INNOCENZO GASPARINI INSTITUTE FOR ECONOMIC RESEARCH

Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

The Design of Ambiguous Mechanisms

Alfredo Di Tillio, Nenad Kos and Matthias Messner

Working Paper n. 446

This Version: July 14, 2012

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

The Design of Ambiguous Mechanisms*

Alfredo Di Tillio † — Bocconi University

Nenad Kos ‡ — Bocconi University

Matthias Messner § — Bocconi University and CESifo

July 14, 2012

First version: April 2011

Abstract

This paper considers the optimal mechanism design problem of an expected revenue maximizing principal who wants to sell a single unit of a good to an agent who is ambiguity averse in the sense of Gilboa and Schmeidler (1989). We show that the optimal static 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. If the type set is composed of N (finite) types, then the (smallest) optimal ambiguous mechanism contains $N-1$ outcome functions.

We show 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 non-atomic 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 Simone Cerreia-Vioglio, Eddie Dekel, Georgy Egorov, Eduardo Faingold, Itzhak Gilboa, Matthew Jackson, Alessandro Lizzeri, George Mailath, Massimo Marinacci, Stephen Morris, Frank Riedel, Larry Samuelson, Martin Schneider, Ilya Segal, seminar participants at Northwestern University, Universitat Autonoma de Barcelona, Universitat Pompeu Fabra, Yale University, participants of UECE Lisbon 2011, Workshop on Ambiguity Bonn 2011, and Arne Ryde Workshop on Dynamic Princing: A Mechanism Design Perspective 2011.

†Dept. of Econ. & IGIER, Bocconi University; e-mail: `alfredo.ditillio@unibocconi.it`

‡Dept. of Econ. & IGIER, Bocconi University; e-mail: `nenad.kos@unibocconi.it`

§Dept. of Econ.& IGIER, Bocconi University; CESifo Munich; e-mail: `matthias.messner@unibocconi.it`

1 Introduction

Starting with the seminal work of Ellsberg (1961), experimental economists have argued that the standard economic model for decision making under uncertainty, namely the Expected Utility Model (henceforth EU model), performs rather poorly in describing individuals' behavior in situations where subjects have very little information regarding the decision problem they are facing. In particular, it has been shown that the overwhelming majority of individuals tends to shy away from alternatives for which they lack the necessary information to form a probabilistic belief about their consequences. It is well known that this aversion against uncertainty/ambiguity is incompatible with the EU model.[1]

This inconsistency between observed decisions and the EU model has stimulated the development of decision theoretic models that are able to accommodate ambiguity aversion.[2] While ambiguity aversion models have been successfully applied in many areas of economics and finance,[3] they have received only limited attention in mechanism design (see the discussion of the literature below). This is surprising not only in light of the fact that the available experimental evidence suggests that most individuals do exhibit at least some degree of ambiguity aversion. But mechanism design is also a subject where issues related to ambiguity aversion cannot simply be ruled out by assuming that the involved parties have sufficiently good information regarding the exogenous elements of the environment (e.g. distributions over type sets). After all, in mechanism design the principal has direct control over the rules of the game (auction format, tax code etc.) and thus also about the information that he wants to release regarding these rules. Since it is up to him whether or not he wants to remain ambiguous vis à vis the agent when he explains the rules of the game, the question whether agents are exposed to uncertainty becomes an endogenous issue. This observation very naturally leads to the question of whether the principal can benefit from specifying the rules in an ambiguous fashion, and if so, how exactly those ambiguous rules should be designed. These are the questions that we address in this paper.

[1] The sense in which ambiguity aversion is incompatible with the EU model is best explained with Ellsberg's famous two urn example. There are two urns, each of which contains one hundred balls. Half of the balls in Urn A are red, the other half is blue. Also Urn B is composed of balls that are either red or blue, but the decision maker has no information about the number of balls of each color. Now consider the following two bets. Bet RA pays one dollar if in a random draw from Urn A a red ball is extracted; bet RB pays one dollar if a random draw from Urn B yields a red ball. When faced with the choice between these two bets the overwhelming majority of subjects picks bet RA. The same they do also when the pair of bets is formulated for the color blue. Within the EU framework it is impossible to rationalize both these decisions: for each possible belief about the composition of Urn B the decision maker should choose the bet on a blue ball from Urn B if and only if between the two bets on red he prefers the one referring to Urn A.

[2] For two recent surveys of the literature on ambiguity aversion and its axiomatic foundations, see Gilboa (2009) and Gilboa and Marinacci (2011).

[3] See for instance Epstein and Schneider (2008) and Castro and Yannelis (2012) for examples of applications of ambiguity aversion in finance and general equilibrium.

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. The agent privately observes his willingness to pay for the good; the principal only knows the distribution from which it has been drawn. 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 static mechanism uses 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 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 'comprehensive' belief is reasonable if it is compatible with the principal being indifferent between all the elements of the ambiguous mechanism if the agent acts optimally with respect to such a belief. We therefore impose that the principal can only propose mechanisms that satisfy this property. That is, we require that all elements of an ambiguous mechanism generate the same expected revenue under the assumption that the agent chooses his strategy based on a comprehensive belief. We refer to ambiguous mechanisms that satisfy this condition as *consistent*.[4] While from a technical point of view we treat consistency as a constraint that limits the feasible actions available to the designer, it should be interpreted as an equilibrium

[4] A formal definition of the consistency condition is given in Section 2.

condition in the interaction between the principal and the agent.

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 (one of) the smallest optimal direct ambiguous mechanism(s) for the case where the set of possible types of the agent is finite.[5] We show that this mechanism is composed of at most $N-1$ elements, where N is the number of types. The n-th outcome function of this ambiguous mechanism assigns the good with probability one to all types $m \neq n, N$ at a price that coincides with the reported type. Thus, every outcome function extracts the entire surplus from $N-2$ types. Also the highest type obtains the good with probability one. Since his transfers are used to guarantee consistency, he typically does not have to pay under every outcome function a price equal to his willingness to pay. The remaining components of the outcome functions (allocations and payments of type $n \leq N-1$ under outcome function n) vary with the details of the type distribution. More specifically, we show that these components depend on the types' so called *adjusted virtual valuations*. Independently of the details of the type distribution, these components satisfy a monotonicity condition: the probability with which type n obtains the good under outcome function n is smaller than or equal to the probability with which outcome function $m > n$ assigns the good to type m.

Using the above described 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 non-atomic type distribution over an interval, the optimal ambiguous mechanism extracts the full surplus from the agent. In the final section of the paper we discuss how this result on full surplus extraction under ambiguity aversion relates to the findings of Matthews (1983), who shows that full surplus extraction is possible if the agent is infinitely risk averse.[6] There we also show that the fact that the number of types is a crucial determinant for the share of the surplus that the principal can extract from the agent, implies that the principal may want to elicit payoff-irrelevant private information from the agent. Since such information is easy to generate, the principal has an incentive to induce the agent to inflate his type set by adding payoff irrelevant elements.

We also consider two dimensions along which our basic model can be extended. First, we argue that the central insight of the paper (the principal can exploit the ambiguity aversion of the agent by offering an ambiguous mechanism) does not hinge on the particular model of ambiguity aversion that we adopt in this paper (MMEU preferences) but remains valid also if alternative models of uncertainty aversion are assumed.[7] Specifically, we provide an example

[5] The term 'smallest' refers to the number of elements of the ambiguous mechanism.

[6] The implications of risk aversion for the design of an optimal mechanism are also studied by Maskin and Riley (1984).

[7] Of course, the details of our characterization of an optimal ambiguous mechanism do depend on the specificities

that shows this for the case of smooth ambiguity aversion. A second dimension in which the core insight of the paper generalizes is the number of agents. While we do not provide a detailed characterization of the optimal ambiguous mechanism for the case where the agents' type sets are finite, we describe the mechanism that extracts the full surplus when the agents' types are drawn from an atomless distribution defined on some interval.

Given the result that in the presence of ambiguity averse agents it is highly beneficial for the designer to use ambiguous mechanisms, one should expect to observe real world mechanisms that exhibit some form of ambiguity. Indeed, we believe that our findings help to rationalize a number of institutions that we observe in real world settings. Probably the most obvious instances of real world mechanisms that can be interpreted as ambiguous mechanisms are auction formats with unknown reservation prices.[8] Even though unknown reservation prices can also be rationalized in the context of a Bayesian setting, it seems to us that the most natural way to interpret the fact that the principal does not want to disclose any information with respect to his reservation price, is that he wants to expose the bidders to uncertainty.[9]

Our results also provide a new way to rationalize laws that leave some discretionary decision power to the executive or judicial branch of government. Arguably, it is very difficult for citizens to predict the behavior of government agencies or judicial courts. In particular, this should be the case for citizens who have only limited experience in dealing with the public administration. Thus, by leaving laws incomplete, and by delegating decisions to the executive agencies or the court system, the parliament exposes its citizens to uncertainty.

Related literature: A number of recent papers consider mechanism design problems with ambiguity averse players. Examples include Bose, Ozdenoren, and Pape (2006), Turocy (2008), Bose and Daripa (2009), Bodoh-Creed (2010), Bose and Renou (2011) and Bergemann and Schlag (2011).[10] The central difference between these papers and ours is that they start from the assumption that the agents (and/or the principal) are uncertain about the other agents' type distribution. That is, the uncertainty in these models refers to an exogenously given variable. The endogenous objects (i.e. the mechanisms) are not allowed to be ambiguous.[11] Instead,

of the Gilboa and Schmeidler (1989) preference model that we use.

[8]For a discussion of real world situations where auctions with unknown reservation prices are applied see, for instance, Elyakime, Laffont, Loisel, and Vuong (1994).

[9] Another prominent instance of ambiguous auction formats are Google's auctions for ad space on its web pages. Apparently auctions with ex ante uncertain auction rules are applied also in the used car market. In these auctions first buyers submit their bids. Upon observing the bids the auctioneer either declares a winner or he calls for a second round of bids and so on. The rule according to which the decision about whether or not to continue is taken, is not known to buyers (and supposedly not easily inferable from previous observations unless the bidder is extremely experienced). We are thankful to Larry Samuelson for pointing us to this example.

[10]Several 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.

[11]Bose and Renou (2011) are an exception to this observation; their work is discussed in more detail in the following paragraph.

these papers characterize the optimal standard (i.e. non-ambiguous) mechanism, where attention is restricted either to direct mechanisms or to simple forms of indirect mechanisms (e.g. standard auction formats).[12]

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.

The paper is also related to the literature on robust mechanism design that originated with the seminal papers by Bergemann and Morris (2005) and Chung and Ely (2007). This literature departs from the standard Bayesian type space framework that has dominated the earlier mechanism design literature and studies what kind of social choice functions are implementable irrespective of the type space that is assumed. Requiring such a form of robustness with respect to the specificities of the type space is similar in spirit to the idea of a designer that is uncertain with respect to the 'correct' type space. Apart from the fact that the 'uncertainty aversion' in the case of this literature is on the side of the designer, the crucial conceptual difference to our work lies in the fact that the family of the relevant type spaces is not an endogenous object (like the ambiguous mechanisms in our work) but is exogenously given.

[12]Similar comments apply both to the literature that considers moral hazard models with ambiguity aversion and the literature that studies models with Knightian uncertainty. For the literature on moral hazard and ambiguity aversion see for instance Kellner (2011) and Szydlowski (2012); on Knightian uncertainty in mechanism design see Lopomo, Rigotti, and Shannon (2009) and Garrett (2011) and the references therein.

2 Framework

Throughout the first part of the paper we consider the mechanism design problem of a principal who wants to sell a single unit of a good to a single 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.

Allocations and preferences. An allocation is a pair $(x, \tau) \in X \times \mathbb{R}$, where $x \in X = [0, 1]$ is the probability with which the good is transferred to the agent and τ represents a monetary transfer (to be paid by the agent).[13] 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 linear utility function,

$$u(x, \tau) = x\theta - \tau.$$

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, $\mathcal{M}$, his utility is given by

$$\inf_{\mu \in \mathcal{M}} \mathbb{E}_{\mu}[x\theta - \tau].$$

The agent's valuation (or type), θ, is only observed by the agent himself. The principal knows that it is distributed according to $p = (p_1, \dots, p_N)$.

The principal is a risk neutral expected utility maximizer. We assume that his objective is to maximize his expected revenue (i.e. the expected transfer payments extracted from the agent). The main results of our paper go through also under the assumption of an ambiguity averse principal. Allowing for that possibility, however, does not add anything of interest to our analysis since the crucial insights that we obtain are driven by the ambiguity aversion of the agent. In order not to obscure this fact it is convenient to assume standard expected utility preferences for the principal.

The principal's strategy set: simple vs. ambiguous mechanisms. 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 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

[13] Instead of interpreting x as the probability with which the indivisible good is transferred one can equivalently assume that the good is perfectly divisible and that x represents the share that is given to the agent.

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 allows for random allocations but not for random transfers: the range of t is $\mathbb{R}$, not the set of probability measures over $\mathbb{R}$. Given that both the principal and the agent are risk neutral, restricting attention to deterministic transfer schemes is without loss of generality. A mechanism with random transfers can be replaced by one with deterministic transfers that specifies for each type report the expected values of the random transfer scheme. Doing so does not alter the two players' expected payoffs for any decision that the agent may take. The same is true for random allocation rules if the good is perfectly divisible.[14]

In 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 particular, if the agent is an expected utility maximizer the principal cannot gain anything from concealing the mechanism to which he committed himself.[15] The central insight of this paper is that this is no longer the case in the presence of an ambiguity averse agent. In fact, we show that it is typically in the principal's best interest not to inform the agent about the exact outcome function to which he commits. Instead, he can benefit from communicating the rules of the mechanisms in an *ambiguous* way.

Formally, we capture the idea of ambiguous rules with the notion of an ambiguous mechanism.

Definition 1. *An* **ambiguous mechanism** *is a pair* (S, Ω)*, where* S *is a set of messages, and* Ω *is a set of simple outcome functions defined on* S*, i.e.* $\Omega \subset X^S \times \mathbb{R}^S$*.*[16] *A generic element of* Ω *is denoted by* (q, t)*, where* $q \in X^S$ *and* $t \in \mathbb{R}^S$*.*

An ambiguous mechanism (S, Ω) *is said to be* **direct** *if* $S = \Theta$*.*

Since in the case of direct ambiguous mechanisms it is clear what the message space is, we identify the direct mechanism (Θ, Ω) with its set of outcome functions Ω. That is, with a slight abuse of terminology we will often refer to Ω alone as an ambiguous mechanism.

Before we continue with our analysis a few remarks on the interpretation and purpose of the concept of an ambiguous mechanism are in order. After choosing a set of possible messages, S, the principal commits to an outcome function $(\hat{q}, \hat{t}) \in X^S \times \mathbb{R}^S$. This commitment may be achieved, say, by depositing $(\hat{q}, \hat{t})$ with an uninterested third party. The agent is not necessarily

[14] If the good is not divisible, then allowing for random allocations expands the set of possible allocations.

[15] Strictly speaking, even in a Bayesian one does not need to assume explicitly that the principal announces the mechanism to the agent. But under every standard equilibrium concept agents would know in equilibrium which mechanism has been chosen by the principal.

[16] As already mentioned in the discussion of simple mechanisms, the restriction to ambiguous mechanisms with sets of deterministic outcome functions is without loss of generality in an environment with risk neutral players.

informed about the exact outcome function that has been chosen. Instead the principal may limit himself to tell the agent that it belongs to a set Ω (i.e. $(\hat{q}, \hat{t}) \in \Omega$). Of course, by announcing an ambiguous outcome rule Ω he exposes the agent to uncertainty about the consequences of his messages. We discuss the principal's motives for doing so in the next section.

The requirement $(\hat{q}, \hat{t}) \in \Omega$ rules out the possibility that the principal completely deceives the agent with regard to $(\hat{q}, \hat{t})$. Thus, an agent who decides to participate in the ambiguous mechanism and announces the message $s \in S$ can be forced to accept the outcome $(\hat{q}(s), \hat{t}(s))$ only if two conditions are satisfied. First, the outcome function $(\hat{q}, \hat{t})$ is the one to which the principal has committed[17] and ii) $(\hat{q}, \hat{t})$ is an element of the announced set of outcome functions, Ω. We stress once more the fact that the principal commits to $(\hat{q}, \hat{t})$ *before* the agent sends his message; therefore the choice of $(\hat{q}, \hat{t})$ cannot be conditioned on the latter.

Agent's beliefs and strategies. Once the designer has specified an ambiguous mechanism, (S, Ω), the agent chooses a message from S. A strategy for the agent is a function σ that maps Θ into S, i.e. $\sigma \in S^{\Theta}$.

Observe that we do not allow the agent to use mixed strategies. It is well known that an ambiguity averse individual who faces the choice between two alternatives with uncertain consequences, may be strictly better off by mixing over the two options than he would be if he chose either one of them with probability one. Through the use of randomizations the agent may be able to hedge against the uncertainty involved in the two alternatives.[18]

This observation notwithstanding, our assumption that the agent may only use pure strategies is without loss of generality. This is so because preventing the agent from using mixed strategies not only benefits the designer, but it is also feasible for him to do so. We discuss the reasons why the principal does not want to allow the agent to send randomized messages in some more detail in Example 1 below. Essentially, by permitting the agent to randomize the principal reduces his own capacity to control the agent's deviation payoffs.

In order to see that the principal can prevent the agent from hedging, observe that the benefits from hedging are gains that the agent perceives before the outcome of his random strategy is observed. Thus, in our mechanism design context the agent can realize such gains only if he has the possibility to commit to the use of an external randomization device that picks the message that he is supposed to report to the principal in his stead.[19] But then the principal just needs to preclude the use of such external randomization devices by imposing that reports have to be communicated directly by the agent himself.

[17]Here we implicitly assume that the principal can credibly disclose this outcome function ex post.

[18]That uncertainty averse agents have an incentive to hedge against uncertainty has already been observed by Raiffa (1961).

[19]No hedging gains can be obtained by 'rolling a die' in his head. Once such a die has been rolled the agent still has the possibility to decide whether or not to communicate the message chosen by the die.

The set of optimal strategies for the agent depends on his beliefs regarding the outcome function $(\hat{q}, \hat{t})$ to which the principal has committed himself. Since the agent's only piece of hard information regarding the chosen outcome function is that $(\hat{q}, \hat{t})$ belongs to Ω, we assume that his (set valued) belief contains all probability measures on Ω, provided that such a belief is not incompatible with the fact that $(\hat{q}, \hat{t})$ is an optimal choice of the principal.

Denote the set of all probability measures over Ω by $\Delta(\Omega)$. We say that the set of beliefs $\Delta(\Omega)$ is consistent with $(\hat{q}, \hat{t})$ being an optimal choice of the principal, if the following condition holds.

Definition 2. *Let $\Sigma^*(S, \Omega) \subset S^\Theta$ denote the agent's set of optimal strategies for the ambiguous mechanism (S, Ω) when his beliefs are described by the set $\Delta(\Omega)$. The ambiguous mechanism (S, Ω) is* **consistent** *with respect to $\sigma \in \Sigma^*(S, \Delta(\Omega))$ if for all $(q', t'), (q'', t'') \in \Omega$*

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

(S, Ω) is **consistent** *if it is consistent with respect to some $\sigma \in \Sigma^*(S, \Omega)$.*

Consistency requires that each element of the ambiguous mechanism Ω delivers the same expected revenue to the principal if the agent bases his choice on the belief set $\Delta(\Omega)$. In order to understand why we impose this condition, consider a situation were it is not satisfied. I.e. assume that the principal proposes an ambiguous mechanism (S, Ω) such that for every $\sigma \in \Sigma^*(S, \Omega)$ there are outcome functions $(\underline{q}, \underline{t}), (\bar{q}, \bar{t}) \in \Omega$ satisfying

$$\mathbb{E}_p[\underline{t}(\sigma(\theta))] < \mathbb{E}_p[\bar{t}(\sigma(\theta))].$$

In this case the agent's assumption that the principal might have chosen any of the elements in Ω leads to the conclusion that the principal strictly prefers some elements of Ω over other elements of Ω, if he correctly predicts the agent's belief and strategy. Consistency serves to rule out such contradictory beliefs.

Essentially, consistency is an equilibrium condition for the two stage game played by the designer and the agent. In equilibrium the agent should not believe in the possibility that a certain outcome function is chosen, if his optimal response to such a belief implies that the outcome function does not maximize the designer's payoff.

3 Optimal ambiguous mechanisms

The designer's objective is to maximize his expected revenue. When designing the (ambiguous) mechanism that he wants to adopt he has to respect two types of constraints: i) consistency and ii) individual rationality. The consistency conditions has been discussed in the preceding section. Individual rationality refers to the agents willingness to participate in the mechanism. We assume that participation is voluntary. Thus, in order to guarantee that no type of the agent refuses to participate, the mechanism must yield to each type a payoff that is no smaller than the value of the type's outside option. We assume that the latter is equal to zero for all types.[20]

Formally, we can state the principal's problem as follows.

$$\begin{aligned} \max_{R\in\mathbb{R}, S\subset\mathbb{R}, \Omega\subset X^S\times\mathbb{R}^S} \quad & R \qquad\qquad (\text{P}) \\ \text{subject to} \quad & \text{there exists } \sigma^* \in \Sigma^*(S,\Omega) \text{ such that} \\ & \text{i) } R = \mathbb{E}_p[t(\sigma^*(\theta))] \quad \text{for all } (q,t)\in\Omega, \text{ and} \\ & \text{ii) } \inf_{(q,t)\in\Omega} \{q(\sigma^*(\theta))\theta - t(\sigma^*(\theta))\} \geq 0 \quad \text{for all } \theta\in\Theta. \end{aligned}$$

Constraint i) is the consistency condition while ii) represents the individual rationality constraint. Since simple mechanisms can be seen as degenerate ambiguous mechanisms, the constraint set of this problem is clearly non-empty.

Before we derive the solution to this problem we show with the help of a simple example why the use of an ambiguous mechanisms allows the principal to extract a larger share of the surplus than a simple mechanism.

Example 1 (Simple vs. Ambiguous mechanisms)**.** Consider an environment where the agent's type is drawn from the set $\Theta = \{1, 2, 4\}$ according to the probability distribution $p = (1/4, 1/4, 1/2)$. The optimal simple mechanism for this setting, excludes the two lowest types and assigns the good only to an agent who reports the highest type at the price of 4. The details of the optimal (direct) outcome function are reported in the following table, where $\hat{\theta}$ denotes the reported type.

$\hat{\theta}$	1	2	4
(q^*, t^*)	$(0, 0)$	$(0, 0)$	$(1, 4)$

The expected revenue generated by this mechanism when the agent reports his type truthfully, is $R^* = \mathbb{E}_p[t^*(\theta)] = 2$.

Now suppose that the planner adopts the (direct) ambiguous mechanism $\Omega = \{(q^1, t^1), (q^2, t^2)\}$, described in the table below.

[20]In what follows individual rationality is defined for general (i.e. direct and non-direct) ambiguous mechanisms. We stress this because it is standard in the literature to define it with respect to truthtelling in direct mechanisms.

$\hat{\theta}$	1	2	4
(q^1, t^1)	$(0, 0)$	$(1, 1)$	$(1, 4)$
(q^2, t^2)	$(1, 1)$	$(0, 0)$	$(1, 4)$

Upon observing Ω the agent knows that the planner has committed himself either to (q^1, t^1) or to (q^2, t^2). Since he does not know which one of the outcome functions the principal has picked he evaluates all possible reports according to the outcome function that yields the worse outcome for that report. Thus, by announcing either 1 or 2 each type of the agent obtains a payoff of 0:

$$\min\{q^1(1)\theta - t^1(1), q^2(1)\theta - t^2(1)\} = \min\{q^1(2)\theta - t^1(2), q^2(2)\theta - t^2(2)\} = \min\{0\times\theta - 0, 1\times\theta - 1\} = 0.$$

Only the highest type gets a zero payoff from announcing 4 ($q^1(4)4 - t^1(4) = q^2(4)4 - t^2(4) = 4 - 4 = 0$); for both of the other types reporting 4 leads to a negative payoff. Combining these observations allows us to conclude that under Ω each type of the agent, by reporting truthfully, obtains a payoff of at least 0, i.e. Ω is individually rational. Moreover, for each type θ it is optimal to report θ. Under truthful reporting both outcome functions yield the expected revenue $R = \mathbb{E}_p[t^1(\theta)] = \mathbb{E}_p[t^2(\theta)] = 9/4 > 2 = R^*$. Therefore, with respect to truth-telling, Ω is not only individually rational but also consistent. We conclude that Ω is an admissible ambiguous mechanism that yields a higher expected revenue than the best simple mechanism.

The above example reveals that the ambiguous mechanism Ω performs better than the optimal simple mechanism (q^*, t^*) because neither element of Ω excludes both lower types: while (q^1, t^1) only excludes type 1, under (q^2, t^2) trade is ruled out only if type 2 is reported. The fact that for each of the two lower types there is one outcome function that specifies that no trade will take place if this type is reported, suffices to deter the highest type from deviating to a non-truthful message. When contemplating the message $\hat{\theta} = 1$ he is afraid of the possibility that the principal has chosen (q^1, t^1); in the case of $\hat{\theta} = 2$, instead, the agent is afraid of (q^2, t^2). Thus, through the use of an ambiguous mechanism it is possible to extract the full surplus from the highest type without having to renounce completely on the possibility to extract some rent also from the lower types.

The preceding arguments are based on the assumption that the agent cannot randomize between different messages. We next show that if given the possibility, the agent would want to hedge against the uncertainty to which the ambiguous mechanism exposes him, by committing to a randomization device for the choice of his report. Thus, truth telling is no longer an optimal strategy for him.

In order to see this consider type 4 and observe that his truth-telling payoff is 0 as is the payoff from either one of the other two available messages. The reason for the low payoff from a pure deviation strategy is of course the ambiguity in the consequences that the deviation will have. If the agent is given the possibility to commit to a randomization device that mixes (with equal chances) between the messages 1 and 2, then by doing so he can eliminate this uncertainty and obtain a strictly positive payoff. Thus, in this case truth-telling would no longer be an optimal strategy.

In particular, suppose the agent mixes half-half between messages 1 and 2. Then both under the outcome function (q^1, t^1) and under the outcome function (q^2, t^2) he gets the good with probability $1/2$ and pays (in expectation) a transfer of $1/2$.[21] His payoff from such a mixed deviation is therefore equal to $(1/2) \times 4 - 1/2 = 3/2 > 0$. □

In what follows we prove a sequence of results that show that Problem P can be substantially simplified. We start our analysis by showing that it is without loss for the principal to restrict his attention to direct ambiguous mechanisms. That is, we prove that in our environment a version of the Revelation Principle holds. Before we do so we introduce the notion of incentive compatibility of direct ambiguous mechanisms.

Definition 3 (Incentive Compatibility)**.** *The direct ambiguous mechanism Ω satisfies the incentive compatibility constraint of type $\theta_n \in \Theta$ with respect to type $\theta_m \in \Theta$ if*[22]

$$\inf_{(q,t)\in\Omega}\{q(\theta_n)\theta_n - t(\theta_n)\} \geq \inf_{(q,t)\in\Omega}\{q(\theta_m)\theta_n - t(\theta_m)\}. \tag{IC(n,m)}$$

The ambiguous mechanism Ω is **downward incentive compatible** *if IC(n,m) holds for all $n, m \in N$ such that $n > m$; it is* **upward incentive compatible** *if 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 it is optimal for the agent to reveal his type truthfully.

Proposition 1 (Revelation Principle)**.** *Let (S, Ω) be an ambiguous mechanism that is consistent with respect to $\sigma \in \Sigma^*(S, \Omega)$. There is a direct ambiguous mechanism (Θ, Ω'), and an optimal strategy for the agent, $\sigma' \in \Sigma^*(\Theta, \Omega')$ such that the following three conditions hold:*

i) (Θ, Ω') *is consistent with respect to σ';*

ii) (Θ, Ω') *is incentive compatible; in particular, $\sigma'(\theta) = \theta$ for every $\theta \in \Theta$;*

[21] Since the agent is risk neutral he evaluates stochastic outcomes according to their expected value.

[22] Since the agent is risk neutral, calculating the infimum of the expected payoffs of the agent with respect to $\Delta(\Omega)$ delivers the same value as the one that is obtained when attention is restricted to Ω.

iii) $(q,t) \in \Omega$ *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. Optimality of σ implies

$$\inf_{(q,t)\in\Omega} q(\sigma(\theta))\theta - t(\sigma(\theta)) \geq \inf_{(q,t)\in\Omega} q(s)\theta - t(s) \qquad \forall s \in S.$$

Let Ω' be defined by iii). By the construction of Ω' we have

$$\inf_{(q',t')\in\Omega'} q'(\theta)\theta - t'(\theta) = \inf_{(q,t)\in\Omega} q(\sigma(\theta))\theta - t(\sigma(\theta)).$$

Similarly,

$$\inf_{(q',t')\in\Omega'} q'(\theta')\theta - t'(\theta') = \inf_{(q,t)\in\Omega} q(s')\theta - t(s')$$

for $s' = \sigma(\theta')$. Combining these three observations yields

$$\inf_{(q',t')\in\Omega'} q'(\theta)\theta - t'(\theta) \geq \inf_{(q',t')\in\Omega} q'(\theta')\theta - t'(\theta') \qquad \forall \theta' \in \Theta,$$

as required by condition ii).

Condition i) follows from the construction of t'.

□

We have formulated the Revelation Principle so that it directly includes the consistency condition. The proposition tells us that for any non-direct ambiguous mechanism (S,Ω) we can find an incentive compatible direct one, Ω', such that, element by element, the two ambiguous mechanisms generate the same allocations and transfers. Therefore, both the principal and the agent obtain the same payoff under the two mechanisms. As a consequence it is without loss for the principal to restrict himself to direct ambiguous mechanisms that are incentive compatible.

Given the restriction to direct ambiguous mechanisms, from now on we write q_n and t_n for respectively $q(\theta_n)$ and $t(\theta_n)$. Next, we formally define the notion of individual rationality for direct ambiguous mechanisms.

Definition 4 (Individual Rationality)**.** *A direct ambiguous mechanism* Ω *is individually rational for type* $\theta_n \in \Theta$ *if*

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

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

The next lemma shows that if downward incentive compatibility holds, then imposing individual rationality for the lowest type implies individual rationality for the other types.

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

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

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

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 also hold for type θ_m, $m = 2, \ldots n$. □

By the Revelation Principle and Lemma 1 we can restrict our attention to direct ambiguous mechanism that satisfy incentive compatibility, consistency and IR(1). In what follows it will be convenient to consider a relaxed version of this problem, where incentive compatibility is substituted by downward incentive compatibility:

$$\max_{R\in\mathbb{R},\Omega\subset X^{\Theta}\times\mathbb{R}^{\Theta}} \quad R \tag{P'}$$

$$\text{s.t.} \quad R = \mathbb{E}_p[t(\theta)] \quad \text{for all } (q,t)\in\Omega,$$

$$\text{IC(n,m)} \quad \text{for all } n, m \in N,\ n > m,$$

$$\text{IR(1)}.$$

Our next objective is to show that the set of admissible ambiguous mechanisms of this problem can be substantially further restricted without reducing the problem's value. Then we will prove that all the elements of the resulting reduced feasible set satisfy upward incentive compatibility. Thus, the solutions to the simplified problem that we obtain, also solve the original problem P.

The first step of our argument is to show that P' admits solutions that exhibit a property to which we refer as **uniformity**. Uniformity requires that for every type, each element of the ambiguous mechanism yields the same truth telling payoff. Moreover, one of the downward incentive constraints must be binding. In order to allow for a compact statement of the result we denote set of P'-feasible ambiguous mechanisms by C.[23] Moreover, for every $\Omega \in C$ we write $R(\Omega)$ for the expected revenue generated by (every element of) Ω.

Lemma 2 (Uniformity)**.** *For every* $\Omega \in C$ *there exists* $\tilde{\Omega} \in C$ *such that* $|\Omega| = |\tilde{\Omega}|$, $R(\Omega) = R(\tilde{\Omega})$,

[23] That is, C is the set of all ambiguous mechanisms, Ω, for which there is an R such that the pair (R, Ω) satisfies the constraints of Problem P'.

and

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

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

Proof. We first show that Ω can always be manipulated - in a revenue and feasibility neutral way, - so that condition (3) holds. IR(1) requires that

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

Suppose that for some $(q', t') \in \Omega$ this inequality is strict, i.e. $q'_1\theta_1 - t'_1 > 0$. Then (q', t') can be replaced by the outcome function (q'', t'') that coincides with (q', t') everywhere except for the transfers of the lowest and the highest type. Those are set as follows:

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

That is, the payment of the lowest type is increased until his truth-telling payoff 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'. This change does not lead to a violation of IR(1). By construction the new ambiguous mechanism that we obtain after the replacement of (q', t') satisfies consistency and generates the same expected revenue as Ω. Furthermore, the change in the transfers can at most relax the downward incentive constraints.

The same type of argument can be applied in order to show that we can obtain (4) without violating either downward incentive compatibility or IR(1) and without changing any of the expected revenues. If Ω contains an element, (q, t), such that the truth telling payoff of type θ_n under this outcome function is strictly larger than the payoff that he can realize from the most attractive downward deviation option, then we can increase t_n until (4) holds. 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 $\theta_m < \theta_n$ or IR(1). The downward incentive constraints for types $\theta_m > \theta_n$ are at most relaxed. □

In the statement of Problem P', Ω is allowed to be of any size/cardinality. The next result shows that it is without loss of generality to restrict attention to ambiguous mechanisms that contain at most $N - 1$ simple outcome functions.

Lemma 3. *For every $\Omega \in C$, there exists $\tilde{\Omega} \in C$ that satisfies the following three conditions*

i) $R(\Omega) = R(\tilde{\Omega})$,

ii) $|\Omega| \leq N - 1$ *and*

iii) *if* $(q, t) \in \tilde{\Omega}$, *then there exists an* $1 \leq m \leq N - 1$ *such that for all* $n = m + 1, \ldots, N$, (q, t) *is the unique element of* $\arg\min_{(q,t)\in\bar{\Omega}}\{q_m\theta_n - t_m\}$.

Proof. Let $\Omega \in C$. For each pair $n, m \in N$, $n > m$, define

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

where $\bar{\Omega}$ is the closure (in the usual Euclidean sense) of Ω. $M_{n,m}$ is the set of outcome functions in Ω that minimize type θ_n's payoff from reporting type θ_m. This set 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. Since the principal can simply give away the good for free and realize zero expect revenue, it cannot possibly make sense for him to offer any transfer scheme that implies a negative expected revenue. Combining this non-negativity constraint on the expected transfers with the upper bound on the transfers of each type, we can also obtain a lower bound on transfers.

Next we show that $M_{n,m} = M_{n',m}$ for all $n, n' > m$; that is, the set of minimizers $M_{n,m}$ does not vary with n. In order to see this, let $(q^*, t^*) \in M_{n,m}$. By the definition of $M_{n,m}$ we have

$$q^*_m\theta_n - t^*_m \leq q_m\theta_n - t_m \qquad \text{for all } (q, t) \in \bar{\Omega}.$$

By Lemma 2 we may assume that

$$q^*_m\theta_m - t^*_m = q_m\theta_m - t_m \qquad \text{for all } (q, t) \in \bar{\Omega}.$$

Combining these two inequalities we obtain

$$(q^*_m - q_m)(\theta_n - \theta_m) \leq 0 \qquad \text{for all } (q, t) \in \bar{\Omega}.$$

That is, (q^*, t^*) is an outcome function in $\bar{\Omega}$ that minimizes the probability with which the good is transferred when the reported type is θ_m. Since this condition must hold for every type $\theta_n > \theta_m$ we can conclude that $M_{n,m} = M_{n',m}$ for all $n, n' > m$. This entitles us to drop the index n from $M_{n,m}$ and simply write M_m for the set of outcome functions that yield the smallest payoff.

We are now ready to construct $\tilde{\Omega}$. In order to do so, pick for each $m = 1, \ldots, N - 1$, an arbitrary element from the set M_m. The thus obtained collection of outcome functions might not be a set for it might contain multiple replica of the same outcome function.[24] If this is the case,

[24]Notice that a given outcome function might belong both to M_m and $M_{m'}$, $1 \leq m, m' < N$.

eliminate all but one of those replica and label the resulting set by Ω. Next, verify if there is any $(q, t) \in \Omega$ such that whenever $(q, t) \in \arg\min_{(q',t')\in\hat{\Omega}} q'_m\theta_n - t'_m$ for some $1 \leq m < n \leq N$, then the set $\arg\min_{(q',t')\in\hat{\Omega}} q'_m\theta_n - t'_m$ contains at least one other element (i.e. (q, t) is never a unique minimizer). If $\hat{\Omega}$ contains such outcome functions pick one of them arbitrarily and eliminate it from $\hat{\Omega}$. Repeat this last step as long as the reduced ambiguous mechanism continues to contain elements that never uniquely determine any downward deviation payoff. Since $\hat{\Omega}$ contains at most $N - 1$ elements, this procedure must end after at most $N - 2$ steps. Denote the resulting set of outcome functions by $\tilde{\Omega}$.

By construction the payoffs associated with downward deviations are the same in $\tilde{\Omega}$ and $\bar{\Omega}$. Moreover, the fact that $\tilde{\Omega}$ is a subset of $\bar{\Omega}$ allows us to conclude that the truth-telling payoffs can be no smaller in $\tilde{\Omega}$ than in $\bar{\Omega}$. Combining these two observations it follows that $\tilde{\Omega}$ is downward incentive compatible and IR(1) if $\bar{\Omega}$ is so. But this is certainly the case since $\bar{\Omega}$ inherits both properties from Ω. $\tilde{\Omega} \subset \bar{\Omega}$ also implies that $\tilde{\Omega}$ is consistent and that $R(\tilde{\Omega}) = R(\bar{\Omega})$ $(= R(\Omega))$. Since $\tilde{\Omega} \subset \hat{\Omega}$ and $|\hat{\Omega}| \leq N - 1$ it follows that $|\tilde{\Omega}| \leq N - 1$. Finally, that $\tilde{\Omega}$ satisfies condition iii) of the lemma follows directly from its construction. □

In the above lemma we have seen that the 'relevant' elements of an ambiguous mechanism Ω are those that define the downward deviation payoffs. Since under uniformity every type θ_n, $n > m$, worries about the same element of Ω when contemplating to report θ_m, there can be at most $N - 1$ simple mechanisms in Ω that matter.

Lemma 3 has immediate implications for the case of a binary type set.

Corollary 1. *If the type set Θ contains only two elements, then the use of ambiguous mechanisms does not allow the principal to achieve a higher expected revenue than the one that he can obtain with an optimal simple mechanims.*

Henceforth, we refer to ambiguous mechanisms that are non-redundant in the sense that they satisfy condition iii) of the Lemma 3, as *minimal* ambiguous mechanism. In the case of minimal ambiguous mechanisms Ω we use the following notational convention: we indicate the element of Ω that defines the payoff for downward deviations to type θ_m by the superscript m, i.e.

$$\{(q^m, t^m)\} = \arg\min_{(q,t)\in\Omega} q_m\theta_n - \theta_m \qquad \text{for } n > m.$$

If a minimal ambiguous mechanisms contains strictly less than $N - 1$ elments, then there must be at least one outcome function, $(\tilde{q}, \tilde{t})$ say, that determines the downward deviation payoffs for at least two types, m and m' (i.e. $(\tilde{q}, \tilde{t}) = \arg\min_{(q,t)\in\Omega} q_m\theta_n - \theta_m$ for $n > m$ and $(\tilde{q}, \tilde{t}) = \arg\min_{(q,t)\in\Omega} q_{m'}\theta_n - \theta_{m'}$ for $n > m'$). Notationally we treat this case by keeping two copies of $(\tilde{q}, \tilde{t})$, one indexed by m ((q^m, t^m)) and one indexed by m' ($(q^{m'}, t^{m'})$). Thus, from now on minimal ambiguous mechanisms are taken to be sets composed of exactly $N-1$ outcome functions, some of which may differ only in their labels.

Next we show that within the set of P'-feasible ambiguous mechanisms that are minimal and uniform, we can concentrate on mechanisms that involve only allocation functions which are equal to 1 in all but perhaps one coordinate. Moreover, the coordinates of the outcome functions that are allowed to differ from the value 1 satisfy a monotonicity condition defined across outcome functions.

Lemma 4 (Monotonicity and maximal slackness)**.** *For every $\Omega \in C$ there exists a minimal $\tilde{\Omega} = \{(\tilde{q}^1, \tilde{t}^1), \ldots, (\tilde{q}^{N-1}, \tilde{t}^{N-1})\} \in C$ such that $R(\tilde{\Omega}) = R(\Omega)$, and*

i) Monotonicity: for all $1 < m < N$, $\qquad \tilde{q}^m_m \geq \tilde{q}^{m-1}_{m-1}$;

ii) Maximal slackness: for all $n, m \in N, n \neq m$, $\tilde{q}^n_m = 1$.

Proof. Let $\Omega \in C$. By Lemmata 2 and 3 we may assume that Ω is minimal and satisfies uniformity.

If Ω does not satisfy maximal slackness we can simply replace it by an ambiguous mechanism that coincides with Ω except where maximal slackness is violated; those elements are set equal to 1 in the new ambiguous mechanism. In order to see that such a change does not affect downward incentive compatibility, remember that (q^m, t^m) defines the payoffs from downward deviations to θ_m, i.e.

$$q^m_m \theta_l - t^m_m \leq q^n_m \theta_l - t^n_m,$$

for all $N > l > m$ and $1 < n < N-1$, $n \neq m$. If for some $1 < n < N-1$, $n \neq m$, q^n_m is increased to 1, the right hand side of this inequality increases while the left hand side does not change. In other words, the downward deviation payoffs remain the same. On the other hand, the truth-telling payoffs can at most increase. Thus, downward incentive compatibility must continue to hold. The resulting ambiguous mechanism must also be consistent since the transfers of the new mechanism are the same as in Ω.

Next we turn to monotonicity. Suppose that Ω does not satisfy condition i) in the statement of the lemma. Let m^* be the smallest index for which this condition is violated. That is, m^* satisfies

$$q^{m^*}_{m^*} < q^{m^*-1}_{m^*-1} \qquad \text{and} \qquad q^m_m \geq q^{m-1}_{m-1} \quad \text{for all } 1 < m < m^*.$$

We will argue that if $1 \leq m < m^*$, then for type $\theta_n > \theta_m$ the most attractive alternative for a downward deviation in the set $\{\theta_1, \ldots, \theta_m\}$ is θ_m. Notice that this means that for types θ_n, $n \leq m^*$, the binding downward incentive constraint is the one with respect to the downward adjacent type, θ_{n-1}.

In order order to see this, notice that for $1 < m < m^*$ and $n > m$ we have

$$
\begin{aligned}
&[q_m^m\theta_n - t_m^m] - [q_{m-1}^{m-1}\theta_n - t_{m-1}^{m-1}] = \\
&[q_m^m(\theta_n - \theta_m) + q_m^m\theta_m - t_m^m] - [q_{m-1}^{m-1}(\theta_n - \theta_m) + q_{m-1}^{m-1}\theta_m - t_{m-1}^{m-1}] = \\
&[q_m^m - q_{m-1}^{m-1}][\theta_n - \theta_m] + \left\{[q_m^m\theta_m - t_m^m] - [q_{m-1}^{m-1}\theta_m - t_{m-1}^{m-1}]\right\} \geq 0.
\end{aligned}
$$

That the inequality sign in the last line must hold can be seen as follows. By assumption we have $q_m^m - q_{m-1}^{m-1} \geq 0$. Since also $\theta_n - \theta_m > 0$, the product of these two terms must be non-negative. The expression in the curly brackets is non-negative by downward incentive compatibility. Therefore type θ_n's payoff from deviating downward to θ_m is increasing in m.

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

$$
\begin{aligned}
&[q_{m^*}^{m^*}\theta_n - t_{m^*}^{m^*}] - [q_{m^*-1}^{m^*-1}\theta_n - t_{m^*-1}^{m^*-1}] = \\
&[q_{m^*}^{m^*} - q_{m^*-1}^{m^*-1}][\theta_n - \theta_{m^*}] + \{[q_{m^*}^{m^*}\theta_{m^*} - t_{m^*}^{m^*}] - [q_{m^*-1}^{m^*-1}\theta_{m^*} - t_{m^*-1}^{m^*-1}]\} = \\
&[q_{m^*}^{m^*} - q_{m^*-1}^{m^*-1}][\theta_n - \theta_{m^*}] < 0,
\end{aligned}
$$

where the inequality in the last line is due to the starting assumption $q_{m^*}^{m^*} < q_{m^*-1}^{m^*-1}$. Thus, for every type $\theta_n > \theta_{m^*}$ reporting θ_{m^*-1} yields a higher payoff than reporting θ_{m^*}.

Consider now the ambiguous mechanism, Ω', that coincides with Ω, except for the value of $q_{m^*}^{m^*}$, which is increased to $q_{m^*-1}^{m^*-1}$. As we have seen above, doing so does not interfere with downward incentive compatibility. Moreover, since $m^* > 1$ this change cannot affect the truth-telling payoff of the lowest type and so IR(1) must continue to hold. Finally, observe that since Ω' exhibits the same transfers as Ω, it must be consistent and the revenue that it generates is the same as the one delivered by Ω.

By iterating on this argument we obtain an ambiguous mechanism $\tilde{\Omega}$ with the desired properties. □

In passing from Problem P to Problem P' we have dropped the upward incentive constraint. The next result justifies this relaxation.

Lemma 5. *Suppose that $\Omega \in C$ is minimal and satisfies the conditions of uniformity, maximal slackness and monotonicity. Moreover, assume that under Ω the downward incentive compatibility constraint of the highest type is binding. Then Ω is upward incentive compatible.*

Proof. Let $N > n > m \geq 1$. Then

$$\begin{aligned}
\min_{(q,t)\in\Omega}\{q_n\theta_m - t_n\} &\leq q_n^n\theta_m - t_n^n \\
&= q_n^n(\theta_m - \theta_n) + q_n^n\theta_n - t_n^n = q_n^n(\theta_m - \theta_n) + q_{n-1}^{n-1}\theta_n - t_{n-1}^{n-1} \\
&= q_n^n(\theta_m - \theta_n) + q_{n-1}^{n-1}(\theta_n - \theta_{n-1}) + q_{n-1}^{n-1}\theta_{n-1} - t_{n-1}^{n-1} \\
&\vdots \qquad\qquad \vdots \\
&= q_n^n(\theta_m - \theta_n) + q_{n-1}^{n-1}(\theta_n - \theta_{n-1}) + \ldots + q_m^m(\theta_{m+1} - \theta_m) + q_m^m\theta_m - t_m^m \\
&= q_m^m\theta_m - t_m^m - \sum_{k=0}^{n-m-1}(q_{n-k}^{n-k} - q_{n-k-1}^{n-k-1})(\theta_{n-k} - \theta_m) \\
&\leq q_m^m\theta_m - t_m^m \\
&= \min_{(q,t)\in\Omega}\{q_m\theta_m - t_m\}.
\end{aligned}$$

The first inequality is definitional, the second inequality follows from the fact that q_n^n is non-decreasing in n and the last equation is implied by uniformity, which requires that all elements of Ω deliver the same truth-telling payoffs to all types, except possibly type θ_N. For the steps between the two inequalities we exploit the fact that the binding downward incentive constraint is the one with respect to the downward adjacent type.

In the preceding argument we do not allow for $n = N$. The reason for this is purely notational. A perfectly analogous argument can be applied in the case $n = N$ by using in the first row instead of (q^n, t^n) the mechanism that minimizes the truth telling payoff of the highest type. □

Lemma 5 allows us to conclude that Problem P' admits solutions that are upward incentive compatible. Clearly, at every solution of Problem P' the downward incentive compatibility constraint of the highest type must be binding.[25] As a consequence, every candidate for a solution of P' that also satisfies the conditions of minimality, uniformity, monotonicity and maximal slackness is upward incentive compatible. This means that if we want to find a solution to the original Problem P, we can restrict attention to the subset of P'-feasible ambiguous mechanisms that satisfies the before mentioned additional properties.

3.1 The optimal ambiguous mechanism

In this subsection we derive an optimal minimal ambiguous mechanism. Exploiting uniformity and maximal slackness we obtain the following result.

[25] If at a P'-feasible ambiguous mechanism Ω the downward incentive constraint of the highest types does not hold with equality, then the highest type's transfers can be increased in every outcome function by some $\varepsilon > 0$ without violating downward incentive compatibility.

Lemma 6. *Problem P admits as solution a direct ambiguous mechanism that is minimal and satisfies the condition*

$$\begin{aligned} t_1^m &= q_1^m \theta_1 && \text{for all } 1 \le m < N, \\ t_{n+1}^m &= (q_{n+1}^m - q_n^n)\theta_{n+1} + t_n^n && \text{for all } 1 \le n \le N-2, \text{ and } 1 \le m < N \\ t_N^m &\le (1 - q_{N-1}^{N-1})\theta_{N-1} + t_{N-1}^{N-1}, && \text{for all } 1 \le m < N. \end{aligned} \tag{5}$$

Proof. Our previous results imply that the solution set of Problem P contains minimal direct ambiguous mechanisms that satisfy the conditions of uniformity and maximal slackness. Uniformity immediately implies the first two lines of condition 5. The third line of the condition is obtained by rewriting the downward incentive constraint IC(N,N-1) and using $q_N^m = 1$ (by maximal slackness). □

Condition 5 can be rewritten as

$$\begin{aligned} t_n^m &= q_n^m \theta_n - \sum_{k=1}^{n-1} q_k^k (\theta_{k+1} - \theta_k) && \text{for all } 1 \le n \le N-1 \text{ and } 1 \le m \le N-1 \\ t_N^m &\le \theta_N - \sum_{k=1}^{N-1} q_k^k (\theta_{k+1} - \theta_k) && \text{for all } 1 \le m \le N-1 \end{aligned} \tag{6}$$

Condition 6 shows that by choosing a vector $(q_1^1, \ldots, q_{N-1}^{N-1})$ one can pin down all transfers, except those of the highest type. For the transfers of the highest type we only have an upper bound.

Calculating the expected value of the right hand side expressions in (6) we obtain the following upper bound on the expected revenue of $(q^m, t^m) \in \Omega$,

$$\bar{R}^m(q_1^1, \ldots, q_{N-1}^{N-1}) \quad = \quad \mathbb{E}_p[\theta] - p_m(1 - q_m^m)\theta_m - \sum_{n=1}^{N-1} q_n^n (1 - P_n)(\theta_{n+1} - \theta_n),$$

where $P_n = \sum_{k=1}^n p_k$. The objective of the designer is to maximize $\min_m \bar{R}^m(q_1^1, \ldots, q_{N-1}^{N-1})$ over the set of admissible vectors $(q_1^1, \ldots, q_{N-1}^{N-1})$. The transfers of the highest type are then chosen so that consistency holds, i.e. such that $R^n = \min_m \bar{R}^m(q_1^1, \ldots, q_{N-1}^{N-1})$ for all $1 \le n \le N-1$.

In what follows we denote the vector $(q_1^1, \ldots, q_{N-1}^{N-1})$ by Q. We say that an ambiguous mechanism $\tilde{\Omega}$ is generated or induced by Q if $\tilde{\Omega}$ is composed of $N-1$ elements that satisfy consistency and maximal slackness and whose allocation components 'along the diagonal', $(\tilde{q}_1^1, \ldots, \tilde{q}_{N-1}^{N-1})$, coincide with Q. For later reference we point out that $R^m(Q)$ is decreasing in q_n^n, $n \neq m$. Given the preceding observations we can restate the principal's problem in the following simple form:

$$\max_{R\in\mathbb{R}, Q\in[0,1]^{N-1}} \quad R \qquad \text{(P”)}$$

$$\text{subject to} \quad R \le \bar{R}^n(Q) \quad \text{for all } 1 \le n \le N-1,$$

$$q_n^n \ge q_{n-1}^{n-1} \quad \text{for all } 1 < n \le N-1.$$

For the presentation of the next results it is convenient to introduce some further notation and terminology. First, we inductively construct the set $\mathcal{M} = \{m_1, \ldots, m_M, m_{M+1}\}$, which is a subset of the index set N. The first element, m_1, is set equal to 1. If for m_{j-1} the set $\{n : N > n > m_{j-1}, p_n\theta_n > p_{m_{j-1}}\theta_{m_{j-1}}\}$ is non-empty, we set $m_j = \min\{n : N > n > m_{j-1}, p_n\theta_n > p_{m_{j-1}}\theta_{m_{j-1}}\}$. Let m_M be the largest index defined in this way and set $m_{M+1} = N$. Observe that if $p_n\theta_n$ is increasing in n, then $\mathcal{M}$ coincides with the set N. Also notice that $p_{m_j}\theta_{m_j}$ is monotonic in $j = 1, \ldots, M$ by construction.

Next we define for all $1 \le j \le M$ the so called *adjusted virtual valuation*, $\bar{v}_{m_j}$:

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

We refer to $\bar{v}_{m_j}$ as adjusted virtual valuation because both its definition and its role are reminiscent of the role of virtual valuations.[26] In particular, in Proposition 2 below we show that the optimal value of Q depends on the signs of the adjusted virtual valuations. In the statement of this result we will exploit the fact that the adjusted virtual valuation can cross the zero only from below. This is shown in the following lemma.

Lemma 7. *If $\bar{v}_{m_j} \le 0$ for $1 < j \le M$, then $\bar{v}_{m_k} \le 0$ for all $1 \le k < j$.*

Proof. In order to see this, we rewrite the virtual valuation $\bar{v}_{m_j}$ in the form

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

The sign of $\bar{v}_{m_j}$ is determined by the expression in the square brackets. It is easy to verify that this term is increasing in j. Thus, if it is negative for a given $1 < j \le M$ then it must be so also for all $1 \le k < j$. □

We are now ready to state the main result of this section in which we characterize one of the solutions of Problem P”.

[26] Strictly speaking the adjusted virtual valuation $\bar{v}_{m_j}$ resembles more the product of the virtual valuation of type θ_{m_j} and its probability p_{m_j} than the virtual valuation itself.

Proposition 2.

i) If $\bar{v}_1 > 0$*, then* $(\hat{q}_1^1, \ldots, \hat{q}_{N-1}^{N-1}) = (1, \ldots, 1)$ *solves Problem P".*

ii) If $\bar{v}_1 \le 0$*, let* $j^* = \max\{j : \bar{v}_{m_j} \le 0\}$ *and let* $\hat{Q} = (\hat{q}_1^1, \ldots, \hat{q}_{N-1}^{N-1})$ *be defined by*

$$\hat{q}_n^n = \begin{cases} 0 & \text{if } n < m_{j^*+1} \\ 1 - \frac{p_{m_{j^*}} \theta_{m_{j^*}}}{p_{m_j} \theta_{m_j}} & \text{if } j^* + 1 \le j \le M \text{ and } m_j \le n < m_{j+1}. \end{cases}$$

$\hat{Q}$ *constitutes a solution of P".*

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

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

In order to see this observe that since for every Q with non-decreasing components, we have $q_n^n \ge q_{m_j}^{m_j}$ it follows that

$$\bar{R}^n(Q) - \bar{R}^{m_j}(Q) = -p_n\theta_n(1 - q_n^n) + p_{m_j}\theta_{m_j}(1 - q_{m_j}^{m_j}) \ge (1 - q_{m_j}^{m_j})(p_{m_j}\theta_{m_j} - p_n\theta_n) \ge 0.$$

That is, there is no admissible Q for which $\bar{R}^n(Q)$ is the (strictly) smallest upper bound on the revenues. But $\bar{R}^n(Q)$ is the only bound that could be increasing in q_n^n. Thus, it is without loss to choose q_n^n as small as possible, i.e. we can set $\hat{q}_n^n = \hat{q}_{n-1}^{n-1}$. Since this argument applies to all $m_j < n < m_{j+1}$ we can conclude that choosing $\hat{q}_n^n = \hat{q}_{m_j}^{m_j}$ for all $m_j < n < m_{j+1}$ is optimal.

Step 2. At the optimum

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

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

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

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

Moreover, rewriting the inequality yields

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

In such a case we can lower $q_{m_{j+1}}^{m_{j+1}}$ without violating the constraint $q_{m_{j+1}}^{m_{j+1}} \ge q_{m_j}^{m_j}$, and thus increase all $\bar{R}^n$, $n \neq m_{j+1}$. Since $\bar{R}^{m_{j+1}}$ is not the smallest bound this means that the minimum of the bounds would increase. But then Q cannot be optimal.

Step 3. If $\bar{v}_1 \leq 0$ then at the optimum

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

if $\bar{v}_1 > 0$ 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

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

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

For the case that j lies strictly between j^* and M (i.e. $j^* < j < M$) assume that we have shown the claim for $s = j + 1, \ldots, M$. 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

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

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 $m_{s-1} < n < m_s$, $s = j+1, \dots, M$,

$$q_n^n = q_{m_s}^{m_s}.$$

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_n^n, $m_{s-1} \le n < m_s$, $s = j+1, \dots, M$, 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 $\bar{R}^{m_j}$ changes at the 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 $\bar{v}_1 \le 0$ then at the optimum $\hat{q}_{m_j}^{m_j} = 0$ for all $j \le j^*$.

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

$$\hat{q}_{m_s}^{m_s} = 1 - \frac{p_{m_{j^*}}\theta_{m_{j^*}}}{p_{m_s}\theta_{m_s}}(1 - \hat{q}_{m_{j^*}}^{m_{j^*}}) \tag{7}$$

holds. Thus, varying $q_{m_{j^*}}^{m_{j^*}}$ implies that we have to change accordingly also all q_n^n, $m_{j^*} < n < N$. In the previous step we have seen that the overall effect that such a change has on $\bar{R}^{m_{j^*}}$ is measured by $\bar{v}_{m_{j^*}}$. Thus, if $\bar{v}_{m_{j^*}} \le 0$, then $\bar{R}^{m_{j^*}}$ is maximized by choosing $q_{m_{j^*}}^{m_{j^*}}$ as small as possible. But that means that we have to set $q_{m_{j^*}}^{m_{j^*}} = q_{m_{j^*-1}}^{m_{j^*-1}}$.

Next, consider the choice of $q_{m_{j^*-1}}^{m_{j^*-1}}$. If $q_{m_{j^*}}^{m_{j^*}} = q_{m_{j^*-1}}^{m_{j^*-1}}$, then

$$\bar{R}^{m_{j^*}}(Q) - \bar{R}^{m_{j^*-1}}(Q) = (1 - q_{m_{j^*-1}}^{m_{j^*-1}})(p_{m_{j^*-1}}\theta_{m_{j^*-1}} - p_{m_j^*}\theta_{m_j^*}). \tag{8}$$

If $q^{m_{j^*-1}}_{m_{j^*-1}} < 1$ this expression is strictly negative, meaning that $\bar{R}^{m_{j^*-1}}$ is not the smallest one of the bounds. Since all other bounds are strictly decreasing in $q^{m_{j^*-1}}_{m_{j^*-1}}$, so must be $\min_j \bar{R}^{m_j}$. If $q^{m_{j^*}}_{m_{j^*}} = q^{m_{j^*-1}}_{m_{j^*-1}} = 1$, then $\bar{R}^{m_{j^*-1}}$ can be increased by a decrease of $q^{m_{j^*-1}}_{m_{j^*-1}}$ that is accompanied with a reduction of all q^n_n, $m_{j^*-1} < n < N$, in accordance with (7). In order to see this notice that by (8) we know that in the initial situation we have $\bar{R}^{m_{j^*}} = \bar{R}^{m_{j^*-1}}$. After the proposed reduction of all q^n_n, $m_{j^*-1} \leq n < N$ instead we have $\bar{R}^{m_{j^*}} < \bar{R}^{m_{j^*-1}}$. By our previous arguments we know that a reduction of $(q^{m_{j^*}}_{m_{j^*}}, \ldots, q^{N-1}_{N-1})$ in accordance with (7) leads to an increase of $\bar{R}^{m_{j^*}}$ and $\min_j \bar{R}^{m_j}$. If in addition also $(q^{m_{j^*-1}}_{m_{j^*-1}}, \ldots, q^{m_{j^*}-1}_{m_{j^*}-1})$ is reduced then certainly $\bar{R}^{m_j}$, $j \neq j^* - 1$, increase further. Moreover, since after the change $\bar{R}^{m_{j^*}} < \bar{R}^{m_{j^*-1}}$ it must be the case that also $\bar{R}^{m_{j^*-1}}$ increases. Combining these arguments we conclude that $q^{m_{j^*-1}}_{m_{j^*-1}}$ must be chosen as small as possible, i.e. $q^{m_{j^*-1}}_{m_{j^*-1}} = q^{m_{j^*-2}}_{m_{j^*-2}}$.

Iterating on the same argument we can show that for all $m_j \leq m_{j^*}$, $q^{m_j}_{m_j}$ 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 $m_j \leq m_{j^*}$.

Step 5. If $\bar{v}_1 > 0$, then at the optimum $q^{m_j}_{m_j} = 1$ for all $1 \leq j \leq M$.

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

$$\hat{q}^{m_j}_{m_j} \leq 1 - \frac{p_{m_{j-1}}\theta_{m_{j-1}}}{p_{m_j}\theta_{m_j}}(1 - \hat{q}^{m_{j-1}}_{m_{j-1}}) \tag{9}$$

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^n_n, $1 < n < N - 1$, must be equal to 1.

□

Proposition 2 gives us a solution for Problem P''. Given $\hat{Q}$ it is straightforward to calculate the problem's optimal value $\hat{R}$. In particular, for all $1 \leq n, m < N$ the optimal transfer $\hat{t}^m_n$ can be obtained from (6). The highest type's transfers are then chosen so that the expected revenue of each of the $N - 1$ outcome functions is equal to the optimal value of Problem P'', $\hat{R} = \min_j \bar{R}^{m_j}(\hat{Q})$. We summarize these observations in the following corollary.

Corollary 2. *Suppose $\hat{Q} = (\hat{q}^1_1, \ldots, \hat{q}^{N-1}_{N-1})$ solves Problem P'' and that $\hat{R}$ is the problem's value. Then the corresponding transfers of the outcome function $(\hat{q}^m, \hat{t}^m)$, $1 \leq m < N$, are given by*

$$\hat{t}^m_n = \begin{cases} \hat{q}^m_n\theta_n - \sum_{k=1}^{n-1} \hat{q}^k_k(\theta_{k+1} - \theta_k) & \text{if } 1 \leq n < N \\ \hat{R} - \sum_{n=1}^{N-1} p_n\hat{t}^m_n & \text{if } n = N. \end{cases}$$

If $\bar{v}_1 > 0$, then the optimal value of the designer's problem is $\hat{R} = \theta_1$. Otherwise, the optimal *expected revenue is*

$$\hat{R} = \bar{R}^{m_{j^*}} = \mathbb{E}_p[\theta] - p_{m_{j^*}}\theta_{m_{j^*}} - \sum_{n=m_{j^*+1}}^{N-1} \hat{q}_n^n(1-P_n)(\theta_{n+1}-\theta_n).$$

Finally, it is interesting to compare the expected revenue of an optimal ambiguous mechanism with the one of the best simple mechanism. Of course, every simple mechanism constitutes a (trivial) ambiguous mechanism. Thus, simple mechanisms cannot possibly deliver a strictly higher revenue than the optimal ambiguous mechanism. But when is it the case that the designer can do strictly better by using a 'truly ambiguous' mechanism?

Consider an arbitrary simple IC and IR mechanism $(\tilde{q}, \tilde{t})$ and define the ambiguous mechanism $\Omega = \{(q^1, t^1), \dots, (q^{N-1}, t^{N-1})\}$ as follows

$$q_n^m = \begin{cases} \tilde{q}_n & \text{if } n = m \\ 1 & \text{else,} \end{cases} \qquad t_n^m = \begin{cases} \tilde{t}_n & \text{if } n = m \\ \tilde{t}_n + (1-\tilde{q}_n)\theta_n & \text{if } n \neq m, N \\ \tilde{t}_N - [p_k\theta_{m_M}(1-\tilde{q}_{m_M}) - p_m\theta_m(1-\tilde{q}_m)]/p_N & \text{if } n = N, \end{cases}$$

where m_M is defined as before, i.e. $m_M \in \arg\max_{n<N} p_n\theta_n$.

Incentive compatibility of the simple mechanism $(\tilde{q}, \tilde{t})$ implies $\tilde{q}_n \leq \tilde{q}_{n+1}$, $1 \leq n < N$. Observe also that for all types θ_n, $n < N$, the truth-telling payoffs are constant across the outcome functions in Ω; moreover they coincide with the truth-telling payoffs under $(\tilde{q}, \tilde{t})$. By construction of Ω the deviation payoffs cannot be larger than the deviation payoffs in $(\tilde{q}, \tilde{t})$. Hence, Ω is both incentive compatible and individually rational. Finally, it can easily be verified that Ω is consistent. The expected revenue generated by $(q^m, t^m) \in \Omega$ is

$$R^m = \sum_{n\leq N} p_n\tilde{t}_n + \sum_{n\neq m_M, N}(1-\tilde{q}_n)\theta_n.$$

Thus Ω delivers a strictly larger expected revenue than $(\tilde{q}, \tilde{t})$ if $\tilde{q}_n < 1$ for some $n < N$.

This means that if the optimal ambiguous mechanism and the optimal simple mechanism yield the same expected revenue, then it must be the case that the optimal simple mechanism assigns the object to every type with probability one. Such a simple mechanism yields the same revenue ($R = \theta_1$) as the ambiguous mechanism induced by $Q = (1, \dots, 1)$; thus the latter must be optimal.

Conversely, assuming that the ambiguous mechanism induced by $Q = (1, \dots, 1)$ is optimal implies that the principal cannot do better by using ambiguous mechanisms than he can by using a simple mechanism.

The following proposition summarizes these observations.

Proposition 3. *An optimal ambiguous mechanism yields a strictly larger revenue than the best simple mechanism if and only if* $\bar{v}_1 < 0$.

We conclude this section with a three-type example that illustrates the above developed results.

Example 2 (Optimal ambiguous mechanisms in the three type case)**.**
Suppose that $\Theta = \{\theta_1, \theta_2, \theta_3\}$. The formula for the optimal Q given in Proposition 2 conditions on the signs of the adjusted virtual valuations. The adjusted virtual valuations in turn depend on the composition of the set $\mathcal{M}$. Remember that $\mathcal{M} = \{m_1, \ldots, m_{M+1}\}$ is a subset of type indices such that $p_{m_j}\theta_{m_j}$ is increasing in j. With three types there are only two possibilities: either i) $p_1\theta_1 > p_2\theta_2$ or ii) $p_1\theta_1 \leq p_2\theta_2$.

i) $p_1\theta_1 > p_2\theta_2$: In this case we have $\mathcal{M} = \{1, 3\}$; i.e. $\mathcal{M}$ does not include 2. Consequently, q_2^2 is always chosen equal to q_1^1 and so we either have $q_1^1 = q_2^2 = 1$ or $q_1^1 = q_2^2 = 0$, depending on whether $\bar{v}_1 > 0$ or $\bar{v}_1 \leq 0$. Notice that $\bar{v}_1$ takes the value

$$\bar{v}_1 = p_1\theta_1 - (1 - p_1)(\theta_2 - \theta_1) - (1 - P_2)(\theta_3 - \theta_2) = \theta_1 - p_2\theta_2 - p_3\theta_3.$$

Hence, $q_1^1 = q_2^2 = 1$ is optimal if θ_1 is larger than the two larger types' contribution to $\mathbb{E}_p[\theta]$. From (6) it follows that all transfers for the two lower types are equal to θ_1. Since by Corollary 2, $\hat{R} = \theta_1$, it follows that also the highest type's transfers are equal to θ_1. $\hat{Q} = (1, 1)$ means that every outcome function of the ambiguous mechanism specifies to give the good with probability one to the agent, irrespective of his message. Incentive compatibility then requires that the transfers do not change either with the reported type. The maximal transfer that is compatible with the lowest type's individual rationality constraint is to have him pay his valuation.

Notice also that $\hat{Q} = (1, 1)$ means that the two outcome functions (q^1, t^1) and (q^2, t^2) coincide. Thus, in this case the designer can achieve the maximal expected revenue by offering a simple mechanism.

If $\bar{v}_1 \leq 0$, then it is optimal to set $q_1^1 = q_2^2 = 0$. That is, each of the two outcome functions excludes one of the two lower types, but neither of them excludes both. According to Corollary 2 the expected revenue in this case is $\hat{R} = p_2\theta_2 + p_3\theta_3$. Using (6) once more we obtain $t_1^1 = 0$, $t_2^1 = \theta_2$, $t_1^2 = \theta_1$ and $t_2^2 = 0$. Finally, consistency implies $t_3^1 = \theta_3$ and $t_3^2 = \theta_3 - [p_1\theta_1 - p_2\theta_2]/p_3$.

Clearly there is no simple mechanism that achieves an expected revenue of $\hat{R} = p_2\theta_2 + p_3\theta_3$. An optimal simple mechanism takes one of the following three forms: i) the good is given to every type with probability one at the price θ_1, ii) the good is given to types two and three at the price θ_2 or iii) it specifies that only the highest type gets the good at the price of his valuation. Option i) generates an expected revenue of θ_1 which is by assumption ($\bar{v}_1 \leq 0$) smaller than $\hat{R}$. The revenues under options ii), $p_2(\theta_2 + \theta_3)$, and iii), $p_3\theta_3$, are clearly smaller than $p_2\theta_2 + p_3\theta_3$.

This confirms the result in Proposition 3.

ii) $p_1\theta_1 \leq p_2\theta_2$: In this case $M = \{1, 2, 3\}$, implying that the choices of both q_1^1 and q_2^2 are non-trivial and depend on the sign of both $\bar{v}_1$ and $\bar{v}_2$. These two variables now take the values

$$\begin{aligned} \bar{v}_1 &= p_1\theta_1 - (1-p_1)(\theta_2-\theta_1) - \frac{p_1\theta_1}{p_2\theta_2} p_3(\theta_3-\theta_2) = \theta_1 - (1-p_1)\theta_2 - \frac{p_1\theta_1}{p_2\theta_2} p_3(\theta_3-\theta_2) \\ \bar{v}_2 &= p_2\theta_2 - (1-P_2)(\theta_3-\theta_2) = (p_2+p_3)\theta_2 - p_3\theta_3. \end{aligned}$$

$\bar{v}_1$ is slightly larger than in case i) (the difference between the two expressions is the smaller the closer $p_1\theta_1$ is to $p_2\theta_2$). $\bar{v}_2$ instead is given by the product of the (regular) virtual valuation of type 2 and his probability.

As in case i) it is optimal to set $q_1^1 = q_2^2 = 1$ if $\bar{v}_1 > 0$. If $\bar{v}_1 \leq 0$ then the optimal value of q_1^1 is 0. But unlike before, $\bar{v}_1 \leq 0$ no longer implies $q_2^2 = 0$. Instead, the optimal value of q_2^2 depends on the sign of $\bar{v}_2$. More specifically, in order for $q_1^1 = q_2^2 = 0$ to be optimal, it must be the case that both $v_1 \leq 0$ and $v_2 \leq 0$. In the remaining case ($v_1 \leq 0$ and $v_2 > 0$) we obtain the solution $q_1^1 = 0$ and $q_2^2 = 1 - p_1\theta_1/p_2\theta_2$.

As for the transfers, we obtain $t_1^1 = 0, t_2^1 = \theta_2, t_3^1 = \theta_2 + p_1\theta_1(\theta_3-\theta_2)/p_2\theta_2$ and $t_1^2 = \theta_1, t_2^2 = (1 - p_1\theta_1/p_2\theta_2)\theta_2, t_3^2 = \theta_2 + p_1\theta_1(\theta_3-\theta_2)/p_2\theta_2$. The expected value of these transfers is $\hat{R} = (p_2+p_3)\theta_2 + p_1p_3\theta_1(\theta_3-\theta_2)/p_2\theta_2$. This revenue exceeds the revenue achieved by the revenue maximizing simple mechanism, $(p_2+p_3)\theta_2$. □

4 Discussion and extensions

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

In the preceding section we have seen that if $\hat{Q}$ solves Problem P" then the ambiguous mechanism that it generates must yield a (weakly) higher expected revenue than the one produced by the ambiguous mechanism that corresponds to $Q = (0, \ldots, 0)$.[27] This latter mechanism takes a particularly simple form: the transfer rule corresponding to the m-th outcome function, t^m, is given by

$$t_n^m = \begin{cases} 0 & \text{if } n = m \\ \theta_n & \text{if } n \neq n, N \\ \theta_N - (p_{m_M}\theta_{m_M} - p_m\theta_m)/p_N & \text{if } n = N. \end{cases}$$

[27] We say that $\Omega = \{(q^1, t^1), \ldots, (q^{N-1}, t^{N-1})\}$ is generated by Q, if Ω satisfies maximal slackness, consistency, $(q_1^1, q_2^2, \ldots, q_{N-1}^{N-1}) = Q$ and condition (5).

The expectation of this transfer is $R = \mathbb{E}_p[\theta] - p_{m_M}\theta_{m_M}$. So the ambiguous mechanism generated by $Q = (0, \ldots, 0)$ extracts all of the agent's (expected) surplus except for type θ_{m_M}'s contribution, $p_{m_M}\theta_{m_M}$. The part of the surplus that is left to the agent, $p_{m_M}\theta_{m_M}$, is small if the probability of each single type (and thus also type θ_{m_M}) is small, as it can be the case in settings with 'large' type sets. These observations suggest, that in environments with large type sets, the designer can essentially extract the full rent from the agent. The following proposition gives a more precise formulation of this insight.

Proposition 4 (Full surplus extraction in the limit)**.** *Let* $\{\Theta^N, p^N\}_N$ *be a sequence of finite environments, such that* $|\Theta^N| = N$. *Assume the limit* $\lim_{N\to\infty} \mathbb{E}_{p^N}[\theta^N]$ *exists. Moreover, let* $\bar{m}_N$ *be such that* $p^N_{\bar{m}_N}\theta^N_{\bar{m}_N} \geq p^N_l \theta^N_l$ *for all* $1 \leq l \leq N-1$ *and write* $\hat{R}^N$ *for the revenue that the designer can generate with the an optimal ambiguous mechanism in the* N*-th environment. If* $p^N_{\bar{m}_N}\theta^N_{\bar{m}_N} \overset{N\to\infty}{\longrightarrow} 0$ *then*

$$\frac{\hat{R}^N}{\mathbb{E}_{p^N}[\theta^N]} \overset{N\to\infty}{\longrightarrow} 1.$$

That is, in the limit the designer is able to extract all of the agent's surplus.

Proof. By our preceding observations for all N we have

$$\mathbb{E}_{p^N}[\theta^N] \geq \hat{R}^N \geq \mathbb{E}_{p^N}[\theta^N] - p^N_{\bar{m}_N}\theta^N_{\bar{m}_N}.$$

Dividing both sides by $\mathbb{E}_{p^N}[\theta^N]$ and taking the limit yields the result. □

In order to get a better intuition for this result, consider again the type of ambiguous mechanism described above. In such a mechanism, for each $n < N$, outcome function (q^n, t^n) assigns the good with probability one to every type except type θ_n, who is excluded from trade (i.e. he receives the good with probability zero). Moreover, under (q^n, t^n) all types, except θ_n and θ_N, are charged their valuations. The fact that under (q^n, t^n) type θ_n does not get the good not only implies that type θ_n himself cannot get a strictly positive payoff from revealing his type, but it also means that no other type can achieve a strictly positive payoff from reporting θ_n either. Thus, the outcome function (q^n, t^n) guarantees that (downward) deviations toward θ_n are unattractive. In the same way each other outcome function (q^m, t^m), $m \neq n$ makes sure that the agent does not have an incentive to report θ_m unless that is his true type. Since each single outcome function in the ambiguous mechanism has to take care of the deviation incentives toward just one type, they can be chosen freely (i.e. unconstrained by incentive considerations) for all other possible reports. In particular, it is feasible to specify that for each other message (except θ_N) the agent gets the good for sure in exchange of a payment that corresponds to his report. The highest type does not necessarily have to pay his valuation since his transfers are used to guarantee consistency across outcome functions.[28]

The downside of a types's exclusion from trade is that no rent can be extracted from him. Since all outcome functions must yield the same expected revenue, all of them can extract only as much as the one that excludes the type with the largest contribution to the expected surplus. If the set of types increases and the likelihood of each single type decreases, the cost of excluding each single type decreases as well.

In a context with a continuum of types and an atomless type distribution, the weight of each single type is exactly zero. For such environments, we obtain the following corollary to Proposition 4.

Corollary 3 (Full surplus extraction)**.** *Suppose that Θ is a compact interval in $\mathbb{R}$ and that the type distribution P is atomless. Then the ambiguous mechanism, $\Omega = \{(q^\theta, t^\theta), \theta \in \Theta\}$, where (q^θ, t^θ) is defined by*

$$q^\theta_{\theta'} = \begin{cases} 0 & \text{if } \theta' = \theta \\ 1 & \text{else} \end{cases} \qquad t^\theta_{\theta'} = \begin{cases} 0 & \text{if } \theta' = \theta \\ \theta' & \text{else,} \end{cases}$$

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

Corollary 3 is important not only because it tells us that the designer can achieve full surplus extraction by using an appropriately constructed ambiguous mechanism. An even more important insight that we can derive from this result is that in situations where type sets are large (i.e. continua) and the type distributions are not too concentrated on single points (i.e. atom less), it is possible to design an ambiguous mechanism that achieves full surplus extraction *without knowing the details of the type distribution*.

4.2 Payoff irrelevant information and the 'splitting' of types

In the preceding (sub-)section we have seen that the share of the surplus that the designer can extract from the agent is the larger 'the more types there are'. In particular, if types are distributed atomless on an interval then full surplus extraction is possible. In this section we use this insight to argue that the principal should not only elicit the agent's payoff types, but that he can benefit also from conditioning outcomes on non-payoff relevant information that the agent may hold.

[28] In the case of simple mechanisms all deviation incentives have to be taken care of by a single outcome function. In order to do so this single outcome function needs to be distorted much more severely than each single element of an ambiguous mechanism.

In order to see this, consider again our basic set up with N payoff relevant types, $\Theta = \{\theta_1, \dots, \theta_N\}$. Now assume that θ is only one component of the agent's type. The second component, v, is payoff irrelevant and takes values in the (finite) set $V = \{v_1, \dots, v_K\}$; for convenience, let $V \subset \mathbb{R}$. Denote the distribution of the (bi-dimensional) type by π and assume that the principal knows this distribution.

Even though the type set of this environment is bi-dimensional the results from the preceding section carry over also to this context if we endow $\Theta \times V$ with the lexicographic order (where payoff relevant types constitute the first criterion). In particular, we can construct an optimal ambiguous mechanism as described in Proposition 2 and Corollary 2.

Using the payoff irrelevant part of a type serves the purpose of 'splitting' payoff types into subtypes. Doing so generates a larger number of types who all have a smaller probability. We have seen in the previous subsection why it is desirable from the designer's perspective to have many types who are all not very likely to occur. The insights from that section do not rest on the assumption that types are different in terms of payoff relevance. Instead they also apply when two types differ only in non-payoff relevant dimensions. We demonstrate this in the next example.

Example 3 (The benefits of eliciting payoff irrelevant information)**.** Consider the following simple environment. The type set is given by $\Theta \times V$, where $\Theta = \{1, 3\}$ and $V = \{L, H\}$. The type distribution is uniform and the type set is endowed with the obvious lexicographic ordering.

If the principal ignores the payoff irrelevant part of the agent's type it is as if facing an agent with only two (equally likely) types, 1 and 3. Remember that by Corollary 1 the best mechanism that the designer can offer is a simple mechanism. It is straightforward to see that the optimal simple mechanism, $(\tilde{q}, \tilde{t})$, is defined by $(\tilde{q}(1), \tilde{t}(1)) = (0, 0), (\tilde{q}(3), \tilde{t}(3)) = (1, 3)$. The expected revenue generate by this mechanism is $\tilde{R} = 3/2$.

Now assume that the designer takes into account also the payoff irrelevant component of the agent's type. Then according to Proposition 2 he should offer the ambiguous mechanism Ω composed by the outcome functions described in the following table.[29]

(θ, v)	$(1, L)$	$(1, H)$	$(3, L)$	$(3, H)$
$(q^{(1,L)}, t^{(1,L)})$	$(0, 0)$	$(1, 1)$	$(1, 3)$	$(1, 3)$
$(q^{(1,H)}, t^{(1,H)})$	$(1, 1)$	$(0, 0)$	$(1, 3)$	$(1, 3)$
$(q^{(3,L)}, t^{(3,L)})$	$(1, 1)$	$(1, 1)$	$(2/3, 2)$	$(1, 3)$

It is easily verified that Ω generates an expected revenue of $7/4 > 3/2$. □

[29]In this case we have $\mathcal{M} = \{(1, L), (3, L), (3, H)\}$, $\bar{v}_{(1,L)} = -3/4$ and $\bar{v}_{(3,L)} = 3/4$.

'Active type splitting' or the artificial 'creation' of types. In the preceding discussion we have seen that the principal can benefit from adopting an ambiguous mechanism that elicits not only payoff relevant information but also payoff irrelevant aspects of the agent's type. But if the principal can take advantage of an agent's payoff irrelevant information, then even if the agent does not have such information to start with, he should induce him to acquire it. A simple way to achieve this is to instruct the agent to take a draw from some distribution. If this distribution is atom free doing so allows the principal to extract the agent's full surplus.

Notice that this 'type creation process' must take place before the revelation game is played. Thus, ambiguous mechanisms that are based on type creation do not belong to the class of static ambiguous mechanisms that we have considered so far. Consequently, the discussion in the preceding paragraph is not in contradiction with our findings in the earlier sections where we have derived the optimal ambiguous mechanism for a *given* finite set of types. Moreover, the possibility of creating types does not reduce the relevance of those findings. On the one hand the analysis for a given type set is by itself of theoretical interest. On the other hand, that analysis constitutes the basis upon which our discussion of the benefits of type splitting rests. Finally, also from a more applied perspective the preceding results retain their importance. Practicality considerations (complexity costs) might well impose limits on increasing the number of outcome functions in the ambiguous mechanism. Whenever that is the case the designer needs to understand the trade off between the costs and benefits of any additional type. Our results allow us to determine with precision the benefits of larger type sets.

4.3 Preferences

The agent's preferences: Throughout our analysis we have assumed that the agent's valuation is (bi-)linear and that his ambiguity aversion can be captured by the Gilboa-Schmeidler model. In this section we comment on the role of these assumptions.

The linearity of the agent's valuation function is crucial in the final steps of the characterization of the optimal ambiguous mechanism (i.e. Proposition 2 relies on this assumption). In all results up to Lemma 5 we have only exploited the increasing difference property that the linear valuation function exhibits. That is, all those results go through for any other valuation function that exhibits increasing differences. The result that under an atomless type distribution the principal can extract the full surplus goes through in even more general settings. If the agent's preferences over allocation-transfer pairs (x, τ)[30] are described by the function $u(x, \tau, \theta)$, then an ambiguous mechanism like the one used in Corollary 3 can be constructed whenever

[30] Here the variable x is to be interpreted as share of the good.

the problem

$$\max_{(q(\theta),t(\theta))\in X^\Theta\times\mathbb{R}^\Theta} \quad \mathbb{E}_p[t(\theta)]$$
$$\text{s.t.} \quad u(q(\theta),t(\theta),\theta) \geq u(0,0,\theta) \qquad \forall\theta\in\Theta,$$

admits a solution.[31] If (q^*, t^*) solves this problem, then the ambiguous mechanism $\Omega = \{(q^\theta, t^\theta), \theta \in \Theta\}$ whose elements are defined by

$$q^\theta_{\theta'} = \begin{cases} 0 & \text{if } \theta' = \theta \\ q^*(\theta') & \text{else} \end{cases} \qquad t^\theta_{\theta'} = \begin{cases} 0 & \text{if } \theta' = \theta \\ t^*(\theta') & \text{else,} \end{cases}$$

extracts the full surplus.

A concern regarding our assumptions on preferences might be the question to what extent our results are driven by the way in which we model ambiguity aversion. The Gilboa-Schmeidler preferences that we use certainly constitute a rather stark model of ambiguity aversion. Our analysis heavily exploits the tractability of MMEU preferences in the derivation of the optimal ambiguous mechanism with finite types. While we do not know how an optimal mechanism would look like for an alternative model of ambiguity aversion, we can say that the basic idea on which the analysis in this paper builds, does generalize. The most fundamental insight of this paper is that a principal who faces an ambiguity averse agent should try to exploit his ambiguity aversion by offering an ambiguous mechanism. In the following example we show that this insight applies also in environments where the agent's attitude toward uncertainty can be captured by a smooth ambiguity model.

Example 4 (Smooth ambiguity aversion)**.** The setup is as in Example 1, except for the agent' attitude towards ambiguity. That is, we have $\Theta = \{1, 2, 4\}$, $p = (1/4, 1/4, 1/2)$, $u(x, \tau, \theta) = x\theta - \tau$.

We now consider three possible assumptions regarding the agent's attitude toward uncertainty. First, suppose the agent is an (ambiguity neutral) expected utility maximizer. In that case the optimal (non-ambiguous) mechanism for this environment prescribes to exclude the two lowest types from trade. An agent who reports the highest type obtains the good with probability one and pays his valuation, 4. The revenue generated by this mechanism is $R = 2$

Next, assume that the agent has MMEU preferences. In Example 1 we have seen that in that case the principal can achieve a revenue of at least $R = 9/4$ by offering an ambiguous mechanism.

Finally, consider the case where the agent's attitude towards uncertainty can be described

[31] We continue to assume that by opting out from the mechanism each type of the agent obtains the allocation-transfer pair $(0, 0)$.

by a model of smooth ambiguity aversion a la Klibanoff, Marinacci, and Mukerji (2005). In particular, assume that given a (direct) ambiguous mechanism Ω, type θ of the agent evaluates messages according to the following procedure. First, he calculates for each message $\hat{\theta} \in \Theta$ and each possible probability $\pi \in \Delta(\Omega)$ his expected utility, i.e.

$$\mathbb{E}_\pi \left[u(q(\hat{\theta}), t(\hat{\theta}), \theta) \right] = \mathbb{E}_\pi \left[q(\hat{\theta})\theta - t(\hat{\theta}) \right].$$

In a second step, he evaluates the thus obtained expected utility values with the increasing and concave function $\phi : \mathbb{R} \to \mathbb{R}$. Finally, the transformed utility indices are integrated with respect to some probability measure μ over $\Delta(\Omega)$. The payoff that type θ of the agent associates with reporting type $\hat{\theta}$ is

$$U(\hat{\theta}, \theta) = \mathbb{E}_\mu \left\{ \phi \left(\mathbb{E}_\pi \left[q(\hat{\theta})\theta - t(\hat{\theta})) \right] \right) \right\}.$$

The function ϕ and the distribution μ describe the agent's attitude towards uncertainty. μ describes the relative weight that the agent assigns to the possible beliefs that he can hold after learning the ambiguous mechanism. The shape of the function ϕ captures the agent's degree of ambiguity aversion. A linear ϕ means that the agent is ambiguity neutral, i.e. exposing him to uncertainty does not generate any cost to him. A strictly concave ϕ instead corresponds to an agent who is strictly ambiguity averse.

In what follows we assume that $\phi(x) = 1 - \exp(-2x)$, i.e. ϕ has the shape of a CARA function. As for μ, we assume that it is uniform over Ω (or the set of degenerate distributions over Ω). This is a natural assumption given that we only allow for consistent ambiguous mechanisms. Consistency means that the designer is indifferent between the different outcome functions of the ambiguous mechanism. Thus there is no reason for the agent to assign asymmetric weights to the different outcome functions.[32]

Returning to our example, consider the ambiguous mechanism $\Omega = \{(q^1, t^1), (q^2, t^2)\}$, described in the following table.

θ	1	2	4
(q^1, t^1)	$(0, 0)$	$(1, 1)$	$(1, 7/2)$
(q^2, t^2)	$(1, 1)$	$(0, 0)$	$(1, 7/2)$.

Under truth-telling the expected revenue of each of to the two outcome functions is 2. Therefore, under truth-telling this ambiguous mechanism yields exactly the same revenue as the optimal simple mechanism.

[32]While the assumption of a uniform μ over Ω is convenient in that it simplifies the presentation of our example, we should point out that it is not an assumption that is necessary for our argument.

The following table shows the expected utility that an agent of type θ obtains if he reports type $\hat{\theta}$; π here denotes the probability of the outcome function (q^1, t^1).

$\theta \setminus \hat{\theta}$	1	2	4
1	$\phi(0)$	$\phi(0)$	$\phi(-5/2)$
2	$\mathbb{E}_\mu[\phi(1-\pi)]$	$\mathbb{E}_\mu[\phi(\pi)]$	$\phi(-3/2)$
4	$\mathbb{E}_\mu[\phi(3(1-\pi))]$	$\mathbb{E}_\mu[\phi(3\pi)]$	$\phi(1/2)$

Observe that a truthful report guarantees each type a payoff that is no smaller than the value of the outside option, $\phi(0)$. Thus, the ambiguous mechanism Ω is individually rational. It is also easily seen that the lowest type's incentive compatibility constraint is satisfied ($\phi(0) > \phi(-5/2)$). For the middle type it is never optimal to report the highest type ($\phi(-3/2) < \phi(\pi)$). Moreover, since μ is uniform (and thus symmetric around 1/2) we have $\mathbb{E}_\mu[\phi(3\pi)] = \mathbb{E}_\mu[\phi(3(1-\pi))]$, meaning that type 2 has no incentive to report type 1 either.

Finally, we have to consider the highest type. Since

$$\phi(1/2) - \mathbb{E}_\mu[\phi(3\pi)] = [1 - \exp(-1)] - [1 - \exp(-6)]/2 \approx 0.13,$$

it follows that also this type's incentive compatibility constraints are satisfied. We conclude that Ω is both incentive compatible and individually rational. Notice that all incentive constraints regarding deviations from or deviations toward the highest type hold with strict inequality. Since the same is true also for the individual rationality constraint of the highest type we can increase $t^1(4)$ and $t^2(4)$ by some $\varepsilon > 0$ without violating any of these constraints. The resulting ambiguous mechanism then yields an expected revenue that strictly exceeds the revenue produced by the optimal simple mechanism. □

The principal's preferences: Throughout the paper we assumed that the principal is ambiguity-neutral. However, our results do not depend on this assumption. This is most obvious in the case of Proposition 1 and Lemmata 1–5 as those results do not refer to the designer's preferences. It is also easily seen that the ambiguous mechanism characterized in Proposition 2 remains the optimal ambiguous mechanism in the class of direct ambiguous mechanisms that are consistent. Thus, allowing for an ambiguity averse principal does not lead to different predictions once attention is restricted to consistent ambiguous mechanisms as we have defined them in Section 2. The question therefore is whether or not allowing for ambiguity aversion on the side of the principal affects the interpretation of the concept of an ambiguous mechanism and the appropriateness of the consistency condition that we impose.

Assuming an ambiguity averse principal introduces an aspect that is not present in the case of an ambiguity neutral designer: in such a framework it is conceivable that the principal has the ability to expose himself to the uncertainty to which he subjects the agent. That is, instead of committing ex ante to a specific element in the ambiguous mechanism that he announces to the agent, he could delegate the task of picking an element from the ambiguous mechanism to an uninterested third party or some mechanical selection device whose functioning neither he nor the agent understands. If for the principal it is not feasible to commit to such uncertain/ambiguous selection devices, then the assumption regarding the principal's attitude towards uncertainty does not affect the analysis in any way; all the findings that we have derived apply not only to the case of an expected utility maximizing principal but also to the case of an ambiguity averse designer.

What are the consequences of allowing for the possibility that the designer can delegate the choice of the outcome function in the above described sense? In situations where it is not the principal who picks the outcome function there is no reason to impose consistency. Remember that the motivation for requiring that condition was that we wanted to rule out the possibility that the agent believes in outcome functions that the principal would not want to choose if the agent does consider them as possible choices of the principal. This motivation disappears if the principal does not have to make the choice. In principle this means that one would have to consider a larger set of possible ambiguous mechanisms (also those that do not satisfy consistency) than the ones that we have considered in this paper. Fortunately, it is not difficult to see that non-consistent mechanisms can be disregarded without loss of generality also when the designer is ambiguity averse.

Remember that Proposition 1 and Lemmata 1–4 state that a given ambiguous mechanism can be replaced by another/simpler ambiguous mechanism without affecting (downward) incentive compatibility, individual rationality and consistency. All these results continue to hold on the larger set of ambiguous mechanisms if the requirement that consistency continues to hold is replaced with the requirement that the expected revenues of all outcome functions that survive the manipulations are not changed in the course of the manipulation. Moreover, everything that can be achieved by an arbitrary downward incentive compatible and individually rational ambiguous mechanisms can be done also with an downward incentive compatible and individually rational ambiguous mechanisms that is composed of at most $N-1$ elements and that satisfies the properties uniformity and monotonicity/maximal slackness. Any such ambiguous mechanism can be transformed into a consistent mechanism without changing the payoff of a principal with MMEU preferences and without affecting its downward incentive and individual rationality.

In order to see this, suppose Ω is an ambiguous mechanism that satisfies the above properties. Assume also that (q^m, t^m) minimizes the expected revenue among the outcome functions in Ω. If (q^n, t^n) is such that $\mathbb{E}_p[t^n] > \mathbb{E}_p[t^m]$, then we can simply replace it by the outcome function $(\tilde{q}^n, \tilde{t}^n)$, which coincides with (q^n, t^n) everywhere except for the payment of the highest type which is defined by

$$\tilde{t}^n_N = t^m_N - [p_m(1 - q^m_m)\theta_m - p_n(1 - q^n_n)\theta_n]/p_N.$$

Since uniformity and maximal slackness imply that

$$\mathbb{E}_p[t^m] - \mathbb{E}_p[t^n] = p_N[t^m_N - t^n_N] - p_m(1 - q^m_m)\theta_m + p_n(1 - q^n_n)\theta_n < 0,$$

it thus follows that

$$\mathbb{E}_p[t^m] - \mathbb{E}_p[\tilde{t}^n] = 0$$

and

$$t^n_N > \tilde{t}^n_N.$$

The preceding equation tells us, that applying this manipulation to all outcome functions we will end up with a consistent ambiguous mechanism. The last inequality implies that the downward incentive compatibility constraint of the highest type cannot be violated by the manipulation, which completes our argument.

4.4 Surplus extraction: ambiguity aversion vs. risk aversion

In Section 4.1 we have shown that with the use of ambiguous mechanisms the principal can extract the entire surplus from the agent provided that the agent has MMEU preferences and his type set is large enough. This result is related to the findings of Matthews (1983) and Maskin and Riley (1984) who have studied mechanism design problems with risk averse agents. Matthews (1983) shows that if the type set is a continuum and the agent has a valuation function that exhibits constant absolute risk aversion, then the share of the surplus that the principal can extract from the agent increases as the agent's coefficient of absolute risk aversion increases; in particular, when the agent becomes infinitely risk averse, the principal can appropriate the entire surplus.

Formally, the case of an ambiguity neutral and risk averse agent with a CARA utility function resembles the case of an agent who is risk neutral and smoothly ambiguity averse with a CARA transformation function ϕ.[33] Moreover, the MMEU preferences à la Gilboa and Schmeidler (1989) that we assume can be seen as the limit of ambiguity averse preferences that are CARA-smooth, when the CARA coefficient goes to ∞. In the light of these observa-

[33] By the term 'transformation function' we mean the function which is applied to transform the expected utility values. It is standard to denote this function by ϕ as we do in Example 4.

tions our full rent extraction result for the case where the type set is a continuum, may seem to be a rather immediate consequence of the results of Matthews (1983).

But the analogy between the analysis of the case of a risk averse agent in Matthews (1983) and our treatment of environments with a smoothly ambiguity averse agent is not quite as close as it appears. In the case of risk aversion the designer is allowed to choose random allocations. From a technical point of view this means that the designer controls the exact distribution over outcomes with respect to which the agent calculates his expected utility. The principal does not have an analogous instrument in the case of smooth ambiguity aversion. While the designer can exploit the agent's uncertainty aversion by offering an ambiguous mechanism, he will typically not be able to control in detail how the agent perceives the uncertainty embedded in the mechanism. In particular, the designer cannot choose the distribution which the agent uses in the aggregation of the expected utility values that he associates with the different beliefs that he holds (i.e. the distribution that in Example 4 has been denoted by μ.) Translated into the risk aversion model this would mean that the principal is not allowed to specify all aspects of the distribution of a random allocation but only its support, while the weighting of the different elements of this support is exogenously given to him as part of the agents preferences/beliefs.

4.5 Environments with multiple agents

In the previous sections we have restricted our attention to optimal mechanism design problems in single agent environments. In this section we show how the full surplus extraction result of Corollary 3 can be extended to a setting with multiple agents.[34] 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 easily be extend to the case with $I > 2$ agents.

Assume that the two agents have 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 (independently) drawn from the atomless distribution p with support $[0, 1]$. We do not need to assume that the designer knows the exact type distribution. For the following result we only need to impose that he knows the support of the distribution. Regarding the two agents' beliefs about each others type distribution we make no assumptions at all.

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

[34]The following full-rent-extraction result for multiple agents implies that ambiguous mechanisms outperform simple mechanisms also in situations with multiple agents. For a characterization of expected revenue maximizing simple mechanisms in general quasi-linear environments see Kos and Messner (2012).

$\theta \in \Theta\}$, be defined by

$$\begin{aligned} q^{\theta}(\hat{\theta}) &= \begin{cases} (0,0) & \textit{if } \theta = \hat{\theta} \\ (1,0) & \textit{if } \left[\hat{\theta}^1 \neq \theta^1 \textit{ and } \hat{\theta}^2 = \theta^2\right] \textit{ or } \left[\hat{\theta}^1 \neq \theta^1, \hat{\theta}^2 \neq \theta^2, \textit{ and } \hat{\theta}^1 \geq \hat{\theta}^2\right] \\ (0,1) & \textit{else} \end{cases} \\ t^{\theta}(\hat{\theta}) &= (q_1^{\theta}(\hat{\theta})\hat{\theta}^1, q_2^{\theta}(\hat{\theta})\hat{\theta}^2). \end{aligned}$$

Under Ω truth-telling is an optimal strategy for the two agents irrespective of their beliefs regarding the other agent's type or play. Moreover, Ω is individually rational and consistent (with respect to truth telling). The expected revenue generated 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 on 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 the two agents, irrespective of what they believe about the other agent's type or play. In order to see this, notice that for every profile of announced types, $\hat{\theta}$, agent i knows that there are simple mechanisms in Ω (all those indexed by a type profile, θ, such that $\hat{\theta}^i = \theta^i$) that specify that he will not receive the object and that he will not have to pay anything. This means that for every $\hat{\theta}$ his payoff is at most zero. On the other hand, by revealing his type truthfully, agent i can never get a strictly negative payoff since every outcome function specifies for every pair of reported types one of two possible outcomes for agent i: either he gets the object with probability one and pays the reported valuation or he does not get the object and pays zero; in either case the resulting payoff is zero. □

5 Conclusion

In this paper we have studied mechanism design problems where the agent is ambiguity averse in the sense of Gilboa and Schmeidler (1989). The central insight of our analysis is the observation that the principal can exploit the agent's ambiguity aversion by offering ambiguous mechanisms. In fact, we find that if the type set is 'large enough' the designer can extract the entire rent from the agent.

While most of our analysis concentrates on the case of a single agent environment, we show that when the type distribution is atomless our result readily generalize to settings with multiple agents. Finally, the core insight of our paper - in order to optimally exploit the uncertainty aversion of the agent the designer should offer ambiguous mechanisms - do not depend on the assumption of MMEU preferences à la Gilboa and Schmeidler (1989). In comparison to other models of ambiguity aversion, MMEU preferences provide important advantages in terms of tractability. In Example 4 we have seen that it is optimal for the principal to use ambiguous mechanisms also if we adopt the less extreme smooth model of ambiguity aversion.

References

Azrieli, Yaron and Roee Teper (2011), "Uncertainty aversion and equilibrium existence in games with incomplete." *Games and Economic Behavior*, forthcoming.

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

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

Bodoh-Creed, Aaron (2010), "Ambiguous beliefs and mechanism design." *Games and Economic Behavior*, forthcoming.

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." Unpublished manuscript.

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

Chung, K.S. and J. Ely (2007), "Foundations of dominant strategy mechanisms." *Review of Economic Studies*, 74, 447–476.

Ellsberg, Daniel (1961), "Risk, ambiguity, and the savage axioms." *Quarterly Journal of Economics*, 75 (4), 643–669.

Elyakime, B., J.J. Laffont, P. Loisel, and Q. Vuong (1994), "First price sealed-bid auctions with secret reservation prices." *Annales d'Économie et de Statistique*, 34, 115–141.

Epstein, Larry and Martin Schneider (2008), "Ambiguity, information quality and asset prices." *Journal of Finance*, 63(1), 197-228., 197–228.

Garrett, D.F. (2011), "Robustness of simple menus of contracts in cost-based procurement." Unpublished manuscript.

Gilboa, Itzhak (2009), *Theory of Decision Under Uncertainty*. Cambridge University Press.

Gilboa, Itzhak and Massimo Marinacci (2011), "Ambiguity and the bayesian paradigm." IGIER Working Paper Nr.

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

Kellner, M. (2011), "Tournaments as response to ambiguity aversion in incentive contracts." Unpublished manuscript.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005), "A smooth model of decision making under ambiguity." *Econometrica*, 73, 1849–1892.

Kos, Nenad and Matthias Messner (2012), "Extremal incentive compatible transfers." *Journal of Economic Theory*, forthcoming.

Lopomo, Giuseppe, Luca Rigotti, and Chris Shannon (2009), "Uncertainty in mechanism design." Working paper.

Maskin, Eric and John Riley (1984), "Optimal auctions with risk averse buyers." *Econometrica*, 52, 1473–1518.

Matthews, Steven (1983), "Selling to risk averse buyers with unobservable tastes." *Journal of Economic Theory*, 30, 370–400.

Raiffa, H. (1961), "Risk, ambiguity, and the savage axioms: Comment." *Quarterly Journal of Economics*, 75, 690–694.

Szydlowski, Martin (2012), "Ambiguity in dynamic contracts." SSRN Working Paper http://ssrn.com/abstract=1986186.

Turocy, Theodore L. (2008), "Auction choice for ambiguity-averse sellers facing strategic uncertainty." *Games and Economic Behavior*, 62, 155–179.