

The Design of Ambiguous Mechanisms*

Alfredo Di Tillio [†]
Bocconi University

Nenad Kos [‡]
Bocconi University

Matthias Messner [§]
Bocconi University and CESifo

November 17, 2011

First version: April 2011

Abstract

This paper considers a mechanism design framework that is standard except for the agents' preferences, which exhibit ambiguity aversion in the sense of [Gilboa and Schmeidler \(1989\)](#). We show that in this framework the optimal strategic form mechanism is an *ambiguous mechanism*. An ambiguous mechanism specifies a message space and a set of outcome functions. After showing that (a version of) the Revelation Principle holds in our environment, we give an exact characterization of the (smallest) optimal ambiguous mechanism for the single agent case. If the type set is composed of N (finite) elements, then the (smallest) optimal ambiguous mechanism contains $N - 1$ outcome functions. We show how the extra revenue that the designer obtains by employing an ambiguous mechanism (as opposed to a standard non-ambiguous mechanism) depends on the parameters of the model. In particular, we prove that the share of the surplus that the designer can extract from the agent increases as the type set becomes larger and the probability of each single type decreases. In the limiting case where the agent's type is drawn from a nonatomic distribution on an interval, the optimal ambiguous mechanism extracts all the rent from the agent.

JEL Code: C72, D44, D82.

Keywords: Optimal mechanism design, Ambiguity aversion, Incentive compatibility, Individual rationality.

*For helpful discussions and comments we thank Eddie Dekel, Eduardo Faingold, Itzhak Gilboa, Matthew Jackson, Alessandro Lizzeri, George Mailath, Massimo Marinacci, Stephen Morris, Frank Riedel, Larry Samuelson, Martin Schneider and Ilya Segal.

[†]Department of Economics and IGIER, Bocconi University, Via Roentgen 1, I-20136 Milano.
E-mail: alfredo.ditillio@unibocconi.it

[‡]Department of Economics and IGIER, Bocconi University, Via Roentgen 1, I-20136 Milano.
E-mail: nenad.kos@unibocconi.it

[§]Department of Economics and IGIER, Bocconi University, Via Roentgen 1, I-20136 Milano.
E-mail: matthias.messner@unibocconi.it

1 Introduction

There is a recent and fast-growing literature that considers mechanism design problems with ambiguity averse players. Examples include [Bose, Ozdenoren, and Pape \(2006\)](#), [Turocy \(2008\)](#), [Bose and Daripa \(2009\)](#), [Bodoh-Creed \(2010\)](#) and [Bose and Renou \(2011\)](#).¹ Typically, the papers in this literature start from the assumption that the agents (and or principal) are uncertain about the other agents' type distribution. They then go on to analyze what implications this has for the optimal design of a mechanism. In doing so, attention is restricted either to direct mechanisms or to simple indirect mechanisms (e.g. standard auction formats). Furthermore, and more importantly, mechanisms are always assumed to be uncertainty free.² That is, by restricting attention to some form of standard Bayesian mechanism, this literature has implicitly ruled out the possibility that the principal exposes the agents to uncertainty by constructing mechanisms that contain some elements of ambiguity.

In this paper we show that in a context with ambiguity-averse agents the restriction to non-ambiguous mechanisms is not without loss of generality. Specifically, we consider a simple screening model—a seller wants to sell a single object to a single buyer—that is standard except for the agent's preferences. Instead of modeling the agent as a classical expected utility maximizer, we assume that his preferences are of the maxmin expected utility (MMEU) type, as in [Gilboa and Schmeidler \(1989\)](#). We show that in this environment the optimal strategic form mechanism comprises an uncertain outcome rule. We refer to mechanisms with this feature as *ambiguous mechanisms*. An uncertain outcome rule is not given by a single (deterministic or stochastic) outcome function. Rather, it consists of a set of outcome functions. When the principal employs an ambiguous mechanism, he does not communicate to the agent the exact outcome function to which he has committed. Instead, he only informs the agent that it belongs to a set of outcome functions (those that belong to the ambiguous mechanism).

By proposing an ambiguous mechanism, the principal exposes the agent to uncertainty regarding the consequences of any given message that he can send. Since the agent has MMEU preferences, it follows that each one of his types associates with each possible message the worst possible outcome that he can obtain under the different outcome functions. Of course, different types evaluate outcomes differently, and hence they may associate different worst case scenarios with a given message. It is exactly this fact that makes the use of ambiguous mechanisms attractive for the principal. By appropriately choosing the ambiguous mechanism, he can (to some extent) manipulate the payoffs that one type of the agent receives by sending a given message, without necessarily changing the payoff that other types associate with that same message. In particular, if the ambiguous mechanism is direct (i.e. if all of its elements are

¹Several models of beliefs and behavior in games that relax the assumption of Bayesian expected-utility maximizing players have been proposed. See e.g. [Azrieli and Teper \(2011\)](#) and the references therein.

²An exception is [Bose and Renou \(2011\)](#). We discuss their paper in more detail below.

direct) this means that the designer can lower the payoff that type θ' associates with sending the message θ'' without necessarily lowering the truth telling payoff of type θ'' .

The arguments in the preceding paragraph presume that the agent believes that the principal might have committed to any of the elements of the ambiguous mechanism. Put differently, it takes for granted that the agent's (set-valued) belief over the set of outcome functions contains at least all degenerate distributions over this set. The assumption that the agent holds such a belief is reasonable if it is compatible with the principal being indifferent between all the elements of the outcome rule. We therefore impose that the principal can only propose mechanisms that satisfy this property. We refer to ambiguous mechanisms that satisfy this condition as *consistent*.

After introducing the concept of ambiguous mechanisms, we formulate and prove a version of the Revelation Principle that is appropriate for our context and objectives. Doing so allows us to restrict attention to direct ambiguous mechanisms. We characterize the smallest optimal direct ambiguous mechanism for the case where the set of possible types of the agent is finite. We show that this mechanism is composed of at most $N - 1$ elements, where N is the number of types, and we provide an exact characterization of these $N - 1$ elements in terms of the agent's type distribution. Furthermore, using this characterization, we prove that the share of surplus that the designer can extract from the agent increases as the type set becomes larger and the probability of each type converges to zero. In the limiting case of a nonatomic type distribution over an interval, the optimal ambiguous mechanism extracts the full surplus from the agent.

To the best of our knowledge, this is the first paper showing that it can be in the designer's interest to introduce uncertainty over outcome functions when agents are ambiguity averse. In a contemporaneous and independent paper, [Bose and Renou \(2011\)](#) also recognize that in such contexts the principal may want to introduce some element of uncertainty into the mechanism that he uses. The two papers are complementary, as they study the impact of ambiguity aversion through quite distinct channels. Unlike in this paper, in their work the uncertainty is not introduced via the outcome functions. Instead, they explore which social choice rules the designer can implement if he engages the agents in a dynamic communication game that he mediates by transforming messages in an ambiguous way. By injecting uncertainty in the exchange of messages between the agents, the principal can manipulate the agents' beliefs about each other's type and hence their behavior. [Bose and Renou \(2011\)](#) remark that the precise extent to which the agents' beliefs can be manipulated depends on the assumed form of (full Bayesian) belief updating. By contrast, our restricting attention to strategic form mechanisms makes the question of what is the most appropriate way to model updating by ambiguity averse individuals—still controversially discussed in the literature—altogether irrelevant in our context. Finally, the ambiguous communication devices in [Bose and Renou \(2011\)](#) serve to manipulate the agents' beliefs over the other agents' types and hence they are ineffective in single

agent environments. Instead, as we show in this paper, (outcome) ambiguous mechanisms have leverage also in the case of a single agent.

2 Framework

Throughout the first part of the paper we consider a mechanism design framework with one agent. The notation and terminology that we introduce below generalizes in the obvious way to the multiple agent setting that we consider in the final parts of the paper.

An allocation is a pair $(x, \tau) \in [0, 1] \times \mathbb{R}$, where x is a (non-monetary) choice variable and τ represents a monetary transfer (to be paid by the agent). With a slight abuse of terminology we will often use the term ‘allocation’ to indicate the non-monetary component x of a pair (x, τ) . The agent’s preferences over $X \times \mathbb{R}$ depend on his (payoff-) type $\theta \in \Theta \subset \mathbb{R}$. More specifically we assume that they can be represented by a utility function of the following form:

$$x\theta - \tau.$$

The agent is ambiguity averse in the sense of [Gilboa and Schmeidler \(1989\)](#). That is, in a situation where his beliefs are described by a family of distributions over allocation-transfer pairs, \mathcal{F} , his utility is given by

$$\inf_{F \in \mathcal{F}} \mathbb{E}_F[x\theta - \tau].$$

The agent’s type is only observed by the agent himself. The principal only knows that the agent’s type is drawn from a distribution P over Θ .

Throughout the first part of the paper we assume that the principal is a risk neutral expected utility maximizer.³ More specifically, we assume that his objective is to maximize the expected transfer that the agent has to pay.

A **simple mechanism** is a triple (S, q, t) , where S is a set of messages for the agent and q and t are functions that map S into X and \mathbb{R} , respectively. That is, for each $s \in S$, q and t specify an allocation-transfer pair. We will therefore refer to q and t respectively as allocation and transfer rules on S . A **direct simple mechanism** is a simple mechanism such that $S = \Theta$. Since all direct mechanisms share the same message space, we drop the latter from the notation and identify the direct mechanism (Θ, q, t) with its outcome function (q, t) .

Notice that our definition of a simple mechanism does not allow for random allocations or transfers; the range of q is X , not the set of probability measures over X . Restricting attention to deterministic mechanisms is without loss of generality in our setting. A random mechanism

³We make this assumption only for convenience. The main insights of our paper extend to situations where also the designer’s preferences are allowed to exhibit ambiguity aversion.

that specifies for each type a distribution over allocations and transfers can be replaced by a deterministic one that assigns to each type the expected values of the corresponding distributions of the random mechanism. Since both the agent and the principal are risk neutral doing so will not alter their expected payoffs.

In the literature on mechanism design it is typically assumed that the principal commits to a specific outcome function and that he announces this outcome function to the agent. In this paper we generalize the second part of this assumption. That is, we continue to assume that the principal makes a commitment to a specific outcome function. But we do not impose that he has to communicate exactly this function to the agent. Instead, we allow for the possibility that he only tells the agent that the outcome function that he has chosen and to which he is committed belongs to a certain set of functions. By announcing a set of outcome functions the planner exposes the agent to uncertainty over outcomes. We refer to a mechanism with outcome uncertainty as **ambiguous mechanism**.

More formally, a (deterministic) ambiguous outcome rule for the message space S is a collection of allocation-transfer functions (q, t) defined on S . If $M(S)$ is the set of all such pairs, then any subset of $M(S)$, is an ambiguous outcome rule. An **ambiguous mechanism** is a pair (S, Ω) , where S is a message space and $\Omega \subset M(S)$ is an ambiguous outcome rule. We say that (S, Ω) is a **direct ambiguous mechanism** if $S = \Theta$. Since in the case of direct mechanisms the message space is always Θ we identify the mechanism (Θ, Ω) with the outcome rule Ω . That is, we will often refer to Ω alone as ambiguous mechanism.

Given an ambiguous mechanism (S, Ω) the agent has to choose an element from S , the set of messages. A strategy of the agent is a function σ that maps Θ into S .

The optimal strategy of the agent depends on his beliefs about the element of Ω to which the principal has committed himself. This belief is allowed to be set valued and we denote by B . More specifically, we assume that B coincides with Ω . Before we discuss this point further it is convenient to introduce some additional notation and terminology.

Given an ambiguous mechanism (S, Ω) , we write $\Sigma^*(S, B)$ for the set of optimal strategies of an agent who holds the belief B . That is, if $\sigma \in \Sigma^*(S, \Omega)$ then for all $\theta \in \Theta$ we have

$$\sigma(\theta) \in \arg \max_{s \in S} \{ \min_{(q,t) \in \Omega} \{q(s)\theta - t(s)\} \}.$$

Definition 1. *The ambiguous mechanism (S, Ω) is **consistent** if there is a $\sigma \in \Sigma^*(S, \Omega)$ such that for all $(q', t'), (q'', t'') \in \Omega$ we have*

$$\mathbb{E}_P[t'(\sigma(\theta))] = \mathbb{E}_P[t''(\sigma(\theta))].$$

Consistency requires that each element of the ambiguous mechanism Ω delivers the same

expected revenue to the principal, if the agent's belief coincides with Ω . Put differently, an ambiguous mechanism is consistent if the principal can be indifferent between all its elements when he assumes that the agent believes that he might have committed to anyone of the ambiguous mechanism's elements.

The reason for imposing consistency is quite straightforward. An agent should not consider all elements of the announced ambiguous outcome rule as real possibilities if under that belief the principal gets a lower payoff from some outcome functions than he gets from others.

The revelation principle

The following result shows that in our model a version of the *revelation principle* holds.

Proposition 1 (Revelation Principle). *For every static ambiguous mechanism (S, Ω) that is consistent and every optimal strategy for the agent, $\sigma \in \Sigma^*(S, \Omega)$, there is a consistent direct ambiguous mechanism (Θ, Ω') , and an optimal strategy for the agent, $\sigma' \in \Sigma^*(\Theta, \Omega')$ such that the following two conditions hold:*

- i) $\sigma'(\theta) = \theta$ for every $\theta \in \Theta$; i.e. under σ' the agent reports his true type;
- ii) $(q, t) \in \Omega$ if and only if $(q', t') \in \Omega'$ where (q', t') is defined by

$$\begin{aligned} q'(\theta) &= q(\sigma(\theta)), \\ t'(\theta) &= t(\sigma(\theta)). \end{aligned}$$

Proof. Omitted. □

The revelation principle as formulated above tells us that if we want to find an optimal ambiguous mechanisms we can restrict ourselves to deterministic and direct ambiguous mechanisms that are chosen such that the agent is willing to reveal his type truthfully. Since in the case of a direct ambiguous mechanism the message space is clear we drop it from the notation and write only Ω for an ambiguous mechanism.

3 Optimal ambiguous mechanisms with finitely many types

In this section we solve the designer's optimal mechanisms design problem for the case where the type set is finite. Thus, throughout this section we assume that the type set is $\Omega = \{\theta_1, \theta_2, \dots, \theta_N\}$, where $\theta_{n+1} > \theta_n$ for all $n = 1, \dots, N - 1$. With a slight abuse of notation we denote the index set by N . p_n indicates the probability of type θ_n . Finally, we write q_n and t_n for $q(\theta_n)$ and $t(\theta_n)$, respectively.

The designer's problem is to choose a revenue maximizing ambiguous mechanism from the set of ambiguous mechanisms that are consistent and in which the agent is willing to participate voluntarily. By the Revelation Principle we know that we can restrict our attention to direct consistent ambiguous mechanisms that are deterministic and for which truth telling is an optimal strategy for the agent.

In order to give a formal statement of the designer's optimization problem it is convenient to introduce the following definitions.

Definition 2 (Incentive Compatibility). *The ambiguous mechanism Ω satisfies the incentive compatibility constraint of type $\theta_n \in \Theta$ with respect to type $\theta_m \in \Theta$ if*

$$\inf_{(q,t) \in \Omega} \{q_n \theta_n - t_n\} \geq \inf_{(q,t) \in \Omega} \{q_m \theta_n - t_m\}. \quad (\text{IC}(n,m))$$

The ambiguous mechanism Ω is **downward incentive compatible** if $\text{IC}(n,m)$ holds for all $n, m \in N$ such that $n > m$; it is **upward incentive compatible** if $\text{IC}(n,m)$ is satisfied for all $n, m \in N$ such that $n < m$. Finally, Ω is **incentive compatible** if it is both downward and upward incentive compatible.

Incentive compatibility of Ω guarantees that for the agent it is optimal to reveal his type truthfully. We also assume that the agent cannot be forced to participate in the mechanism. If we assume that the value of not participating in the mechanism is constant across types and normalized to zero, then this means that the following condition must hold.

Definition 3 (Individual Rationality). *An ambiguous mechanism Ω is individually rational for type $\theta_n \in \Theta$ if*

$$\inf_{(q,t) \in \Omega} \{q_n \theta_n - t_n\} \geq 0. \quad (\text{IR}(n))$$

Ω is **individually rational** if condition $\text{IR}(n)$ holds for all $n \in N$.

Exploiting the Revelation Principle and using the above definitions, the designer's problem can be written as follows.

$$\begin{aligned} \max_{R \in \mathbb{R}, \Omega \subset M(\Theta)} \quad & R & (\text{P}) \\ \text{s.t.} \quad & R = \mathbb{E}[t(\theta)] \quad \text{for all } (q, t) \in \Omega \\ & \text{IR}(n) \quad \text{for all } n \in N \\ & \text{IC}(n,m) \quad \text{for } n, m \in N. \end{aligned}$$

The first constraint in Problem (P) requires that the solution of the designer's problem must

satisfy consistency. Through the second and third constraint we impose individual rationality and incentive compatibility.

The first result shows that whenever incentive compatibility holds, then individual rationality is satisfied for all types if and only if it holds for the lowest type.

Lemma 1. *If Ω satisfies $IC(n, n-1)$ for all $n = 2, \dots, N$ and $IR(1)$, then Ω also satisfies $IR(n)$.*

Proof. Suppose that Ω satisfies $IR(1)$ and $IC(n, n-1)$ for $n = 2, \dots, N$. Then, for all $n \geq 2$ we have

$$\inf_{(q,t) \in \Omega} \{q_n \theta_n - t_n\} \stackrel{IC(n, n-1)}{\geq} \inf_{(q,t) \in \Omega} \{q_{n-1} \theta_n - t_{n-1}\} \geq \inf_{(q,t) \in \Omega} \{q_{n-1} \theta_{n-1} - t_{n-1}\}.$$

But if the truth-telling payoff is weakly increasing in the agent's type and the lowest type's individual rationality constraint is satisfied, then individual rationality must hold for type θ_m too, $m = 2, \dots, n$. \square

In the following we show that an ambiguous mechanism that satisfies individual rationality and downward incentive compatibility can be transformed into an ambiguous mechanism that is fully incentive compatible ambiguous mechanism without any loss in the expected revenue that it generates. We also show that for the designer it is without loss if he restricts himself to ambiguous mechanisms that contain at most $N - 1$ elements. In order to arrive at these results we proceed in a number of intermediate steps.

We start by showing that we can restrict attention to ambiguous mechanisms that are finite.

Lemma 2. *For every downward incentive compatible and individually rational ambiguous mechanism Ω there is a downward incentive compatible and individually rational ambiguous mechanism Ω' that contains at most $N(N - 1)/2$ elements and that generates the same revenue as Ω .*

Proof. Define Ω' as follows. For each pair $n, m \in N$, $n \neq m$, Ω' contains exactly one element from the set

$$M_{n,m} = \arg \min_{(q,t) \in \bar{\Omega}} \{q_m \theta_n - t_m\},$$

where $\bar{\Omega}$ is the closure of Ω . Notice that $M_{n,m}$ is non-empty since $\bar{\Omega}$ is not only closed but also bounded. Boundedness is obvious in the dimension of the allocations which belong to $[0, 1]^N$. Transfers are bounded above by the individual rationality condition. Moreover, by consistency all transfers in Ω have the same expected value. But if transfers are bounded above and have the same expected value, then they must also be bounded below.

Since Ω' is a subset of the closure of Ω it is consistent and generates the same revenue as Ω . Moreover, by construction the deviation payoffs in Ω' and Ω are the same. But since the

truth telling payoffs must be weakly higher in Ω' . Thus, Ω' is both incentive compatible and individually rational. \square

The next result shows that we can restrict attention to ambiguous mechanisms that exhibit the characteristic that for each type each element of the ambiguous mechanism yields the same truth telling payoff. We will refer to this property as **uniformity**.

Lemma 3 (Uniformity). *For every downward incentive compatible and individually rational ambiguous mechanism Ω there exists a downward incentive compatible and individually rational ambiguous mechanism $\tilde{\Omega}$ that delivers the same revenue as Ω and that satisfies the conditions*

$$q_1\theta_1 - t_1 = 0 \quad \text{for all } (q, t) \in \tilde{\Omega} \quad (3)$$

$$q_n\theta_n - t_n = \max_{m < n} \left\{ \min_{(q, t) \in \tilde{\Omega}} q_m\theta_m - t_m \right\} \quad \text{for all } (q, t) \in \tilde{\Omega} \text{ and } N > n > 1. \quad (4)$$

Proof. Individual rationality requires that

$$q_1\theta_1 - t_1 \geq 0 \quad \text{for all } (q, t) \in \Omega.$$

If Ω contains an element (q', t') for which $q'_1\theta_1 - t'_1 > 0$, then we can replace it by (q'', t'') that coincides with (q', t') in all components except for the transfers of the lowest and the highest type that are set as follow:

$$\begin{aligned} t''_1 &= q'_1\theta_1 \\ t''_N &= t'_N - \frac{p_1}{p_N}[t''_1 - t'_1]. \end{aligned}$$

The payment of the lowest type is increased so that the truth telling payoff of this type is equal to zero. The transfer of the highest type is lowered so that the expected value of t'' coincides with the expected value of t' . Thus the new ambiguous mechanism that we obtain after the replacement of (q', t') also satisfies consistency and generates the same expected revenue as Ω . Furthermore, the change in the transfers can at most relax the downward incentive constraints. If necessary, we can repeat the procedure described in the previous paragraph for other mechanisms in Ω for which the truth telling payoff of the lowest type is different from zero.

We can apply the same reasoning also to $N > n > 1$: If Ω contains an element, (q, t) , such that the truth telling payoff of type n is strictly larger than the payoff that he can realize from any downward deviation, then we can increase t_n without violating his (downward) incentive constraints. In order to keep the expected revenue constant we lower the transfer of the highest type accordingly. Neither of these two changes affects the downward incentive constraints for types $\theta_m < \theta_n$. The downward incentive constraints for types $\theta_m > \theta_n$ are at

most relaxed. Finally, since also the individual rationality constraint of the lowest type is not affected it follows by Lemma 1 that all individual rationality constraints must continue to hold after the modification. \square

Lemma 3 will be used in the proof of the next result. Lemma 4 tells us that everything that can be done (in terms of revenue) with downward incentive compatible and individually rational ambiguous mechanisms can be achieved by ambiguous mechanisms that contain at most $N - 1$ simple mechanisms. Thus, this result further strengthens Lemma 2.

Lemma 4. *For every downward incentive compatible and individually rational ambiguous mechanism, Ω , there exists a downward incentive compatible and individually rational ambiguous mechanism, $\tilde{\Omega}$, that delivers the same expected revenue as Ω and that is composed of at most $N - 1$ elements.*

Proof. By Lemma 2 we know that it is without loss of generality to assume that Ω is composed of at most $N(N - 1)/2$ elements. Moreover, by Lemma 3 we can assume that truth telling payoffs are uniform across Ω and either equal to the highest payoff that can be achieved by announcing a lower type (this holds for types θ_n , $n > 1$) or equal to the value of the outside option (for type θ_1).

Recall that $M_{j,k}$ denotes the set of solutions of $\min_{(q,t) \in \Omega} \{q_k \theta_j - t_k\}$. Let n, m, l be such that $n, m > l$. We will show now that $(q', t') \in M_{n,l}$ and $(q'', t'') \in M_{m,l}$ implies

$$(q'_l, t'_l) = (q''_l, t''_l).$$

That is, the simple mechanisms (q', t') and (q'', t'') coincide in their l -th component. But this means that if there is no other pair, $(j, k) \neq (n, l)$, $j > k$, such that (q', t') is the unique element of $M_{j,k}$, then we can eliminate (q', t') from Ω without altering the fact that the agent has no incentive to under-report his type. Since a similar observation applies to all other elements of Ω it follows that Ω can be reduced to at most $N - 1$ elements without violating downward incentive compatibility or individual rationality.

By the definition of (q', t') and (q'', t'') we have that

$$\begin{aligned} q'_l \theta_n - t'_l &\leq q''_l \theta_n - t''_l \\ q''_l \theta_m - t''_l &\leq q'_l \theta_m - t'_l. \end{aligned} \tag{5}$$

We also know that

$$q'_l \theta_l - t'_l = q''_l \theta_l - t''_l. \tag{6}$$

Using (6) we can rewrite the inequalities (5) as follows

$$\begin{aligned}(q'_l - q''_l)(\theta_n - \theta_l) &\leq 0, \\ (q''_l - q'_l)(\theta_m - \theta_l) &\leq 0.\end{aligned}$$

Clearly, these two conditions can hold simultaneously only if $q'_l = q''_l$. By (6) it then follows that we must also have $t'_l = t''_l$. \square

In what follows we write (q^l, t^l) for the mechanism which defines the payoffs for downward deviations toward type θ_l .

Lemma 5 (Monotonicity). *For every downward incentive compatible and individually rational ambiguous mechanism $\tilde{\Omega}$ there exists a downward incentive compatible and individually rational ambiguous mechanism Ω that yields the same revenue as $\tilde{\Omega}$, is composed of at most $N - 1$ elements and satisfies the following two conditions:*

- i) *Monotonicity: For all $1 < n < N$, $q_n^n \geq q_{n-1}^{n-1}$;*
- ii) *Maximal slackness: For all $n, m \in N, n \neq m, q_n^m = 1$.*

Proof. By Lemma 4 we know that there must be a downward incentive compatible and individually rational ambiguous mechanism, Ω that is composed of at most $N - 1$ elements and that yields the same revenue as $\tilde{\Omega}$. Moreover, by Lemma 3 we can assume that for each type θ_n , $n < N$, the truth telling payoffs for all elements of Ω are either all equal to the payoff generated by the most profitable downward deviation ($1 < n < N$) or they are all equal to the outside option payoff ($n = 1$).

Suppose that this Ω does not satisfy the monotonicity condition. Let n be the smallest index for which the monotonicity condition is violated, i.e. for which we have

$$q_n^n < q_{n-1}^{n-1}.$$

We will argue that this implies that if $l \leq n - 1$, then for every type $\theta_m > \theta_l$, θ_l is the most tempting option for a deviation in the set $\{\theta_1, \theta_2, \dots, \theta_l\}$. This in particular implies that for all types $l \leq n$ the binding downward incentive constraint of type θ_l must be the one with respect to the next lowest type, θ_{l-1} .

In order to see this notice that for every $l \leq n - 1$ and $m > l$ we have

$$\begin{aligned} & [q_l^l \theta_m - t_l^l] - [q_{l-1}^{l-1} \theta_m - t_{l-1}^{l-1}] = \\ & [q_l^l (\theta_m - \theta_l) + q_l^l \theta_l - t_l^l] - [q_{l-1}^{l-1} (\theta_m - \theta_l) + q_{l-1}^{l-1} \theta_l - t_{l-1}^{l-1}] = \\ & [q_l^l - q_{l-1}^{l-1}] [\theta_m - \theta_l] + \{ [q_l^l \theta_l - t_l^l] - [q_{l-1}^{l-1} \theta_l - t_{l-1}^{l-1}] \} \geq 0. \end{aligned}$$

That the inequality sign in the last line must hold can be seen as follows. By assumption we have $q_l^l - q_{l-1}^{l-1} \geq 0$. Since $m > l$ we also have $\theta_m - \theta_l > 0$ and thus the product of these two expressions must be non-negative. The expression in the curly brackets is non-negative by incentive compatibility.

Next consider the downward deviation incentives of type θ_m , $m > n$, with respect to θ_n and θ_{n-1} . Since for type θ_n the binding downward incentive constraint is IC(n,n-1) we have

$$\begin{aligned} & [q_n^n \theta_m - t_n^n] - [q_{n-1}^{n-1} \theta_m - t_{n-1}^{n-1}] = \\ & [q_n^n - q_{n-1}^{n-1}] [\theta_m - \theta_n] + \{ [q_n^n \theta_n - t_n^n] - [q_{n-1}^{n-1} \theta_n - t_{n-1}^{n-1}] \} = \\ & [q_n^n - q_{n-1}^{n-1}] [\theta_m - \theta_n]. \end{aligned}$$

Thus, IC(m,n) cannot be the binding incentive constraint for θ_m , $m > n$, as long as $q_n^n - q_{n-1}^{n-1} < 0$.

Consider now the ambiguous mechanism, Ω' , that coincides with Ω , except for the value of q_n^n , which is set equal to q_{n-1}^{n-1} . Clearly, since Ω' has the same transfers as Ω , it must be consistent and it must generate exactly the same revenue as Ω . Furthermore, the increase of q_n^n can at most increase the agent's payoff and so Ω' must also be individually rational. As for incentive compatibility, notice that by construction under Ω' every type $\theta_m > \theta_n$ obtains the same payoff from announcing θ_n as he gets by announcing θ_{n-1} . Since the incentive to announce lower types different from θ_n are the same in Ω and Ω' it thus follows that Ω' is downward incentive compatible.

By iterating on this argument we obtain a downward incentive compatible and individually rational ambiguous mechanism that satisfies monotonicity.

As for the maximal slackness condition notice that an increase of q_n^n does not affect the deviation payoffs of the ambiguous mechanism. Also, since truth telling payoffs are assumed to be uniform across the elements of Ω , increasing q_n^n does not alter the truth telling payoffs either. Thus, we can assume that Ω satisfies maximal slackness. \square

The next result justifies our initial decision to ignore the upward incentive compatibility conditions.

Lemma 6. *Let $\Omega = \{(q^1, t^1), \dots, (q^{N-1}, t^{N-1})\}$ be a consistent, downward incentive compatible and individually rational ambiguous mechanism, where (q^l, t^l) , $1 \leq l \leq N - 1$, denotes the*

the condition

$$\begin{aligned}
t_1 &= q_1 \theta_1 && \text{for all } (q, t) \in \Omega, \\
t_{l+1} &= (q_{l+1} - q_l^l) \theta_{l+1} + t_l^l && \text{for all } l = 1, \dots, N-2, \text{ and } (q, t) \in \Omega \\
t_N &\leq (1 - q_{N-1}^{N-1}) \theta_{N-1} + t_{N-1}^{N-1}
\end{aligned} \tag{7}$$

or equivalently,

$$\begin{aligned}
t_l &= q_l \theta_l - \sum_{k=1}^{l-1} q_k^k (\theta_{k+1} - \theta_k) && \text{for all } l = 1, \dots, N-1, \text{ and } (q, t) \in \Omega \\
t_N &\leq \theta_N - \sum_{k=1}^{N-1} q_k^k (\theta_{k+1} - \theta_k)
\end{aligned} \tag{8}$$

Then the expected revenue from the l -th element of Ω is bounded above by

$$T^l(Q) = \mathbb{E}[\theta] - p_l (1 - q_l^l) \theta_l - \sum_{k=1}^{N-1} q_k^k (1 - P_k) (\theta_{k+1} - \theta_k),$$

where $Q = (q_1^1, \dots, q_{N-1}^{N-1})$ and $P_k = \sum_{j=1}^k p_j$. For later reference we note here that T^l is decreasing in q_k^k , $k \neq l$.

Define

$$T(Q) = \min_{1 \leq l \leq N-1} T^l(Q).$$

Since we have to satisfy consistency, the maximal expected revenue cannot exceed the maximum of T . We therefore now analyze which Q solves the problem

$$\hat{T} := \max_{Q \in \mathcal{Q}} T(Q), \tag{9}$$

where \mathcal{Q} is the set of all $Q \in [0, 1]^{N-1}$ with increasing components.

For what follows it is convenient to introduce some further notation and terminology. Let the set $M = \{m_1, \dots, m_M, m_{M+1}\}$ be such that the following two conditions are satisfied

- i) $m_{M+1} = N$ and $m_1 = 1$.
- ii) For all $1 < j \leq M$, m_j is the smallest index l , $m_{j-1} < l \leq N-1$, which satisfies the condition $p_{m_{j-1}} \theta_{m_{j-1}} < p_{m_j} \theta_{m_j}$.

Notice that if p_n is increasing in n , then the set M coincides with N .

Next recall that the *virtual valuation* of type l , v_l is defined as

$$v_l = \theta_l - (\theta_{l+1} - \theta_l) \frac{1 - P_l}{p_l}.$$

For all $1 \leq j \leq M$ we define the *adjusted virtual valuation*, \bar{v}_{m_j} , as follows

$$\bar{v}_{m_j} = p_{m_j} \theta_{m_j} - \sum_{s=j}^M \frac{p_{m_j} \theta_{m_j}}{p_{m_s} \theta_{m_s}} \sum_{i=m_s}^{m_{s+1}-1} (1 - P_i) (\theta_{i+1} - \theta_i).$$

We will explain the meaning and usefulness of the adjusted virtual valuation in the course of the proof of the following proposition, which constitutes the main result of the current section.

Proposition 2. *Let K be the set of all elements of M for which $\bar{v}_{m_j} \leq 0$. If K is non-empty, denote its largest element by m_k .*

i) *If $1 \leq m_k \leq M$, then T is maximized by $\hat{Q} = (\hat{q}_1^1, \dots, \hat{q}_{N-1}^{N-1})$, where*

$$\hat{q}_l^l = \begin{cases} 0 & \text{if } l < m_{k+1} \\ 1 - \frac{p_{m_k} \theta_{m_k}}{p_{m_j} \theta_{m_j}} & \text{if } m_j \leq l < m_{j+1} \text{ and } k+1 \leq j \leq M. \end{cases}$$

ii) *If K is empty, then $(\hat{q}_1^1, \dots, \hat{q}_{N-1}^{N-1}) = (1, \dots, 1)$ is optimal.*

Proof. In order to prove this we proceed in several steps. In the first step we show that the problem of choosing $(q_1^1, \dots, q_{N-1}^{N-1})$ can easily be reduced to a problem where only $(q_{m_1}^{m_1}, \dots, q_{m_M}^{m_M})$ have to be chosen.

Step 1. If $m_j < l < m_{j+1}$, $1 \leq j \leq M$, then at the optimum $\hat{q}_l^l = \hat{q}_{m_j}^{m_j}$.

In order to see this observe that since for every $Q \in \mathcal{Q}$ we have $q_l^l \geq q_{m_j}^{m_j}$ it follows that

$$T^l(Q) - T^{m_j}(Q) = -p_l \theta_l (1 - q_l^l) + p_{m_j} \theta_{m_j} (1 - q_{m_j}^{m_j}) \geq (1 - q_{m_j}^{m_j}) (p_{m_j} \theta_{m_j} - p_l \theta_l) \geq 0.$$

That is, for no admissible Q is T^l the unique minimal bound that defines T . But T^l is the only bound that could be increasing in q_l^l . Thus, it is optimal to choose q_l^l as small as possible, i.e. $\hat{q}_l^l = \hat{q}_{l-1}^{l-1}$. Since this applies to all $m_j < l < m_{j+1}$ we can conclude that $\hat{q}_l^l = \hat{q}_{m_j}^{m_j}$ for all $m_j < l < m_{j+1}$.

Step 2. At the optimum

$$\hat{q}_{m_{j+1}}^{m_{j+1}} \leq 1 - \frac{p_{m_j} \theta_{m_j}}{p_{m_{j+1}} \theta_{m_{j+1}}} (1 - \hat{q}_{m_j}^{m_j})$$

for all $1 \leq j \leq M - 1$.

In order to see this, notice that for every Q such that

$$q_{m_{j+1}}^{m_{j+1}} > 1 - \frac{p_{m_j} \theta_{m_j}}{p_{m_{j+1}} \theta_{m_{j+1}}} (1 - q_{m_j}^{m_j})$$

we have

$$T^{m_{j+1}}(Q) - T^{m_j}(Q) > 0.$$

Moreover, rewriting the inequality we get

$$q_{m_{j+1}}^{m_{j+1}} - q_{m_j}^{m_j} > \left(1 - \frac{p_{m_j} \theta_{m_j}}{p_{m_{j+1}} \theta_{m_{j+1}}}\right) (1 - q_{m_j}^{m_j}) \geq 0.$$

In such a case we could lower $q_{m_{j+1}}^{m_{j+1}}$ without violating the constraint $q_{m_{j+1}}^{m_{j+1}} \geq q_{m_j}^{m_j}$, and thus increase all T^l , $l \neq m_{j+1}$. Since $T^{m_{j+1}}$ is not the smallest bound this means that T would increase. But then Q cannot be optimal.

Step 3. If K is non-empty then at the optimum

$$\hat{q}_{m_j}^{m_j} = 1 - \frac{p_{m_{j-1}} \theta_{m_{j-1}}}{p_{m_j} \theta_{m_j}} (1 - \hat{q}_{m_{j-1}}^{m_{j-1}}) \quad \text{for all } k < j \leq M$$

If K is empty then this condition holds for all $1 < j \leq M$.

By Step 2 we know that at the optimum

$$\hat{q}_{m_j}^{m_j} \leq 1 - \frac{p_{m_{j-1}} \theta_{m_{j-1}}}{p_{m_j} \theta_{m_j}} (1 - \hat{q}_{m_{j-1}}^{m_{j-1}}),$$

for all $1 < j \leq M$ or equivalently

$$T^{m_j}(\hat{Q}) \leq T^{m_{j-1}}(\hat{Q}).$$

Now suppose that \hat{Q} is such that this condition holds with strict inequality for $j = M$. Then, $T^{m_M}(\hat{Q})$ is strictly smaller than any other bound. If \hat{Q} is optimal then it should not be possible to increase T^{m_M} . An increase of T^{m_M} can be achieved only if $q_{m_M}^{m_M}$ is increased. On the other hand, since for all $m_M < l < N$ we have $\hat{q}_l^l = \hat{q}_{m_M}^{m_M}$, $q_{m_M}^{m_M}$ can be increased without violating monotonicity only if at the same time we also increase q_l^l . The impact of a uniform increase of $(q_{m_M}^{m_M}, \dots, q_{N-1}^{N-1})$ on T^{m_M} is

$$p_{m_M} \theta_{m_M} - \sum_{i=m_M}^{N-1} (1 - P_i)(\theta_{i+1} - \theta_i) = \bar{v}_{m_M}.$$

Thus \hat{Q} cannot be optimal if $\bar{v}_{m_M} > 0$. This proves the claim for $k < j = M$.

For the case $k < j < M$. Assume, that we have shown the claim for $s = j + 1, \dots, M$. If \hat{Q} is such that

$$\hat{q}_{m_j}^{m_j} < 1 - \frac{p_{m_{j-1}}\theta_{m_{j-1}}}{p_{m_j}\theta_{m_j}}(1 - \hat{q}_{m_{j-1}}^{m_{j-1}})$$

then

$$T^{m_M}(\hat{Q}) = \dots = T^{m_{j+1}}(\hat{Q}) = T^{m_j}(\hat{Q}) < T^{m_{j-1}}(\hat{Q}) \leq \dots \leq T^{m_1}(\hat{Q}).$$

The assumption that the claim holds for $s = j + 1, \dots, M$ implies that

$$\begin{aligned} q_{m_s}^{m_s} &= 1 - \frac{p_{m_{s-1}}\theta_{m_{s-1}}}{p_{m_s}\theta_{m_s}}(1 - q_{m_{s-1}}^{m_{s-1}}) = 1 - \frac{p_{m_{s-1}}\theta_{m_{s-1}}}{p_{m_s}\theta_{m_s}} \left[1 - \left(1 - \frac{p_{m_{s-2}}\theta_{m_{s-2}}}{p_{m_{s-1}}\theta_{m_{s-1}}} (1 - q_{m_{s-2}}^{m_{s-2}}) \right) \right] \\ &= \frac{p_{m_{s-2}}\theta_{m_{s-2}}}{p_{m_s}\theta_{m_s}} (1 - q_{m_{s-2}}^{m_{s-2}}) = \dots \\ &= 1 - \frac{p_{m_j}\theta_{m_j}}{p_{m_s}\theta_{m_s}} (1 - q_{m_j}^{m_j}) \end{aligned}$$

Moreover, by Step 1 we know that for $m_{s-1} < l < m_s$, $s = j + 1, \dots, M$,

$$q_l^l = q_{m_s}^{m_s}.$$

Thus, if starting from \hat{Q} we want to increase $q_{m_j}^{m_j}$ then monotonicity combined with the fact that the claim holds for all $s = j + 1, \dots, M$ implies that we must increase q_l^l , $m_{s-1} \leq l < m_s$, $s = j + 1, \dots, M$, at the rate

$$\frac{p_{m_j}\theta_{m_j}}{p_{m_{s-1}}\theta_{m_{s-1}}}.$$

If $(q_{m_j}^{m_j}, \dots, q_{N-1}^{N-1})$ is increased in this way then T^{m_j} changes at the following rate

$$p_{m_j}\theta_{m_j} - \sum_{s=j}^M \frac{p_{m_j}\theta_{m_j}}{p_{m_s}\theta_{m_s}} \sum_{i=m_s}^{m_{s+1}-1} (1 - P_i)(\theta_{i+1} - \theta_i) = \bar{v}_{m_j}.$$

Thus, if $\bar{v}_{m_j} > 0$ then \hat{Q} cannot be optimal.

Step 4. If K is non-empty, then at the optimum $\hat{q}_{m_j}^{m_j} = 0$ for all $j \leq k$.

Consider first the case $j = k$. From Step 3 we know that for all $s = k + 1, \dots, M$ the condition

$$\hat{q}_{m_s}^{m_s} = 1 - \frac{p_{m_k}\theta_{m_k}}{p_{m_s}\theta_{m_s}}(1 - \hat{q}_{m_k}^{m_k}) \quad (10)$$

holds. Thus, varying $q_{m_k}^{m_k}$ implies that we have to change accordingly also all q_l^l , $m_k < l < N$. In the previous step we have seen that the overall effect of such a change on T^{m_k} is measured

by \bar{v}_{m_k} . Thus, if $\bar{v}_{m_k} \leq 0$, then T^{m_k} is maximized by choosing $q_{m_k}^{m_k}$ as small as possible. But that means that we have to set $q_{m_k}^{m_k} = q_{m_{k-1}}^{m_{k-1}}$.

Next, observe that if at Q we have $q_{m_k}^{m_k} = q_{m_{k-1}}^{m_{k-1}}$, then

$$T^{m_k}(Q) - T^{m_{k-1}}(Q) = (1 - q_{m_{k-1}}^{m_{k-1}})(p_{m_{k-1}}\theta_{m_{k-1}} - p_{m_k}\theta_{m_k}) < 0.$$

That is, $T^{m_{k-1}}$ cannot be the smallest bound and so if at Q we have $q_{m_{k-1}}^{m_{k-1}} > q_{m_{k-2}}^{m_{k-2}}$ lowering $q_{m_{k-1}}^{m_{k-1}}$ would lead to an increase of T .

Iterating on the same argument we can show that for all $m_j \leq m_k$, $q_{m_j}^{m_j}$ must be chosen as small as possible. Since for m_1 this means $q_{m_1}^{m_1} = 0$ we thus get $q_{m_j}^{m_j} = 0$ for all $m_j \leq m_k$.

Step 5. If K is empty, then at the optimum we must have $q_{m_j}^{m_j} = 1$ for all $1 \leq j \leq M$.

In Step 3 we have seen that if $\bar{v}_{m_j} > 0$ for all $k < j \leq M$ then each $q_{m_j}^{m_j}$ has to be chosen as large as the constraint

$$\hat{q}_{m_j}^{m_j} \leq 1 - \frac{p_{m_{j-1}}\theta_{m_{j-1}}}{p_{m_j}\theta_{m_j}}(1 - \hat{q}_{m_{j-1}}^{m_{j-1}}) \quad (11)$$

allows. Since there is no such constraint for $j = 1$ it follows that $q_{m_1}^{m_1}$ must be optimally set equal to 1. Monotonicity then requires that also q_l^l , $1 < l < N - 1$, must be equal to 1.

□

Proposition 2 gives us a solution for the problem

$$\hat{T} = \max_{Q \in \mathcal{Q}} T(Q).$$

As for the value of this problem Proposition 2 implies that $\hat{T} = T^{m_k}(\hat{Q})$, whenever m_k is well defined, and $\hat{T} = \theta_1$ if the adjusted virtual valuations of all types θ_{m_j} , $1 \leq j \leq M$ are all strictly positive. More precisely, we have

$$\hat{T} = \begin{cases} \mathbb{E}[\theta] - p_{m_k}\theta_{m_k} - \sum_{s=m_{k+1}}^{N-1} q_s^s(1 - P_s)(\theta_{s+1} - \theta_s) & \text{if } 1 \leq m_k \leq M \\ \theta_1 & \text{else.} \end{cases}$$

Individual transfers can be recovered from (8), with the exception of the highest type's payments. The transfers of type N can be obtained as residual between the \hat{T} and the sum of the transfers of the other types, i.e.

$$t_N^l = \hat{T} - \sum_{n=1}^{N-1} t_n^l.$$

4 Discussion

4.1 Surplus extraction: Ambiguous mechanisms vs simple mechanisms

Since the designer does have the option to pick a non-ambiguous simple mechanism, it follows that he will choose an ambiguous mechanism only if doing so does not make him worse off. That is, the optimal ambiguous mechanism must allow him to extract at least as much revenue as does the best simple mechanism.

At the end of the preceding section we have observed that if $\bar{v}_{m_j} > 0$ for all $1 \leq j \leq M$, then $q_l^l = 1$ for all $1 \leq l \leq N$. Since every optimal ambiguous mechanism also satisfies the condition $q_n^l = 1$ for all $n, l \leq N - 1, n \neq l$, it follows that in this case all elements of $\hat{\Omega}$ are identical. That is, $\hat{\Omega}$ is a singleton set. So in this case the possibility of introducing ambiguity does not allow the mechanism designer to do better than with the best simple mechanism. Of course, the condition $\bar{v}_{m_j} > 0$ for all $1 \leq j \leq M$, can be relaxed to $\bar{v}_{m_j} \geq 0$ for all $1 \leq j \leq M$.

Conversely, we have seen that each time the condition $\bar{v}_{m_j} > 0$ for all $1 \leq j \leq M$, does not hold, the optimal ambiguous mechanism derived in Proposition 2 is not a singleton. Thus, it is natural to conjecture that in this case the designer can obtain a higher revenue by employing an ambiguous mechanism than what he can achieve with simple mechanisms. The following proposition shows that this conjecture is correct provided \bar{v}_{m_k} is strictly negative.

Proposition 3. *An optimal ambiguous mechanism yields a strictly larger revenue than the best simple mechanism if and only if it is not optimal to set $Q = 1$.*

Proof. Necessity: If $Q = 1$ then all elements of the ambiguous mechanism are identical. Thus, saying that $Q = 1$ is optimal amounts to saying that the revenue maximum can be achieved by a simple mechanism.

Sufficiency: If $Q = 1$ is not optimal, we distinguish two cases. First, assume that $(q, t) = (1, \theta_1)$ is an optimal simple mechanism. Then the optimal ambiguous mechanism must yield a strictly larger revenue than $(q, t) = (1, \theta_1)$, because $(q, t) = (1, \theta_1)$ can be replicated within the set of ambiguous mechanism that we consider.

So it remains to be shown that sufficiency also holds when $(q, t) = (1, \theta_1)$ is not an optimal simple mechanism. Let (q^*, t^*) be an optimal simple mechanism and let k be the largest element of $\{n : q_n^* < 1\}$. Moreover, assume $m \leq k$ is such that $p_m \theta_m (1 - q_m^*) \leq p_n \theta_n (1 - q_n^*)$ for all $n = 1, \dots, k$. Next consider the ambiguous mechanism, $\hat{\Omega}$, that is composed of k elements that

are defined as follows.

$$q_n^l = \begin{cases} q_n^* & \text{if } n = l \\ 1 & \text{else,} \end{cases} \quad t_n^l = \begin{cases} t_n^* & \text{if } l = n \text{ or } [l = m, n = N] \\ \theta_n^* + (1 - q_n^*)\theta_n & \text{if } [l \neq n, n < N] \\ t_n^m - [p_l\theta_l(1 - q_l^*) - p_m\theta_m(1 - q_m^*)]/p_N & \text{if } l \neq m, n = N \end{cases}$$

This ambiguous mechanism is consistent by construction. Also, for every type $n < N$ the payoff from truth telling is constant across all simple mechanisms in $\hat{\Omega}$ and equal to the truth telling payoff under the simple mechanism (q^*, t^*) . For $n = N$ the m -th mechanism yields the same truth telling payoff as does (q^*, t^*) ; all other elements of $\hat{\Omega}$ instead yield a (weakly larger) truth telling payoff to the highest type. These observations imply that $\hat{\Omega}$ is individually rational and incentive compatible. The expected revenue generated by $\hat{\Omega}$ is

$$T = \sum_{n=1}^N p_n t_n^* + p_m \theta_m (1 - q_m^*) > \sum_n p_n t_n^*.$$

□

4.2 Increasing the number of types: Full surplus extraction in the limit

In this section we show that under reasonable assumptions the share of the surplus that the designer can extract increases as the number of types increases.

Proposition 4 (Full surplus extraction in the limit). *Let $\{\Theta^N, P^N\}_N$ be a sequence of finite environments, such that $|\Theta^N| = N$. Moreover, let M_N be such that $p_{M_N}^N \theta_{M_N}^N \geq p_l^N \theta_l^N$ for all $1 \leq l \leq N - 1$ and write \hat{T}^N for the revenue that the designer can generate with the best ambiguous mechanism in the N -th environment. If $p_{M_N}^N \theta_{M_N}^N \xrightarrow{N \rightarrow \infty} 0$ then*

$$\frac{\hat{T}^N}{\mathbb{E}[\theta^N]} \xrightarrow{N \rightarrow \infty} 1.$$

That is, in the limit the designer is able to extract all the surplus from the agent.

Proof. Fix N and consider the ambiguous mechanism, Ω^N , defined as follows. Ω^N is composed of $N - 1$ simple mechanisms. The corresponding allocation matrix has 1's everywhere except along its 'diagonal', where all elements are equal to zero. Also the diagonal elements of the transfer matrix are equal to zero. The off-diagonal elements of the n -th column of the transfer matrix (vector of possible transfers for type θ_n^N) are equal to θ_n^N if $n < N$. The transfers for the highest type instead are defined by

$$\theta_N^N - (p_{M_N}^N \theta_{M_N}^N - p_l^N \theta_l^N) / p_N^N.$$

The expected transfer of the l -th element of this ambiguous mechanism is

$$T^{N,l} = \sum_{n \neq l, N} p_n^N \theta_n^N + p_N^N [\theta_N^N - (p_{M_N}^N \theta_{M_N}^N - p_l^N \theta_l^N) / p_N^N] = \mathbb{E}[\theta^N] - p_{M_N}^N \theta_{M_N}^N.$$

Since this expression does not depend on l it follows that Ω^N is consistent.

Next observe that for every type θ_n^N , $n < N$, the truth telling payoff are equal to zero under all simple mechanisms in Ω^N . In each mechanism $1 \leq l \leq N - 1$, type θ_n^N , $1 \leq n < N$, either gets the object with probability one and makes a transfer equal to his valuation, or he does not get the object and does not pay anything. The same is also true for the highest type θ_N^N under the M_N -th simple mechanism in Ω^N . This means that Ω^N is individually rational.

The fact that for every $l < N$ there is a simple mechanism that specifies that the agent does not get the object and does not have to pay anything if he announces θ_l^N means every type expects to get at most zero from announcing θ_l^N . Since under the M_N -th simple mechanism the highest type obtains a zero payoff from truth telling it follows that every other type must get a negative payoff from announcing θ_N^N under this mechanism. Combining these observations we can conclude that Ω^N is also incentive compatible.

The result now follows immediately from the observations that $\mathbb{E}[\theta^N] \geq \hat{T}^N \geq T^N$ for all N and

$$\frac{T^N}{\mathbb{E}[\theta^N]} = 1 - \frac{p_{M_N}^N \theta_{M_N}^N}{\mathbb{E}[\theta^N]} \xrightarrow{N \rightarrow \infty} = 1.$$

□

The proof of the preceding result relies on the insight that in environments with a large number of types, each single type's weight in the expected value over types can be very close to zero. Of course in an environment with a continuum of types and an atomless type distribution each single type's weight is exactly zero. Thus for such environments we obtain the following corollary to Proposition 4.

Corollary 1 (Full surplus extraction). *Suppose that Θ is a compact interval in \mathbb{R} and that P is atomless. Then the ambiguous mechanism, $\Omega = \{(q^\theta, t^\theta), \theta \in \Theta\}$, defined by*

$$q_{\theta'}^\theta = \begin{cases} 0 & \text{if } \theta = \theta' \\ 1 & \text{else} \end{cases} \quad t_{\theta'}^\theta = \begin{cases} 0 & \text{if } \theta = \theta' \\ \theta' & \text{else.} \end{cases}$$

is individually rational, incentive compatible and consistent. Moreover, Ω extracts the full surplus from the agent, that is $T = \mathbb{E}[\theta]$.

Proposition 4 and Corollary 1 are important not only because they tell us that the designer can achieve (almost) full surplus extraction by using an appropriately constructed ambiguous

mechanism. An even more important insight that we can derive from these results is that in situations where type sets are large and the type distributions are not too concentrated, it is possible to design an ambiguous mechanism that achieves (almost) full surplus extraction without knowing the details of the type distribution. In particular, these results tell us that in ‘large’ environments the assumptions that i) the planner is ambiguity neutral and ii) that he knows the distribution from which the agents type is drawn can be substantially relaxed. The only assumption that we need to maintain is that the true type distribution is atomless and that the planner believes or knows this.

5 Multiple agents

In the previous sections we have seen how the designer can benefit from offering an ambiguous mechanism in a screening environment. In particular, in Corollary 1 we have shown that the designer can extract the entire surplus from the agent if the set of types is a continuum. The objective of this section is to show how this result can be extended to the case of multiple agents. More specifically, in what follows we consider a setting with two agents whose types are drawn from an atomless distribution. The assumption of two agents is made for notational convenience only. All the arguments that we apply can easily be extended (and in obvious way so) to the case with $I \geq 2$ agents. We will show that if the agents are ambiguity averse then we can construct an ambiguous mechanism such that i) truth telling is an optimal strategy for each agent independently of their beliefs about the other agent’s type or play and ii) full surplus extraction is achieved.

So assume that there are 2 agents with preferences as in the previous sections. We denote the type set of agent i , $i = 1, 2$, by Θ^i . For a generic element of this set we write θ^i ; generic type profiles in $\Theta = \times_i \Theta^i$, are indicated by θ . We assume that the agents’ type are drawn from the atomless distribution P , whose support is $[0, 1]$. We make no specific assumption as to what the designer knows (believes) about P , except for the fact that they know (believe) that P is atomless. Similarly, we make no assumption about the agents’ beliefs about each other’s type.

Proposition 5 (Full surplus extraction with multiple agents). *Consider a two agent setting as described in the preceding paragraph. Moreover, let the ambiguous mechanism $\Omega = \{(q_\theta, t_\theta) : \theta \in \Theta\}$, be defined by*

$$q_\theta(\hat{\theta}) = \begin{cases} (0, 0) & \text{if } \theta = \hat{\theta} \\ (1, 0) & \text{if } \hat{\theta}^1 \geq \hat{\theta}^2, \hat{\theta}^1 \neq \theta^1, \hat{\theta}^2 \neq \theta^2 \\ (0, 1) & \text{if } \hat{\theta}^2 \geq \hat{\theta}^1, \hat{\theta}^1 \neq \theta^1, \hat{\theta}^2 \neq \theta^2, \end{cases} \quad t_\theta(\hat{\theta}) = (q_\theta^1(\hat{\theta})\hat{\theta}^1, q_\theta^2(\hat{\theta})\hat{\theta}^2).$$

Ω is incentive compatible, individually rational and consistent. The expected revenue gener-

ated by each element of Ω is $T = \mathbb{E}[\max\{\theta^1, \theta^2\}]$. That is, Ω achieves full surplus extraction.

Proof. Consistency follows from the fact that any two simple mechanisms in Ω differ only for a set of types with zero probability. Notice also that every simple mechanism almost always awards the object to the agent with the higher announced type at a price that is equal to the announced type. Thus under truth telling each simple mechanism in Ω generates a revenue of $T = \mathbb{E}[\max\{\theta^1, \theta^2\}]$.

As for individual rationality observe that the ambiguous mechanism never specifies a payment for an agent unless he receives the object. When an agent receives the object then he has to make a payment that corresponds to his announced valuation. Thus, truth telling always guarantees a non-negative payoff.

Finally, we have to argue that under Ω truth telling is an optimal strategy for both agents, irrespective of their beliefs over the other agents type or play. This simply follows from the fact that the ambiguous mechanism is designed such that both agents get the same payoff irrespective of their own actions and of the actions of the other agents. In order to see this, notice that for every profile of announced types, $\hat{\theta}$, agent $i = 1, 2$ knows that there are simple mechanisms in Ω (all those indexed by a type profile, θ , such that $\hat{\theta}_1 = \theta_1$) that specify that he will not receive the object and that he will not have to pay anything. Thus, for every $\hat{\theta}$ his payoff is 0. \square

6 Conclusion

In this paper we have assumed that the principal has ambiguity-neutral preferences. However, our results do not depend on this assumption. In particular, natural analogues of Proposition 1 and Lemmas 1–6 hold with an ambiguity-averse principal, even without assuming knowledge of the distribution P over the agent's types. In effect, the analysis is even simpler in this case: consistency need not be imposed, as the principal would always focus on the simple mechanism that delivers the lowest profit. Moreover, if we assume that the principal knows P , then Proposition 2 delivers a lower bound on what the principal will be able to achieve. Finally, immediate extensions of Proposition 1 and Lemmas 1–6 can be proved for the case of any finite number of agents (and ambiguity-averse principal).

The assumption that the agent has MMEU preferences, which we hold throughout the paper, is certainly extreme. A natural extension of our analysis would be to consider alternative models of preferences exhibiting ambiguity aversion. We conjecture that our results can be extended to the case of smooth ambiguity and also to other models. We plan to develop a full analysis of several alternative models in future work.

Another important extension regards moral hazard problems. In ongoing work we are analyzing the design of ambiguous contracts in the moral hazard setup. In this framework it is not hard to show that when the agent has only two actions available, then the principal cannot obtain a higher payoff by offering an ambiguous reward scheme. Moreover, again paralleling the results obtained in the screening framework of this paper, when more than two actions are available, the principal does stand to gain by offering an ambiguous contract.

References

- Azrieli, Yaron and Roe Teper (2011), “Uncertainty aversion and equilibrium existence in games with incomplete.” *Games and Economic Behavior*, forthcoming.
- Bodoh-Creed, Aaron (2010), “Ambiguous beliefs and mechanism design.” Working paper.
- Bose, Subir and Arup Daripa (2009), “A dynamic mechanism and surplus extraction under ambiguity.” *Journal of Economic Theory*, 144(5), 2084–2114.
- Bose, Subir, Emre Ozdenoren, and Andreas Pape (2006), “Optimal auction with ambiguity.” *Theoretical Economics*, 1, 411–438.
- Bose, Subir and Ludovic Renou (2011), “Mechanism design with ambiguous communication devices.” Mimeo.
- Gilboa, Itzhak and David Schmeidler (1989), “Maxmin expected utility with a non-unique prior.” *Journal of Mathematical Economics*, 18, 141–153.
- Turocy, Theodore L. (2008), “Auction choice for ambiguity-averse sellers facing strategic uncertainty.” *Games and Economic Behavior*, 62, 155–179.