University *of* Leicester

DEPARTMENT OF ECONOMICS

Bilateral Bargaining in an Ambiguous Environment

Subir Bose, University of Leicester, UK
Suresh Mutuswami, University of Leicester, UK

Working Paper No. 12/10
May 2012

BILATERAL BARGAINING IN AN AMBIGUOUS ENVIRONMENT

SUBIR BOSE AND SURESH MUTUSWAMI

Abstract. We perturb the bilateral bargaining model by introducing small ambiguity (via the epsilon contamination model) about the agents' types. We assume that the preferences are characterized by ambiguity aversion (Gilboa-Schmeidler). The rest of the setup is exactly as in Myerson and Satterthwaite [10]. And we show that in this environment, it is possible to design a mechanism that generates almost-efficient trade. Crucially, the mechanism has to be extensive-form; standard (normal form) direct revelation mechanism can only generate outcome that is continuous with respect to the amount of ambiguity.

JEL: C78, D8

Key words: Ambiguity Aversion, Mechanism Design, Bilateral Bargaining, Myerson Satterthwaite.

1. Introduction

In a seminal paper, Myerson and Satterthwaite [10] analyzed a bilateral bargaining model and established that there are no individually rational, balanced budget, and incentive compatible mechanism that is also efficient. In this paper, we show that we get a more positive result if the Myerson-Satterthwaite model is "perturbed" in a small way. In particular, we consider a setup where the buyer does not know the precise distribution from which the seller's valuation is drawn. There are various ways of modeling such ambiguity; in this paper we use the ε-*contamination* model. Our central result is that even when the amount of ambiguity is small, there exists an incentive compatible, individually rational and budget balanced mechanism which is almost efficient in a sense made precise below.

There is by now a small but growing literature on ambiguity on mechanism design.[1] The papers most closely related to ours are Bodoh-Creed [1], and Castro and Yannelis [5] and Bose and Daripa [3]. Both

Date: March, 2012.

[1]See for example, [1], [2], [3], [4], [5].

Bodoh-Creed and Castro and Yannelis consider the bilateral bargaining problem when agents have ambiguity averse preferences as represented by the MMEU (MaxMin Expected Utility) model; see below for more details on the MMEU model. Bodoh-Creed makes use of the revenue maximizing transfer scheme and shows that (i) the ex-ante budget deficit needed to implement the efficient outcome is smaller, the higher is the amount of agent ambiguity and (ii) when ex-ante budget balanced mechanisms are used, the extent of efficient trade increases as the amount of agent ambiguity is increased. Castro and Yannelis consider a special form of MMEU preferences where each individual considers the *worst case scenario* for each action when choosing an action and show that in this setting of complete ambiguity, efficient trade is possible using an incentive compatible, individually rational and budget balanced mechanism. Both these papers thus show the role of ambiguity aversion towards achieving efficiency when the level of ambiguity is large (limiting result in Bodoh-Creed) whereas we are mainly concerned with exploitation of ambiguity aversion for efficient trade when the level of ambiguity is small. Thus we view our work to be complementary to these two papers.

Our work is also closely related to Bose and Daripa [3]. The dynamic mechanism we use is a modification of the mechanism used there. Bose and Daripa consider an auction environment and show that with epsilon contamination preferences (see below for more on this preference) the seller (auctioneer) can extract almost full surplus from almost all buyer types even when the level of ambiguity (as represented by the parameter epsilon) is small. The insight we use is that the type of mechanism used there can also be used in other situations when the mechanism designer has objective that is different from revenue maximization. Specifically, in the present construction, one side (here the seller) is given all the gains from trade (in other words the seller is made the residual claimant of all the surplus arising from trade) which "in effect" completely relaxes the seller's incentive compatibility constraint. (We need to make clear that since we do not use a direct revelation game, we use the term incentive compatibility a bit loosely; what we mean by the preceding sentence is that in the (indirect) mechanism we use, the seller does not have any incentive to say no to trade *whenever* the price is higher than the seller's valuation of the item). As for the buyer, the mechanism is - essentially - characterized by a decreasing sequence of price offers and we show, as in Bose and Daripa, that in equilibrium, buyer types who trade do so at prices that are arbitrarily close to their true valuation. [The proofs and results in subsection 4.2 simply mimic steps in many of the proofs in Bose and Daripa]. Crucially, what this results in is buyer types agreeing to trade in almost all situations where trade is desirable on grounds of efficiency. This along with the fact that almost all buyer types trade in equilibrium then gives us our almost-efficient trade result.

Note that in our case, the seller also extracts almost all surplus from almost all buyer types (with whom it is efficient that the seller trades). An important implication of our result is that the Bose and Daripa construction can be applied in situations where the mechanism designer has objectives other than surplus extraction *provided* the objective does not *conflict* with extracting almost all surplus from some agent. [For example, a setup that causes problem with the Bose and Daripa construction is when the designer, in addition to desiring efficient outcome *also* wants a fair division of gains from trade, a division that depends on agent's privately known types.]

The rest of the paper is organized as follows. Section 2 describes the basic, non-ambiguous setup. The notion of ambiguity aversion, as well as that of the updating rule are discussed in section 3. Section 4 contains all the main results. Section 5 contains a discussion on the important topic as to why normal form mechanisms are in general inadequate for our purpose. Section 6 concludes with a brief discussion of some interesting future research.

2. The Basic, Non-ambiguous Setup

In this section we describe the standard setup as in Myerson and Satterthwaite (for more details, see [10]); ambiguity is introduced into the model in Section 4. There is one seller and one potential buyer with the seller having one unit of an indivisible item to sell. The seller's valuation (opportunity cost) of the object is c while the buyer's valuation is v. Assume that $v \in V$ and $c \in C$ where V and C are compact, convex subsets of $\mathbb{R}_+$ such that there exists a non-degenerate interval $[a, b]$ such that $[a, b] \subset V \cap C$.[2] As is standard, we assume that the valuations are private information. A crucial element of the standard setup - and the one from which we depart in section 4 when we introduce ambiguity into the model - is the assumption that it is commonly known that v and c are determined according to independent draws from the distributions $G(v)$ and $F(c)$. It is assumed that the distributions have continuous densities $g(v) > 0$ and $f(v) > 0$.[3]

The preferences of the buyer and seller are as follows. Let x_b denote the probability with which the buyer obtains the item and x_s the probability with which the seller gives up the item.[4] Similarly, let t_b and t_s denote (monetary) payments received by the buyer and the seller respectively. Then the utilities of the buyer, U_b and the seller, U_s of types v and c respectively are given by $U_b = x_b v + t_b$ and $U_s = -x_s c + t_s$.

A feasible *mechanism* consists of message spaces M_b, M_s with generic elements m_b and m_s respectively, an allocation rule $x(m_b, m_s)$ which is the probability with which the item is transferred from the seller to the buyer and monetary payments, $t_b(m_b, m_s)$ and $t_s(m_b, m_s)$. A direct mechanism is where the message spaces $M_b = V$ and $M_s = C$. The interpretation is that in the direct mechanism the buyer and the seller are asked to report their type to the mechanism. By a standard invocation of the *revelation principle*, when searching for the optimal mechanism it suffices to restrict the search to direct mechanisms. Let $\sigma_b(v)$ and $\sigma_s(c)$ be the strategies of the buyer and the seller in the game (form) resulting from the (direct) mechanism. Then the mechanism is *incentive compatible* if there exists an equilibrium of the direct game where the equilibrium strategies of the buyer and the seller are $\sigma_b(v) = v$ and $\sigma_s(c) = c$; in other words the buyer and seller report their types truthfully to the mechanism. According to the revelation principle, it is sufficient to consider incentive compatible direct mechanisms when searching for the optimal mechanism.

A mechanism is *efficient* if it results in trade whenever $v > c$; in other words $x(v, c) = 1$ when $v > c$ and $x(v, c) = 0$ when $v < c$. [5]

[2]This condition is referred to as *overlapping support* in the literature.
[3]As is well-known, efficiency is possible if the assumption on densities is violated.
[4]Currently, the basic preferences are being defined. Later when we describe mechanisms, we will impose feasibility constraints requiring x_b to be equal to x_s.
[5]It is irrelevant whether trade takes place when $v = c$.

It is *individually rational* if $\int_C \{p(v,c)v + t_b(v,c)\}\, dF(c) \geq 0$ for all $v \in V$ and $\int_V \{-p(v,c)c + t_s(v,c)\}\, dG(v) \geq 0$ for all $c \in C$. It is *budget balanced* if $t_s(v,c) + t_b(v,c) = 0$ for all $(v,c) \in V \times C$. Note that while efficiency and budget balance are *ex post*, the individual rationality is *interim* individual rationality.

Myerson and Satterthwaite showed that in the above setting, there is no efficient mechanism satisfying *individual rationality* (IR henceforth), *incentive compatibility* (IC henceforth) and *budget balance.* (BB henceforth).[6]

In this paper, we perturb this environment by introducing (a small amount of) ambiguity and study the effect this has on designing of mechanisms. Before we discuss how specifically we introduce ambiguity into the bargaining framework, we first describe briefly the notion of ambiguity and the behavioral notion of ambiguity aversion that we use in our work. We also discuss the updating rule agents use.

3. Ambiguity Aversion

In this section we briefly review the model of ambiguity aversion that we use; for more details please refer to the original articles cited. Let Ω be a set of states and let Σ be a σ-field on Ω. Let X be a set of outcomes. An act is a function $f : \Omega \to X$ which is measurable.

Let $\mathcal{P}$ be the set of possible priors. In the standard expected utility model $\mathcal{P}$ is a singleton set. In a model with ambiguity $\mathcal{P}$ typically consists of more than one element. The question in this extended model is how an agent ranks various acts. The maxmin expected utility (MMEU, for short) proposed by Gilboa and Schmeidler [7] says that an act f is preferred by an agent to an act g if and only if

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

The Gilboa-Schmeidler setting is a static one and hence, it does not contain a method of updating beliefs following the arrival of new information. Our mechanism, however, is dynamic (as will become clear). There are various ways of extending the Gilboa-Schmeidler

[6] There is actually a stronger result: there exists no mechanism satisfying IR, IC and even *ex ante* BB where the latter requires $t_s(v,c)$ and $t_b(v,c)$ to be such that $\int_V \int_C \{t_s(v,c) + t_b(v,c)\} dF(c)dG(v) = 0$.

model to allow for the updating of beliefs following the arrival of new information.

The updating rule that we use is called the "full Bayesian" (or "prior-by-prior") rule.[7] In words, on arrival of some new information, the agent updates each prior in $\mathcal{P}$ using Bayes rule. This will give a new set of priors, say $\mathcal{P}'$. She then chooses the action f such that

$$f = \text{argmax}_{g:\Omega\to X} \min_{p\in\mathcal{P}'} \int u(g)dp$$

It is worth observing at this point that there is a potential source of dynamic inconsistency in the above procedure.[8] In particular, consider a set of priors $\mathcal{P}$, and a filtration $\mathcal{F}$. When $\mathcal{P}$ is not a singleton, it is possible that the set of updated priors $\mathcal{P}'$ is such that the act g is preferred to the act f (in the sense that the minimum expected utility from g is higher than minimum expected utility from f) for every event $E \in \mathcal{F}$, and yet the minimum expected utility of f is higher than minimum expected utility of g when the set of priors is $\mathcal{P}$. Hence in this situation f is the unconditional preferred act even though g is preferred to f conditional on *every* event E of the filtration $\mathcal{F}$.

Notice that for such a "preference reversal" (over acts f and g) to happen, it is necessary that there be at least one event $E \in \mathcal{F}$ where the minimizing updated prior p' is *not* the update of the unconditional minimizing prior $p \in \mathcal{P}$. Such "dynamic inconsistency" is therefore not present in the standard expected utility framework where there is a single prior. As we shall see, the dynamic inconsistency associated with ambiguity averse preferences plays a crucial role in our results.

4. Bilateral Bargaining with Ambiguity

We now introduce ambiguity into the environment by supposing that there exists some amount of ambiguity regarding the type of the

[7]See Kopylov [8] for an axiomatization of this rule. See also Pires [11], Siniscalchi [12].

[8]It is possible to have a set of priors $\mathcal{P}$ and a filtration $\mathcal{F}$ of Ω leading to the set of posteriors $\mathcal{P}'$, such that no such dynamic inconsistency arises when the prior-by-prior updating rule is used. These are called rectangular priors in the literature. See Epstein and Schneider [6] for more on this.

agents[9]. First, however, (solely) for ease of exposition, we assume that both V and C are $[0,1]$; we discuss generalizations later (in subsection 4.5). The important point of departure from the standard setup is with respect to an agent's beliefs regarding the other agent's type. Specifically, instead of it being common knowledge that the types are drawn from the distributions G and F, we allow the agents to entertain (small) doubt regarding the distributions from which the valuations are drawn. This doubt is represented, parametrically, through the ϵ-contamination model. Formally, letting $\mathcal{P}$ denote the set of distributions over $[0,1]$, the ambiguous beliefs of the buyer regarding the seller's valuation is represented by set of priors $\mathcal{H}$ where there exists an $\varepsilon \in (0,1]$ such that every $H \in \mathcal{H}$ is represented as

$$H(c) = (1-\varepsilon)F(c) + \varepsilon P(c);\;\; P \in \mathcal{P}$$

In what follows, it would be useful (though strictly speaking, not necessary) to think of ε to be small for the results that follow.

We can also describe the ambiguous beliefs of the seller regarding the buyer's type as well. However, since this will play no role in the formal analysis that follows, we refrain from introducing unnecessary notation.

4.1. **The Mechanism.** The mechanism is an extensive game. It involves a strictly decreasing sequence of prices $\mathbf{p} = \{p_0, p_1, \cdots, p_n\}$, for some finite positive integer n. A third party (this could be the mechanism designer; in any case we will call the third party designer henceforth) starts by asking the seller whether she is willing to sell at the highest price.[10] If the seller says yes, then the designer asks the buyer whether he is willing to buy at that price. If the buyer accepts the offer, then trade takes place at that price. If the buyer declines then the designer asks the seller whether she is willing to sell at the second highest price and so on. This process continues till one of three things happen: (i) buyer and seller both say yes to some price p_k; the outcome then is trade taking place at that price (ii) after the seller having said

[9]It will be clear, after going through the model and the result, that it is enough to introduce ambiguity about one agent's (the seller's, for the mechanism described in this paper) type. See below for a more detailed discussion on this.

[10]We will see shortly that it will not affect anything if instead the first asking price is p_1, which is the first "non-trivial" price.

yes while the buyer having said no to all prices $p_0, ..., p_k$, the seller says no to the asking price p_{k+1}; in this case the game is over with no trade taking place, and (iii) the buyer says no to the last offered price p_n as well; again the game ends with no trade taking place.

The role of the price sequence $\mathbf{p} = \{p_0, p_1, \cdots, p_n\}$ is crucial. Letting $\Delta_k = p_k - p_{k+1}$ denote the "price gaps", it is an important requirement that Δ_k be a decreasing sequence as well. More specifically, the prices p_k are as follows. Fix $\delta \in (0, 1)$. (The role of this will be clarified shortly). The price sequence is given by:

$$p_k = \begin{cases} 1 & \text{if } k = 0, \\ (1-\delta)^k \alpha^{k-1} & \text{if } k \geq 1 \end{cases}$$

where

$$\frac{1}{1-\delta+\varepsilon\delta} \leq \alpha < \frac{1}{1-\delta}.$$

The sequence of price gaps, denoted by $\Delta_k \equiv p_k - p_{k+1}$, are then given by:

$$\Delta_k = \begin{cases} \delta & \text{if } k = 0, \\ (1-\delta)^k \alpha^k \left[\frac{1}{\alpha} - (1-\delta)\right] & \text{if } k \geq 1. \end{cases}$$

Note that since $(1-\delta)\alpha < 1$, both p_k and Δ_k are strictly decreasing in k. Next, we record a couple of fairly simple, but very useful, properties of the price sequence in the following two lemmas.

Lemma 1. *For all $k \geq 1$, $\Delta_k \leq \delta\varepsilon$.*

Proof. For $k \geq 1$, since $(1-\delta)\alpha < 1$,

$$\Delta_k \leq \left[\frac{1}{\alpha} - (1-\delta)\right] \leq 1 - \delta + \varepsilon\delta - (1-\delta) = \varepsilon\delta.$$

where the second inequality follows because $\alpha \geq \frac{1}{1-\delta+\varepsilon\delta}$. □

Lemma 2. *Given any $\lambda \in (0, 1)$, there exists a positive integer T such that $\sum_{k=1}^{T} \Delta_k \geqslant 1 - \lambda$.*

Proof omitted since it is simply a straightforward application of the fact that $\lim_{n\to\infty} \sum_{k=0}^{n} \Delta_k = 1$.

For any $\lambda \in (0, 1)$, 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. Let us now indicate the role of the other parameter, δ, in the price sequence. We show below that in an equilibrium of this game, the strategy of a seller of type c involves saying "YES" to all prices $p_0, ..., p_\ell$, and saying "NO" to price $p_{\ell+1}$ where $p_{\ell+1} < c \leq p_\ell$.

More interestingly, for small $\delta > 0$, the strategy of a buyer of type v is to say "NO" to all prices $p_0, ..., p_{k-1}$, and to say "YES" to all prices $p_k, ..., p_n$ where $v - p_k$ is no greater than δ. Hence, when trade takes place, it does so at prices such that the buyer obtains an ex-post surplus that is no larger than δ. As we will see shortly, this will be crucial for obtaining the result that the mechanism is almost efficient (made precise in Proposition (3) below). Notice that for any given λ, a decrease in ε or δ results in an increase in the value of T_*.

We proceed as follows. We first fix the seller's strategy as described above and show that for δ sufficiently small, all buyer types above λ plan to purchase at a price that is at most δ less than their true valuation. We then show that given the mechanism, and the buyer's best response, the optimal strategy of the seller is to indeed follow the seller-strategy described above. These will then be used to show that trade is almost efficient.

4.2. **Buyer's best response.** In this section we show that for δ sufficiently small, the buyer's best response can be characterized by a monotone sequence of cutoffs $\{v_k\}_{k=1,2,...}$ such that $v_k > v_{k+1}$ and for each price p_k, type v_k is the lowest type that buys at price p_k; moreover, the maximum (ex-post) surplus any buyer-type obtains is equal to δ.

Consider a type v who (has not bought yet and) is offered a price p. If $p > v$, then it is clear that the buyer-type v should not accept the offer. If on the other hand $v > p$, then the buyer-type can accept and get a sure payoff of $v-p$. Suppose instead of accepting the current offer buyer-type v decides to wait for the next offer and - provided the next offer is made by the seller - plans to accept that offer. Letting $p - \Delta$ denote the next price offer, the expected payoff [11] from this alternative strategy (of waiting) is equal to

$$(v - p + \Delta)(1 - \varepsilon)\left(\frac{F(p - \Delta)}{F(p)}\right)$$

[11] The expectation is taken with respect to the updated minimizing distribution. Hence, we should write "updated maxmin expected payoff"; however for the sake of brevity, we will simply write "expected" payoff since it will always be the case that the phrase expected payoff would mean updated maxmin expected payoff.

To see why this is the expected payoff, note that given that the seller makes the offer p it must be that the seller's valuation is below p. The next offer will be made if the seller's valuation c is less than $p-\Delta$.Since the updated minimizing prior puts the weight ε on the event $c \in (p - \Delta,\, p)$, we obtain the above expression for the expected payoff for the waiting strategy.

We can deduce from the above the buyer's optimal decision when the buyer faces the following two-period problem: buy now or wait and buy next period. We will use this to characterize the buyer's optimal response. Since the buyer can, however, choose to wait more than one period - or follow some other more complicated strategy - we will need to show that the buyer's optimal decision can in fact be reduced to a suitable two-period problem. [12] More specifically, we show that the equilibrium, and hence the buyer's best response strategy, has the following structure: For all buyer-types $v \geq p_n$, there exists a cut-off price $p(v)$ such that $p(v)$ is the first (that is the highest) price at which type v plans to buy. Further, type v plans to accept the seller's offer for *all* prices $p \leq p(v)$. These results are shown through a series of lemmas. The first of them shows that there is an interval of types of buyers that plan to buy at price p_n (but not at higher price).

Lemma 3. *There is an interval of types that plan to buy at price p_n but at no higher price.*

Proof. Since p_n is the last offer, all types $v > p_n$ should accept the offer of p_n. Since the buyer's surplus from not buying is 0, no type $v \in (p_{n-1}, p_n]$ should however, plan to buy at price $p_k \geq p_{n-1}$. □

We next show that there are types who plan to buy at price p_1.

It is convenient to first define the notation $H_{k,k+\ell}$ and the function $G_{k,k+\ell}(v)$.

Definition 1. *Let $H_{k,k+\ell}$ be defined as:*

$$H_{k,k+\ell} \equiv \frac{F(p_{k+\ell})}{F(p_k)}$$

[12] Using dynamic programming terminology, we show that the optimal strategy has a "no one-step-deviation" property

In other words, $H_{k,k+\ell}$ is the non-ambiguous (i.e., under the distribution F) conditional probability that the seller's valuation is below $p_{k+\ell}$, given that it is below p_k.

Definition 2. *Let the function $G_{k,k+\ell}(v)$ be defined as*

$$G_{k,k+\ell}(v) \equiv (v - p_k) - (1-\varepsilon)(v - p_{k+\ell})H_{k,k+\ell}$$

The function $G_{k,k+\ell}(v)$ is the difference between two payoffs for type v. The first is the *certain*, in particular*, unambiguous*, payoff from accepting the offer p_k at stage k. The other is the *ambiguous* expected payoff (evaluated in stage k) from waiting to buy at stage $k+\ell$ at the lower price $p_{k+\ell}$. This payoff is uncertain since the item may be withdrawn before stage $k+\ell$ is reached and in this environment, the expected payoff is ambiguous as well. Note that under the latter, the buyer type v plans to purchase at stage $k+\ell$ but not at any earlier stage. Note also that $G_{k,k+\ell}(v)$ is strictly increasing in v; this will play a crucial role in what follows.

Lemma 4. *There is an interval of types $[\widehat{v}_1, 1]$, with $\widehat{v}_1 < 1$, that plan to buy at price p_1.*

Proof. Suppose to the contrary that no type plans to buy at price p_1. Let p_m, where m is an integer strictly greater than 1, be the first price which is accepted by any type of the buyer. Note that such a m must exist according to Lemma 3. But this means, in particular that all buyer types would reject offer p_{m-1} (and accept p_m). Now, consider $G_{m-1,m}(v)$, for some integer $m \geq 2$, evaluated at $v = 1$. We have,

$$\begin{aligned}
G_{m-1,m}(1) &= (1-p_{m-1}) - (1-\varepsilon)(1-p_m)H_{m-1,m} \\
&> (1-p_{m-1}) - (1-\varepsilon)(1-p_m) \\
&= (1-p_{m-1}) - (1-\varepsilon)(1-p_{m-1}+\Delta_{m-1}) \\
&= (1-p_{m-1})(\varepsilon) - (1-\varepsilon)(\Delta_{m-1}) \\
&\geq \delta\varepsilon - (1-\varepsilon)(\Delta_{m-1}) \\
&> 0
\end{aligned}$$

where the first inequality follows since $H_{m-1,m} < 1$, the second inequality follows since $1 - p_{m-1} \geq 1 - p_1 = \delta$, and the final inequality uses

Lemma 1.But this means when facing the price offer p_{m-1}, type $v = 1$, and by continuity, some interval of types $(\tilde{v}, 1]$, will find accepting price p_{m-1} preferable to waiting for the offer p_m. Contradiction. □

The next lemma shows that if a type v plans to accept a price p then the optimal strategy involves accepting all lower prices. This result is crucial since, as mentioned earlier, it allows us to simplify the buyer's dynamic optimization problem into a two-step (buy now or at the next stage) problems.

Lemma 5. *Let $p(v)$ be the highest price that type v plans to accept. Then for δ sufficiently small, the optimal response of the buyer-type v is to accept all lower prices.*

Proof. Recall that $G_{k,k+\ell}(v)$ is strictly increasing in v. Hence, if a type, say v', is indifferent between accepting a current offer p_k and waiting till period $k + \ell$ to accept offer $p_{k+\ell}$, then all types $v > v'$ strictly prefer to accept the offer p_k when faced with the same choice.

Consider now a price p_k and suppose to the contrary that there is a set of types, $\widetilde{V}_k$, such that for all $v \in \widetilde{V}_k$, price p_k is the highest price that is acceptable, and furthermore, the next acceptable price is $p_{k+\ell}$ for some integer $\ell > 1$. [Such a price $p_{k+\ell}$ must exist according to Lemma 3]. Hence, types in $\widetilde{V}_k$ do not plan to purchase at any price $\{p_{k+1}, \dots, p_{k+\ell-1}\}$. Let v_k be the lowest type in $\widetilde{V}_k$.[13]

Since,

$$v_k - p_k = (1 - \varepsilon)(v_k - p_{k+\ell})H_{k,k+\ell}$$

and by definition,

$$p_{k+\ell} = p_k - \sum_{i=0}^{\ell-1} \Delta_{k+i}$$

we have,

$$v_k - p_k = \frac{(1 - \varepsilon)\left(\sum_{i=0}^{\ell-1} \Delta_{k+i}\right) H_{k,k+\ell}}{1 - (1 - \varepsilon)H_{k,k+\ell}} \tag{1}$$

[13]Such a v_k must exist. By definition, $G_{k,k+\ell}(v) \geq 0$ for all $v \in \widetilde{V}_k$. Since $G_{k,k+\ell}(v = p_{k+\ell}) < 0$, it is clear that all $v \in \widetilde{V}_k$ are strictly greater than $p_{k+\ell}$. Since $G_{k,k+\ell}(v)$ is continuous in v, it is impossible to go from $G_{k,k+\ell}(v) > 0$ to $G_{k,k+\ell}(v) < 0$ without having a v, which we call v_k, such that $G_{k,k+\ell}(v_k) = 0$

Now consider the term $G_{k+\ell-1,k+\ell}(v_k)$ where recall that

$$G_{k+\ell-1,k+\ell}(v_k) = v_k - p_{k+\ell-1} - (1-\varepsilon)(v_k - p_{k+\ell})H_{k+\ell-1,k+\ell}$$

We have,

$$\begin{aligned} & v_k - p_{k+\ell-1} - (1-\varepsilon)(v_k - p_{k+\ell})H_{k+\ell-1,k+\ell} \\ = & \left(v_k - p_k + \sum_{i=0}^{\ell-1}\Delta_{k+i}\right)(1-(1-\varepsilon)H_{k+\ell-1,k+\ell}) - \Delta_{k+\ell-1} \\ = & \left(\sum_{i=0}^{\ell-1}\Delta_{k+i}\right)\left(\frac{1-(1-\varepsilon)H_{k+\ell-1,k+\ell}}{1-(1-\varepsilon)H_{k,k+\ell}}\right) - \Delta_{k+\ell-1} \end{aligned}$$

where the last equality follows from substituting the expression $v_k - p_k$ from equation (1)

We want to show that the above is positive when δ is sufficiently small. We need to consider two sub-cases. The first sub-case deals with the situation where type $v \in \widetilde{V}$ has a strategy such that it does not plan to accept the next $\ell - 1$ prices after price p_k. In the second case, type v plans to reject all offers unless the price *gap* between p_k and the next acceptable price offer is at least as large as some number $L > 0$.[14] Note that in the first case ℓ is finite and hence as $\delta \to 0$, $p_k - p_{k+\ell} \to 0$. In the second case, as $\delta \to 0$, we have $\ell \to \infty$.

Sub-case 1: ℓ finite.

In this case the ratio $\frac{H_{k+\ell-1,k+\ell}}{H_{k,k+\ell}} \to 1$ as $\delta \to 0$. Since $\sum_{i=0}^{\ell-1}\Delta_{k+i} > \ell\Delta_{k+\ell-1}$, and $\ell \geq 2$, we have

$$\begin{aligned} & \left(\sum_{i=0}^{\ell-1}\Delta_{k+i}\right)\left(\frac{1-(1-\varepsilon)H_{k+\ell-1,k+\ell}}{1-(1-\varepsilon)H_{k,k+\ell}}\right) - \Delta_{k+\ell-1} \\ > & \ [\Delta_{k+\ell-1}]\left[\ell\frac{1-(1-\varepsilon)H_{k+\ell-1,k+\ell}}{1-(1-\varepsilon)H_{k,k+\ell}} - 1\right] \end{aligned}$$

and when δ is sufficiently small, $\frac{1-(1-\varepsilon)H_{k+\ell-1,k+\ell}}{1-(1-\varepsilon)H_{k,k+\ell}} > \frac{1}{\ell}$. Hence when δ is sufficiently small, $G_{k+\ell-1,k+\ell}(v_k) > 0$.

Sub-case 2: There exists $L > 0$ such that $p_k - p_{k+\ell} \geq L$ as $\delta \to 0$.

Note that in this case, $\ell \to \infty$, as $\delta \to 0$. [15]

[14] Since we know that p_n will be accepted, there is thus a natural bound on how large L can be depending on the price p_k.

[15] Hence, unlike in sub-case 1, $\frac{H_{k+\ell-1,k+\ell}}{H_{k,k+\ell}}$ does not necessarily tend to 1 as $\delta \to 0$.

Note however that the ratio $\left(\frac{1-(1-\varepsilon)H_{k+\ell-1,k+\ell}}{1-(1-\varepsilon)H_{k,k+\ell}}\right)$ is at least as large as ε. Hence the first term, $\left(\sum_{i=0}^{\ell-1}\Delta_{k+i}\right)\left(\frac{1-(1-\varepsilon)H_{k+\ell-1,k+\ell}}{1-(1-\varepsilon)H_{k,k+\ell}}\right)$, remains bounded by $L\varepsilon$ as $\delta \to 0$. Since $\Delta_{k+\ell-1} \to 0$ as $\delta \to 0$, again $G_{k+\ell-1,k+\ell}(v_k) > 0$ when δ is sufficiently small. Contradiction. □

The next lemma shows the useful result that whenever δ is sufficiently small, there is a non-degenerate interval of types of buyers who plan to purchase at every price.[16]

Lemma 6. *For δ sufficiently small, for every price p_k, there exists an interval $[v_k, v_{k-1})$ with $v_k < v_{k-1}$ and $v_0 = 1$ such that the optimal strategy of a type $v \in [v_k, v_{k-1})$ involves rejecting all offers $p_j > p_k$ and accepting all offers $p_j \leq p_k$.*

Proof omitted [17]. We do need to point out here that while the main inductive arguments in the proof follow closely the steps used in the proof of Lemma 5 and both lemmas essentially characterize the same buyer behavior, the reason the above result is placed as a separate lemma is that Lemma 6 also shows the important result that there are no "price-gaps". In other words, for every price p_k there is an interval of buyer types who buy at that price. This along with Proposition 1 that characterizes the buyer types who buy at any price p_k is then used to show the main result, Proposition 3, that trade is almost-efficient.

We are now ready to show the main result of this subsection.

Proposition 1. *In equilibrium all buyer types who plan to buy do so at a price such that their ex post surplus is at most δ.*

Proof. The highest type who buys at price p_k is the type (that is just below) v_{k-1}. Hence, the maximum surplus from buying at price p_k is $v_{k-1} - p_k$. This is equal to δ for $k = 1$. For $k > 1$, we have,

$$v_{k-1} - p_{k-1} = (1-\varepsilon)(v_{k-1} - p_k)H_{k-1}$$

[16]Formally, the purchase price in the phrase "plan to purchase at a price" refers to the highest price at which the buyer-type's optimal strategy involves choosing the action "accept" (and hence the optimal action is to choose "reject" for all higher prices)

[17]Available upon request from the authors

which, given that $\Delta_{k-1} = p_{k-1} - p_k$ can be rewritten as

$$v_{k-1} - p_{k-1} = \Delta_{k-1} \frac{(1-\varepsilon)H_{k-1}}{1-(1-\varepsilon)H_{k-1}}$$

Therefore, since

$$\begin{aligned} & v_{k-1} - p_k \\ = \; & v_{k-1} - p_{k-1} + p_{k-1} - p_k \end{aligned}$$

we have

$$\begin{aligned} & v_{k-1} - p_k \\ = \; & \Delta_{k-1} \frac{(1-\varepsilon)H_{k-1}}{1-(1-\varepsilon)H_{k-1}} + \Delta_{k-1} \\ = \; & \Delta_{k-1} \left[\frac{(1-\varepsilon)H_{k-1}}{1-(1-\varepsilon)H_{k-1}} + 1 \right] \\ < \; & \Delta_{k-1} \left[1 + \frac{1-\varepsilon}{\varepsilon} \right] \\ = \; & \frac{\Delta_{k-1}}{\varepsilon} \\ \leq \; & \delta \end{aligned}$$

where the first inequality follows since $H_{k-1} < 1$ and the last inequality follows from Lemma 1. □

4.3. **Seller's best response.** Given the buyer's best response, we next show the (very intuitive) result that the seller indeed follows the strategy described above.

Proposition 2. *For all types $c > p_n$, let $k(c)$ be the integer such that $p_{k(c)} \geq c > p_{k(c)+1}$; for all $c \leq p_n$ let $p_{k(c)} = p_n$. The seller's optimal strategy is as follows: Type c plans to accept all offer to sell till price $p_{k(c)}$ is reached; in other words, if the buyer has not accepted already, type c withdraws the item from sale if (and only if) the buyer rejects the price offer $p_{k(c)}$.*

Proof. Since $p_{k(c)+\ell} < c$ for all $\ell \geq 1$, clearly the seller type c should have no intention of accepting any offers to sell beyond price $p_{k(c)}$. By rejecting the price offer $p_{k(c)-m}$ for some integer $m \geq 1$, the seller obtains a payoff equal to zero. By accepting it, since there is some chance that the buyer will accept at prices $p_{k(c)-m}$ through $p_{k(c)}$, the expected payoff is positive.[18] □

4.4. **Main Result.** We are now ready to state the main result that with our mechanism it is possible to achieve almost-efficient trade even when the extent of ambiguity (as measured by ε) is small.

Proposition 3. *$\forall \varepsilon > 0$, $\exists \overline{\delta} > 0$ such that for all $\delta < \overline{\delta}$, there exists a mechanism with an equilibrium where trade takes place whenever $v - c > \delta$.*

In other words, presence of even a very small amount of ambiguity allows designing of a mechanism such that inefficiency from absence of trade (when trade is desirable) can be reduced to a negligibly small amount.

Proof. There are two sources of inefficiency. First, buyer types $v \leq p_n$ do not trade at all and whenever seller type c is such that $p_n \geq v > c$, this results in loss of efficiency. Second, even for types $v \geq p_n$, there are situations when v, c are such that $v > c$ but these types do not trade in equilibrium.

For the first type of inefficiency, notice that for any $v, c \in [p_n, 0]$, such that $v > c$, it must be that $v - c < p_n$. Since for any given δ and ε, we have $\lim_{n \to 0} p_n = 0$, hence for any $c \in [p_n, 0]$, the extent of the first source of inefficiency can be made arbitrarily small by increasing the value of n.

As for the second source, note that for any $c > p_n$, the mass of buyer types for which the second type of inefficiency happens is equal

[18]Expected payoff, as usual, means maxmin expected payoff since the seller may have ambiguity about the buyer's type as well. If the seller's ambiguous beliefs are such that the worst distribution puts zero weight on the buyer types who would accept prices $p_{k(c)-m}$ through $p_{k(c)}$ then the seller would be indifferent between accepting and rejecting. Under epsilon contamination, however, and given that the distribution $G(v)$ is such that $g(v) > 0$ for all $v > 0$, it is strictly better for the seller to continue rather than to stop at price $p_{k(c)-m}$.

to $v_{k(c)} - c$. Furthermore, since $c > p_{k(c)+1}$

$$\begin{aligned} & v_{k(c)} - c \\ < \; & v_{k(c)} - p_{k(c)+1} \\ = \; & v_{k(c)} - p_{k(c)} + p_{k(c)} - p_{k(c)+1} \\ \leq \; & \delta + \Delta_{k(c)} \\ \leq \; & 2\delta \end{aligned}$$

which goes to zero as $\delta \to 0$.

Since both sources of inefficiency can be made arbitrarily small, the result follows. □

4.5. **General V and C.** Before we leave this section, let us note that it is straightforward extending the results to more general intervals V and C rather than $[0, 1]$ (but with the intersection containing an interval). So, let us suppose now that V is the interval $[\underline{v}, \overline{v}]$ and C is the interval $[\underline{c}, \overline{c}]$ where $\overline{v} \geq \overline{c} \geq \underline{v} \geq \underline{c} \geq 0$. If the inequalities above are strict, then there are many different price sequences that will result in efficient trade (but with somewhat different implications for surplus extraction from the buyer, depending on what p_n is chosen, that is how far below $\underline{v}$ is p_n.) To be more specific, suppose all the above inequalities are strict (crucially, $\underline{v} > \underline{c}$). First we choose a number, $a \in (\underline{v}, \underline{c})$. Note that we can choose a to be very close to $\underline{v}$. Now define the price sequence as:

$$p_k = \begin{cases} \overline{c} & \text{if } k = 0, \\ (\overline{c} - a - \delta)^k \left(\frac{1}{\overline{c}-a-\delta+\delta\varepsilon}\right)^{k-1} + a & \text{if } k \geq 1. \end{cases}$$

so that we have

$$\Delta_k = \begin{cases} p_0 - p_1 & \text{if } k = 0, \\ p_k - p_{k+1} = \left(\frac{\overline{c}-a-\delta}{\overline{c}-a-\delta+\delta\varepsilon}\right)^k \delta\varepsilon. \end{cases}$$

For given ε and δ, we choose n to be the smallest integer such that $p_n \leq \underline{v}$ and proceed as before[19]. Note that in this setting, the first source of inefficiency is absent since now all types v have a chance to

[19]Since

$$\sum_{k=0}^{\infty} \Delta_k = \overline{c} - a$$

and $a < \underline{v}$, such a n exists.

trade. Also, types $v \in (\bar{c}, \bar{v}]$ may be getting ex post surplus that is larger than δ.

5. Inadequacy of normal form mechanisms

The mechanism proposed by us is an extensive game (form). We argue now that mechanisms that are normal form will in general be inferior to our proposed mechanism. This might seem puzzling since in the standard expected utility case, the direct revelation game – which is atemporal – is sufficient to characterize the optimal mechanism. Put differently, in the standard expected utility framework, the mechanism designer cannot obtain a higher welfare by using some extensive form indirect mechanism as compared to the outcomes obtainable as equilibria of the direct revelation game. The resolution of this puzzle is the observation – which may be of some independent interest – that in the standard setting, it is the revelation principle that guarantees that concentrating on the standard (atemporal) direct revelation game is sufficient when searching for the optimal mechanism. However, in our setting, the standard revelation game is able to pick out only the optimal normal form indirect mechanism. In other words, in the setting with non-EU preferences that we have, the standard revelation principle fails.

We can understand the reason for the inadequacy of normal form mechanisms in a different way as follows. When preferences are not dynamically consistent, it is as if the agent (buyer in our case) has "multiple selves" at various stages. The preferences of these multiple selves are not perfectly aligned and a dynamic mechanism allows the designer to exploit this tension. A standard direct revelation game, which is essentially normal form, amounts to asking *only* the "initial self" to report his type. In essence therefore, it allows the decision maker to *commit* to future course of action that would not be possible if the mechanism was truly dynamic. For dynamically consistent preferences (for example, expected utility preferences) such power to commit has no value and an *atemporal* direct mechanism can be used instead of a truly dynamic one. When preferences are not dynamically consistent, using a normal form mechanism may reduce the set of outcomes that the mechanism designer can implement by using a dynamic mechanism. [20]

6. Conclusion

We see the results obtained in this paper as indicative. They show that in some circumstances, a mechanism designer can exploit the ambiguity aversion on the part of agents even when the ambiguity is in some sense small. However, the results also raise some new, difficult questions that we hope to answer in future research.

Firstly, there is the question of *optimal* mechanisms. We have identified *one* mechanism that is almost efficient. However, we can ask whether this is the best that can be done or whether there are mechanisms that do better. To answer this question, we need a revelation principle for dynamically inconsistent preferences. While there is no such general result, (as far as we are aware), see Bose and Renou [4] for some progress on this question that shows a revelation principle for a class of mechanisms (termed ambiguous mechanisms) that include as a strict subset all the classic static and dynamic indirect mechanisms considered in the literature so far.

Secondly, we would like to understand better the class of environments where the particular type of mechanism used in this paper can be utilized to exploit ambiguity aversion of the agents. Let us indicate briefly why an extension (to more general environments) so not seem very straightforward. As an example, consider the generalization of the Myerson-Satterthwaite model with continuous quantities (see McAfee [9]. When quantity is fixed (for example, as in this paper, when a unit of indivisible item is to be traded), under the mechanism considered in this paper, at each stage the buyer faces a choice between an *unambiguous payoff* (accept the good at the offered price) and an *ambiguous payoff* (reject the offered price, and wait for a lower price with the risk that the good may not be available at all). With continuous quantities, the buyer faces a choice between two ambiguous payoffs because even accepting the current offer essentially means facing an ambiguous payoff (arising from the fact that the quantity traded and hence the

[20] For a related discussion, see [3].

payoff will depend on the seller's type. While one can work out suitable mechanisms for certain parametric specifications of the model, we have so far been unable to work out a general result as the one presented here incorporating a single indivisible item of trade.

In summary, our results point to the need for more work to understand the implications of ambiguity aversion in mechanism design.

References

[1] A. Bodoh-Creed, Ambiguous Beliefs and Mechanism Design, mimeo, 2012

[2] S. Bose, E. Ozdenoren, and A. Pappe, Optimal Auctions with Ambiguity, *Theoretical Economics* 2006, Vol 1, Issue 4, 411-438

[3] S. Bose, and A. Daripa, A Dynamic Mechanism and Surplus Extraction Under Ambiguity, *Journal of Economic Theory*, 2009, Vol 144, Issue 5, 2084-2115

[4] S. Bose and L. Renou, Mechanism Design with Ambiguous Communication Devices, mimeo, 2012

[5] L. De Castro, and N.C. Yannelis, Ambiguity aversion solves the conflict between efficiency and incentive compatibility, mimeo, 2012.

[6] L. Epstein and M. Schneider, Recursive Multiple Priors, *Journal of Economic Theory* 2003, Vol 113, issue 1, 1-31

[7] I. Gilboa and D. Schmeidler, Maximin expected utility with a non-unique prior, *Journal of Mathematical Economics*, 18, 1989, 141–153.

[8] I. Kopylov, Subjective probability and confidence, mimeo, 2009.

[9] R. Preston McAfee, Efficient allocation with continuous quantities, *Journal of Economic Theory*, 53, 1991, 51-74.

[10] R. Myerson and M. Satterthwaite, Efficient mechanisms for bilateral trading, *Journal of Economic Theory*, 29, 1983, 265-281.

[11] C. P. Pires, A Rule For Updating Ambiguous Beliefs, *Theory and Decision* 2002, Vol 52, No 2, 137-152

[12] M. Siniscalchi, Dynamic choice under ambiguity, *Theoretical Economics*, 2011, vol 6, No 3, 379-421

Department of Economics, University of Leicester, United Kingdom
E-mail address: sb345@le.ac.uk

Department of Economics, University of Leicester, United Kingdom
E-mail address: sm403@le.ac.uk