decrease, the term $H^j_{k-\ell-1}$, which is the conditional probability (under distribution F) that j obtains the item in period $k+1$ if he passes in period $k-\ell-1$, is approximately equal to H^j_k, the conditional probability (again, under F) that j obtains the item in period $k+1$ if he passes in period k.

The above argument is used to derive a contradiction. In the proposed equilibrium, a type just below $v^j_{k-\ell-1}$ is approximately indifferent between $p_{k-\ell-1}$ and p_{k+1} (and in particular is not supposed to buy at prices $p_{k-\ell}$ through p_k). However, given that $H^j_{k-\ell-1} \approx H^j_k$, and $p_k < p_{k-\ell-1}$, in period k when the seller actually offers the price p_k, such a type strictly prefers to buy at p_k rather than wait till period $k+1$, contradicting the supposed equilibrium behavior. It is essentially this argument that rules out any gaps in the strategies adopted by either player in equilibrium, proving the stated result. While this is the basic intuition, the formal proof has to carefully check several cases, and is somewhat lengthy. We have relegated it to the appendix.

Recall that our definition of an interior cut-off strategy above did not impose any out-of-equilibrium restrictions. We now impose such restrictions and define a perfect cut-off strategy. For a strategy to be part of a perfect equilibrium, it must specify behavior that is optimal at every information set given the (correct) assessment of the behavior of the other player as well as one's own behavior at every continuation information set. Specifically, if, say $p_k(v)$ is the highest acceptable price for type v, it must be better for v to accept $p_k(v)$ than to reject and act optimally at every (off-equilibrium-path) future occasion if asked by the seller. The following result shows that for the appropriately chosen price sequence, such optimality simply implies that v must accept all subsequent (off-equilibrium-path) offers by the seller as well.

Lemma 2 *Let $p_k(v)$ be the highest acceptable price for type v of buyer i, $i \in \{1,2\}$. For $\delta < \overline{\delta}$, optimal behavior at any subsequent information set requires that type v also accepts all prices lower than $p_k(v)$.*

Proof: From Proposition 1 we know that for δ low enough equilibrium strategies are interior cut-off strategies. Hence for every price, there are types of positive measure who plan to buy at that price.

Next, suppose p_k is the highest acceptable price for a type v. Suppose, to the contrary, that type v does not find some price $p_{k+\ell}$, $\ell \geqslant 1$, acceptable. From Proposition 1, there is some type v' for whom $p_{k+\ell}$ is the highest acceptable price. The monotonicity property then immediately gives a contradiction. If $v' > v$, then the higher price p_k cannot be acceptable to v. On the other hand, if $v' < v$, then since v' finds it optimal to accept when offered $p_{k+\ell}$, the same must be true of the higher type v.

Thus, if p_k is the highest price type v accepts in equilibrium, then the (off-equilibrium-path) strategy of type v is to accept every lower price as well.$\|$

Therefore, a **perfect cut-off strategy** is an interior cut-off strategy with the additional requirement that if a type accepts any price, it must also accept all subsequent prices.

From definition 1, the highest price accepted by a non-degenerate interval of types $(v^i_{k-1}, v^i_k]$ is p_k. It follows that for a perfect cut-off strategy, v^i_k is the lowest type of i who buys at p_k, and is indifferent between accepting p_k or continuing for just one more period and accepting the next available price p_{k+1}.

We now use the results above to characterize perfect cut-off strategies. Since any equilibrium involves such strategies, this characterizes all equilibria.

Proposition 2 *For $\delta < \overline{\delta}$, in any equilibrium the strategy of any bidder i is a perfect cut-off strategy $\mathbf{v}^i = (v^i_1, \ldots, v^i_n)$ where $v_n = p_n$. Further, for $1 \leqslant k \leqslant (n-1)$, $v^i_k \in (p_k, v^i_{k-1})$, where $v_0 \equiv 1$, and v^i_k is given by*

$$v^i_k = p_k + \Delta_k \frac{(1-\varepsilon)H^i_k}{1-(1-\varepsilon)H^i_k}$$

where H^i_k is given by equation (3.4). *For any given $\mathbf{v}^j$, v^i_k is unique.*

Proof: For δ small, it follows from Lemma 2 that in any equilibrium buyers must use a perfect cut-off strategy. From Lemma 3 (in Appendix **A**.2), we have $v^i_n = p_n$. From Lemma 1 we know that if the strategy of j gives rise to the cut-off vector $\mathbf{v}^j = (v^j_1, \ldots, v^j_n)$, then for any type v of i the difference in payoff from buying immediately versus waiting one period to buy at price p_{k+1} is given by $G^i_k(v)$. Since the type v^i_k is the lowest type that buys at k, it must be that v^i_k is determined by solving $G^i_k(v) = 0$ for v.

Now, clearly, $G^i_k(p_k) < 0$. Therefore $v^i_k > p_k$. Since (as shown by Proposition 1) a positive measure of types of i plan to buy at each price, we also have $v^i_k < v^i_{k-1}$. Thus it must be that $G^i_k(v_{k-1}) > 0$. Further, $G^i_k(v)$ is strictly increasing and continuous in v. Therefore if an equilibrium $(\mathbf{v}^i, \mathbf{v}^j)$ exists, for any given $\mathbf{v}^j$ there exists a unique $v^i_k \in (p_k, v^i_{k-1})$ such that $G^i_k(v^i_k) = 0$. Finally, $G^i_k(v^i_k) = 0$ implies (from equation (3.3)) $v^i_k - p_k = (1-\varepsilon)(v^i_k - p_{k+1})H^i_k = (1-\varepsilon)(v^i_k - p_k + \Delta_k)H^i_k$. Solving, we get the stated equation.‖

The result above characterizes all equilibria. Finally, the next result proves existence. The proof is essentially an application of Brouwer's fixed point theorem and has been relegated to the appendix.

Proposition 3 *There is $\underline{\delta} > 0$ such that for any $\delta < \underline{\delta}$, a symmetric equilibrium exists.*

4 The Main Result

We now present the main result of the paper which follows directly from the characterization results derived in the last section. For any preference parameter $\varepsilon > 0$, the seller can design an MDM such that the object is sold to almost all types and the types who buy pay a price that is arbitrarily close to their valuation of the item. More specifically, for any given $\varepsilon > 0$, there is $\delta^*(\varepsilon)$ such that for any chosen $\delta \in (0, \delta^*(\varepsilon))$ and $\eta > 0$, the reserve type is no greater than η (i.e. the item is sold if at least one buyer's valuation is greater than η) and no buyer type obtains (an ex post) surplus greater than δ. (Of course, the types that do not buy get zero surplus. However, the seller makes zero revenue from them as well and so an important point of the result is that while extracting almost all surplus from the types that buy, the mass of non-buying types can be made to be arbitrarily small.) Since the set of types who are excluded are at most $[0, \eta]$ and the ex post surplus of the types who buy is at most δ, and since both δ and η can be arbitrarily small, the result follows.

Proposition 4 *For any preference parameter $\varepsilon > 0$, there exists $\delta^*(\varepsilon) > 0$ such that for any $\delta < \delta^*(\varepsilon)$, and $\eta > 0$, there is an MDM such that in any equilibrium of the game induced by the MDM, the item is sold if at least one buyer has valuation greater than η and no type obtains an ex post surplus greater than δ.*

Proof: The results in the previous section show that for any $\varepsilon > 0$, there is $\delta^*(\varepsilon) > 0$ such that whenever $\delta < \delta^*(\varepsilon)$, an equilibrium exists, and all equilibria can be characterized as in Proposition 2. Further, as noted in section 3.1, for any $\eta \in (0, 1)$, there exists an integer T such by choosing $n = T$, the price sequence (which consists of n prices) of the MDM covers at least a fraction $(1 - \eta)$ of types so that the item is not sold to at most types in $[0, \eta]$. Thus, it only remains to show that no type that buys gets an ex post surplus greater than δ.

Now, since types in $[v_k, v_{k-1})$ buy at price p_k, the ex post surplus of any type buying at p_k is at most $v_{k-1} - p_k$, which is bounded above by δ as follows: From the necessary conditions for equilibrium presented in Proposition 2, we have

$$\begin{aligned}
v_{k-1} - p_k &= p_{k-1} - p_k + \Delta_{k-1}\frac{(1-\varepsilon)H_{k-1}}{1-(1-\varepsilon)H_{k-1}} \\
&< \Delta_{k-1} + \Delta_{k-1}\frac{(1-\varepsilon)}{\varepsilon} \\
&= \frac{\Delta_{k-1}}{\varepsilon} \quad = \quad \frac{\delta}{2}\left(\frac{1-\delta}{1-\delta+\delta\varepsilon/2}\right)^{k-1} \quad < \quad \delta
\end{aligned}$$

where the second step follows from the fact that $H_{k-1}(v_{k-1}) < 1$. The final inequality follows from the fact that the coefficient of δ is less than 1 for any $\varepsilon > 0$. This completes the proof.‖

As mentioned in the introduction, the basic intuition for the result is that for any v and faced with any price p where $p < v$, a buyer can get a sure payoff $v - p$ from buying now, or wait to be offered a lower price but face the prospect that the other buyer accepts it before this can happen. As explained before, the loss from waiting is at least $(v-p)\varepsilon$ whereas the gain from waiting is of the order Δ_p. For any given ε, by making Δ_p successively small, the gain from waiting can be made arbitrarily small. Further, the price sequence is *constructed* so that $v - p$ is at most δ, and even when δ is small, the loss from waiting is still larger than the gain from waiting.

It is also worth reiterating that this cannot happen in the standard unique prior model. Roughly speaking, in the absence of ambiguity, given that F is smooth, the expected gain and expected loss from waiting shrink at the same rate as the price gap becomes smaller.

5 A Numerical Example

Suppose F is the uniform distribution on the unit interval. For any $k < n$ the equation for v_k is

$$v_k = p_k + \Delta_k \frac{(1-\varepsilon)H_k}{1-(1-\varepsilon)H_k}$$

where

$$H_k = \begin{cases} (v_1+v_2)/(1+v_1) & \text{for } k=1\text{, and} \\ (v_k+v_{k+1})/(v_{k-1}+v_k) & \text{for } 2 \leqslant k \leqslant (n-1) \end{cases}$$

Given $v_n = p_n$, the equations can be solved for any given n. It can be directly verified (as well as already noted in Proposition 3) that there is a unique positive solution for any v_k.

The following table shows a few steps ($n = 7$) for $\delta = 0.05$, and $\varepsilon = 0.2$. In this case we extract a surplus of at least 0.95 of value from the top 8% types. The prices p_k and cut-offs v_k are as shown. The right hand column shows the maximum rent obtained by any type. The rent obtained by any type $v \in [v_{(k+1)}, v_k)$ is given by $v - p_{(k+1)} \leqslant v_k - p_{(k+1)}$, which is the maximum rent.

Price	V_k	Maximum Rent
0.950	0.968	0.0500
0.945	0.964	0.0232
0.940	0.959	0.0243
0.935	0.954	0.0241
0.930	0.949	0.0240
0.925	0.943	0.0238
0.920	0.920	0.0228

Continuing in this fashion (i.e. by increasing n beyond 7), it is possible to extract a rent of at least a fraction 0.95 of value from any fraction of types less than 1.

It is interesting to compare this with the outcome of the static optimal mechanism. Bose et al. (2006) show that the optimal (static) direct revelation mechanism is a full insurance mechanism in which the reserve type v_* is such that $v_* - (1-\varepsilon)\frac{1-F(v_*)}{f(v_*)} = 0$, and the (expected) surplus of type $v \geqslant v_*$ is $(1-\varepsilon)\int_{v_*}^{v} F(y)dy$. If F is uniform on the unit interval, the reserve type is (approximately) 0.44, and the surplus is approximately 0.32 for $v = 1$. To have a direct comparison with the calculations above, suppose the seller were to choose 0.92 as the reserve type. Then type $v = 1$ gets a surplus equal to $(.8)\int_{0.92}^{1} ydy$ which is about 0.06.In contrast, in the MDM, type $v = 1$ gets a surplus of exactly δ which in this numerical example is 0.05 (and, in general, can be made arbitrarily small). Further, under the optimal static mechanism the surplus of type $v = 1$ increases further as the reserve price falls, but never exceeds δ under the MDM.

6 Discussion

Maxmin preferences, and in particular the epsilon contamination formulation, have been used extensively in the literature to represent ambiguity averse behavior. Our results show the effect that a dynamic mechanism can have in a setting which is IPV in all aspects other than the ambiguity aversion of the buyers. In this section we discuss various aspects of our model to explain their role in delivering the main result. Some of these issues have been briefly mentioned earlier.

6.1 Dynamic Consistency

Preferences satisfy dynamic consistency if an optimal plan based on prior preferences coincides with the sequentially optimal plan in a decision tree, and vice versa. This is unproblematic in the expected utility paradigm,(13) but does not arise naturally under ambiguity. It is well known that with ambiguity sensitive (and in general, non-expected utility) preferences, well known updating rules (including the one we use) can give rise to dynamic inconsistency. One response, adopted by several authors, is to impose dynamic consistency as an added axiom.(14) However, the literature also points out that in certain situations it makes more intuitive sense to allow for preferences that violate dynamic consistency. In particular, we refer the reader to Epstein and Schneider (2003) for an excellent discussion in an Ellsberg type setting which shows that when there are intuitive choices for *different* periods, ambiguity may result in dynamic consistency being problematic.

We explore this issue in the following parts. Using a simple example, we explore in section 6.1.1 the role of a dynamic mechanism when preferences are not dynamically consistent over Savage acts. In particular, we discuss how a dynamic mechanism differs from the optimal direct revelation mechanism. A second objective is to highlight

(13)In the expected utility paradigm, assuming that preferences satisfy dynamic consistency is not a problem if one assumes that the updating follows Bayes rule. It is well known (see Epstein and Schneider (2003)) that if the conditional preferences at every time-event pair satisfy expected utility theory, they satisfy dynamic consistency if and only if the updating is done using Bayes Rule.

(14)See, for example, Epstein and Schneider (2003), Maccheroni, Marinacci, and Rustichini (2006), Klibanoff, Marinacci, and Mukerji (2006). Alternatively, one can have a dynamically consistent updating rule but give up consequentialism. For such an approach, see Hanany and Klibanoff (2006).

the general insight regarding the role of dynamic mechanisms under ambiguity that extends beyond the specific settings used in the paper. We construct the example using the general maxmin model rather than the epsilon contamination specification.

While section 6.1.1 discusses the role of dynamic mechanisms in the presence of dynamic inconsistency over acts, we should emphasize that our formal model does not suffer from any inconsistency. Section 6.1.2 clarifies this latter point. It also points out why defining preferences over acts might be inadequate for the setting we consider.

6.1.1 Dynamic Versus Static Mechanisms and the Role of the Revelation Principle

Consider the following 2 period decision problem. A buyer has value v for an object. Initially, nature moves and either offers or does not offer the object at a price p. We assume that $v > p$. If the offer is made, the buyer chooses to accept or reject. If the buyer rejects, nature moves again and either offers or does not offer a price $p - \Delta$, for some $\Delta > 0$. If the offer is made, the buyer accepts. (Note that the mechanism in the paper has a similar structure of choice at each stage for each buyer. Since in this example the game ends in 2 periods, there is no loss of generality in simplifying the second period actions.)

Let us translate the above decision tree into states and acts. Let E_1 denote the state that "nature does not offer p." The complement of E_1 is broken up into two further states: E_2 (nature offers p, does not offer $p - \Delta$) and E_3 (nature offers p, offers $p - \Delta$). Hence, the state space is $E_1 \cup E_2 \cup E_3$. Since the buyer accepts the offer $p - \Delta$ if he hasn't already accepted the previous offer of p, it suffices to consider two acts: a which denotes "accept p" and r, denoting "reject p." The acts are defined as:

$$a(E_1) = 0, \;\; a(E_2) = a(E_3) = v - p$$
$$r(E_1) = 0, \;\; r(E_2) = 0, \;\; r(E_3) = v - p + \Delta$$

Suppose $\Pr[E_2] = \alpha, \Pr[E_3] = \beta$, and $\Pr[E_1] = 1 - \alpha - \beta$. By varying α and β we can generate a suitable set of distributions. To avoid the degenerate case we assume $\alpha > 0, \beta > 0$ and $\alpha + \beta < 1$.

Initially, the buyer compares $\min\{(\alpha + \beta)(v - p)\}$ to $\min\{\beta(v - p + \Delta)\}$. However, suppose E_1 did not happen and the buyer now faces the offer p. Now the comparison is between the sure payoff of $v - p$ and (using the full Bayesian updating rule)

$\min\left\{\left(\frac{\beta}{\alpha+\beta}\right)(v-p+\Delta)\right\}$. It is easy to find numbers such that even though $\min\{\beta(v-p+\Delta)\} > \min\{(\alpha+\beta)(v-p)\}$ (implying r is preferred to a initially), we have $v-p > \min\left\{\left(\frac{\beta}{\alpha+\beta}\right)(v-p+\Delta)\right\}$ (implying a is preferred to r conditional on E_1 not occurring). Note that unlike in an EU setting, such a "switch" can occur here because in calculating the conditional expected payoff under ambiguity, the updated minimizing distribution is not the update of the minimizing distribution used at the initial stage.

For example, suppose $v-p = 10$, $\Delta = 5$, and there are only two possible priors as shown below (where the entries are probabilities):

	E_1	E_2	E_3
Prior 1	0.89	0.01	0.1
Prior 2	0.8	0.1	0.1

The initial payoff from r is $0.1 \times 15 = 1.5$ and that from a is $\min\{.2, .11\} \times 10 = 1.1$ Therefore r is preferred to a. However the conditional payoffs (conditioning on E_1 not occurring) are: 10 from a and $\min\left\{\frac{0.1}{0.11} \times 15, \frac{0.1}{0.2} \times 15\right\} = 7.5$ from r. Therefore a is now preferred to r. Thus a plan based on the initial preferences does not coincide with the sequentially optimal plan for this decision problem.

An important objective of this exercise is to point out that even though the *full* surplus extraction is special (and as we indicate later, depends on specific preferences and properties of sets of priors), in a general setting with ambiguity averse dynamically inconsistent preferences, a dynamic mechanism can exploit conditional preferences and thus have extra opportunities for surplus extraction compared to a static mechanism that are not available in the standard EU model.

An important related point is that the standard direct revelation game essentially carries out the mechanism design exercise in terms of the "initial preferences." If the actual (indirect) mechanism is also static, this procedure is without loss of generality. However the mechanism from the standard direct revelation game extends to dynamic mechanisms with agents of the type we consider *only if* agents were to have the *power to commit* to the decisions formed from the initial preferences.[15] In contexts where

[15]The surplus earned by the buyers if they face the same price sequence as in the MDM but could somehow commit to a strategy, would be at least as much as the surplus they earn in the optimal static mechanism in Bose et al. (2006). Hence one could view the difference between an agent's surplus in their model and ours as the maximum price the agent is willing to pay to access a commitment device. We thank Peter Klibanoff for this observation.

such additional power is unavailable, the mechanism derived using the revelation game corresponds to the optimal static mechanism and use of a dynamic mechanism can produce results different from the optimal static one.

6.1.2 Consistent Plans

A coherent theory of dynamic choice must allow for consistent planning. Further, the theory must provide a consistent way of reconciling preferences at different time periods.

In this paper, we do this by adopting Siniscalchi (2006) as the foundation for dynamic behavior in our model. First, we use the idea that agents are forward looking. Formally, the axiom of sophistication allows for agents to form plans and to carry them out consistently: a sophisticated decision maker correctly anticipates his future preferences while making current plans, hence precluding the problem of future deviations from the current plan. Second, considering preferences defined over (decision) *trees* rather than over *acts* (with conditional preferences defined over sub-trees), results in a coherent theory of dynamic choice that does not need to appeal to the notion of dynamic consistency.(16),(17) In fact, as we now argue, the example from the previous section suggests that in settings like ours, acts may not be rich enough to capture all aspects of a problem that a decision maker cares about.

Consider a variation of the example discussed in the previous section. Now, the buyer moves first and chooses between one of two actions, a or r. Following each choice of the buyer, nature chooses between E_1, E_2 and E_3. The corresponding payoffs are $(0, v-p, v-p)$ if a is chosen, and $(0, 0, v-p+\Delta)$ if r is chosen.

This alternative tree gives rise to exactly the same acts as in the previous section. Hence if preferences are defined over acts only, the two situations should not give rise to different outcomes. However, this second decision tree corresponds to a static choice model, or alternatively, a situation where the buyer facing the choice problem de-

(16)See in particular section 4 of Siniscalchi for an axiomatization for the maxmin model and full Bayesian updating.

(17)It is indeed the case that buyers in our model have preferences that are sophisticated (so that they can carry out a plan consistently) and dynamically inconsistent (over Savage acts). The formal model however, does not refer to the property of dynamic consistency.

scribed in the previous section can somehow access a commitment device and hence can commit to choices based on the initial preferences.[18] From a formal point of view, this alternative decision tree is different from that of the previous section, and therefore when preferences are defined over trees, there is no inconsistency if the buyer's choice of actions differ across these two trees. Note finally, that the buyer, being sophisticated, will anticipate correctly that he chooses r in the decision tree described here and that he chooses a in the decision problem faced in the previous section.

6.2 Updating Rules

We use the full Bayesian updating rule where the decision maker uses Bayes rule to update all distributions (except those under which the observed event would be impossible), and the payoff is equal to the minimum expected utility calculated by considering this entire set of updated distributions. While this is one of the most well used rules in the literature, other rules have been proposed as well. For example, a second well known rule with an axiomatic foundation is the generalized maximum likelihood rule (Gilboa and Schmeidler, 1993). Under this, the retained (and updated) distributions are those that give maximum likelihood to the event known to have occurred. Our result holds under this alternative updating rule as well. Faced with an offer p_k in period k, a buyer knows that the object remained unsold in previous rounds, i.e. the other buyer's type is in $[0, v_{k-1})$. Therefore, when calculating the payoff from rejecting p_k, the buyer's contaminating set admits only distributions that are "most favorable" in terms of the event $[0, v_{k-1})$. However, within this set, the worst distribution is still the one the puts the entire weight on the event that the current buyer will not obtain the item if he waits. Therefore the minimum expected payoff is the same as under the full Bayesian rule.

Note that the general issue with respect to updating rules is that the results in the paper would hold as long as the rule does not throw away these worst distributions. And we are not aware of any general argument that would require removal of these worst distributions as the game progresses. To be specific, consider an event $[v_k, v_{k-1})$ and note that the set of distributions that the buyers consider initially has amongst it (at least) one distribution $\widetilde{F}$ which puts an epsilon weight on this event. Now, suppose

[18] See footnote (15).

the buyer is told that the event $[v_{k-1}, 1]$ has not occurred. There is no obvious reason to suggest that this extra information should now make $\widetilde{F}$ irrelevant.

6.3 Set of Priors

The version of epsilon contamination we use takes the set of all distributions to be the set of contaminating distributions. It has often been pointed out, especially in the statistics literature, that this may be too general, and a "reasonable" modification of the model might involve requiring each element of this set to satisfy certain properties. For example, suppose F is differentiable and satisfies the monotone hazard rate property; it might then be considered desirable to require elements of the set of contaminating distributions to satisfy these two properties as well.[19] However, such modifications do not automatically invalidate the main result of this paper. To see why, suppose the contaminating distributions are of the form $L_n(v) = v^n$, $n = 1, 2, ...$Each distribution in this family is differentiable and satisfies the monotone hazard rate property.[20] Since the conditional probability of the event $[v_k, v_{k-1})$, given the event $[0, v_{k-1})$ is equal to $1 - \left(\frac{v_k}{v_{k-1}}\right)^n$, and since $\inf_n \left(\frac{v_k}{v_{k-1}}\right)^n = 0$, it is as if, for all practical purposes, there is a contaminating distribution that puts the entire mass on $[v_k, v_{k-1})$ (which is precisely what is done in the formal analysis).

On the other hand, if the above example was changed to include the set of distributions of the form $L_n(v) = v^{\frac{1}{n}}, n = 1, 2, ...,$ the full surplus extraction result clearly does not hold. Similarly, if the set contains only finitely many distributions, again full surplus extraction is not possible.

The above discussion illustrates that even though we have taken the set of contaminating distributions to be the set of all distributions over $[0, 1]$, our result should hold for many–though not all–subsets of distributions as well. Letting $\mathcal{P}$ denote the set of contaminating distributions, a sufficient condition is that $\mathcal{P}$ contains all the distributions

[19]The interpretation is that as before the decision maker puts a weight ε on the true distribution not being F. Now, however, he has confidence that the true distribution, even if not F, has certain properties similar to F.

[20]Of course this is only a particular family of distributions having the property of being differentiable and having monotone hazard rate. However, since the minimum cannot increase if sets are made bigger, considering only a particular family like this suffices to illustrate our point.

such that[21]

$$\inf_{L\in\mathcal{P}}[L(x)\equiv \Pr\{v\leqslant x\} \text{ according to } L]=0 \text{ for all } x\in[0,1]$$

As noted before, the crucial idea is that when an agent decides whether to accept the current price p_k (and get a certain payoff), his gain from waiting for a further drop in price is made small compared to the loss from waiting. The contaminating set of distributions must be such that the updated distribution - used to calculate the (minimum) expected payoff from waiting - allows this to happen and it can be seen that the sufficient condition guarantees exactly this.

6.4 (Seemingly) minor differences in the design of the mechanism

In the MDM, in any round k, the price p_k is offered to the buyers *sequentially* which allows each buyer to obtain an unambiguous payoff from accepting the offer. To make the mechanism symmetric, we have chosen a formulation where in each round the seller randomly chooses the order in which buyers are offered the price for that round. We can also allow a variation in which the seller chooses, at the start of the game, a buyer who is asked first in each period. This encumbers the algebra since the two buyers have slightly different problems to contemplate, but the main result holds unchanged.

However, consider now a variation where at each stage k, the price p_k is offered to the two buyers *simultaneously*, and the winner is chosen randomly if both accept. In all other respects this alternative mechanism is similar to the MDM.

In a standard Bayesian setting, the two formats would clearly not yield significantly different outcomes. However, with ambiguity aversion, this seemingly minor difference produces a dramatically different outcome. Under the alternative specification, both accepting and rejecting an offer produce ambiguous outcomes, which is why the full surplus extraction result fails.

From the perspective of mechanism design theory this implies that minor differences in the construction of mechanisms - differences that might be outcome *irrelevant* in the standard EU framework - can have a drastic impact under ambiguity aversion.

[21]We thank an anonymous referee for this observation as well as the one about the role of seemingly minor design differences noted in the next section.

7 Conclusion

Evidence (experimental and otherwise) suggests that it is important for economic models to explore the consequence of non-expected utility preferences. The fairly large (and growing) literature in this area has given us many valuable insights.

In this paper, we consider a private values auction model with ambiguity and buyers with ambiguity averse preferences. In the standard setting with a unique prior, the optimal mechanism leaves all but the lowest participating type with information rent. Previous work shows that even under ambiguity aversion, the optimal static mechanism leaves buyer types with rent. In contrast, we show that in the latter environment, dynamic mechanisms have more power, and using the epsilon contamination specification to model ambiguity aversion, we construct a very simple dynamic mechanism that extracts almost all surplus.

We view the contribution of our work as providing an example of the non-standard effects that ambiguity aversion can have on mechanism design. Our formal model uses the epsilon contamination specification and clearly our result of full surplus extraction is related to this setting. Nevertheless, the idea that in auction like settings, dynamic mechanisms can extract a greater surplus than static ones by exploiting ambiguity aversion is a more general one. By showing that the equivalence between static and dynamic mechanisms (standard under the unique prior model) need not extend to a setting with ambiguity, our results strike a cautionary note for working in the non-unique prior environment. This is further highlighted by the contrast between our results and those in Bose et al. (2006) who study optimal static auctions under ambiguity, and leads us to conclude that a straightforward application of the revelation principle has its limitations when preferences are no longer characterized by subjective expected utility. Understanding the proper scope of the revelation principle with such "non-probabilistically sophisticated" preferences is an interesting question that we hope to address in future research.

8 Appendix: Proofs

A.1 Some Conditional Probabilities

This section derives some conditional probabilities that are used repeatedly in the analysis.

Let H^i_k denote the probability under the distribution F (i.e. if there were no ambiguity) that i obtains the item the item at p_{k+1} given that he refuses the current offer of p_k. This can be calculated in two parts.

First, let ϕ^i_k denote the probability under the distribution F that i obtains the item at p_{k+1} conditional on the item not being sold at p_k. Second, let π^i_k denote the probability (again, this is the probability under F) that if i refuses the current offer p_k the object remains unsold till the next price p_{k+1}. Then we have $H^i_k = \pi^i_k \phi^i_k$.

Calculating $\boldsymbol{\phi^i_k}$: ϕ^i_k can be derived is as follows. If buyer i is asked first in period $k+1$ (which happens with probability $1/2$), he obtains the item for sure. If j is asked first (probability $1/2$), i obtains the item only if j passes. Given that the object is unsold at p_k, we know that the type of j is lower than v^j_k. Therefore the probability that j will refuse p_{k+1} given that he has refused p_k is given by $Prob(v^j < v^j_{k+1} | v^j < v^j_k) = \dfrac{F(v^j_{k+1})}{F(v^j_k)}$. Therefore

$$\phi^i_k = \frac{1}{2} + \frac{1}{2}\frac{F(v^j_{k+1})}{F(v^j_k)} \tag{A.1}$$

Calculating $\boldsymbol{\pi^i_k}$: Next, π^i_k can be derived as follows.

First, we need to work out the probability that a buyer is being asked first given that he is asked whether he wants to buy at p_k. The conditioning on being asked is important since the fact that a buyer is asked whether he wants to buy at p_k conveys information about whether he is first or second. Let $q^i \in \{1, 2\}$ denote the position (1st or 2nd) of buyer i in any period. Further, let A^i denote the event that "buyer i is asked whether he wants to buy at p_k." We want to determine $Prob(q^i = 1|A^i)$.

$$
\begin{aligned}
Prob(q^i = 1|A^i) &= \frac{Prob(q^i = 1)Prob(A^i|q^i = 1)}{Prob(q^i = 1)Prob(A^i|q^i = 1) + Prob(q^i = 2)Prob(A^i|q^i = 2)} \\
&= \frac{\frac{1}{2}}{\frac{1}{2} + \frac{1}{2}\frac{F(v_k^j)}{F(v_{k-1}^j)}} = \frac{F(v_{k-1}^j)}{F(v_{k-1}^j) + F(v_k^j)}
\end{aligned}
$$

where $v_0^j \equiv 1$. Next, $Prob(q^i = 2|A^i) = 1 - Prob(q^i = 1|A^i) = \dfrac{F(v_k^j)}{F(v_{k-1}^j) + F(v_k^j)}$.

We are now ready to derive π_k^i. Note that given i refuses p_k, the probability of the object being unsold if i is second ($q^i = 2$) is 1, and the probability of the object being unsold if i is first ($q^i = 1$) is $\dfrac{F(v_k^j)}{F(v_{k-1}^j)}$. Therefore

$$
\begin{aligned}
\pi_k^i &= Prob(q^i = 1|A^i)\frac{F(v_k^j)}{F(v_{k-1}^j)} + Prob(q^i = 2|A^i)(1) \\
&= \frac{2F(v_k^j)}{F(v_{k-1}^j) + F(v_k^j)} \qquad \text{(A.2)}
\end{aligned}
$$

where $v_0^j \equiv 1$. Finally, using equations (A.1) and (A.2), we get

$$
H_k^i = \pi_k^i\, \phi_k^i = \frac{F(v_k^j) + F(v_{k+1}^j)}{F(v_k^j) + F(v_{k-1}^j)}
$$

where $v_0^j \equiv 1$.

A.2 Proof of Proposition 1

In this section we prove Proposition 1. The basic outline of our argument is as follows. In Lemma 3 we show that in any equilibrium, for both buyers, the cut-off type for price p_n is in fact p_n and that there are types of positive measure who plan to buy at price p_n. Next, we show in Lemma 4 that for both buyers, there are types of positive measure that plan to buy at p_1. Lemma 6 is crucial, it shows that whenever δ is sufficiently small, given that a positive measure of types of both buyers buy at prices p_1 and p_n, there must be a positive measure of types of both buyers who buy at price p_{n-1}. Proposition 1 now follows from a recursive argument: provided types of positive measure plan to buy at prices p_1 and $p_{k+1}, ..., p_n$, there must be types of positive measure who plan to buy at p_k as well.

We remind the reader that the term v_k^i, is used to denote the lowest type of buyer i who plans to buy at price p_k. Also, to avoid confusion with respect to superscripts versus exponents, in the rest of this Appendix, we refer to the two buyers as i and j instead of 1 and 2.

Lemma 3 *In any equilibrium, $v_n^i = v_n^j = p_n$. Further, a positive measure of types of both buyers plan to buy at price p_n but not at any earlier price.*

Proof: Consider a type $v \in (p_n, p_{n-1})$ of either buyer. Buying at any price greater than p_{n-1} is dominated by not buying at all. Further, while the surplus from not buying is zero, that from buying at price p_n is $v - p_n > 0$. Hence types of positive measure (p_n, p_{n-1}) must plan to buy at p_n but not at any earlier price. Furthermore, the lowest type (the type that is indifferent between buying at price p_n and not buying at all) that buys at p_n is p_n, so that $v_n^i = v_n^j = p_n$. $\|$

In what follows, we use the word "probability" to mean probability with respect to the distribution F.

Before we proceed, a comment on the notation we use in considering out-of-equilibrium cases where no type of a buyer plans to buy at some prices. For $k > 1$, let p_k through p_{k+t} be prices such that no types of a buyer buy at these prices; however, there are types who buy at price p_{k-1}. In this case as $v_{k-1} = v_k = ... = v_{k+t}$ and let p_{k-1} be the price at which types $[v_{k+t}, v_{k-1})$ plan to buy. For $k = 1$, we still denote $1 = v_0 = v_1 = ... = v_{k+t}$, but now p_{k+t+1} is the price at which types $[v_{k+t+1}, 1]$ plan to buy.

Lemma 4 *In equilibrium a positive measure of types of each buyer plan to buy at p_1.*

Proof: Suppose buyer j does not plan to buy at prices $p_1, \ldots, p_k$ for $1 \leqslant k < n$, and p_{k+1} is the first price at which j buys. (This is denoted as $v_1^j = \ldots = v_k^j = 1$ and $v_{k+1}^j < 1$.) Clearly, the best response of i is not to buy at prices $p_1, \ldots, p_{k-1}$. If i refuses p_k, the probability that the game reaches p_{k+1} is 1. Thus $\pi_k^i = 1$. Therefore $H_k^i = \pi_k^i \phi_k^i = \phi_k^i = 1/2 + (1/2)F(v_{k+1}^j)$.[22] Further, if i refuses p_k, he is asked first next period with probability $1/2$ and gets the unambiguous payoff of $(v - p_{k+1})$. Therefore, the payoff from refusing p_k is $\left(1/2 + 1/2(1-\varepsilon)F(v_{k+1}^j)\right)(v - p_{k+1})$.

Define the following function.

$$\widehat{G}_k^i(v) \equiv v - p_k - \left(1/2 + 1/2(1-\varepsilon)F(v_{k+1}^j)\right)(v - p_{k+1}) \tag{A.3}$$

$\widehat{G}_k^i(v)$ can be rewritten as $\frac{1}{2}(v - p_k)(1 - (1-\varepsilon)F(v_{k+1}^j)) - \frac{1}{2}(1 + (1-\varepsilon)F(v_{k+1}^j))\Delta_k$. Note that

$$\begin{aligned} 2\widehat{G}_k^i(1) &= (1-p_k)(1-(1-\varepsilon)F(v_{k+1}^j)) - (1+(1-\varepsilon)F(v_{k+1}^j))\Delta_k \\ &> \delta\varepsilon - (2-\varepsilon)\Delta_k \;\geqslant\; \delta\varepsilon - (2-\varepsilon)\Delta_1 \;=\; \frac{\delta\varepsilon^2}{2(1-\delta)+\delta\varepsilon} \;>\; 0 \end{aligned}$$

where the second step follows from the fact that $(1 - p_k) \geqslant (1 - p_1) = \delta$, and the fact that $F(v_{k+1}^j) < 1$, and the third step uses $\Delta_1 \geqslant \Delta_k$.

Since $\widehat{G}_k^i(v)$ is continuous, increasing in v, and negative at $v = p_k$, there exists v_k^i such that $\widehat{G}_k^i(v) > 0$ for $v > v_k^i$ and $\widehat{G}_k^i(v_k^i) = 0$. Since we know that i does not plan to buy at any earlier price than p_k, it must be that types $[v_k^i, 1]$ of buyer i plan to buy at p_k.

Now consider buyer j. If j refuses p_k, the probability (under F) that the object is sold at p_k is strictly positive. Therefore, the minimizing distribution for j puts the entire weight on the type of i being such that the object is sold at p_k. Let

$$G_k^j(v) \equiv v - p_k - (1-\varepsilon)(v - p_{k+1})H_k^j$$

(22)Note that this is the same formula as in equation (A.1), since here $F(v_k^j) = F(v_{k-1}^j) = 1$.

We have

$$\begin{aligned}
G_k^j(1) &= 1 - p_k - (1-\varepsilon)(1-p_{k+1})H_k^j \;=\; (1-p_k)(1-(1-\varepsilon)H_k^j) - (1-\varepsilon)\Delta_k H_k^j \\
&> \delta\varepsilon - (1-\varepsilon)\Delta_k \;\geqslant\; \delta\varepsilon - (1-\varepsilon)\Delta_1 \;=\; \delta\varepsilon\left(\frac{1-\delta+\varepsilon}{2(1-\delta)+\delta\varepsilon}\right) \;>\; 0
\end{aligned}$$

where the first inequality follows since $H_k^j < 1$, and $(1-p_k) \geqslant (1-p_1) = \delta$ and the second one follows since $\Delta_k \leqslant \Delta_1$. Since $G_k^j(v)$ is increasing and continuous, there are types of j of positive measure near 1 who would deviate and buy at p_k. Contradiction.‖

We now derive an inequality in the next lemma that is useful for later proofs. The reader can skip the proof of the lemma without losing the thread of the argument.

Lemma 5 $v^i_{n-\ell-t} - v^i_n < \delta(\ell + t)$.

Proof: Case 1: Some types of j buy at least some price in $\{p_{n-\ell-t}, \ldots, p_{n-\ell-1}\}$, $t \geqslant 1$.

In this case, if i refuses $p_{n-\ell-t}$, it is possible that the game ends before $p_{n-\ell}$ is offered. We know that $v^i_{n-\ell-t}$ is given by $G^i_{n-\ell-t}(v) = 0$, i.e.

$$\begin{aligned}
v^i_{n-\ell-t} - p_{n-\ell-t} &= (1-\varepsilon)(v^i_{n-\ell-t} - p_{n-\ell})H^i_{n-\ell-t} \\
&= (1-\varepsilon)(v^i_{n-\ell-t} - p_{n-\ell-t} + \Delta_{n-\ell-t} + \ldots + \Delta_{n-\ell-1})H^i_{n-\ell-t}
\end{aligned}$$

Solving,

$$\begin{aligned}
v^i_{n-\ell-t} - p_{n-\ell-t} &= (\Delta_{n-\ell-t} + \ldots + \Delta_{n-\ell-1})\frac{(1-\varepsilon)H^i_{n-\ell-1}}{1-(1-\varepsilon)H^i_{n-\ell-1}} \\
&< (\Delta_{n-\ell-t} + \ldots + \Delta_{n-\ell-1})\frac{(1-\varepsilon)}{\varepsilon} \qquad \text{(A.4)}
\end{aligned}$$

Let $\alpha \equiv \frac{1-\delta}{1-\delta+\delta\varepsilon/2}$. Note that $\alpha < 1$. From equation (3.2), we have $\Delta_k = \frac{1}{2}\delta\varepsilon\alpha^k < \frac{1}{2}\delta\varepsilon$. Therefore

$$\begin{aligned}
v^i_{n-\ell-t} - p_n &= v^i_{n-\ell-t} - p_{n-\ell-t} + p_{n-\ell-t} - p_n \\
&= v^i_{n-\ell-t} - p_{n-\ell-t} + \Delta_{n-\ell-t} + \ldots + \Delta_{n-1} \\
&< (\Delta_{n-\ell-t} + \ldots + \Delta_{n-\ell-1})\frac{(1-\varepsilon)}{\varepsilon} + \Delta_{n-\ell-1} + \Delta_{n-\ell} + \ldots + \Delta_{n-1} \\
&= (\Delta_{n-\ell-t} + \ldots + \Delta_{n-\ell-2})\frac{(1-\varepsilon)}{\varepsilon} + \frac{\Delta_{n-\ell-1}}{\varepsilon} + \Delta_{n-\ell} + \ldots + \Delta_{n-1} \\
&< \frac{1}{2}\left(\delta(1-\varepsilon)(t-1) + \delta + \delta\varepsilon\ell\right) \\
&< \delta(\ell + t) \qquad \text{(A.5)}
\end{aligned}$$

where the third step follows from the inequality (A.4) above, and the last step follows from the facts that $0 < \varepsilon < 1$.

Finally, since $v_n^i = p_n$, $v_{n-\ell-t}^i - v_n^i < \delta(\ell + t)$.

Case 2: No type of j buys at prices $\{p_{n-\ell-t}, \ldots, p_{n-\ell-1}\}$, $t \geqslant 1$.

We know that types of i buy at $p_{n-\ell-t}$ and at $p_{n-\ell}$ but not at the prices in between. If $t > 1$, any type of i who buys at $p_{n-\ell-t}$ can deviate profitably and buy at $p_{n-\ell-1}$ instead. Contradiction. Therefore in this case the only possibility is $t = 1$.

So it remains to prove the inequality when $t = 1$ and no type of j buys at $p_{n-\ell-1}$. In this case, analogously with (A.3), $v_{n-\ell-1}^i$ is given by $\widehat{G}_{n-\ell-1}^i(v) = 0$, where

$$\widehat{G}_{n-\ell-1}^i(v) \equiv v - p_{n-\ell-1} - \left(1/2 + 1/2(1-\varepsilon)R_{n-\ell-1}^i\right)(v - p_{n-\ell})$$

where $R_{n-\ell-1}^i$ is the conditional probability that j rejects $p_{n-\ell}$.[23] Using the fact that $v - p_{n-\ell} = v - p_{n-\ell-1} + \Delta_{n-\ell-1}$, and solving,

$$v_{n-\ell-1}^i - p_{n-\ell-1} = \Delta_{n-\ell-1} \frac{1 + (1-\varepsilon)R_{n-\ell-1}^i}{1 - (1-\varepsilon)R_{n-\ell-1}^i} < \frac{2-\varepsilon}{\varepsilon}$$

Proceeding as in (A.5),

$$\begin{aligned} v_{n-\ell-1}^i - p_n &= v_{n-\ell-t}^i - p_{n-\ell-1} + \Delta_{n-\ell-1} + \ldots + \Delta_{n-1} \\ &< \frac{2}{\varepsilon}\Delta_{n-\ell-1} + \Delta_{n-\ell} + \ldots + \Delta_{n-1} \; < \; \delta + \frac{\delta\varepsilon}{2}\ell \; < \; \delta(\ell+1) \end{aligned}$$

This completes the proof.‖

To continue with the proof of the proposition, let us now show that both buyers have types who plan to buy at price p_{n-1}.

We use the following observation repeatedly in the following proofs. Suppose buyer j does not plan to buy at prices $\{p_{n-\ell+1}, \ldots, p_{n-1}\}$ where $2 \leqslant \ell \leqslant n-1$, but plans to buy at $p_{n-\ell}$ (and of course at p_n). Then the best response of buyer i involves not planning to buy at prices $\{p_{n-\ell+1}, \ldots, p_{n-2}\}$ whenever $\ell > 2$. Further, there must be types of i who plan to buy at $p_{n-\ell}$. (Otherwise types of j buying at $p_{n-\ell}$ can profitably deviate to, say, p_{n-2}. This contradicts the assumption that j buys at $p_{n-\ell}$). Armed with these facts, let us now show the result.

[23] Suppose the lowest price higher than $p_{n-\ell}$ at which some types of j buy is $p_{n-\ell-1-s}$ for $s \geqslant 1$. Then $R_{n-\ell-1-s}^i = \frac{F(v_{n-\ell}^j)}{F(v_{n-\ell-1-s}^j)}$.

Lemma 6 *There is $\overline{\delta} > 0$ such that for $\delta < \overline{\delta}$ there are types (of positive measure) of j who buy at p_{n-1}.*

Proof: In the proposed equilibrium, types $v \geqslant v^j_{n-\ell}$ of j buy at prices $p \geqslant p_{n-\ell}$, with type $v^j_{n-\ell}$ and some types just above buying at price $p_{n-\ell}$. But since j does not buy at prices $\{p_{n-\ell+1}, \ldots, p_{n-1}\}$, types just below $v^j_{n-\ell}$ must buy at p_n and not before. Therefore, in the proposed equilibrium, it must be that $v^j_{n-\ell}$ is indifferent between buying at $p_{n-\ell}$ or p_n. So we have, for buyer j,

$$v^j_{n-\ell} - p_{n-\ell} = (1-\varepsilon)(v^j_{n-\ell} - p_n)H^j_{n-\ell} \tag{A.6}$$

where $H^j_{n-\ell} = \pi^j_{n-\ell}\ \widehat{\pi}^j_{n-1}\ \phi^j_{n-1}$, where $\pi^j_{n-\ell}$ is the probability that the object is unsold at $p_{n-\ell}$ given that j refuses the current offer of $p_{n-\ell}$, $\widehat{\pi}^j_{n-1}$ is the probability that the object will remain unsold at p_{n-1}, and ϕ^j_{n-1} is the probability that j obtains the item at price p_n.

We know some types of i buy at price $p_{n-\ell}$. Let $p_{n-\ell-t}$, $t \geqslant 1$, be the price before $p_{n-\ell}$ at which some types of i buy in equilibrium. We have

$$\pi^j_{n-\ell} = \frac{2F(v^i_{n-\ell})}{F(v^i_{n-\ell}) + F(v^i_{n-\ell-t})} \quad \text{and} \quad \widehat{\pi}^j_{n-1} = \frac{F(v^i_{n-1})}{F(v^i_{n-\ell})}$$

Note that if there are no types of i who buy at p_{n-1}, then $F(v^i_{n-1}) = F(v^i_{n-\ell})$, and $\widehat{\pi}^j_{n-1} = 1$. Otherwise $\widehat{\pi}^j_{n-1}$ is less than 1.

Finally $\phi^j_{n-1} = \dfrac{1}{2} + \dfrac{1}{2}\dfrac{F(v^i_n)}{F(v^i_{n-1})}$, where, again, if there are no types of i who buy at p_{n-1}, then $F(v^i_{n-1}) = F(v^i_{n-\ell})$. From the above,

$$H^j_{n-\ell} = \frac{F(v^i_{n-1}) + F(v^i_n)}{F(v^i_{n-\ell}) + F(v^i_{n-\ell-t})} \tag{A.7}$$

Now, we can rewrite equation (A.6) above as

$$v^j_{n-\ell} - p_{n-\ell} = \frac{(1-\varepsilon)(p_{n-\ell} - p_n)H^j_{n-\ell}}{1 - (1-\varepsilon)H^j_{n-\ell}} \tag{A.8}$$

Let

$$G^j_{n-1}(v) \equiv v - p_{n-1} - (1-\varepsilon)(v - p_n)\ H^j_{n-1}$$

where $H^j_{n-1} = \pi^j_{n-1}\ \phi^j_{n-1}$, where ϕ^j_{n-1} is as given above, and π^j_{n-1} is the probability that the object remains unsold at p_{n-1} given that j refuses the current offer of p_{n-1}. Note that $\pi^j_{n-1} = 1$ if no types of i buy at price p_{n-1}, otherwise it is equal to $\frac{2F(v^i_{n-1})}{F(v^i_{n-1})+F(v^i_n)}$. In either case, since $\phi^j_{n-1} < 1$, we have $H^j_{n-1} < 1$ as well. To establish that contrary to what has been supposed, there are types of j who will in fact want to buy at price p_{n-1}, it is useful to break up the analysis into several cases.

Case 1: ℓ and t are fixed positive integers.

Intuitively, this is the case where both i and j follow strategies where they do not buy for some finite number of prices. Note that in this case $\delta(\ell + t) \to 0$, as $\delta \to 0$

We use the fact that $v^i_{n-\ell-t} - v^i_n < \delta(\ell + t)$ (shown in Lemma 5) in the proof below. We must consider two subcases: the case in which some types of i buy at p_{n-1} and the complementary case.

Case 1.1: Some types of i buy at p_{n-1}.

Now, since there are no types of j who buy at p_{n-1}, it must be that $G^j_{n-1}(v)$ is not strictly positive for any $v \in [p_{n-1}, v^j_{n-\ell}]$. Consider the value of $G^j_{n-1}(\cdot)$ at $v^j_{n-\ell}$. We have

$$\begin{aligned}
G^j_{n-1}(v^j_{n-\ell}) &= v^j_{n-\ell} - p_{n-1} - (1-\varepsilon)(v^j_{n-\ell} - p_n)\ H^j_{n-1} \\
&= (v^j_{n-\ell} - p_{n-\ell}) + (p_{n-\ell} - p_n) - \Delta_{n-1} \\
&\qquad\qquad - (1-\varepsilon)\left[(v^j_{n-\ell} - p_{n-\ell}) + (p_{n-\ell} - p_n)\right] H^j_{n-1} \\
&= (p_{n-\ell} - p_n)\left(\frac{1-(1-\varepsilon)\ H^j_{n-1}}{1-(1-\varepsilon)\ H^j_{n-\ell}}\right) - \Delta_{n-1} \\
&> \left[2\left(\frac{1-(1-\varepsilon)\ H^j_{n-1}}{1-(1-\varepsilon)\ H^j_{n-\ell}}\right) - 1\right]\Delta_{n-1}
\end{aligned}$$

where the second step follows from equation (A.8), and the third step follows from the fact that $p_{n-\ell} - p_n \geqslant p_{n-2} - p_n = \Delta_{n-2} + \Delta_{n-1} > 2\Delta_{n-1}$.

Now, $H^j_{n-\ell}$ is given by (A.7), and $H^j_{n-1} = \dfrac{F(v^i_{n-1}) + F(v^i_n)}{F(v^i_{n-1}) + F(v^i_{n-\ell})}$. Therefore

$$\frac{H^j_{n-1}}{H^j_{n-\ell}} = \frac{F(v^i_{n-\ell}) + F(v^i_{n-\ell-t})}{F(v^i_{n-\ell}) + F(v^i_{n-1})}$$

From Lemma 5, $v^i_{n-\ell-t} - v^i_{n-1} < v^i_{n-\ell-t} - v^i_n < \delta(\ell + t)$. Therefore, as $\delta \to 0$, the ratio $\frac{H^j_{n-1}}{H^j_{n-\ell}}$ converges to 1. Hence for sufficiently small δ, the term $\frac{1-(1-\varepsilon)\ H^j_{n-1}}{1-(1-\varepsilon)\ H^j_{n-\ell}}$ is greater than $\frac{1}{2}$ and we have $G^j_{n-1}(v^j_{n-\ell}) > 0$.

Case 1.2: No type of i buys at p_{n-1}.

In this case, if buyer j refuses p_{n-1}, he knows that the game proceeds to the next stage and with probability $1/2$ he gets the first offer next period. Analogously to equation (A.3), define

$$\widehat{G}^j_{n-1}(v) \equiv v - p_{n-1} - \left(\frac{1}{2} + \frac{1}{2}(1-\varepsilon)\frac{F(v^i_n)}{F(v^i_{n-1})}\right)(v - p_n)$$

Let $R^j_{n-1} \equiv \dfrac{F(v^i_n)}{F(v^i_{n-1})}$. It follows that

$$\begin{aligned}
2\widehat{G}^j_{n-1}(v^j_{n-\ell}) &= (v^j_{n-\ell} - p_{n-1})(1-(1-\varepsilon)R^j_{n-1}) - \Delta_{n-1}(1+(1-\varepsilon)R^j_{n-1}) \\
&= (p_{n-\ell} - p_n)\left(\frac{1-(1-\varepsilon)\ R^j_{n-1}}{1-(1-\varepsilon)\ H^j_{n-\ell}}\right) - \Delta_{n-1}(1+(1-\varepsilon)R^j_{n-1}) \\
&> \left[2\left(\frac{1-(1-\varepsilon)\ R^j_{n-1}}{1-(1-\varepsilon)\ H^j_{n-\ell}}\right) - (2-\varepsilon)\right]\Delta_{n-1}
\end{aligned}$$

where the final inequality follows, as before, from the fact that $p_{n-\ell} - p_n > 2\Delta_{n-1}$. From Lemma 5, $v^i_{n-\ell-t} - v^i_n < \delta(\ell + t)$. Therefore, as $\delta \to 0$, the ratio $\frac{R^j_{n-1}}{H^j_{n-\ell}}$ converges to 1. Hence for sufficiently small δ, $\widehat{G}^j_{n-1}(v^j_{n-\ell}) > 0$.

Case 2: <u>t is arbitrary and ℓ varies with n.</u>

This is the case when the gap $p_{n-\ell} - p_{n-1}$ does not vanish as $\delta \to 0$.

As $\delta \to 0$, since $(p_{n-\ell} - p_n)$ does not vanish, and since for any given $\eta > 0$, $H^j_{n-\ell}$ is bounded away from zero, it follows from equation (A.8) that $v^j_{n-\ell} - p_{n-\ell}$ does not vanish. Therefore, $v^j_{n-\ell} - p_{n-1}$ does not vanish. However, $p_n - p_{n-1} \to 0$, and $(1-\varepsilon)H^j_{n-1} < 1$ (for case 1.1) and $(1-\varepsilon)R^j_{n-1} < 1$ (for case 1.2). Therefore for δ small enough, $G^j_{n-1}(v^j_{n-\ell}) > 0$ (case 1.1) and $\widehat{G}^j_{n-1}(v^j_{n-\ell}) > 0$ (case 1.2).

In cases 1 and 2 above, we have shown that $G^j_{n-1}(v^j_{n-\ell}) > 0$ (and $\widehat{G}^j_{n-1}(v^j_{n-\ell}) > 0$). But since $G^j_{n-1}(\cdot)$ (and $\widehat{G}^j_{n-1}(\cdot)$) is strictly increasing, continuous, and negative at p_{n-1}, this implies that there is $v^j_{n-1} \in (p_{n-1}, v^j_{n-\ell})$ such that $G^j_{n-1}(v) > 0$ (and $\widehat{G}^j_{n-1}(v) > 0$) for $v \in (v^j_{n-1}, v^j_{n-\ell})$. Since types below $v^j_{n-\ell}$ do not buy at any price greater than or equal to $p_{n-\ell}$, these types (of positive measure) strictly prefer to stop at p_{n-1} rather than wait till p_n. This contradicts the supposition that there are no types of j who buy at p_{n-1}.

We need to consider a third possibility in order to complete the Lemma.

Case 3: ℓ is a fixed integer and t varies with n.

This is the case when as $\delta \to 0$, $\delta(\ell + t - 1)$ does not go to zero because t (and n) becomes arbitrarily large as δ becomes small. However, this is analogous to a case we have analyzed before with i and j roles switched. We know that in equilibrium, both buyers have types who plan to buy at price $p_{n-\ell}$. If i plans to buy at prices $p_{n-\ell-t}$ and $p_{n-\ell}$, but not to buy at prices $\{p_{n-\ell-t+1}, ..., p_{n-\ell-1}\}$, the best response of j should involve not buying at prices $\{p_{n-\ell-t+1}, ..., p_{n-\ell-2}\}$. If $p_{n-\ell-t} - p_{n-\ell-1}$ does not go to zero, we can use the arguments in case 2 above to argue that contrary to what is being supposed, for small δ, buyer i will in fact have some types of positive measure who buy at $p_{n-\ell-1}$ rather than waiting till $p_{n-\ell}$.

This completes the proof of the lemma.$\|$

To continue now with the proof of the Proposition, suppose both buyers have a positive measure of types buying at prices p_{n-k} to p_n, where $1 \leqslant k \leqslant n-2$. By exactly the same argument as above we can establish that both buyers must also buy at p_{n-k-1}. This, combined with the previous steps complete the proof of Proposition 1.$\|$

A.3 Proof of Proposition 3

Define $A^k = [p_k, 1]$ for $k = 1, 2, \ldots, n-1$. Let A be the cartesian product of A^k. A vector $\mathbf{x} \in A$ is of the form: $\mathbf{x} = \{x_1, \ldots, x_{n-1}\}$, such that $x_k \in [p_k, 1]$. Note that A is closed and bounded and hence compact, and it is also convex.

Let Z be the cartesian product of $[0, 1]$ taken $n-1$ times. Let C be the subset of Z such that

$$C = \left\{ \mathbf{x} \in [\mathbf{0}, \mathbf{1}]^{n-1} \ \ |x_1 \geqslant x_2 \geqslant \ldots \geqslant x_{n-1} \right\}$$

Note that C is compact and convex. Let $D \equiv C \cap A$. Since C and A are both finite dimensional compact and convex sets, D is also compact and convex.

Finally, we define the set of cut-off vectors E:

$$E = \left\{ \mathbf{v} \in [\mathbf{0}, \mathbf{1}]^{n+1} \ \ |v_0 = 1, \{v_1, \ldots, v_{n-1}\} \in D, v_n = p_n \right\}$$

E is the set of cut-off vectors D with each vector augmented by an initial and final element, which are fixed at 1 and p_n, respectively.

Throughout the proof we assume that δ is small enough so that all previous results hold. The following definitions are used throughout the proof.

Any vector $(v_1, \ldots, v_{n-1})$ is said to be in the interior of D if $v_k > v_{k+1}$ for all $k \in \{1, \ldots, n-2\}$, and any vector in D not in the interior of D is said to be in the border of D. Any vector $\mathbf{v}$ is said to be in the **interior (border)** of E if $(v_1, \ldots, v_{n-1})$ is in the interior (border) of D.

Let E_B denote the border of B, and let E_I denote the interior of E. Clearly, $E = E_B \cup E_I$.

Next, similar to the term H_k^i in Lemma 1 (as well as in appendix **A**.1), let $H_k(\mathbf{v})$ denote the probability that the a buyer can buy at p_{k+1} conditional on passing at p_k. As before, this is given by $H_k(\mathbf{v}) = \dfrac{F(v_k) + F(v_{k+1})}{F(v_k) + F(v_{k-1})}$.

Let $\mathbf{y}(\mathbf{v}) \equiv \{y_0(\mathbf{v}), \ldots, y_n(\mathbf{v})\}$ denote the best response to any $\mathbf{v} \in E$.

STEP 1: First, consider the set of vectors in E_I (the interior of E). From Proposition 2, we know that the best response is unique, continuous, and given by $y_0(\mathbf{v}) = 1$, $y_n(\mathbf{v}) = p_n$, and for $0 < k < n$:

$$y_k(\mathbf{v}) = p_k + \Delta_k \frac{(1-\varepsilon) H_k(\mathbf{v})}{1 - (1-\varepsilon) H_k(\mathbf{v})} \tag{A.9}$$

STEP 2: Next consider $\mathbf{v} \in E_B$. Let us first show that the best response mapping can be discontinuous at border vectors. For example, consider any border vector where $v_{k-1} = v_k = v_{k+1}$ for some k. For any such vector, $H_k = 1$, and the best response involves not buying at p_k, implying $y_k = y_{k-1}$.

Note that for any $\mathbf{v} \in E_I$, $H_k < 1$. The argument in step 1 shows that for any $\mathbf{v} \in E_I$, $y_k < y_{k-1}$. Further, $\lim_{H_k \to 1} y_k < y_{k-1}$. But if $H_k = 1$, $y_k = y_{k-1}$. Thus the best response mapping could be discontinuous at any $\mathbf{v} \in E_B$.

To solve the problem we proceed as follows. We set up a "pseudo best response" function as follows. In taking a best response to any $\mathbf{v} \in E_I$, a buyer behaves according to equation (A.9). Thus the pseudo best response coincides with the true best response on E_I. For any $\mathbf{v} \in E_B$, there is at least one k for which $v_{k-1} = v_k$ (i.e. no type of the other buyer buys at p_k). Faced with any degenerate interval $[v_k, v_{k-1})$, the pseudo best response corresponds to the best response of the (fictitious) situation where the buyer believes that if he does not buy at p_k, then with probability ε the object is sold to the other buyer before the game reaches p_{k+1}.(24)

Thus faced with any border vector, when offered any price p_k, the pseudo best response compares $v - p_k$ to $(1-\varepsilon)(v - p_{k+1})H_k$ for all values of $H_k \leqslant 1$ (i.e. *even when* $H_k = 1$).

For any $\mathbf{v} \in E_B$ let $\widehat{y}_k(\mathbf{v})$ denote the pseudo best response. From the above, this is exactly similar to equation (A.9) for interior points, and is given by

$$\widehat{y}_k(\mathbf{v}) = p_k + \Delta_k \frac{(1-\varepsilon)H_k(\mathbf{v})}{1-(1-\varepsilon)H_k(\mathbf{v})}$$

for $H_k(\mathbf{v}) \leqslant 1$.

Consider any $\widetilde{\mathbf{v}} \in E_B$. From equation (A.9), clearly $\lim_{\mathbf{v} \to \widetilde{\mathbf{v}}} y_k(\mathbf{v}) = \widehat{y}_k(\widetilde{\mathbf{v}})$. Thus replacing y by $\widehat{y}$ on E_B preserves continuity of the best response mapping.

With this specification, the calculations in Proposition 2 can be retraced and it can be easily seen that all conclusions are exactly the same (even with $H_k = 1$ we preserve the factor $(1-\varepsilon)$, and none of the results require $H_k < 1$). In particular, note that $\widehat{y}_k < \widehat{y}_{k-1}$ for all $k \in \{1, \ldots, n\}$, and therefore the pseudo best response vector belongs in the interior of E.

(24) In other words, while the true best response to any such border point would assume correctly that refusing p_k would mean that the game reaches the next stage with probability 1, the pseudo best response in effect assumes that the interval $[v_k, v_{k-1})$ is not degenerate, but there are some types of the other buyer who do buy at p_k.

STEP 3: Finally, define the mapping $\Psi : E \to E$ such that

$$\begin{aligned}
\Psi_0(\mathbf{v}) &= 1 \\
\Psi_k(\mathbf{v}) &= \begin{cases} y_k(\mathbf{v}) & \text{if } \mathbf{v} \in E_I \\ \widehat{y}_k(\mathbf{v}) & \text{if } \mathbf{v} \in E_B \end{cases} \\
\Psi_n(\mathbf{v}) &= p_n
\end{aligned}$$

Since Ψ maps E continuously to itself, by Brouwer's fixed point theorem, there exists a fixed point of Ψ, i.e. there exists $\mathbf{v}^*$ such that $\Psi(\mathbf{v}^*) = \mathbf{v}^*$.

We know from Proposition 2 that for any $\mathbf{v} \in E_I$, $\Psi(\mathbf{v})$ belongs to the interior of E. As noted at the end of step 2, the same is true for vectors in E_B. Thus the range of Ψ is a subset of E_I. Therefore any fixed point must be in E_I. But $\Psi(\mathbf{v})$ is the *true best response* for any $\mathbf{v} \in E_I$. It follows that any fixed point must be a true mutual best response, and therefore a symmetric equilibrium.$\|$

References

AUGUSTIN, T. (2003): "On the suboptimality of the generalized Bayes Rule and robust Bayesian procedures from the decision theoretic point of view: A cautionary note on updating imprecise priors," In 3rd Int. Symp. Imprecise Probabilities and Their Applications, pp 31–45.

BOSE, S., E. OZDENOREN, AND A. PAPE (2006): "Optimal Auctions with Ambiguity," *Theoretical Economics*, 1(4), 411–438.

CAMERER, C. AND M. WEBER (1992): "Recent developments in modeling preferences: uncertainty and ambiguity," *Journal of Risk and Uncertainty*, 5, 325–370.

CHEN, Z. AND L. G. EPSTEIN (2002): "Ambiguity, Risk, and Asset Returns in Continuous Time," *Econometrica*, 70, 1403–1443.

CHU, C. AND W. LIU (2002): "Knightian uncertainty and the voting paradox," Mimeo, University of Washington.

CRAWFORD, V. (1990): "Equilibrium without Independence," *Journal of Economic Theory*, 50, 127–154.

CRÉMER, J. AND R. P. MCLEAN (1985): "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, 53, 345–61.

——— (1988): "Full extraction of the surplus in Bayesian and dominant strategy auctions," *Econometrica*, 56, 1247–1257.

ELLSBERG, D. (1961): "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics*, 75, 643–669.

EPSTEIN, L. AND M. SCHNEIDER (2003): "Recursive Multiple Priors," *Journal of Economic Theory*, 113, 1–31.

GILBOA, I. AND D. SCHMEIDLER (1989): "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 18, 141–153.

——— (1993): "Updating Ambiguous Beliefs," *Journal of Economic Theory*, 59, 33–49.

GRUNWALD, P. D. AND J. Y. HALPERN (2004): "When Ignorance is Bliss," Proceedings of the Twenthieth Conference on Uncertainty in AI. ACM Int. Conf. Proc. 70, 226–234.

HANANY, E. AND P. KLIBANOFF (2006): "Updating Preferences with Multiple Priors," Mimeo, Northwestern University.

HUBER, P. (1973): "The use of Choquet capacities in statistics," *Bulletin of International Statistics Institute*, 45, 181–191.

JAFFRAY, J.-Y. (1994): "Dynamic decision making with belief functions," in *Advances in the Dempster-Shafer Theory of Evidence*, ed. by R. R. Yager, J. Kacprzyk, and M. Fedrizzi, Wiley.

KARNI, E. (1988): "On the equivalence between descending bid auctions and first price sealed bid auctions," *Theory and Decision*, 25, 211–217.

KARNI, E. AND Z. SAFRA (1986): "Vickrey auctions in the theory of expected utility with rank-dependent probabilities," *Economics Letters*, 20, 15–18.

——— (1989a): "Ascending bid auctions with behaviorally consistent bidders," *Annals of Operations Research*, 19, 435–446.

——— (1989b): "Dynamic consistency, revelations in auctions and the structure of preferences," *Review of Economic Studies*, 56, 421–433.

KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2006): "Recursive Smooth Ambiguity Preferences," Mimeo.

KOPYLOV, I. (2008): "Subjective Probability and Confidence," Working Paper, UC Irvine.

LO, K. C. (1998): "Sealed bid auctions with uncertainty averse bidders," *Economic Theory*, 12, 1–20.

MACCHERONI, L., M. MARINACCI, AND A. RUSTICHINI (2006): "Dynamic Variational Preferences," *Journal of Economic Theory*, 128, 4–44.

MASKIN, E. S. AND J. G. RILEY (1984): "Optimal Auctions with Risk Averse Buyers," *Econometrica*, 52, 1473–1518.

MATTHEWS, S. (1983): "Selling to Risk Averse Buyers with Unobservable Tastes," *Journal of Economic Theory*, 30, 370–400.

MCAFEE, R. P. AND P. J. RENY (1992): "Correlated information and mechanism design," *Econometrica*, 60(2), 395–421.

MUKERJI, S. (1998): "Ambiguity Aversion and Incompleteness of Contractual Form," *American Economic Review*, 88(5), 1207–1231.

MUKERJI, S. AND J.-M. TALLON (2004): "Ambiguity Aversion and the Absence of Wage Indexation," *Journal of Monetary Economics*, 51, 653–670.

MYERSON, R. (1981): "Optimal Auctions," *Mathematics of Operations Research*, 6, 58–63.

NAKAJIMA, D. (2004): "Allais Paradox Bidders in First-Price Auctions, Dutch Auctions and their Variants," Working Paper, Princeton University.

NISHIMURA, K. G. AND H. OZAKI (2004): "Search and Knightian uncertainty," *Journal of Economic Theory*, 119, 299–333.

——— (2006): "An axiomatic approach to ε-contamination," *Economic Theory*, 27, 333–340.

OZDENOREN, E. (2002): "Auctions and bargaining with a set of priors," Unpublished manuscript, University of Michigan.

PIRES, C. P. (2002): "A rule for updating ambiguous beliefs," *Theory and Decision*, 53(2), 137–152.

RILEY, J. G. AND W. F. SAMUELSON (1981): "Optimal Auctions," *American Economic Review*, 71, 381–392.

RIOS, I. D. AND F. RUGGERI, eds. (2000): *Robust Bayesian Analysis*, vol. 152 of *Lecture Notes in Statistics*, Springer Verlag.

ROBERT, J. (1991): "Continuity in auction design," *Journal of Economic Theory*, 55(1), 169–179.

SEIDENFELD, T. (2004): "A contrast between two decision rules for use with (convex) sets of probabilities: $\gamma-$maximin versus E-admissibility," *Synthese*, 140, 69–88.

SINISCALCHI, M. (2006): "Dynamic Choice under Ambiguity," Mimeo, Northwestern University.

VOLIJ, O. (2002): "Payoff equivalence in sealed bid auctions and the dual theory of choice under risk," *Economics Letters*, 76, 231–237.

WALLEY, P. (1991): *Statistical Reasoning with Imprecise Probabilities*, vol. 42 of *Monographs on Statistics and Applied Probability*, Chapman and Hall.