Working Papers

www.cesifo.org/wp

The Design of Ambiguous Mechanisms

Alfredo Di Tillio
Nenad Kos
Matthias Messner

CESifo Working Paper No. 4949
Category 13: Behavioural Economics
August 2014

CESifo *Working Paper No. 4949*

The Design of Ambiguous Mechanisms

Abstract

This paper explores the sale of an object to an ambiguity averse buyer. We show that the seller can increase his profit by using an ambiguous mechanism. That is, the seller can benefit from hiding certain features of the mechanism that he has committed to from the agent. We then characterize the profit maximizing mechanisms for the seller and characterize the conditions under which the seller can gain by employing an ambiguous mechanism.

JEL-Code: C720, D440, D820.

Keywords: optimal mechanism design, ambiguity aversion, incentive compatibility, individual rationality.

Alfredo Di Tillio
Department of Economics & IGIER
Bocconi University
Via Roentgen 1
Italy – 20136 Milan
alfredo.ditillio@unibocconi.it

Nenad Kos
Department of Economics & IGIER
Bocconi University
Via Roentgen 1
Italy – 20136 Milan
nenad.kos@unibocconi.it

Matthias Messner
Department of Economics & IGIER
Bocconi University
Via Roentgen 1
Italy – 20136 Milan
matthias.messner@unibocconi.it

July 26, 2014
We thank the editor and three anonymous referees for their insightful and constructive comments that have helped us to greatly improve the paper. For helpful discussions and comments we also 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.

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. For two recent surveys of the literature on ambiguity aversion and its axiomatic foundations, see Gilboa (2009) and Gilboa and Marinacci (2011). While ambiguity aversion models have been successfully applied in many areas of economics and finance,[2] they have received only limited attention in mechanism design (see the discussion of the literature below).

We consider a screening model in which a seller is selling an object to a single ambiguity averse buyer. For most of the paper we assume that the agent's preferences can be described by the maxmin expected utility model (MMEU) proposed by Gilboa and Schmeidler (1989). The agent privately observes his willingness to pay for the good, while the principal only knows the distribution from which it has been drawn. We introduce the concept of an *ambiguous mechanism*, i.e. a mechanism where the principal announces a set of possible standard mechanisms (henceforth simple mechanisms), and commits to one of them without revealing to the buyer which one he has chosen.

We then proceed to show that a seller who faces an ambiguity averse agent can strictly benefit from using such ambiguous mechanisms. This result has wide ranging consequences. It implies that in any mechanism design environment with ambiguity averse agents—be it auctions, bilateral trade, optimal taxation, unemployment insurance, or some other setting—the

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

analysis is not without loss of generality unless ambiguous mechanisms are considered.

Mechanisms where the buyer is not fully informed about all details of the rules, and which can therefore be considered as ambiguous, are commonly observed. Probably the most prominent examples are auctions with secret reserve prices.[3] Bergemann and Horner (2010) discuss the case of Google's sponsored search auctions where the algorithm to pick the winner is unknown.[4] Our result also provides a new way to rationalize laws that leave some discretionary decision power to the executive or judicial branch of government. It is reasonable to assume that for citizens it is much more difficult to predict the behavior of government agencies or judicial courts with proper decision power then it is to predict the behavior of agencies who just implement fully specified legal rules. 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 and/or the court system, the parliament exposes its citizens to ambiguity.

Through the use of an ambiguous mechanism, the principal exposes the agent to ambiguity regarding the consequences of his report. Since the agent has MMEU preferences, he associates with each possible report the worst possible outcome that he can obtain under all the simple mechanisms that compose the ambiguous mechanism. Different types evaluate outcomes differently, and hence they may associate different worst case scenarios with a given report. It is precisely this feature that makes the use of ambiguous mechanisms attractive for the principal: the principal can design the ambiguous mechanism in such a way that each outcome function that it contains deters the agent from a subset of his deviation possibilities. In this way each of the simple mechanisms that compose the ambiguous mechanism needs to be less distorted than the outcome function of a simple mechanism has to be, since the latter has to prevent the agent from *all* his possible deviations.

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, provided that the agent acts optimally with respect to such a belief. We therefore require that all elements of an ambiguous mechanism

[3] Ashenfelter (1989) documents the existence of secret reserve prices at the famous auction houses as Christie's and Sotheby's, Hendricks, Porter, and Spady (1989) in auctions for off-shore oil, Elyakime, Laffont, Loisel, and Vuong (1994) in timber auctions in France, and Bajari and Hortacsu (2003) on eBay, to name a few.

[4] Auctions with ex ante uncertain auction rules are also applied 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.

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*. A formal definition of the consistency condition is presented in Section 3. 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 condition in the interaction between the principal and the agent.

In the remainder of the paper we characterize the profit maximizing static mechanisms in the above described environment. First, 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. The highest type also obtains the good with probability one. However, since his transfers are used to guarantee consistency, he typically does not have to pay 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 with the risk aversion growing towards infinity the seller's rents grows towards full surplus extraction.[6] We then proceed to show 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.

[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).

We provide two robustness checks. First, we argue that the central insight of the paper—that the principal can exploit the ambiguity aversion of the agent by offering an ambiguous mechanism—does not depend on the specific model of ambiguity aversion we adopt in this paper (MMEU preferences) but remains valid under alternative models of uncertainty aversion.[7] More specifically, we provide an example 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 atomfree distribution defined on some interval.

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 (2014), Bergemann and Schlag (2011), Auster (2013) and Wolitzky (2013).[8] 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.[9] Instead, 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).[10]

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 (2014) 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

[7] The details of our characterization of an optimal ambiguous mechanism do depend on the assumption of MMEU preferences.

[8] 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. Moreover, ambiguity aversion has also been applied in environments with moral hazard; see Lang and Wambach (2013).

[9] Bose and Renou (2014) are an exception to this observation; their work is discussed in more detail in the following paragraph.

[10] 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.

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 (2014) 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, restricting attention to strategic form mechanisms makes the question of what is the most appropriate way to model updating by ambiguity averse individuals— an issue still controversially discussed in the literature—altogether irrelevant in our context. Finally, the ambiguous communication devices in Bose and Renou (2014) 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.

2 Motivating Example

A principal is selling an object to an ambiguity averse and risk neutral buyer whose preferences can be represented by maxmin expected utility and whose valuation for the object, θ, is 1 with probability $1/4$, 2 with probability $1/4$ and 4 with probability $1/2$. The seller's objective is to maximize expected revenue.

The optimal standard mechanism in this setting is a take-it-or-leave-it offer at the price of 4. The corresponding direct mechanism asks the agent to report his type and awards him the object at the price of 4 if the agent announces $\theta = 4$. For any other report the seller keeps the object and no transfers take place. The expected revenue generated by this mechanism is 2. It will prove convenient to represent the described direct mechanism in the following table form, where (q^*, t^*) denotes the outcome function (probability with which the good is transferred and transfer to be paid) and $\hat{\theta}$ denotes the reported type.

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

Table 1: The optimal non-ambiguous direct mechanism

Suppose that instead of offering the above standard mechanism (henceforth, we will refer to such mechanisms also as *simple* mechanisms) the seller proceeds as follows. Before he asks the buyer to communicate his valuation of the good he informs him that he has committed to a simple (direct) mechanism. But instead of letting the buyer know to which simple mechanism he has committed he only tells him that this simple mechanism is an element of some set of simple mechanisms that he reveals to the buyer. By not providing the buyer with any further information about the simple mechanism to which he has committed, the seller exposes the buyer to ambiguity about the consequences of his messages. We therefore refer to the set of simple mechanisms that is communicated to the buyer as an *ambiguous mechanism*.

For the sake of concreteness, suppose the seller offers an ambiguous mechanism that contains two (direct) simple mechanisms, denoted by (q^1, t^1) and (q^2, t^2), respectively.[11] Assume that the first outcome function, (q^1, t^1), specifies that upon a message $\hat{\theta} = 1$ the object remains with the seller and there are no transfers. If the agent reports $\hat{\theta} = 2$, he obtains the object at the price of 1. Finally, in case the agent's message is $\hat{\theta} = 4$, he obtains the object and pays price 4. The second outcome function, (q^2, t^2), awards the object at the price of 1 to the agent if he reports the type $\hat{\theta} = 1$, does not award the object to the agent, and no transfers take place, if the agent reports type $\hat{\theta} = 2$. If the agent sends the message $\hat{\theta} = 4$, he receives the object with certainty and pays 4. We denote the ambiguous mechanism that is composed of these two simple mechanisms by Ω. The details of Ω are summarized in the following table.

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

Table 2: The ambiguous (direct) mechanism Ω

We now turn to the question of how the buyer should behave when he is offered the mechanism Ω. We have assumed that the buyer's preferences are of the max-min expected utility type. This means that whenever he is faced with a decision problem under ambiguity he associates with each action that he may take the payoff that this action yields in the (action specific)

[11] Throughout the paper we slightly abuse terminology by identifying direct mechanisms with their outcome functions.

worst case scenario. In the context of our example the buyer is exposed to ambiguity through the mechanism Ω: the consequences of any message that he can send to the seller depend on the outcome function to which the seller has committed and he has no explicit or even implicit information regarding the seller's choice.

If the buyer does consider it possible that the principal might have committed to either of the two outcome functions in Ω, then the payoff that type θ associates with the message $\hat{\theta}$ is $\min\{q^1(\hat{\theta})\theta - t^1(\hat{\theta}), q^2(\hat{\theta})\theta - t^2(\hat{\theta})\}$.[12] We will now argue that truthful reporting is an optimal strategy for the buyer. To see this, suppose first that the buyer's type is $\theta = 4$. If he reports this truthfully he obtains the object with certainty and pays the price 4 (both simple mechanisms specify this outcome in case of the message $\hat{\theta} = 4$). His payoff in that case is therefore 0. If he reports type 2 he gets the object at the price of 1 if he is facing the simple mechanism (q^1, t^1); the corresponding payoff is $1 \times 4 - 1 = 3$. On the other hand, if he is faced with (q^2, t^2) he does not get the object, and does not have to pay anything. Consequently, his payoff in that case is 0. So the worst-case payoff both for truth-telling and for reporting type 2 is equal to 0, meaning that reporting type 2 when the true type is 4 does not represent a profitable deviation. The same reasoning can be applied to show that type 4 cannot do better by reporting 1 instead of 4. In fact, the 'symmetry' in the outcomes after reporting either type 1 or 2 implies that all three types are indifferent between these two reports. Thus, in order to complete our argument we just have to show that neither type 1 nor type 2 can do better by reporting 4 than by telling the truth. This follows since reporting 4 implies a payment of 4 which exceeds the valuations of both lower types.

Under truthful reporting both outcome functions in Ω generate an expected revenue of 9/4:

$$\begin{aligned} (q^1, t^1): &\qquad 0 \times \frac{1}{4} + 1 \times \frac{1}{4} + 4 \times \frac{1}{2} = \frac{9}{4} \\ (q^2, t^2): &\qquad 1 \times \frac{1}{4} + 0 \times \frac{1}{4} + 4 \times \frac{1}{2} = \frac{9}{4}. \end{aligned}$$

This expected revenue exceeds the expected revenue of 2 which is achieved by the optimal standard mechanism (which is a take it or leave it offer at the price of 4). Hence we have shown that in the context of the current example the seller can do strictly better by adopting an ambiguous mechanism rather than limiting himself to a standard mechanism.

The above analysis relies on two crucial assumptions on which we want to comment further. First, in order for the ambiguous mechanism Ω to generate a higher expected revenue than the best simple mechanism it is crucial that the seller has the possibility to commit to one of the two simple mechanisms in Ω *before* the agent makes his choice. It is straightforward to see that if the principal does not have the possibility to commit but must make his choice between

[12] Remember that we have also imposed risk neutrality.

(q^1, t^1) and (q^2, t^2) ex post, then truthtelling would no longer be optimal. Without commitment the principal would always choose the first outcome function after receiving the message 2 and the second one upon getting message 1. But if that is the case, then type 4 of the buyer would never want to report his type truthfully since by choosing either of the other two messages he could get the object with the same probability but at a lower price. Thus, in the absence of commitment Ω generates a lower expected revenue than the optimal simple mechanism. In fact, without commitment no ambiguous mechanism could do better than the best simple mechanism.

Given this observation, it is important to understand that the commitment ability that we need for our results is no stronger than the kind of commitment that is imposed in most of the mechanism design literature. The only difference is that here we assume that the principal can commit to something that the agent cannot observe before he makes his choice. But the ex ante observability of the mechanism for the agent should not be a crucial determinant for the principal's ability to commit. After all we can always assume that the simple mechanism to which the principal commits is described in some document that is stored in some place the access to which is jointly controlled by the principal and the agent (so that ex post they can verify together which allocation-transfer pair should be implemented). If anything, what should matter is the verifiability of the mechanism and of the messages vis-a-vis a third party who can guarantee the correct implementation of the mechanism.

The second important assumption on which our analysis of the above example builds is the assumption that the buyer does consider it possible that the seller might have committed to either of the two outcome functions in Ω. On the one hand this is certainly a reasonable assumption in light of the fact that the seller is not providing the buyer with any explicit information whatsoever as to which outcome function he might have chosen. The question remains whether the choice of Ω itself might be carrying relevant information. In particular, it would seem reasonable that a buyer asks himself what the expected revenue is that each element of the ambiguous mechanism generates, and that he would not believe in the possibility that the principal might have chosen a simple mechanism that generates a lower expected revenue than other simple mechanisms in Ω.

The possibility to indirectly infer the simple mechanism chosen by the principal is incompatible with the idea that the buyer perceives the ambiguous mechanism that he is presented with as ambiguous. We therefore impose the requirement that the principal can offer only ambiguous mechanisms with the property that all their elements generate the same expected revenue provided the agent acts optimally based on the belief that the principal might have committed to any of the simple mechanisms in the ambiguous mechanism. We refer to this condition as *consistency* condition. Under consistency the principal is indifferent between all the simple mechanisms in the ambiguous mechanism if he assumes that the agent considers all the elements of the ambiguous mechanism as 'real' possibilities. Conversely, for the agent there exists no reason to deem one of the elements of the ambiguous mechanism as impossible if he is aware of the principal's indifference. As an ambiguity averter he should therefore not exclude any of them from his considerations.[13] It is due to these considerations that we have chosen Ω in such a way that both its simple mechanisms generate the same expected revenue.

3 Framework

Throughout the first part of the paper we consider the mechanism design problem of a principal selling a single unit of a good to a single agent. The notation and terminology that we introduce below generalize in the obvious way to the case of multiple agents, which we consider later in the paper.

Allocations and preferences. An allocation is a pair $(x,\tau) \in X \times \mathbb{R}$, where $x \in X = [0,1]$ denotes the probability with which the good is transferred to the agent and τ the monetary transfer he has to pay to the principal.[14] With a slight abuse of terminology we will typically 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 type $\theta \in \Theta \subset \mathbb{R}$. More specifically, we assume that they are represented by the linear utility function

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

The agent's valuation of the good, that is, his type θ, is his private information. Throughout the first part of the paper we assume that Θ is a finite set with N elements and we index types so that θ_n is increasing in n. The principal's beliefs regarding the agent's type are described the probability distribution $p = (p_1, \dots, p_N)$.

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, Λ, his utility is given by

$$\inf_{\lambda \in \Lambda} \mathbb{E}_\lambda[x\theta - \tau].$$

The principal is risk and ambiguity neutral. His objective is to maximize expected revenue, that is, the expected transfer payments paid by the agent. We show in Section 5.3 that our main results go through also under the assumption of an ambiguity averse principal. Allowing for

[13] The possibility that the buyer would choose to arbitrarily disregard some of the elements of the ambiguous mechanism would be difficult to reconcile with the spirit of models of ambiguity aversion.

[14] 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 given to the agent.

this possibility, however, does not add anything of interest to our analysis as the central insights that we obtain are driven by the agent's ambiguity aversion.

Simple vs. ambiguous mechanisms. A *simple mechanism* is a triple (S, q, t). S is a set of messages that the agent may send to the principal. The functions q and t map S into X and $\mathbb{R}$, respectively. $q(s)$ is the probability with which the good is transferred to the agent if he sends message s, while $t(s)$ is the corresponding transfer that the agent has to pay to the principal.[15] We refer to q and t, respectively, as allocation and transfer rules on S; the pair (q, t) is the outcome function of the mechanism. A *direct simple mechanism* is a simple mechanism such that $S = \Theta$. Since all direct mechanisms share the same message space, we drop the latter from the notation and identify the direct mechanism (Θ, q, t) with its outcome function (q, t).

Mechanism design models typically assume that besides committing to a particular outcome function the principal also fully and credibly reveals it to the agent. In effect, if the agent has standard expected utility preferences then the latter part of this assumption is innocuous, as the principal cannot gain anything from concealing this information.[16] The central insight of this paper is that this is no longer true with an ambiguity averse agent. Indeed, we show that it is typically in the principal's best interest not to inform the agent about the exact outcome function he commits to. Instead, he can benefit from communicating the rules of the mechanisms in an *ambiguous* way, by only announcing that it belongs to a certain set. The notion of an *ambiguous mechanism* captures the idea of such ambiguous rules.

Definition 1 (Ambiguous mechanism)**.** An **ambiguous mechanism** is a pair (S, Ω), where S is a set of messages, and Ω is a set of outcome functions defined on S, i.e. $\Omega \subset X^S \times \mathbb{R}^S$.[17] A generic element of Ω is denoted by (q, t), where $q \in X^S$ and $t \in \mathbb{R}^S$.

Before we go on, 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 $(\hat{q}, \hat{t})$. This commitment may be achieved, say, by depositing $(\hat{q}, \hat{t})$ with an uninterested third party. The agent is not fully informed about the chosen outcome function. Instead the principal limits himself to telling the agent that it belongs to a set Ω. Of

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

[16]Under any standard equilibrium concept the agent would know in equilibrium which function has been chosen by the principal.

[17]As already argued earlier, the restriction to ambiguous mechanisms with (sets of) deterministic outcome functions is without loss of generality in an environment with risk neutral players.

course, by announcing such ambiguous rules he exposes the agent to uncertainty about the consequences of his messages and 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 $(\hat{q},\hat{t})$. 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 message.

Agent's strategies and beliefs. 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. $\sigma \in S^{\Theta}$.

We assume that the agent cannot use mixed strategies. This assumption, which is commonly adopted in the ambiguity literature, has some bite, as an ambiguity averse individual facing two alternatives with uncertain consequences may strictly prefer mixing over the alternatives to *each* of the two.[18] However, we maintain that besides being pervasive in the literature, the assumption is especially weak in our context. Unlike an expected utility maximizer, an ambiguity averter may ex ante wish to randomize even over alternatives that he is not indifferent over. But if the individual has strict preferences over the alternatives he is randomizing over, then the strong ex ante incentives to mix conflicts with the individuals ex post incentives to implement the outcome of the randomization. In this case, allowing for mixed strategies may therefore matter only if the agent can commit to obeying the recommendation of some randomizing device. Making such a commitment in a mechanism design context is difficult, because the designer can do better by declining reports generated by such devices, that is, by requiring reports to be made directly by the agent.[19]

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. The agent's only piece of hard information in this respect is that the function belongs to Ω. On the other hand, the agent knows that the principal seeks to maximize his revenue. Given the agent's ambiguity aversion, it thus seems appropriate to assume that his belief set contains the entire family $\Delta(\Omega)$ of probability measures on Ω, provided that such a belief set is not incompatible with the principal's optimizing behavior in a sense that we formalize next.

[18] Ever since Raiffa (1961), it is well known that randomization may help the agent to hedge against the uncertainty involved in the two alternatives. Recently Saito (2013) provided an axiomatization of ambiguity aversion which does not give rise to a hedging motive.

[19] While this argument only refers to randomizations over alternatives that the individual does not consider as equivalent, it is sufficient for our purpose. In our context, the agent will have to decide which type to report to the principal. The optimal direct mechanism that we will derive can be arbitrarily closely approximated by a mechanism with the property that no type of the agent is indifferent between any two messages.

For any ambiguous mechanism (S,Ω), let $\Sigma^*(S,\Omega)$ designate the corresponding set of optimal strategies for the agent, when his beliefs are given by $\Delta(\Omega)$. Since the agent is risk neutral, calculating the infimum of his expected payoffs with respect to $\Delta(\Omega)$ delivers the same value as the one obtained when attention is restricted to Ω. Thus, the set $\Sigma^*(S,\Omega)$ is the set of all $\sigma \in S^\Theta$ such that for each $\theta \in \Theta$,

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

Definition 2 (Consistency)**.** An ambiguous mechanism (S,Ω) is *consistent with respect to* $\sigma \in \Sigma^*(S,\Omega)$ if under σ all outcome functions in Ω yield the same expected revenue to the principal, i.e. if for all $(q,t),(q',t') \in \Omega$

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

The ambiguous mechanism (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 $\Delta(\Omega)$. To shed further light on this condition, consider a situation were it is not satisfied. Thus, suppose that the principal proposes an ambiguous mechanism (S,Ω) such that for every $\sigma \in \Sigma^*(S,\Omega)$ there exist $(q,t),(q',t') \in \Omega$ with $\mathbb{E}_p[t(\sigma(\theta))] < \mathbb{E}_p[t'(\sigma(\theta))]$. 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 rules out such contradictory beliefs. It is essentially an equilibrium condition for the two stage game played by the two parties. In equilibrium the agent should not entertain the possibility that a certain outcome function is chosen, if the strategy that he intends to implement in response to this belief implies that the outcome function does not maximize the designer's payoff.[20]

Finally, we remark that requiring the seller to be indifferent before the buyer reports his type rather than after, is to the seller's benefit. If the seller could choose his preferred simple mechanism in the ambiguous mechanism after the report, he would choose a mechanism with the highest transfer given the report. The buyer would foresee this and calculate his payoffs accordingly. In particular, every type of the agent would associate with any given report the same worst case scenario, namely the outcome function in Ω with the smallest probability of trade (for that report) among the simple mechanisms that specify the highest transfer (for that report). But then the agent's behavior vis-a-vis the ambiguous mechanism Ω would be exactly

[20] In Section 5.3 we show that the assumption of consistency is without loss of generality when also the seller has maxmin expected utility preferences.

the same as the agent's behavior vis-a-vis a simple mechanism that specifies only these worst case allocations and transfers.[21] Consequently, the expected revenue of the designer would also be exactly the same. Thus, if the designer does not commit ex ante to an element in Ω he cannot do any better than by using only simple mechanisms.

4 Optimal ambiguous mechanisms

In designing the optimal ambiguous mechanism the principal has to take into account two types of constraints. First, he must respect the consistency condition that we have discussed in the preceding section. Second, since we assume that the buyer's participation in the mechanism is voluntary, the principal must make sure that the mechanism allows each type of the agent to earn at least his outside option. We assume that the latter is equal to zero for every type. Thus, the principal's, problem is to choose among all ambiguous mechanisms (S,Ω) for which there exists some $\sigma \in \Sigma^*(S,\Omega)$ satisfying the conditions

$$\mathbb{E}_p[t(\sigma(\theta))] = \mathbb{E}_p[t'(\sigma(\theta))] \qquad \text{for all } (q,t),(q',t') \in \Omega, \tag{1}$$

$$\inf_{(q,t)\in\Omega}\{q(\sigma(\theta))\theta - t(\sigma(\theta))\} \geq 0 \qquad \text{for all } \theta \in \Theta, \tag{2}$$

the one that delivers the highest expected revenue.

In what follows we show that the principal's problem can be substantially simplified.

4.1 The Revelation Principle

First we prove a version of the *Revelation Principle* that applies to our environment, by showing that the principal can without loss of generality offer the agent an ambiguous mechanism that (i) asks the agent to report his type, and (ii) is constructed such that the agent is willing to do so in a truthful manner.

Definition 3 (Incentive compatibility)**.** An ambiguous mechanism (S,Ω) is direct if $S = \Theta$, in which case we identify the mechanism with its set of outcome functions Ω, and for all $(q,t) \in \Omega$ and $1 \leq n \leq N$ we write q_n and t_n for $q(\theta_n)$ and $t(\theta_n)$, respectively. A direct ambiguous mechanism Ω is *downward incentive compatible* if

$$\inf_{(q,t)\in\Omega}\{q_n\theta_n - t_n\} \geq \inf_{(q,t)\in\Omega}\{q_m\theta_n - t_m\} \qquad \text{for all } 1 \leq m < n \leq N, \tag{DIC}$$

[21] Formally, this simple mechanism specifies the same messages, S, and its outcome function (q,t) is such that for all s, $t(s) = \sup_{(q',t')\in\Omega}\{\tau \in \mathbb{R} : t'(s) = \tau\}$ and $q(s) = \inf_{(q',t')\in\Omega}\{x \in X : q'(s) = x \text{ and } t'(s) = t(s)\}$.

upward incentive compatible if

$$\inf_{(q,t)\in\Omega}\{q_n\theta_n - t_n\} \geq \inf_{(q,t)\in\Omega}\{q_m\theta_n - t_m\} \qquad \text{for all } 1 \leq n < m \leq N \tag{UIC}$$

and *incentive compatible* if it is both downward and upward incentive compatible.

Proposition 1 (Revelation Principle)**.** *Let (S,Ω) be an ambiguous mechanism that is consistent with respect to $\sigma \in \Sigma^*(S,\Omega)$. The direct ambiguous mechanism*

$$\Omega' = \left\{(q',t') \in X^\Theta \times \mathbb{R}^\Theta : q' = q\circ\sigma, t' = t\circ\sigma \textit{ for some } (q,t) \in \Omega\right\}$$

is incentive compatible and consistent with respect to truthtelling.

Proof. See the Appendix. ∎

The Revelation Principle guarantees that given any consistent ambiguous mechanism, (S,Ω), we can find a direct ambiguous mechanism, Ω', satisfying incentive compatibility and such that, element by element, (S,Ω) and Ω' generate the same allocations and transfers, and hence give both the principal and the agent the same payoff. As a consequence the principal can restrict himself to direct ambiguous mechanisms that satisfy incentive compatibility, consistency with respect to truth-telling and individual rationality (condition (2)). In the case of direct ambiguous mechanisms Ω, the latter may be rewritten as

$$\inf_{(q,t)\in\Omega} q_n\theta_n - t_n \geq 0 \qquad \text{for all } 1 \leq n \leq N. \tag{IR}$$

Thus, the problem of the principal can be written as follows:

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

$$\text{s.t.} \qquad R = \sum_{n=1}^{N} p_n t_n \quad \text{for all } (q,t) \in \Omega, \tag{C}$$

(DIC), (UIC) and (IR).

We now consider the relaxed version of this problem where constraint (UIC) is removed, and show through a sequence of lemmata that the set of feasible mechanisms for the relaxed problem can be restricted while leaving the problem's value unchanged. Finally, we prove that all optimal mechanisms for the relaxed problem in fact satisfy (UIC), and are therefore also optimal for the original problem (P). Given a direct ambiguous mechanism Ω that satisfies (C), in what follows we write $R(\Omega)$ for the expected revenue associated to every simple mechanism in the ambiguous mechanism Ω, so that

$$R(\Omega) = \sum_{n=1}^{N} p_n t_n.$$

4.2 Uniform, minimal and monotonic ambiguous mechanisms

We first show that the relaxed version of Problem (P) always admits solutions that do not expose truthfully reporting types to ambiguity, except possibly the highest type. That is, the truthtelling payoff of every type θ_n, $n < N$, is constant across the outcome functions of the optimal ambiguous mechanism. Moreover, at the optimum for each type one of the downward deviation constraints must be binding. Thus, the truthtelling payoffs coincide with the payoff that types can obtain from the most attractive downward deviation. In what follows we refer to these properties as *uniformity*.

Definition 4 (Uniformity)**.** A direct ambiguous mechanism Ω is *uniform* if

$$\begin{aligned} & q_1\theta_1 - t_1 = 0, && \text{for all } (q,t) \in \Omega, \\ & q_n\theta_n - t_n = \max_{1<m<n} \inf_{(q',t')\in\Omega} \{q'_m\theta_n - t'_m\}, && \text{for all } (q,t) \in \Omega, \\ & q_N\theta_N - t_N = \max_{1<m<N} \inf_{(q',t')\in\Omega} \{q'_m\theta_N - t'_m\}, && \text{for some } (q,t) \in \Omega \end{aligned} \tag{Uni}$$

Note that uniformity implies both downward incentive compatibility ((DIC)) and individual rationality ((IR)). This is immediate to see in the case of (DIC). As for (IR) observe that for all $(q,t) \in \Omega$ and all $1 < n \leq N$ we have

$$q_1\theta_n - t_1 \geq q_1\theta_1 - t_1.$$

That is, type θ_n, $n > 1$, cannot obtain a lower payoff from reporting θ_1 than type θ_1 himself. By the first condition in the definition of uniformity the lowest type's payoff from truth-telling is zero. But then, the second and third condition of uniformity can be satisfied only if the truth-telling payoffs of all types θ_n, $n > 1$, are (weakly) larger than zero too.

Lemma 1. *For every direct ambiguous mechanism Ω satisfying (C), (DIC) and (IR) there is a direct ambiguous mechanism Ω' satisfying satisfying (C), (Uni) and $R(\Omega') \geq R(\Omega)$.*

Proof. See the Appendix. ∎

The fact that imposing the uniformity condition (Uni) is without loss of generality – in the relaxed problem where (UIC) is removed but, as we argue later, in problem (P) as well – resembles the standard result from mechanism design, stating that at the optimum the downward incentive compatibility constraints and the individual rationality constraint of the lowest type are binding. Ambiguity aversion and the consistency requirement, however, demand special attention in establishing this fact. In the proof of the lemma we show that if any simple mechanism, (q,t), in the ambiguous mechanism Ω gives the lowest type a strictly positive payoff, then it can be changed by increasing t_1 to $q_1\theta_1$ and decreasing the transfer t_N in a way that leaves the simple mechanism's expected revenue unaltered. Given that the expected revenue remains constant, this modification is neutral with respect to the consistency condition (C). Moreover, increasing the lowest type's transfer and decreasing the one of the highest type cannot possibly lead to a violation of any downward incentive compatibility condition.

Using similar arguments, we show that the value of the designer's problem is not affected if he only considers ambiguous mechanisms Ω such that each $(q,t)\in\Omega$ satisfies the property that truthful reporting of the type θ_n, $n\le N$, yields the same payoff as the most attractive misreport of a lower type (not necessarily the downward adjacent one).

In the statement of Problem P, ambiguous mechanisms are allowed to be of any size. The next result shows that the problem can be substantially simplified since attention can be restricted to ambiguous mechanisms that are both 'small' and have a simple structure. In particular, Lemma 2 shows that there is always a solution of the relaxed version of Problem P that contains $N-1$ (not necessarily distinct) simple mechanisms.[22] Each one of these outcome functions serves the purpose to deter downward deviations towards one particular report. We will henceforth refer to mechanisms with this property as *minimal* mechanisms.

Definition 5 (Minimality)**.** An ambiguous mechanism Ω is *minimal* if $\Omega=\{(q^1,t^1),\dots,(q^{N-1},t^{N-1})\}$, with

$$q^m_m\theta_n - t^m_m \le q^\ell_m\theta_n - t^\ell_m \qquad \text{for all } 1\le \ell, m < n \le N. \tag{Min}$$

Lemma 2. *For every direct ambiguous mechanism Ω, satisfying (C), (Uni) there is a direct ambiguous mechanism Ω' satisfying (C), (Uni), (Min), and $R(\Omega)=R(\Omega')$.*

Proof. See the Appendix. ∎

The fact that each outcome function, (q^m,t^m), of a minimal mechanism has to dissuade the agent only from reporting θ_m when his true type is higher, provides the central intuition for why the seller can do better with an ambiguous mechanism than with a simple mechanism. With multiple outcome functions the designer has more instruments to take care of the incentive constraints. Each single simple mechanism, in the ambiguous mechanism, takes care of only a

[22] By allowing for the possibility that minimal ambiguous mechanisms contain multiple copies of one and the same outcome function we slightly abuse the meaning of the term 'set' that we are using when referring to ambiguous mechanisms. The reasons for adopting this convention are purely notational.

subset of all incentive compatibility constraints. While the outcome function (q^m, t^m) guarantees that no type $\theta_n > \theta_m$ wishes to report θ_m, another outcome function, $(q^{m'}, t^{m'})$, performs the same task with respect to report $\theta_{m'}$. Each simple mechanism in the ambiguous mechanism is therefore less distorted than the optimal non-ambiguous mechanism which has to take care of all the incentive compatibility conditions.

The principal's ability to limit himself to minimal mechanisms has immediate consequences for the case of a binary type set. In this case, having multiple simple mechanisms in the ambiguous mechanism does not provide any advantage in handling the incentive constraints. Indeed, when there are are only two types Lemma (5) readily implies that the seller cannot do better with an ambiguous mechanism containing multiple simple mechanism than with a standard non-ambiguous mechanism.

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 non-ambiguous mechanism.*

Finally, we show that within the set of ambiguous mechanisms that are minimal and uniform, we only have to consider ambiguous mechanisms that exhibit allocation rules that have a particularly simple structure. More specifically, attention can be limited to mechanisms with allocation rules that are are equal to 1 for all except possibly one report. Moreover, the coordinates of the outcome functions that are allowed to differ from the value 1 can be assumed to satisfy a monotonicity condition defined across outcome functions.

Definition 6 (Monotonicity)**.** A minimal direct ambiguous mechanism $\Omega = \{(q^1, t^1), \ldots, (q^{N-1}_{N-1}, t^{N-1}_{N-1})\}$ is *monotonic* if

$$\begin{aligned} q^m_n &= 1 \qquad \text{for all } 1 \le m < N, 1 \le n \le N, n \ne m \quad \text{and} \\ q^m_m &\le q^n_n \qquad \text{for all } 1 \le m < N, m \le n \le N-1. \end{aligned} \tag{Mon}$$

Lemma 3. *For every direct ambiguous mechanism Ω, satisfying (C), (Uni) and (Min), there is a direct ambiguous mechanism Ω' satisfying (C), (Uni), (Min), (Mon) and $R(\Omega') \ge R(\Omega)$.*

Proof. See the Appendix. ∎

The above lemma shows that one can restrict attention to minimal ambiguous mechanisms in which every simple mechanism awards the object with probability one to the agent after all but (possibly) one report. The intuition for the fact that in the simple mechanism (q^m, t^m) only the allocation q^m_m needs to be left unrestricted is rather straightforward. We have observed earlier, that the purpose of (q^m, t^m) is to prevent the agent from reporting θ_m when he is of a higher type. Since by uniformity the truthtelling payoff of type θ_m has to be constant across outcome functions, it follows that the payoff of higher types who value the good more, must be minimized by the outcome function that awards the object with the lowest probability. Thus, $q_m^m \leq q_m^{m'}$ for all $m' \neq m$. Finally, if (q^m, t^m) takes care of the downward deviation constraints toward θ_m, then an increase of $q_{m'}^m$ could neither affect the (downward) incentive compatibility of the mechanism nor its individual rationality. Consequently, $q_{m'}^m$ can be set equal to one.

The second part of (Mon), $q_m^m \geq q_{m-1}^{m-1}$ for all m strictly between 1 and N, parallels the monotonicity result in a standard mechanism design problem with ambiguity neutral agents, where an allocation rule is implementable if and only if it is monotonic. This property translates in a natural way into our setting with an ambiguity averse agent.

In the proof of the lemma we show that this is the case because under a uniform ambiguous mechanism the most attractive (downward) deviation option for type θ_n is the report that guarantees the largest worst case allocation, i.e., the report that guarantees $\max_{m<n} \min_{1\leq \ell < m} q_m^l$. Thus, if $1 \leq m' < m < N$, and $q_{m'}^{m'} > q_m^m$, then there is no type θ_n, $n > m$, for whom the incentive constraint with respect to θ_m is binding. Consequently, by an increase of q_m^m up to $q_{m'}^{m'}$ that is accompanied by a corresponding increase of t_m^m (so that the truth-telling payoff of type θ_m remains unchanged) no downward incentive constraints of any type $\theta_n > \theta_m$ is violated. Since such an increase of q_m^m (and the associated increase of t_m^m) does not affect the downward incentive constraints of types θ_n, $n \leq m$ it follows that the assumed non-monotonicity can be eliminated without affecting downward incentive compatibility. Through appropriate adjustments of the transfers of the highest type (C) can be reestablished and (DIC) can be strengthened into (Uni). None of these modifications affects (Min).

The three results above show that in solving the relaxed version of Problem (P), where (UIC) is removed, one can restrict attention to mechanisms that satisfy condition (C), uniformity, minimality and monotonicity. We now show that properties (C), (Uni), (Min) and (Mon) are actually sufficient for feasibility in the *original* Problem (P) where (UIC) is present.

Lemma 4. *If a direct ambiguous mechanism Ω satisfies (Uni), (Min) and (Mon), then it also satisfies (UIC).*

Proof. See the Appendix. ∎

The main result of this section now follows.

Proposition 2. *Every solution to problem*

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

$$\textit{subject to} \quad (C), (Uni), (Min) \textit{ and } (Mon),$$

is also a solution to Problem (P).

Proof. See the Appendix. ■

4.3 The optimal ambiguous mechanism

We are now ready to describe explicitly an optimal ambiguous mechanism. First, we provide a useful characterization of the constraint set of Problem (P').

Lemma 5. *A minimal and monotonic direct ambiguous mechanism* $\Omega = \{(q^1, t^1), \ldots, (q^{N-1}, t^{N-1})\}$ *satisfies (Uni) if and only if the following hold:*

$$t_n^m = q_n^m \theta_n - \sum_{k=1}^{n-1} q_k^k (\theta_{k+1} - \theta_k) \qquad \textit{for all } 1 \le m, n \le N-1, \tag{3}$$

$$\max_{1 \le m < N} t_N^m = \theta_N - \sum_{k=1}^{N-1} q_k^k (\theta_{k+1} - \theta_k) \tag{4}$$

Proof. See the Appendix. ■

Lemma 5 shows that the transfers of minimal, and monotonic mechanisms that are also uniform can be expressed in terms of the allocation vector $(q_1^1, \ldots, q_{N-1}^{N-1})$ only. The only exception to this rule are the transfers of the highest type. Those are bounded above by an expression that only depends on $(q_1^1, \ldots, q_{N-1}^{N-1})$ (condition (4)) . Conversely, any minimal and monotonic mechanism whose transfers satisfy conditions (3) and (4) is also uniform. Thus, solving Problem P' amounts to optimally choosing allocations $q_1^1 \le \ldots \le q_{N-1}^{N-1}$ and transfers $t_N^1, \ldots, t_N^{N-1}$. All other allocations are equal to one and all other transfers are determined via (3) through the choice of $(q_1^1, \ldots, q_{N-1}^{N-1})$. The two constraints to be respected are condition (4) and consistency; i.e. $t_N^1, \ldots, t_N^{N-1}$ together with the transfers that are determined through (3) must be such that $\sum_{n=1}^N p_n t_n^m$ is constant in m. Given these observations, in what follows we will say that the vector of allocations $(\bar{q}_1^1, \ldots, \bar{q}_{N-1}^{N-1})$ generates or induces the mechanism $\Omega = \{(q^1, t^1), \ldots, (q_{N-1}^{N-1}, t_{N-1}^{N-1})\}$, if Ω satisfies all constraints of Problem P' and $q_m^m = \bar{q}_m^m$ for all $1 \le m \le N-1$.

In what follows we outline how to compute the expected revenue of the mechanism generated by $(q_1^1, \ldots, q_{N-1}^{N-1})$. The problem is that one does not know for which m the maximum in (4) is attained. However, the right-hand sides of (3) and (4) can be used to compute an upper boundary on the expected transfer of each outcome function of the ambiguous mechanism that is generated by $(q_1^1, \ldots, q_{N-1}^{N-1})$. Since the generated ambiguous mechanism needs to be such that all the simple mechanisms yield the same expected transfer the relevant upper boundary is the lowest one. More precisely, let $\bar{R}^m(q_1^1, \ldots, q_{N-1}^{N-1})$ be the expected value of the sum of the terms in the right hand sides of (3) and (4), that is,

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

where $P_n = \sum_{k=1}^{n} p_k$. If the designer chooses the ambiguous mechanism that is generated by the vector of allocations $(q_1^1,\ldots,q_{N-1}^{N-1})$, his expected revenue under the outcome function (q^m, t^m) cannot exceed $\bar{R}^m(q_1^1,\ldots,q_{N-1}^{N-1})$. In fact, since we require that $\sum_{n=1}^{N} p_n t_n^m$ is constant in m it must be the case that transfers in Ω are such that for each $1 \le m \le N-1$ we have

$$R^m(\Omega) = \min_{1\le l\le N-1} \bar{R}^l(q_1^1,\ldots,q_{N-1}^{N-1}). \tag{5}$$

That is, the lowest upper boundary on the expected revenue is binding and thus yields the expected revenue of the ambiguous mechanism generated by $(q_1^1,\ldots,q_{N-1}^{N-1})$.

Since the seller is maximizing his expected revenue, an optimal choice of $(q_1^1,\ldots,q_{N-1}^{N-1})$ must solve the problem

$$\max_{(q_1^1,\ldots,q_{N-1}^{N-1})\in Q} \min_{1\le m\le N-1} \bar{R}^m(q_1^1,\ldots,q_{N-1}^{N-1}), \tag{P”}$$

where Q is the set of all vectors $Q = (q_1^1,\ldots,q_{N-1}^{N-1}) \in [0,1]^{N-1}$ whose components are weakly increasing. The corresponding optimal transfers for the highest type, $(t_N^1,\ldots,t_N^{N-1})$, are then determined by condition (5), i.e. they are chosen so that condition (C) holds.

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.[23] In particular, in Proposition 3 below we show that

[23] 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.

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

Proof. See the Appendix. ∎

We are now ready to state the main result of this section in which we characterize a solution of Problem P".

Proposition 3.

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

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

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

Proof. See the Appendix. ∎

Proposition 3 yields a solution to 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 $\hat{t}_n^m$ can be obtained from (3). The highest type's transfers are then chosen so that the expected revenue of each of the $N-1$ simple mechanisms 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. Moreover, write $(\hat{q}^m, \hat{t}^m)$, $m = 1, \ldots, N-1$ for the m-th element of the (optimal) ambiguous mechanism generated by $\hat{Q}$. Then, $\hat{t}^m$ is given by*

$$\hat{t}_n^m = \begin{cases} \hat{q}_n^m \theta_n - \sum_{k=1}^{n-1} \hat{q}_k^k (\theta_{k+1} - \theta_k) & \text{if } 1 \leq n < N \\ (\hat{R} - \sum_{n=1}^{N-1} p_n \hat{t}_n^m)/p_N & \text{if } n = N, \end{cases}$$

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

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

In our environment the buyer values the good more than the seller. Allocative efficiency would therefore require that the good always be allocated to the buyer. According to Proposition 3 this is typically not the case in the revenue maximizing ambiguous mechanism. The seller might distort the allocative efficiency to increase the revenue, much like it is done in revenue maximizing simple mechanisms. In Section 5.1 we will see that unlike in the case of optimal simple mechanisms these distortions tend to vanish in environments with large type sets. An additional source of inefficiency is introduced through the ambiguity the agent faces when presented with an ambiguous mechanism. This inefficiency though regards only the highest type, for only his truth-telling payoffs vary across outcome functions. For all other types, the uncertainty embedded in the optimal ambiguous mechanism regards only the payoffs from deviations, which are never realized.

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

Certainly, this cannot be the case whenever $Q = (1, \ldots, 1)$ solves Problem P”. The ambiguous mechanism that is generated by the allocation vector $Q = (1 \ldots, 1)$, yields an expected revenue of $\hat{R} = \theta_1$, which is the same as the revenue obtained from a a simple mechanism that prescribes that the object is transferred with probability one at the price of θ_1, irrespective of which message the buyer sends (take-it-or-leave-it offer at the price θ_1).

Therefore a necessary condition for the optimal ambiguous mechanisms to yield a higher expected revenue than the best simple mechanism is that $Q = (1, \ldots, 1)$ is not a solution to Problem P”. We will argue now that this condition is also sufficient. Towards that, assume that $Q = (1, \ldots, 1)$ is not a solution of P”. There are two possibilities regarding the optimal simple mechanism $(\tilde{q}, \tilde{t})$: either $\tilde{q} = (1, \ldots, 1)$ or $\tilde{q} \neq (1, \ldots, 1)$. In the first case, we have that $\tilde{R} = \theta_1$[24] and thus $\hat{R} > \tilde{R}$. The latter inequality is implied by the following two facts: a) the revenue of the ambiguous mechanism generated by $Q = (1, \ldots, 1)$ is equal to θ_1 and b) $Q = (1, \ldots, 1)$ is not optimal. This leaves us with the case $\tilde{q} \neq (1, \ldots, 1)$.

Let $(\tilde{q}, \tilde{t})$ be some optimal simple mechanism, and assume that no $(\tilde{q}', \tilde{t}')$ with $\tilde{q}' = (1, \ldots, 1)$ is optimal among simple mechanism. Define the ambiguous mechanism $\Omega = \{(q^1, t^1), \ldots, (q^{N-1}, t^{N-1})\}$

[24] In what follows all variables with a tilde refer to the optimal simple mechanism $(\tilde{q}, \tilde{t})$. For instance $\tilde{R}$ is the expected revenue generated by $(\tilde{q}, \tilde{t})$.

as follows

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

where m_M is defined as before, i.e. $m_M \in \arg\max_{n<N} p_n\theta_n$. In words, we take the optimal simple mechanism $(\tilde{q},\tilde{t})$ and construct an ambiguous mechanism the following way. The ambiguous mechanism consists of $N-1$ simple mechanisms. In the simple mechanism (q^m,t^m), $m = 1,\ldots,N-1$, q_n^m is set equal to 1 whenever $n \neq m$; q_m^m instead is set equal to $\tilde{q}_m$. The transfers are such that $t_m^m = \tilde{t}_m$, and $t_n^m = \tilde{t}_n + (1-\tilde{q}_n)\theta_n$ for $n \neq m$. t_N^m is defined so that consistency is satisfied.

Incentive compatibility of the optimal 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 simple mechanisms in Ω. In particular, 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 under $(\tilde{q},\tilde{t})$; we verify this in the proof of Proposition 4. 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

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

The above equation shows that the expected revenue of the ambiguous mechanism Ω, constructed from the simple mechanism $(\tilde{q},\tilde{t})$, can be written as the expected revenue from the simple mechanism (the first term) plus a term that depends on the allocation in the simple mechanism $(\tilde{q},\tilde{t})$. This second term is strictly larger than zero if and only if there is a $n \neq m_M, N$, such that $\tilde{q}_n < 1$.

Thus, whenever $(\tilde{q},\tilde{t})$ is such that $\tilde{q}_n < 1$ for some $n \neq m_M, N$, then Ω is an ambiguous mechanism that delivers a strictly larger expected revenue than $(\tilde{q},\tilde{t})$. What if, $\tilde{q}_n = 1$ for all $n \neq m_M, N$? If this condition holds and $m_M \neq 1$ then $\tilde{q}_1 = 1$; by monotonicity of $\tilde{q}$ it then follows that $\tilde{q} = (1\ldots,1)$, a case that we have already considered before. So the only case that is left to be considered is $\tilde{q} = (0,1,\ldots,1)$.[25] If the object is assigned to all types but the lowest one, then the expected revenue is $\tilde{R} = (1-p_1)\theta_2$. But when $m_M = 1$, then the ambiguous mechanism that

[25] If $\tilde{q}_{m_M} = \tilde{q}_1 \neq 1$ then we can assume that $\tilde{q}_1 = 0$. If choosing $q = (q_1,1,\ldots,1)$ with $0 < q_1$ was optimal then so would be $q = (1,\ldots,1)$. But we have already dealt with that case.

is generated by $Q = (0, \dots, 0)$ yields an expected revenue of

$$R = \mathbb{E}_p[\theta] - p_1\theta_1 = (1 - p_1)\theta_2 + \sum_{n=3}^{N} p_n(\theta_n - \theta_2) > \tilde{R},$$

meaning that also in this case there exists an ambiguous mechanism that does strictly better than the optimal simple mechanism.

We summarize the preceding observations in the following proposition.

Proposition 4. *An optimal ambiguous mechanism yields a strictly larger revenue than the best simple mechanism if and only if $\bar{v}_1 < 0$. That is, the use of ambiguous mechanisms is strictly beneficial for the principal if and only if $Q = (1, \dots, 1)$ is not a solution of Problem P”.*

Proof. See the Appendix. ∎

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

Example 1 (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 3 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 $\mathcal{M} = \{m_1, \dots, m_{M+1}\}$ 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) $p_1\theta_1 \le p_2\theta_2$.

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 $\bar{v}_1 \le 0$. 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 $\mathbb{E}_p[\theta]$. (3) implies 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 that the agent is awarded the good with probability one, 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, 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 (**??**) 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 4.

ii) $p_1\theta_1 \leq p_2\theta_2$: In this case $\mathcal{M} = \{1, 2, 3\}$, 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, $(p_2 + p_3)\theta_2$. ■

5 Discussion and extensions

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

In this section we focus on the ambiguous mechanism generated by $Q = (0,\dots,0)$.[26] This mechanism takes a particularly simple form: the transfer rule corresponding to the m-th outcome function, t^m, is given by

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

The expectation of this transfer is $R = \mathbb{E}_p[\theta] - p_{m_M}\theta_{m_M}$. So the ambiguous mechanism generated by $Q = (0,\dots,0)$ 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 the one of type θ_{m_M}) is small, as it can be the case in settings with 'large' type sets. The optimal ambiguous mechanism then leaves to the agent at most this much surplus. 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 5 (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 $n < N$, the outcome function (q^n, t^n)

[26] Remember that we say that $\Omega = \{(q^1,t^1),\dots,(q^{N-1},t^{N-1})\}$ is generated by Q, if Ω satisfies the properties (C), (Uni), (Min), (Mon) and $(q_1^1, q_2^2, \dots, q_{N-1}^{N-1}) = Q$.

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. 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.

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 than each single element of an ambiguous mechanism.

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 5.

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*. Moreover, in this case the mechanism is ex post efficient with probability one. That is, each simple mechanism transfers the good to the agent with probability one.

5.2 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.

In order to see this, consider again our basic set up with N payoff relevant types, $\Theta = \{\theta_1, \ldots, \theta_N\}$. 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 $V = \{v_1, \ldots, v_K\}$; 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 3 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 2 (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 3 he should offer the ambiguous mechanism Ω composed by the outcome functions described in the following table.[27]

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

Table 3: An ambiguous mechanism conditioning on payoff-irrelevant type dimensions

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

'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.

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

Our results allow us to determine with precision the benefits of larger type sets.

5.3 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—risk neutrality—is crucial in the final steps of the characterization of the optimal ambiguous mechanism (i.e. Proposition 3 relies on this assumption). In all results up to Lemma 4 we have only exploited the increasing difference property of the linear valuation function. That is, all those results go through for valuation functions that exhibit 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) are described by the function $u(x,\tau,\theta)$, where x is a fraction of the good, then an ambiguous mechanism like the one used in Corollary 3 can be constructed whenever the problem

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

admits a solution.[28] 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. MMEU certainly constitutes a rather stark model of ambiguity aversion. Our analysis heavily exploits the tractability of these 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 might be able to exploit his ambiguity aversion by offering an ambiguous mechanism. In the following example we show that this insight applies also in environments

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

where the agent's attitude toward uncertainty is represented by a model of smooth ambiguity aversion.

Example 3 (Smooth ambiguity aversion)**.** The setup is as in the example considered in Section 2, except for the agent's attitude towards ambiguity. That is, we have $\Theta = \{1,2,4\}$, $p = (1/4,1/4,1/2)$, $u(x,\tau,\theta) = x\theta - \tau$.

The only difference with respect to the example in Section 2 lies in the agent's attitude towards ambiguity. Here we consider the case of an agent who is smoothly ambiguity averse in the sense of Klibanoff, Marinacci, and Mukerji (2005). In particular, we assume that when faced with 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 capture 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.

For the sake of concreteness, in what follows we assume that $\phi(x) = 1 - \exp(-7x)$, 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 seems 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 treat the different outcome functions asymmetrically.[29]

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

[29] 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.

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

Table 4: An ambiguous mechanism with $R = t/2 + 1/4$ (under truth-telling)

It is straightforward to verify that under truth-telling the expected revenue of both outcome functions is $R = t/2 + 1/4$. We will now solve for the largest t such that this mechanism is incentive compatible. The following table shows the payoffs that each type θ obtains from the available messages $\hat{\theta}$.[30]

$\theta \setminus \hat{\theta}$	1	2	4
1	$\phi(0)$	$\phi(0)$	$\phi(1-t)$
2	$\phi(1)/2$	$\phi(1)/2$	$\phi(2-t)$
4	$\phi(3)/2$	$\phi(3)/2$	$\phi(4-t)$

Table 5: Payoffs for the ambiguous mechanism in the preceding table

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 two lowest type's incentive compatibility constraints are satisfied if $t > 2$ ($\phi(0) > \phi(1-t)$ and $\phi(1)/2 > \phi(2-t)$). The highest type has no incentive to deviate if $\phi(4-t) \geq \phi(3)/2$. The largest t which satisfies this condition is approximately $t = 3.9$.

With $t = 3.9$ the ambiguous mechanism generates an expected revenue of $R = t/2+1/4 = 2.2$ which exceeds the revenue of the best simple mechanism by 0.2.

■

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 1–4 as those results do not refer to the designer's preferences. It is also easily seen that the ambiguous mechanism characterized in Proposition 3 remains the optimal ambiguous mechanism in the class of direct ambiguous mechanisms that are consistent when the seller is risk neutral and his preferences are represented by max-min expected utility. 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

[30] Remember that the agent's belief is described by the uniform distribution over the degenerate distributions on Ω.

them in Section 3. 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 indeed introduces an aspect that is not present in the case of an ambiguity neutral designer: in such a framework the principal may 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.

In situations where the principal is MMEU and a third party picks the outcome function there is no need to impose consistency. Consistency has been introduced in order to make sure that the buyer's beliefs as to which outcome function in the ambiguous mechanism he is facing take into account the seller's motives in choosing that outcome function. 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, including those that do not satisfy consistency, than we have considered in this paper so far. Fortunately, it is not difficult to see that non-consistent mechanisms can be disregarded without loss of generality when the designer is ambiguity averse.

In order to see this notice first that dropping consistency would not upset the validity of the Revelation Principle. Without the requirement that every outcome function in an ambiguous mechanism needs to yield the same expected revenue the relaxed v Problem P becomes

$$\max_{R\in\mathbb{R},\Omega\subset X^{\Theta}\times\mathbb{R}^{\Theta}} \quad R \tag{P-2}$$

$$\text{s.t.} \quad R \leq \min_{(q,t)\in\Omega} \mathbb{E}_p[t(\theta)], \tag{C'}$$

$$\text{(DIC)}, \text{(UIC)}, \text{(IR)}.$$

Note that the only difference between this problem and Problem P is that the constraint (C) in the latter is replaced by the constraint (C') in Problem P-2. (C') represents the fact that the payoff of a seller who is MMEU ambiguity averse and who delegates the choice of the outcome function that will be implemented to an external 'ambiguity device' (uninterested third party), is given by the minimum of the expected revenues that the outcome functions in his ambiguous mechanism generate.

The following result tells us that replacing (C) with (C') is inconsequential for the values of the optimization problems.

Proposition 6. *The value of Problem P-2 coincides with the value of Problem P.*

Proof. See the Appendix. ■

The intuition for this result is rather straightforward: given that the principal's payoff is determined by the outcome function that delivers the lowest expected revenue, he might as well choose to start with a mechanism that contains only outcome functions that yield the same expected revenue. Suppose we are given a mechanism Ω that does not satisfy this condition. Let R denote the minimum of the expected revenues of the outcome functions in Ω. By appropriately lowering the transfers of the highest type in all outcome functions who do not produce an expected revenue of R one can obtains a mechanism that is still downward incentive compatible and individually rational. Applying Lemmata 1 through 3 this mechanism can be further transformed to obtain a new mechanism that is also upward incentive compatible and generates an expected revenue of at least R.

5.4 Surplus extraction: ambiguity aversion vs. risk aversion

In Section 5.1 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.[31] 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 tends to infinity. In the light of these observations our full rent extraction result for the case where the type set is a continuum, may seem to be similar to the results in Matthews (1983).

But the analogy between risk aversion and stochastic mechanisms on the one hand and (smooth) ambiguity aversion and ambiguous mechanisms on the other hand is not quite as close as it appears at first sight. The central difference between the two cases lies in the fact that in the case of stochastic mechanisms the distribution of the outcomes is determined by the designer. That is, the distribution of the outcomes is given by an objective probabilistic

[31] 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 3.

distribution that is a choice variable of the principal. There is no analogous instrument in the case of smooth ambiguity aversion.

Even though in that model uncertainty is described by a distribution (μ) over (distributions of) outcomes, the higher level distribution (μ)—which is the mathematical analog of the probability distributions over outcomes in a stochastic mechanism—does not allow for an objective interpretation.[32] In particular, one cannot think of it as a variable that the principal can choose and can commit to in the way he chooses and commits to the distributions in a stochastic mechanism. Instead, the higher level distribution (μ) is only a description of the uncertainty that the agent perceives and – by the very idea that underlies the concept of uncertainty – there is no sense in which this perception could be given an objective interpretation as can be done in a setting with risk aversion by considering only objective probabilistic distribution of outcomes. Thus, unlike in the case of a stochastic mechanism which fully pins down the perceptions of the risk averse agent by specifying in an objective way all aspects of the distribution of the outcomes, an ambiguous mechanism can always just determine the support of the distribution that describes the agent's perceptions. All remaining aspects of the agent's uncertainty perception are necessarily of a purely subjective nature.

While on the one hand the fact that the higher order distribution in the smooth ambiguity model describes the subjective uncertainty perceptions of the agent means that it is outside the direct control of the principal, on the other hand it also implies that it should be considered as endogenous with respect to the principal's choices. In particular, this endogeneity of the agent's perceptions would call for the use of the consistency concept in an environment with a smoothly ambiguity averse agent for the same reasons we have adopted it in our setting with MMEU preferences. Notice that there is no analogous constraint that has to be considered in a setting with risk aversion and stochastic mechanisms. Thus, the conceptual differences between the two cases also translate into important differences in their mathematical treatment.

5.5 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.[33] 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 extend to the case with more than two agents.

[32] See the discussion of this issue in Klibanoff, Marinacci, and Mukerji (2012).

[33] 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 (2013).

Assume that the two agents have preferences as in the previous sections. We denote the type set of agent, $i = 1,2$, 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' types 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 7 (Full surplus extraction with multiple agents)**.** *Consider a two agent setting as described in the preceding paragraph. Moreover, let the ambiguous mechanism* $\Omega = \{(q_\theta, t_\theta) : \theta \in \Theta\}$*, be defined by*

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

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.*

The ambiguous mechanisms presented in the above result can be constructed as follows. For every profile of types $\tilde\theta = (\tilde\theta^1, \tilde\theta^2)$ we add to the ambiguous mechanism a simple mechanism $(q^{\tilde\theta}, t^{\tilde\theta})$ with the following property. If both agents' reports coincide with the corresponding component of the label of the simple mechanism, i.e. if for every i we have $\hat\theta^i = \tilde\theta^i$, then the seller keeps the object and there are no transfers. If, neither agent's report coincides with label, i.e. if for every i, $\theta^i \neq \tilde\theta^i$, then the agent with the higher report receives the object at the price he reported. Finally, if only one agent's report coincides with the label of the simple mechanism, then the agent whose report does not coincide receives the object at the price equal to the value he announced. Since the ambiguous mechanism contains all such simple mechanisms, there is a simple mechanism for each report of each agent such that the agent does not receive the object if he reports that type. This bounds the agent's expected payoff above by 0. By reporting truthfully the agent either obtains the object and pays the reported, and therefore true, value or does not receive it and pays zero. In either case the agent's payoff is zero, which is also the before established upper bound. But then the agent has no incentives to deviate from truthful reporting. This establishes that truthful reporting is an equilibrium of the proposed mechanism. On the other hand, assuming that the agents do report truthfully and that each profile of types occurs with probability zero, each simple mechanism yields for the seller the expected surplus

$\mathbb{E}[\max\{\theta^1, \theta^2\}]$. With other words, the seller extracts the full surplus.

6 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 3 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.

A Appendix

Proof of Proposition 1. Optimality of σ implies

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

Consider the direct ambiguous mechanism Ω' defined in the proposition. 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

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

and so we can conclude that Ω' is incentive compatible.

That Ω' is consistent with respect to thruthtelling follows immediately from the fact that Ω is consistent with respect to σ. ■

Proof of Lemma 1. Let Ω be a mechanism that satisfies (C), (DIC) and (IR). Let Ω_1 be the set of all simple mechanisms of the form (q,t'), where $(q,t)\in\Omega$, $t'_n = t_n$ for all $1 < n < N$, and

$$t'_1 = q_1\theta_1 \qquad \text{and} \qquad t'_N = t_N - \frac{p_1}{p_N}[t'_1 - t_1].$$

Since Ω satisfies (IR), $t'_1 \geq t_1$ and $t'_N \leq t_N$. Thus, in passing from t to t', the transfer of type θ_1 is increased until his truth-telling payoff is zero, while that of θ_N is lowered so that t and t' have the same expected value. Since Ω satisfies (C) so does Ω_1; in particular, $R(\Omega) = R(\Omega_1)$. By construction Ω_1 satisfies the individual rationality constraint of the lowest type with equality. Moreover, it is also downward incentive compatible. To see this observe that the truth-telling payoffs of all types θ_n, $1 < n < N$, are the same under Ω and Ω_1. The highest type's truth-telling payoff instead is (weakly) larger in Ω_1 than in Ω. Regarding the payoffs from downward deviations, observe that the only downward deviation report which may deliver different payoffs in Ω than Ω_1 is θ_1. But whenever that is the case it is lower in Ω_1 than in Ω. Given that the truth-telling payoffs of types who have downward deviation opportunities are at most higher in Ω_1 than in Ω, it follows that Ω_1 must indeed satisfy (DIC) given that Ω does.

Ω_1 constitutes the base case for our inductive argument. We next show the inductive step. Suppose we have defined the mechanisms $\Omega_1, \ldots, \Omega_{n-1}$ for some $1 < n < N-1$. Proceeding in similar fashion as before, we define Ω_n as the set of all simple mechanisms of the form (q,t'), where $(q,t)\in\Omega_{n-1}$, and the transfer rule t' coincides with t except for the transfers of types θ_n and θ_N, which are

$$t'_n = q_n\theta_n - \max_{1<m<n}\inf_{(\tilde{q},\tilde{t})\in\Omega_{n-1}} \tilde{q}_m\theta_n - \tilde{t}_m \qquad \text{and} \qquad t'_N = t_N - \frac{p_n}{p_N}[t'_n - t_n].$$

Since Ω_{n-1} is satisfies (DIC), it follows that $t'_n \geq t_n$ and $t'_N \leq t_N$. The above transformation increases the transfer to be paid upon reporting type θ_n in all outcome functions until type θ_n's truth-telling payoff from each outcome function becomes equal to the payoff from his most attractive downward deviation. As before, Ω_n inherits the properties (C), (DIC) and (IR) from Ω_{n-1}; moreover, by construction $R(\Omega_n) = R(\Omega_{n-1})$.

Finally, define Ω_N as the set of all outcome functions of the form (q,t'), where $(q,t)\in\Omega_{N-1}$ and the transfers t' coincide with t except for the transfers of the highest type which are set equal to

$$t'_N = t_N + \inf_{(\tilde{q},\tilde{t})\in\Omega_{N-1}}\{\tilde{q}_N\theta_N - \tilde{t}_N\} - \max_{1\leq m<N}\inf_{(\tilde{q},\tilde{t})\in\Omega_{N-1}}\{\tilde{q}_m\theta_N - \tilde{t}_m\}.$$

Since Ω_{N-1} is downward incentive compatible it follows that the difference

$$\inf_{(\tilde{q},\tilde{t})\in\Omega_{N-1}} \{\tilde{q}_N\theta_N - \tilde{t}_N\} - \max_{1\leq m<N} \inf_{(\tilde{q},\tilde{t})\in\Omega_{N-1}} \{\tilde{q}_m\theta_N - \tilde{t}_m\}$$

is non-negative. This means that in passing from Ω_{N-1} to Ω_N the transfers of the highest type are increased by the same amount in all outcome functions. Thus, property (C) is preserved and $R(\Omega_N) \geq R(\Omega_{N-1})$.

In the first step we have shown that the transfers of the lowest type can be increased so that each outcome function gives payoff zero to the lowest type when he reports truthfully. The inductive step then takes an ambiguous mechanism in which the downward incentive compatibility with respect to types θ_1 through θ_{n-1} are binding and shows that one can increase the transfers of type θ_n and decrease the transfers of the type θ_N so that also type θ_n's downward incentive compatibility constraint becomes binding and the expected transfer of each of the outcome functions does not change. In the last step the transfers t_N in *all* outcome functions are uniformly increased until θ_N's downward incentive compatibility constraint becomes binding. Ω_N is, therefore, an ambiguous mechanism that satisfies all conditions of (Uni). Thus, setting $\Omega' = \Omega_N$ proves the lemma.

■

Proof of Lemma 2. Let Ω be an ambiguous mechanism that satisfies (C) and (Uni). Denote its closure (in the usual Euclidean sense) by $\bar{\Omega}$. Clearly, $\bar{\Omega}$ inherits the properties (C) and (Uni) from Ω. Moreover, $R(\bar{\Omega}) = R(\Omega)$. For each $1 \leq m < N$, let Ω_m be the set of ambiguous mechanisms in $\bar{\Omega}$ that minimize the probability of the allocation for the report θ_m, that is,

$$\Omega_m = \{(q,t) \in \bar{\Omega} : q_m \leq q'_m \text{ for all } (q',t') \in \bar{\Omega}\}. \tag{6}$$

Let $m < n \leq N$. Since $\theta_n > \theta_m$, an outcome function $(q,t) \in \bar{\Omega}$ belongs to Ω_m if and only if $(q_m - q'_m)(\theta_n - \theta_m) \leq 0$ for all $(q',t') \in \Omega$. By (Uni), $q_m\theta_m - t_m = q'_m\theta_m - t'_m$, hence the above inequality can be written as $q_m\theta_n - t_m \leq q'_m\theta_n - t'_m$. It follows that Ω_m is also the set of outcome functions that minimizes the payoff of any type $\theta_n > \theta_m$ when he untruthfully reports θ_m, that is, for all $m < n \leq N$

$$\Omega_m = \{(q,t) \in \bar{\Omega} : q_m\theta_n - t_m \leq q'_m\theta_n - t'_m \text{ for all } (q',t') \in \bar{\Omega}\}. \tag{7}$$

The mechanism Ω' is now obtained as follows: for each $1 \leq m < N$, let (q^m, t^m) be an arbitrarily chosen element from Ω_m. Since, Ω' is composed of simple mechanisms that where picked from $\bar{\Omega}$, it immediately follows that Ω' satisfies (C) and $R(\Omega') = R(\bar{\Omega})$. Moreover, uniformity of $\bar{\Omega}$ implies that the truth-telling payoffs in Ω' must be the same as in $\bar{\Omega}$. Finally, Ω' is composed exactly of those elements of $\bar{\Omega}$ which define the downward deviation payoffs in $\bar{\Omega}$. Therefore the downward deviation payoffs under both mechanisms coincide as well. Since neither the truth-telling payoffs nor the downward deviation payoffs change in passing from $\bar{\Omega}$ to Ω', we can conclude that Ω' inherits the property (Uni) from $\bar{\Omega}$. ■

Proof of Lemma 3. Let $\Omega = \{(q^1,t^1),\dots,(q^{N-1},t^{N-1})\}$ be a mechanism satisfying (C), (Uni) and (Min).

We first show that Ω may be changed so that it satisfies the first part of (Mon) while still satisfying (C), (Uni) and (Min). Consider the mechanism $\bar{\Omega} = \{(\bar{q}^1,t^1),\dots,(\bar{q}^{N-1},t^{N-1})\}$, where for each $1 \le m < N$ the allocation rule $\bar{q}^m$ is defined as follows. For every $1 \le n \le N$, $\bar{q}^m_n = q^m_n$ if $n = m$ and $\bar{q}^m_n = 1$ otherwise. By construction, $\bar{\Omega}$ satisfies the first part of (Mon). Moreover, $\bar{\Omega}$ has the same transfer rules as Ω. Given that the latter satisfies (C) so must $\bar{\Omega}$. For the same reason we also have $R(\bar{\Omega}) = R(\Omega)$. Next, observe that the downward deviation payoffs in Ω and $\bar{\Omega}$ are the same. Clearly, these payoffs cannot decrease due to the increase in the allocations that occurs when passing from Ω to $\bar{\Omega}$ (transfers do not change). That they cannot increase follows from the construction of $\bar{\Omega}$ and the fact that Ω satisfies (Min). Taken together these properties imply that for all $1 \le m < n \le N$ we have

$$\min_{1\le\ell<N}\{q^\ell_m\theta_n - t^\ell_n\} = q^m_m\theta_n - t^m_m = \bar{q}^m_m\theta_n - t^m_m = \min_{1\le\ell<N}\{\bar{q}^\ell_m\theta_n - t^\ell_n\}. \tag{8}$$

Since also the truth-telling payoffs can at most increase, it follows that $\bar{\Omega}$ inherits from Ω the properties (DIC) and (IR). Applying Lemma 1 to $\bar{\Omega}$ delivers an ambiguous mechanism $\tilde{\Omega} = \{(\tilde{q}^1,\tilde{t}^1),\dots,(\tilde{q}^{N-1},\tilde{t}^{N-1})\}$ that satisfies (Uni). Since this last step does not involve any changes in the allocation rules it follows that $\tilde{\Omega}$ satisfies the first part of (Mon), so that for all $1 \le m \le N-1$ we have $\tilde{q}^m_m \le \tilde{q}^\ell_m$ for all $1 \le \ell \le N-1$. In the proof of Lemma 2 we have seen that this implies that $(\tilde{q}^m,\tilde{t}^m)$ is the outcome function that defines the payoff from downward deviations towards θ_m. We can therefore conclude that $\tilde{\Omega}$ satisfies (Min) as well.

We now come to the second part of (Mon). Let $\Omega = \{(q^1,t^1),\dots,(q^{N-1},t^{N-1})\}$ be a mechanism satisfying (C), (Uni), (Min) and the first part of (Mon). Observe first that if for $1 \le m < m' < N$ we have $q^m_m \le q^{m'}_{m'}$ then for every type θ_n, $n > m'$, a deviation to θ_m can never be more attractive than a deviation to $\theta_{m'}$. In order to see this, remember that (Uni) implies (DIC) and so type $\theta_{m'}$ must be better off by reporting truthfully than by reporting θ_m:

$$q^m_m\theta_{m'} - t^m_m \le q^{m'}_{m'}\theta_{m'} - t^{m'}_{m'}. \tag{9}$$

Then

$$t_{m'}^{m'} - t_m^m \leq q_{m'}^{m'}\theta_{m'} - q_m^m\theta_{m'} \leq (q_{m'}^{m'} - q_m^m)\theta_n$$

for $m' < n \leq N$. Therefore

$$q_m^m\theta_n - t_m^m \leq q_{m'}^{m'}\theta_n - t_{m'}^{m'} \qquad \text{for all } m' < n \leq N. \tag{10}$$

Now assume that Ω does not satisfy the second part of (Mon), and let m' be the smallest $1 \leq m < N$ for which the condition is violated. Thus,

$$q_{m'}^{m'} < q_{m'-1}^{m'-1} \quad \text{and} \quad q_m^m \geq q_{m-1}^{m-1} \qquad \forall\ 1 < m < m'. \tag{11}$$

By our previous observation for every type θ_n, $n > m' - 1$ the most attractive deviation in the set $\{\theta_1, \ldots, \theta_{m'-1}\}$ is $\theta_{m'-1}$. Since Ω is uniform this means that the payoff of type $\theta_{m'}$ from reporting truthfully and from reporting $\theta_{m'-1}$ is the same, i.e.

$$q_{m'-1}^{m'-1}\theta_{m'} - t_{m'-1}^{m'-1} = q_{m'}^{m'}\theta_{m'} - t_{m'}^{m'}. \tag{12}$$

But if type $\theta_{m'}$ is indifferent between the reports $\theta_{m'}$ ant $\theta_{m'-1}$, then for each $n > m'$, type θ_n must strictly prefer reporting $\theta_{m'-1}$ over reporting θ_m. Indeed, (12) implies

$$q_{m'-1}^{m'-1}\theta_n - t_{m'-1}^{m'-1} > q_{m'}^{m'}\theta_n - t_{m'}^{m'} \qquad \text{for all } m' < n \leq N. \tag{13}$$

Thus, the downward deviation constraint with respect to $\theta_{m'}^{m'}$ cannot be binding for any type θ_n, $n > m'$.

Consider now the mechanism $\bar{\Omega} = \{(\bar{q}^1, \bar{t}^1), \ldots, (\bar{q}^{N-1}, \bar{t}^{N-1})\}$ that coincides with Ω except for the values of $\bar{q}_{m'}^{m'}$, $\bar{t}_{m'}^{m'}$ and $\bar{t}_N^{m'}$. $\bar{q}_{m'}^{m'}$ is increased to $q_{m'-1}^{m'-1}$, $\bar{t}_{m'}^{m'}$ is increased so that the payoff that type $\theta_{m'}$ gets under the outcome function $(\bar{q}^{m'}, \bar{t}^{m'})$ when reporting truthfully is the same one that he gets under the outcome function $(q^{m'}, t^{m'})$, i.e.

$$\bar{q}_{m'}^{m'}\theta_{m'} - \bar{t}_{m'}^{m'} = q_{m'}^{m'}\theta_{m'} - t_{m'}^{m'}.$$

Finally, $\bar{t}_N$ is chosen such that the expected values of $\bar{t}^{m'}$ and $t^{m'}$ coincide, i.e.

$$\bar{t}_N^{m'} = t_N^{m'} - \frac{p_{m'}}{p_N}\left[\bar{t}_{m'}^{m'} - t_{m'}^{m'}\right].$$

Since $\bar{t}_m^{m'} \geq t_m^{m'}$, it follows that $\bar{t}_N^{m'} \leq t_N^{m'}$. Notice that in passing from Ω to $\bar{\Omega}$ only the consequences of reporting $\theta_{m'}$ and θ_N under outcome function $(q^{m'}, t^{m'})$ are affected. We will argue now that $\bar{\Omega}$ satisfies all desired properties except possibly the third condition in (Uni) (a binding downward deviation incentive constraint of type θ_N). Transfers have been modified in a way such that the expected value of $t^{m'}$ remains unchanged. Thus $\bar{\Omega}$ satisfies (C) with $R(\bar{\Omega}) = R(\Omega)$.

The truth-telling payoff of type $\theta_{m'}$ under $(\bar{q}^{m'}, \bar{t}^{m'})$ does not change with respect to $(q^{m'}, t^{m'})$. Thus, also in $\bar{\Omega}$ type $\theta_{m'}$ gets the same truth-telling payoff from all outcome functions. From Lemma 2 we know that in combination with $\bar{q}^{m'}_{m'} \leq \bar{q}^{\ell}_{m'}$, $1 \leq \ell \leq N-1$, this implies that $(\bar{q}^{m'}, \bar{t}^{m'})$ defines the payoff from downward deviations toward $\theta_{m'}$ by types θ_n, $n > m'$. Since no other downward deviation payoffs could be affected when passing from Ω to $\bar{\Omega}$, we can conclude that the latter must satisfy (Min).

The payoffs from downward deviations toward $\theta_{m'}$ increase, but they cannot exceed those from deviations toward $\theta_{m'-1}$. To see this, notice that since Ω satisfies (Uni) we have

$$\bar{q}^{m'}_{m'}\theta_{m'} - \bar{t}^{m'}_{m'} = q^{m'}_{m'}\theta_{m'} - t^{m'}_{m'} = q^{m'-1}_{m'-1}\theta_{m'} - t^{m'-1}_{m'-1} = \bar{q}^{m'-1}_{m'-1}\theta_{m'} - \bar{t}^{m'-1}_{m'-1}.$$

Combining this with $q^{m'}_{m'} = \bar{q}^{m'-1}_{m'-1}$ we get $\bar{t}^{m'}_{m'} = \bar{t}^{m'-1}_{m'-1}$. Thus, the payoff that any type can get from report $\theta_{m'}$ under outcome function $(\bar{q}^{m'}, \bar{t}^{m'})$ coincides with the payoff that he gets from report $\theta_{m'-1}$ under outcome function $(\bar{q}^{m'-1}, \bar{t}^{m'-1})$. Since, the latter defines the payoffs from downward deviations to $\theta_{m'-1}$ the claim follows. We can therefore conclude that $\bar{\Omega}$ satisfies the first two conditions in (Uni).

As for the third requirement of (Uni) (a binding incentive constraint for downward deviations by the highest type) observe that the truth-telling payoff of type θ_N may increase since $t^{m'}_N$ decreases. If this is not the case, then $\bar{\Omega}$ satisfies all conditions of (Uni) and so we are done by setting $\Omega' = \bar{\Omega}$.

If instead the highest type's truth-telling payoff is higher in $\bar{\Omega}$ than in Ω, then consider the mechanism Ω' which coincides with $\bar{\Omega}$ everywhere except for the transfers of the highest type. The latter are chosen as follows: for each $1 \leq m < N$, set $t'^{m}_N = \bar{t}^m_N + \varepsilon$, where ε is given by

$$\varepsilon = \min_{1\leq\ell<N}\{\bar{q}^{\ell}_N\theta_N - \bar{t}^{\ell}_N\} - \max_{1\leq m<N}\min_{1\leq\ell<N}\{\bar{q}^{\ell}_m\theta_m - \bar{t}^{\ell}_m\}.$$

Since in passing from $\bar{\Omega}$ to Ω' the highest type's transfer is increased uniformly across all outcome functions it follows that Ω' satisfies (C) with $R(\Omega') > R(\Omega)$. Moreover, the fact that by the switch from $\bar{\Omega}$ to Ω' only the transfers for the highest type are affected, implies that Ω' inherits both (Min) and the first two parts of (Uni) from $\bar{\Omega}$. Finally, in Ω' the transfers of the highest are chosen exactly such that the highest type's downward deviation constraint is binding. Consequently, Ω' satisfies all parts of (Uni) and so the proof is complete. ■

Proof of Lemma 4. Let Ω be a mechanism that satisfies (Uni), (Min) and (Mon). 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 equality is implied by (Uni). The equalities between the two inequalities follow from the fact that for each $1 < n \leq N$, the binding downward incentive constraint of type θ_n, is the one with respect to the adjacent lower type θ_{n-1}. This has been shown in the proof of Lemma 3, where we have seen that in a uniform ambiguous mechanism (one of) the binding incentive constraint for downward deviations for type θ_n, $1 < n < N$, is the one with respect to type θ_{m_n}, where

$$m_n \in \arg\max_{1\leq m<n}\left\{\min_{1\leq\ell<N}\{q_m^\ell\}\right\}.$$

For a mechanism that is also minimal and monotonic it thus follows that $m_n = n-1$.

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. ■

Proof of Proposition 2. By Lemma (4) we know that every mechanism that is feasible in Problem P' also satisfies (UIC). It also satisfies (DIC) and (IR) because it satisfies (Uni); see the paragraph after Definition 4. Therefore every such mechanism is feasible in Problem P and thus the value of Problem P cannot be smaller than the value of Problem P'.

On the other hand, Lemmata 1 through 3 imply that for every mechanism that is feasible in Problem P there exists a mechanism with at least as high an expected revenue for the seller, that is feasible in Problem P'. But then the value of Problem P' must be at least as large as the value of Problem P. ■

Proof of Lemma 5. If Ω satisfies (Uni), (Min) and (Mon), then for each $1 < n \leq N$ the binding downward incentive constraint of type θ_n is the one with respect to θ_{n-1} (see the proof of Lemma 3). Thus,

$$
\begin{aligned}
t_1^m &= q_1^m \theta_1 && \text{for all } 1 \leq m < N, && (14)\\
t_{n+1}^m &= (q_{n+1}^m - q_n^n)\theta_{n+1} + t_n^n && \text{for all } 1 \leq n \leq N-2 \text{ and } 1 \leq m < N,\\
t_N^m &\leq (1 - q_{N-1}^{N-1})\theta_N + t_{N-1}^{N-1}, && \text{for all } 1 \leq m < N.
\end{aligned}
$$

From this it is straightforward to derive (3) and (4) by recursively substituting the expressions t_n^n into the formula for t_{n+1}^m.

Conversely, suppose we are given a minimal and monotonic mechanism Ω the transfers of which satisfy (3) and (4). It is then easily verified that for each $1 \leq n < N$ the truth-telling payoffs of type θ_n are constant across outcome functions. This in turn implies that the payoff type θ_n, $n \leq N$, obtains from a deviation to θ_m, $1 \leq m < n$, is determined by the outcome function (q^m, t^m) (since this is the outcome function that minimizes the allocation probability after report θ_m). Using this it is straightforward to show that the payoff of type θ_n from deviating θ_m, $m < n$ is increasing in m and is equal to the truthtelling payoff for $m = n-1$. With other words, type θ_n's downward adjacent IC constraint is binding.

In order to see this, consider the difference in the payoff of type θ_n from reporting θ_m and θ_{m-1}, $m \leq n$.

$$
\min_{1 \leq l < N} \{q_m^l \theta_n - t_m^l\} - \min_{1 \leq l < N} \{q_{m-1}^l \theta_n - t_{m-1}^l\} = q_m^m \theta_n - t_m^m - q_{m-1}^{m-1}\theta_n - t_{m-1}^{m-1} \tag{15}
$$

$$
= q_m^m \theta_n - q_m^m \theta_m + \sum_{k=1}^{m-1} q_k^k (\theta_{k+1} - \theta_k) - q_{m-1}^{m-1}\theta_n + q_{m-1}^{m-1}\theta_{m-1} - \sum_{k=1}^{m-2} q_k^k (\theta_{k+1} - \theta_k) \tag{16}
$$

$$
= (\theta_n - \theta_m)(q_m^m - q_{m-1}^{m-1}) \geq 0. \tag{17}
$$

Given that $q_m^m \geq q_{m-1}^{m-1}$ this difference is non-negative. Moreover, it becomes zero for $m = n$. Hence, Ω satisfies all downward incentive constraints and the one with respect to the adjacent lower type is binding. Combining all these observations we can conclude that Ω satisfies (Uni). ∎

Proof of Lemma 6. 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 $1 < j \leq M$ then it must be so also for all $1 \leq k < j$. ■

Proof of Proposition 3. 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 $m_j < n < m_{j+1}$, $1 \leq j \leq M$, 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 $q_n^n \geq q_{m_j}^{m_j}$ it follows that

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

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

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

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

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

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

we have

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

Moreover, rewriting the inequality yields

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

In such a case we can lower $q_{m_{j+1}}^{m_{j+1}}$ without violating the constraint $q_{m_{j+1}}^{m_{j+1}} \geq q_{m_j}^{m_j}$, 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 $s = j+1, \ldots, M$ implies that

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

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

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

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

Step 4. If $\bar{\nu}_1 \le 0$ 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

$$\hat{q}_{m_s}^{m_s} = 1 - \frac{p_{m_{j^*}}\theta_{m_{j^*}}}{p_{m_s}\theta_{m_s}}(1-\hat{q}_{m_{j^*}}^{m_{j^*}}) \tag{18}$$

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 $\bar{\nu}_{m_{j^*}}$. Thus, if $\bar{\nu}_{m_{j^*}} \le 0$, 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

$$\bar{R}^{m_{j^*}}(Q) - \bar{R}^{m_{j^*-1}}(Q) = (1-q_{m_{j^*-1}}^{m_{j^*-1}})(p_{m_{j^*-1}}\theta_{m_{j^*-1}} - p_{m_j^*}\theta_{m_j^*}). \tag{19}$$

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 (18). In order to see this notice that by (19) 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, $m_{j^*-1} \le n < N$ instead we have $\bar{R}^{m_{j^*}} < \bar{R}^{m_{j^*-1}}$. By our previous arguments we know that a reduction of $(q_{m_{j^*}}^{m_{j^*}},\dots,q_{N-1}^{N-1})$ in accordance with (18) leads to an increase of $\bar{R}^{m_{j^*}}$ and $\min_j \bar{R}^{m_j}$. If in addition also $(q_{m_{j^*-1}}^{m_{j^*-1}},\dots,q_{m_{j^*}-1}^{m_{j^*}-1})$ 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

$$\hat{q}_{m_j}^{m_j} \leq 1 - \frac{p_{m_{j-1}}\theta_{m_{j-1}}}{p_{m_j}\theta_{m_j}}(1 - \hat{q}_{m_{j-1}}^{m_{j-1}}) \tag{20}$$

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.

■

Proof of Proposition 4 . In the main text it has been shown that a necessary and sufficient condition for the existence of an ambiguous mechanisms that does strictly better than the best simple mechanism(s) is that $Q = (1, \ldots, 1)$ is not a solution of Problem P''. The two parts of the proof that have not been argued in the text, are i) the incentive compatibility and individual rationality of the ambiguous mechanism Ω constructed from the optimal simple mechanism $(\tilde{q}, \tilde{t})$, and ii) the equivalence result showing that $Q = (1, \ldots, 1)$ is not a solution to Problem P'' if and only if $\bar{v}_1 < 0$. We do so in the following.

As we have anticipated in the text, we will show that the agent's payoff from truthful reporting is no smaller than the truth-telling payoff under the optimal simple mechanism $(\tilde{q}, \tilde{t})$. At the same time the deviation payoffs are bounded above by the corresponding deviation payoffs under $(\tilde{q}, \tilde{t})$. Both individual rationality and incentive compatibility of the ambiguous mechanism Ω therefore follow from the individual rationality and incentive compatibility of $(\tilde{q}, \tilde{t})$.

We start with the truth-telling payoffs. Consider first an agent of type θ_n, $n < N$. Truthful reporting yields the payoff $\tilde{q}_n\theta_n - \tilde{t}_n$ if the simple mechanism is (q^n, t^n), and $q_n^m\theta_n - t_n^m = \theta_n - (1 - \tilde{q}_n)\theta_n - \tilde{t}_n = \tilde{q}_n\theta_n - \tilde{t}_n$ if the simple mechanism (q^m, t^m), $m \neq n$, is chosen. That is, for all the simple mechanism in Ω the agent's payoff from truthfully reporting his type is the same, and in particular, equal to $\tilde{q}_n\theta_n - \tilde{t}_n$.

For the highest type the payoff from truth-telling is no smaller $\theta_N - \tilde{t}_N$ which in turn cannot be smaller than the truth-telling payoff under $(\tilde{q}, \tilde{t})$, $\tilde{q}_N\theta_N - \tilde{t}_N$.

If the agent reports a type $\theta_l < \theta_n$, then his payoff is at most

$$\tilde{q}_l \theta_n - \tilde{t}_l,$$

which is the payoff that agent would get under the simple mechanisms (q^l, t^l). Since $(\tilde{q}, \tilde{t})$ is incentive compatible, it follows that the agent does not have any incentive to deviate downwards. By a perfectly analogous argument it follows that the agent has no profitable upward deviation either.

In order to complete the proof we need to argue that $Q = (1, \ldots, 1)$ not being a solution of P" is equivalent to $\bar{v}_1 < 0$. Proposition 3 implies that when $\bar{v}_1 > 0$, $Q = (1, \ldots, 1)$ is a solution to P". On the other hand, step 4 of the proof of Proposition 3 states that if $\bar{v}_1 \leq 0$, then at the optimum $\hat{q}_{m_j}^{m_j} = 0$ for all $j \leq j^*$. It is easy to see from the proof that when $\bar{v}_1 < 0$, $Q = (1, \ldots, 1)$ cannot be a solution to P". However, when $\bar{v}_1 = 0$, there can be other solutions beside the one in the statement of Proposition 3. In that case $\bar{R}^1(Q)$, as used in the proof, is constant when one varies q_1^1 and adjusts the other $q_{m_j}^{m_j}$ as in the said proof. Thus setting Q to $(1, \ldots, 1)$ does not affect $\bar{R}^1(Q)$ and provides an upper bound on the expected profit for the seller. In addition, for $Q = (1, \ldots, 1)$, $\bar{R}^{m_j}$ is constant in m_j, and in particular equal to θ_1. Thus the upper bound is attained, and $Q = (1, \ldots, 1)$ is also a solution. To conclude, $Q = (1, \ldots, 1)$ is not a solution to P" if and only if $\bar{v}_1 < 0$.

■

Proof of Proposition 6. Consider the relaxed versions of Problem P-2 from which (UIC) has been dropped. Suppose that Ω is a feasible mechanism of that problem and that $(\bar{q}, \bar{t})$ is one of the simple mechanisms in Ω that generates the smallest expected revenue, i.e. $(\bar{q}, \bar{t}) \in \arg\min_{(q,t) \in \Omega} \mathbb{E}_p[t(\theta)]$.

Let Ω' be the set of all outcome functions of the form (q, t'), where $(q, t) \in \Omega$ and the transfer rule t' coincides with t except for the transfers of the highest type, which are equal to

$$t'_N = t_N - \left[\mathbb{E}_p[t(\theta)] - \mathbb{E}_p[\bar{t}(\theta)]\right] / p_N.$$

Ω' inherits from Ω the properties (IR) and (DIC). This follows from the fact that in passing from Ω to Ω' only the highest type's transfers are lowered. Thus, both individual rationality and downward incentive compatibility can at most be relaxed. Moreover, by construction we have that

$$\mathbb{E}_p[t(\theta)] = \mathbb{E}_p[\bar{t}(\theta)] \qquad \text{for all } (q,t) \in \Omega'.$$

Thus, Ω' satisfies (C) and generates the same value as Ω. But that means that the value of the relaxed version of Problem P-2 does not change if the constraint (C') is replaced by (C). Doing so yields the relaxed version of Problem P. We already know that the latter has the same value as Problem P itself. ■

Proof of Proposition 7. 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. ■

References

Ashenfelter, Orley (1989), "How auctions work for wine and art." *The Journal of Economic Perspectives*, 23–36.

Auster, Sarah (2013), "Adverse selection under ambiguity." Working Paper.

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

Bajari, Patrick and Ali Hortacsu (2003), "The winner's curse, reserve prices, and endogenous entry: empirical insights from ebay auctions." *RAND Journal of Economics*, 329–355.

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.

Bergemann, Dirk and Johannes Horner (2010), "Should auctions be transparent?" Working Paper.

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 (2014), "Mechanism design with ambiguous communication devices." *Econometrica*. Forthcoming.

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.

Hendricks, Kenneth, Robert H Porter, and Richard H Spady (1989), "Random reservation prices and bidding behavior in ocs drainage auctions." *JL & Econ. S83*, 32.

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.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2012), "On the smooth ambiguity model: a reply." *Econometrica*, 80, 1303–1321.

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

Lang, Matthias and Achim Wambach (2013), "The fog of fraud - mitigating fraud by strategic ambiguity." *Games and Economic Behavior*, 81, 255–275.

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.

Saito, Kota (2013), "Preference for flexibility and preference for randomization under ambiguity." Working paper.

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.

Wolitzky, Alexander (2013), "Bilateral trading with maxmin agents." Working Paper.