What is behind ambiguity aversion?

Abstract

The results of an experiment extending Ellsberg's setup demonstrate that attitudes towards ambiguity and compound uncertainty are closely related. However, this association is much stronger when the second layer of uncertainty is subjective than when it is objective. Provided that the compound probabilities are simple enough, we find that most subjects, consisting of both students and policy makers, (1) reduce compound objective probabilities, (2) do *not* reduce compound subjective probabilities, and (3) are ambiguity non-neutral. By decomposing ambiguity into risk and model uncertainty, and jointly eliciting the attitudes individuals manifest towards these two types of uncertainty, we characterize individuals' degree of ambiguity aversion. Our data provides evidence of decreasing absolute ambiguity aversion and constant relative ambiguity aversion.

Keywords: Ambiguity aversion, model uncertainty, reduction of compound lotteries, non-expected utility, subjective probabilities, decreasing absolute ambiguity aversion

JEL Classification: D81

"Ambiguity may be high [...] when there are questions of reliability and relevance of information, and particularly where there is *conflicting* opinion and evidence."

Ellsberg (1961, p. 659)

1 Introduction

We report the results of an experiment aiming at disentangling the mechanisms behind ambiguity aversion. By distinguishing preferences for objective and subjective probabilities, and showing that individuals generally do not consider the two in the same way, we are able to detect a strong association between preferences for subjective probabilities and those for ambiguity, and bring our collective understanding on what originates ambiguity aversion a step forward.

Slightly more than fifty years ago, Ellsberg (1961) proposed a series of experiments whose results suggest that people prefer situations in which they perfectly know the probabilities of events occurrence (risk) to situations in which these probabilities are unknown (ambiguity). This seminal paper has given rise to a large body of literature in economics, exploring both the theoretical and experimental sides of the problem. Ambiguity aversion has since been subject to heated debates among scholars, questioning whether preferences emerging from observed behaviors in these experiments should be considered as a deviation from rationality, or instead, they could be seen as a rational way to cope with uncertainty, which could therefore claim a normative status. For instance, Halevy (2007) proposed a series of experiments extending Ellsberg's setup, from which he suggested that attitudes towards ambiguity and towards compound lotteries are tightly associated. The implication of this result, which has been replicated in several other experiments, may have an important impact on the way attitude to ambiguity is perceived and treated in economic models. In particular, if one sees the violation of independence in risky choices as a departure from rationality, and if subjects who are ambiguity averse are also less likely to reduce compound risk, then this weakens the potential for ambiguity aversion to claim a normative status. We test this association in a simple context of decision making to rule out other potential confounding factors, and measure whether ambiguity remains strongly associated with compound risk, or whether it is the nature of probabilities that plays a central role. In particular, if preferences over compound uncertainty with subjective probabilities correlate better with preferences over ambiguity, then it would shed new light on what are the intrinsic behavioral components driving ambiguity attitudes.

Ambiguity (also known as deep uncertainty or Knightian uncertainty) characterizes situations in which a decision maker does not have sufficient information to quantify through a single probability distribution the nature of the problem she is facing. It is distinct from the notion of *risk*, which refers to situations in which probabilities of a random event are perfectly known. Ambiguity is present in most real life situations involving uncertainty. It plays a major role in many economic problems and directly affects the process of decision making. A noteworthy example concerns the decision to mitigate greenhouse gas emissions in the presence of climate change (Berger et al., 2016).

Economics has traditionally treated those situations in which no objectively known or commonly agreed upon probabilities are given by following the approach proposed by the Subjective Expected Utility (SEU) theory (Savage, 1954). In line with the Bayesian tradition, the idea in this framework is that any source of uncertainty can be quantified in probabilistic terms and considered as a *risk.* The prior probabilities an individual has over the different states of the world are subjective, and the decision is made by maximizing the expected utility, given the individual's prior. Following the broad evidence –initiated by Ellsberg (1961)– that most individuals treat ambiguity differently than objective risk (and as a consequence the difficulty for Savage's axioms to be fulfilled), several lines of research have been followed and alternative decision models have been proposed. These models do not treat objective and subjective probabilities in the same way. For example, Gilboa and Schmeidler (1989) gave a behavioral foundation to the influential maxmin expected utility (MEU) model, in which an individual's utility is given by the minimal expected utility over a set of multiple priors that the individual might have. Ghirardato et al. (2004) then axiomatized an extension of this model by considering combinations of minimal and maximal expected utilities over the individual's set of priors. Other models such as the ones proposed by Klibanoff et al. (2005); Nau (2006); Ergin and Gul (2009); Seo (2009) assume that individuals have both a set of first order priors and a second order probability over them, and are expected utility maximizers over both the first and the second layer of uncertainty (two-stage models).

Given the pervasiveness of ambiguity in all fields of economics and its relevance in the process of real-life decision making, we feel that it is crucial (1) to better understand what drives the observed behavior of ambiguity aversion; and (2) to quantify more precisely the extent to which ambiguity aversion exists and characterize its main properties.[1] The aim of this paper is to address this twofold objective. To do so, we follow the decomposition of uncertainty into two distinct layers of analysis, as proposed by Hansen (2014) and Marinacci (2015) building on Arrow's (1951) work. The first layer, commonly referred to as *aleatory uncertainty,* features the probability measure associated with the randomness of an event. It refers to the physical quantification of uncertainty by means of a probabilistic model, and refers to a notion –central in classical statistics– which is generally called *risk* in economics. The second layer, referred to as *epistemic uncertainty*, characterizes the fact that the correct probabilistic model (or risk) is itself unknown.[2] Rather multiple models may exist, each of them associated with a subjective probability representing the decision maker's (DM) degree of belief in that particular model. These subjective probabilities quantify what we will be referring to as *model uncertainty.* We feel this distinction can help us analyze in practice the vast majority of decision problems under uncertainty. Through the lens of this two layer model, ambiguity can be interpreted as the combination of risk and model uncertainty.

In a controlled experimental environment which extends both Ellsberg's (1961) and Halevy's

[1]While the vast majority of the experimental effort has been devoted to finding evidence of ambiguity aversion (Trautmann and van de Kuilen, 2014), few studies have attempted to actually quantify the strength of this effect (notable exceptions, using various models, are Abdellaoui et al. (2011); Dimmock et al. (2015)).

[2]The term *"epistemic"* derives from the Ancient Greek ἐπιστήμη, which means "knowledge", while the term *"aleatory",* which originates from the Latin *alea*, refers to any game of chance involving dice.

(2007) setups, we confront our subjects with four different types of uncertain situations, represented by urns filled with balls that may be either red or black. The two extreme situations are the standard Ellsberg ones, in which the number of red and black balls –and therefore probabilities– are either objectively known (risk) or unknown (ambiguity). In the spirit of Halevy (2007), a fair coin determines the number of red and black balls in the third urn, which therefore presents objective probabilities in two distinct layers (compound risk). Finally in the last situation, perfect ignorance is achieved by letting both the number and the color of balls be unknown before two "experts" provide their assessment of the urn composition. This latter situation also presents uncertainty in two layers, but probabilities in the second layer are subjectively determined (model uncertainty). In particular, model uncertainty is achieved by providing the subjects with two possible models (represented by the two experts' assessments[3]). We focus our research on the distinction between objective and subjective probabilities in a context characterized by two layers of uncertainty. Building on setups used by Holt and Laury (2002) and Andersen et al. (2008), we propose a design that enables us to test the association between ambiguity non-neutrality and, respectively, violation of reduction of compound lottery (ROCL, when probabilities are objective) and violation of reduction of compound uncertainty (ROCU, when probabilities are subjective). We run our within-subjects experiment on both a panel of university students and a panel of policy makers. The first experiment took place in a laboratory at Bocconi University, and we carried out the second as an artefactual field experiment during the 2015 United Nations Climate Change Conference, COP 21.

There are three main findings emerging from our analysis. First, attitudes towards ambiguity and uncertain situations presented in two layers are closely related. However, this association is much stronger when the second layer of uncertainty is subjective than when it is objective. Provided that the compound probabilities are simple enough, most of our subjects behave according to the ROCL and violate the ROCU. Second, subjects tend to be both risk and model uncertainty averse, and furthermore exhibit stronger aversion towards model uncertainty than towards risk. This behavioral characteristic is interpreted as evidence of ambiguity aversion, which is elicited via a joint estimation procedure. Third, and analogously to what has been previously reported for risk aversion (Holt and Laury, 2002, 2005), we find that model uncertainty aversion is decreasing when considered in absolute terms, and increasing when considered in relative terms. In terms of ambiguity attitude, we find evidence of decreasing absolute ambiguity aversion (DAAA) and constant relative ambiguity aversion (CRAA).

2 Experimental procedures

The experiment consists of a sequence of nine tasks, divided in two sets each employing a different elicitation procedure. In the first part, subjects are confronted with a random lottery pair (RLP) procedure, in which they face pairs of uncertain alternatives and are asked to pick one of two. This set of choices enable us to test the predictions of expected utility theory (EUT)

[3] We refer to experts as those individuals or entities who presumably have more information and/or expertise than the DM, and who are acting as advisors by providing information (Budescu et al., 2003).

and detect potential deviations from it in situations of uncertainty. In the second part of the experiment, we use a double price list (PL) procedure to jointly elicit risk and model uncertainty attitudes. In this part, each subject is confronted with a series of binary choices, presented in the form of ordered tables, as popularized by Holt and Laury (2002).

2.1 The choice situations

Our subjects may be confronted with four different uncertain situations across the various tasks. These situations are represented by urns containing balls that can either be red or black. Each urn describes a particular type of uncertainty. The urns are characterized as follows:

- *Urn 1 (risk)*: the number of red and black balls is perfectly known;
- *Urn 2 (compound risk)*: the number of red and black balls is determined by flipping a fair coin in the air;
- *Urn 3 (model uncertainty)*: the number of red, black and the total number of balls in the urn are unknown (to reflect complete ignorance), but information is provided by two "experts", each giving her own assessment of the composition of the urn;
- *Urn 4 (Ambiguity à la Ellsberg)*: the total number of balls in the urn is given, but the exact composition of the urn is unknown.

In the RLP task for example, the urn compositions (red balls, black balls) are as follows: Urn 1 is (50,50); Urn 2 is either (100,0) or (0,100) (flipping a fair coin determines which of the two); Urn 3's composition is unknown but Expert 1's assessment is that there are only red balls, while Expert 2's assessment is that there are only black balls; Urn 4 is composed of any possible combination of red and black balls between 0 and 100. In the first two urns, the probability of drawing a red ball $P(r)$ is objectively known to be 1/2, the only difference between the two being that Urn 1 corresponds to a simple risk situation, while Urn 2 is presented as a compound risk. In Urns 3 and 4 the probabilities are unknown. However, subjects are still given some information taking the form of the experts' beliefs about the urn's composition (in Urn 3), or the total number of balls (in Urn 4). As is the case in Urn 2, the information is presented in two layers, with the difference being that the probabilities associated with the different compositions of the urn are not objective. These urns are presented two by two in a randomized sequence.[4] In each decision, subjects are required to place a bet on the color of the ball drawn from each urn (Red or Black), and to decide on which urn to place their bet (allowing for indifference). The bet may win the subject €15 and entails no losses otherwise.

In the second part of the experiment, we specifically focus on urns of type 1 and 3. Subjects are confronted with three risky and five model uncertainty tasks. The first of the risky tasks takes the common form of a certainty equivalent (CE) task in which subjects are asked to choose between a binary lottery and a sure amount of money. Specifically, by letting O denote the set of monetary outcomes, and $\bar{o}_p\underline{o}$ the binary lottery yielding $\bar{o} \geq \underline{o} \in O$ with probability p and $\underline{o} \in O$

[4] In order to replicate results previously obtained in the literature while introducing the model uncertainty framework, the risky Urn 1 is kept as a reference and systematically paired with the other three urns.

otherwise, subjects are asked to make a series of ten choices between $\bar{o}_p\underline{o}$ and different values of $o \in O$ ordered from $\bar{o}$ to $\underline{o}$.[5] This task allows us to characterize the certainty equivalent, which is defined as the payoff that would leave the subject indifferent between the sure amount and the lottery. The design of the other two risky tasks is also standard and follows Holt and Laury's (2002) PL procedure. The first model uncertainty task is analogous to the certainty equivalent task. In this case however, the binary lottery $\bar{o}_p\underline{o}$ is replaced by the uncertain situation denoted $\bar{o}_{\hat{p}_1\hat{p}_2}\underline{o}$, in which the subject is only given information on two experts' assessed probabilities $\hat{p}_1$ and $\hat{p}_2$. In the four remaining uncertain tasks, subjects make a series of choices between risky situations $\bar{o}_p\underline{o}$ and situations of model uncertainty $\bar{o}_{\hat{p}_1\hat{p}_2}\underline{o}$.

The PL procedure is one of the most commonly employed elicitation methods to represent choices between gambles (Andersen et al., 2006). It is widely considered as a transparent procedure that rarely confuses subjects about the incentives to respond truthfully (Harrison and Rutström, 2008). However, one of the main disadvantages of this method is that subjects typically have the possibility to switch freely between the two options as they progress down the decision tables. They may therefore make inconsistent choices either by switching more than once, or by making reverse choices (Charness et al., 2013). While we recognize these inconsistent behaviors raise additional difficulties –given that they are difficult to rationalize under standard assumptions on preferences, and that the estimation technique and inference of risk and model uncertainty attitudes require a unique switching point– we decided not to enforce consistent choices in this experiment.[6] Rather, we view such behavior as indicative of failing to understand the instructions correctly, or of confusion on the part of the subjects, and discard this inconsistent data from the analysis of our results.

We aimed at a design that, at the same time, remains extremely simple and yet emphasizes the difference between compound risk and model uncertainty. As in any experiment, some features had to be left out for pragmatic purposes. The design may therefore present imperfections. First it should be clear that, while two expert's assessments were provided in Urn 3, the experts were not physicaly present in the room. Although this situation is close to many real life situations in which expert opinions are presented without the physical presence of the experts themselves, this feature may lead subjects to downplay experts' role and to (partially) ignore them, making the situation closer to that of full ignorance. To minimize the potential bias, we gave particular attention to the way we presented the experts and their assessments.[7] We are confident in the effectiveness of our efforts as the information provided via the experts is clearly reflected in the subjects' responses. Indeed, using the results obtained in the various PL tasks, we can show that subjects' choices monotonically follow the information provided by the two experts, suggesting that our subjects incorporated the information provided. Second, the first layer of uncertainty in some of the tasks (for example the RLP task) is structured so that

[5] In the CE tasks of our experiment we used the following values: $\bar{o} = 25$, $\underline{o} = 4$, $p = 0.5$, and $o \in \{25, 18, 15, 14, 13, 12, 10, 8, 6, 4\}$.

[6] Several techniques have been proposed in the literature to enforce consistency in the subjects' choices (see for example Andersen et al. (2006)), but with the major drawback that they may significantly bias the results (Charness et al., 2013).

[7] For that purpose, we specifically mentioned the following in the instructions: "These experts are the best we could find for this situation. They are both experienced and both have excellent track records".

one of the two events is associated with a 100% probability. While one could argue that such a situation is too extreme to be realistic or that the compound lottery –being degenerate– can no longer be considered compound, we argue that the potential bias introduced is kept minimal. We design the RLP task with first layer probabilities being either 0 or 100% exactly because we want the problem to be computationally simple. The objective is to isolate the effect of model uncertainty alone and to minimize the potential role played by cognitive skills in the reduction of compound lotteries. While risk has *stricto sensu* disappeared in the second stage, we argue that subjects still need to make *some* computational effort to find out the final probability of winning the prize. This is reinforced by the fact that our subjects are given the choice of color on which to bet, which requires them to consider the different possibilities depending on the outcome of the coin toss. We are not the first to consider compound risk with a degenerate second stage. For example, Halevy (2007) also uses this specific form to test for reduction of compound lotteries and finds differences in the way subjects value it relative to simple risk, rejecting therefore any systematic reduction. To make the "model uncertainty" situation fully comparable with the objective compound risk, experts in the RLP task are presented as being dogmatic (in the sense that they both assess a 100% probability to one particular event). This allows us to isolate directly the impact of model uncertainty aversion from risk aversion (see below). Overall, any concern associated with the incompleteness of the extremely simplified setup should then equally influence both the subjective and the objective compound situations.

2.2 The randomness device

Since one of our main goals in the experiment is to characterize the way individuals behave in the presence of model uncertainty –i.e. in situations where the only source of information is the one given by the experts– we need to make sure that, in the absence of experts' information, subjects are indeed in a situation of perfect ignorance.[8] This is different from the canonical Ellsberg examples in which (objective) information is given concerning the total number of balls, thus enabling the decision maker to posit a restricted set of possible objective models M.[9] To mimic the situation of perfect ignorance (and compel subjects to consider all the possible probabilities in $[0, 1]$), we construct Urn 3 in such a way that the total number of balls in the urn is itself unknown, and comprised between 1 and 100. We call this modification of Ellsberg's canonical experiment, that reduces the information bias due to the peculiarity of the urn representation, the *randomness device*. In such a situation, the total number of potential objective models is 3045, which is the cardinality of the Farey sequence of order 100.[10] To see this, consider Table 1 below. It presents the sets of potential models when the *maximum* number of balls N is known to be between 1 (first row) and 8 (last row). As can be seen, when the maximum number of balls is $N = 1$, the set of models M consists of two elements:

[8]In other words, they should a priori consider the continuum of probabilities between 0 and 1. In the context of our experiment, this situation would correspond to an urn containing an infinite number of balls.

[9]In particular, $M = \{P(r) \in \{\frac{0}{100}, \frac{1}{100}, \dots, \frac{100}{100}\}\}$ in Ellsberg's two urn example. This makes the standard Ellsberg case itself a situation of model uncertainty.

[10]A Farey sequence of order N, denoted $\mathfrak{F}_N$, is the ascending series of irreducible fractions between 0 and 1 whose denominators do not exceed N (Hardy et al., 1979).

Table 1: Sets of models and their corresponding cardinality when the maximum number of balls in the urn is N

N	Set of possible models: $M_N = \{P(r)\}$																							$\lvert M_N \rvert$
1	$\frac{0}{1}$																						$\frac{1}{1}$	2
2	$\frac{0}{1}$											$\frac{1}{2}$											$\frac{1}{1}$	3
3	$\frac{0}{1}$							$\frac{1}{3}$				$\frac{1}{2}$				$\frac{2}{3}$							$\frac{1}{1}$	5
4	$\frac{0}{1}$					$\frac{1}{4}$		$\frac{1}{3}$				$\frac{1}{2}$				$\frac{2}{3}$		$\frac{3}{4}$					$\frac{1}{1}$	7
5	$\frac{0}{1}$				$\frac{1}{5}$	$\frac{1}{4}$		$\frac{1}{3}$		$\frac{2}{5}$		$\frac{1}{2}$		$\frac{3}{5}$		$\frac{2}{3}$		$\frac{3}{4}$	$\frac{4}{5}$				$\frac{1}{1}$	11
6	$\frac{0}{1}$			$\frac{1}{6}$	$\frac{1}{5}$	$\frac{1}{4}$		$\frac{1}{3}$		$\frac{2}{5}$		$\frac{1}{2}$		$\frac{3}{5}$		$\frac{2}{3}$		$\frac{3}{4}$	$\frac{4}{5}$	$\frac{5}{6}$			$\frac{1}{1}$	13
7	$\frac{0}{1}$		$\frac{1}{7}$	$\frac{1}{6}$	$\frac{1}{5}$	$\frac{1}{4}$	$\frac{2}{7}$	$\frac{1}{3}$		$\frac{2}{5}$	$\frac{3}{7}$	$\frac{1}{2}$	$\frac{4}{7}$	$\frac{3}{5}$		$\frac{2}{3}$	$\frac{5}{7}$	$\frac{3}{4}$	$\frac{4}{5}$	$\frac{5}{6}$	$\frac{6}{7}$		$\frac{1}{1}$	19
8	$\frac{0}{1}$	$\frac{1}{8}$	$\frac{1}{7}$	$\frac{1}{6}$	$\frac{1}{5}$	$\frac{1}{4}$	$\frac{2}{7}$	$\frac{1}{3}$	$\frac{3}{8}$	$\frac{2}{5}$	$\frac{3}{7}$	$\frac{1}{2}$	$\frac{4}{7}$	$\frac{3}{5}$	$\frac{5}{8}$	$\frac{2}{3}$	$\frac{5}{7}$	$\frac{3}{4}$	$\frac{4}{5}$	$\frac{5}{6}$	$\frac{6}{7}$	$\frac{7}{8}$	$\frac{1}{1}$	23

$M_1 = \{P(r) \in \{0, 1\}\}$, where $P(r)$ denotes the probability of drawing a red ball. When $N = 2$, the cardinality of M increases to $|M| = 3$, such that $M_2 = \{P(r) \in \{0, \frac{1}{2}, 1\}\}$. We view this device sufficiently complicated to prevent subjects from doing any calculation of probability distribution over the possible compositions of the urn and their corresponding weights.[11] In that sense, absent of any additional information from the experts, subjects will –most likely– be unable to compute the set of possible objective models, and end up in a situation close to one of perfect ignorance :

$$M_{100} = \{P(r) \in \{\mathfrak{F}_{100}\}\} \sim \{P(r) \in [0, 1]\}. \tag{1}$$

We feel that such a setup, that emphasizes Frank Knight's (1921) original distinction between "measurable" and "unmeasurable" uncertainty (which cannot be represented by numerical probabilities) better reflects the actual state of individuals facing complex problems. Such is the case in a large fraction of modern science problems for example, where the level of abstraction and mathematical requirement to understand processes are such that individuals cannot have a mental construct of the problem they are facing.

2.3 Recruitment and administration

The laboratory experiment took place in Bocconi University (Italy) in April 2015. 189 subjects were recruited through an internal experimental economics recruitment system. Each subject was authorized to participate only once and had to sign up in advance for a particular time slot. The experiment was organized into 12 sessions taking place over four days. Each session lasted approximately 75 minutes, and comprised of 13 to 19 subjects. Subjects were provided with paper, pen and a calculator. A session typically started with silent reading of

[11] Note that even in the case in which all the possible total numbers of balls, and all their possible compositions are assumed to be equally probable, the possible models are not weighted uniformly. E.g., for a total number of balls comprised between 1 and 3, five different models exist: $M_3 = \{P(r) \in \{0, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, 1\}\}$. Assuming the number of balls in the urn is uniformly distributed between 1 and 3, and that for each case the different models are weighted equally, we end up with weights attached to the possible models that are respectively $q = \{\frac{13}{36}, \frac{1}{12}, \frac{1}{9}, \frac{1}{12}, \frac{13}{36}\}$.

general instructions which were printed and provided to each subject in the cubicle to which she/he was assigned. The experimenter then read once more the instructions aloud and made sure everything was clear, before the subjects started a computerized training session that introduced them to the concepts of risky and uncertain urns, and decision tables. The experiment was then performed on computers, with the order of tasks being randomized. Overall, the nine tasks constituting our experiment were associated with a random incentive system to determine the final payoff. Once all subjects had answered all the questions, they were asked to fill in a short socio-economic questionnaire before being told their payoffs (i.e. which of their decisions had been randomly selected, what was the color of the ball drawn from the urn they chose (if any), and what was the corresponding amount they won). Subjects were then paid in cash a 5 euro participation fee, and the additional amount (up to €35) won on the basis of the choices they made. The average gain was about €18.50 per subject. The lab experiment was programmed and conducted with the experiment software z-Tree (Fischbacher, 2007). Details of the experimental procedure, instructions and demographic data are provided in the Supplemental Material available online. To confirm the results of the lab experiment, we also conducted an artefactual field experiment (Harrison and List, 2004) at the 2015 United Nations Climate Change Conference, held in Paris in December, 2015. The 91 subjects who participated in this robustness round were policy makers and climate negotiators originating from 52 different countries. In individual in-person interviews, we confronted respondents who volunteered for the study with the RLP task (we did not include the PL task as time was a binding constraint in the field). Other procedural differences are that the experiment was conducted with pen and paper, the payoff reached €50 if the bet was correct, and subjects did not have access to calculators.

3 Theoretical predictions and literature review

In this section, we present our theoretical predictions. Since the main part of our experiment enables us to measure attitudes directly from behavior, we start by expressing the two first predictions from a very general point of view (i.e. abstracting from any specific model of choice under uncertainty). In the second part, we add some structure by using a specific model of choice which has been proposed by Marinacci (2015) building on Klibanoff et al. (2005). This model, known as the *smooth model*, is fairly general. It enriches and encompasses many of the recent theories of choice under uncertainty in cases where Waldean information is incorporated in the decision problem, while allowing for a distinction between objective and subjective probabilities.

3.1 The setting

The decision maker in our setting evaluates acts (or bets) f whose outcome depends on the realization of an observable state. In the experiment a state is the quartet $\{c_1, c_2, c_3, c_4\}$, where $c_i \in \{R, B\}$ $\forall i = 1, \ldots, 4$ is the color of the ball (red R, or black B) drawn from Urn i. The state space $S = \{R, B\}^4$ is therefore made up of 16 states, but there are only two events $c \in E$ of importance for each bet on a specific Urn i: either a red ball is drawn ($c = r_i$) or a black ball is drawn ($c = b_i$). In this context, each ball draw may be seen as the realization of a random variable that can be described by a specific objective *model.*[12] The uncertainty about the outcome of a given model is of the aleatory type and generally called *risk.* This risk is directly relevant to the DM since it determines the probability with which each event realizes. Probabilities of the different events can, in this case, be defined as *objective* (they refer to a physical concept, represented by a specific composition of the urn). As is the case in the vast majority of decision problems, it may however happen that the DM does not know exactly which probability model generates the observations. In such a situation, a second layer of uncertainty adds onto the first layer of risk. This second layer of uncertainty, which concerns the possible compositions of the urn, may have different natures. It may be a second layer of risk, in which case the uncertain situation is simply an instance of compound risk. Or it may be characterized by *epistemic* uncertainty, if multiple compositions of the urn are possible, but the DM does not know how likely each of them is. The probabilities in this second layer of uncertainty represent the DM's degree of belief in each potential model, and are then *subjective.*[13] This situation is referred to as *model uncertainty.* As in Cerreia-Vioglio et al. (2013b) and Marinacci (2015), it is assumed that the DM knows the possible alternative models belong to a subset M of Δ, the collection of all probability measures on the state space S. In our case, this is the information given to subjects that allows them to posit this subset. Elements of M are seen as possible compositions of the urn that are consistent with the available information, and that could hence be selected by nature to generate observations. In accordance with Wald (1950), the set M is assumed to be taken as a datum of the decision problem.[14]

3.2 General predictions

Hypothesis 1. We expect subjects to be both risk and model uncertainty averse, in the sense that they generally prefer the degenerate lottery, giving $\sum_{c\in\{r,b\}} \bar{P}(c)\, f(c)$ with certainty, to any uncertain situation in which an act f yields $f(c)$ with (expected) probability $\bar{P}(c)$ $\forall c \in \{r, b\}$. By letting Ci denote the certainty equivalent for Urn i, and C_0 the sure amount corresponding to the expected gain of the uncertain bet, we expect to observe:

$$C_0 \geq C1, \tag{2}$$

$$C_0 \geq C3. \tag{3}$$

While the first result is trivial (Holt and Laury, 2002, 2005; Andersen et al., 2008), it is necessary to study what we are ultimately testing, which is whether ambiguity aversion is related to a stronger aversion towards model uncertainty than towards risk.

[12]The term *"model"* here refers to a probability model (or distribution). In our experiment, a model corresponds to a possible composition of the urn.

[13]We here follow the definition of Schmeidler (1989) who interprets subjective probabilities of an event as the number used in calculating the expectation of a random variable. Remark that this definition includes objective probabilities as a special case where we know exactly which number to use.

[14]Cerreia-Vioglio et al. (2013a) call the models in M "objectively rational beliefs". It is analogous to what Ellsberg (1961) calls "reasonable" distributions in his subjective setting. Note that in general incompleteness of information makes the set M non-singleton, contrary to what is assumed in the standard subjective expected utility theory. The true model is assumed to belong to M, abstracting therefore from misspecification issues.

Hypothesis 2. Following Ellsberg's (1961) seminal results and the subsequent experimental literature on ambiguity aversion (see Trautmann and van de Kuilen (2014) for a survey), we predict our subjects to generally prefer to be confronted with risk (Urn 1) rather than with ambiguity (Urn 4). Considering the decomposition of ambiguity into model uncertainty and risk, we also expect this behavior to be related to the fact that they prefer risk (Urn 1) to model uncertainty (Urn 3). Moreover, we predict that when the compound risk is sufficiently easy to reduce, subjects are indifferent between a simple risk (Urn 1) and a compound one (Urn 2). Finally, we predict that the degrees of model uncertainty aversion and of ambiguity aversion are finite. In other words, we expect subjects not to behave according to a Waldean maxmin criterion. In terms of certainty equivalents, these predictions may be written (under the assumption of equal expected values) as:

$$C1 = C2 \geq C3 > \underline{C}, \tag{4}$$

$$C1 \geq C4 > \underline{C}. \tag{5}$$

where $\underline{C}$ corresponds to the certainty equivalent obtained under the worst possible model. Our central hypothesis is the association between expressions (4) and (5). The first equality of statement (4) says that different layers of objective sources of uncertainty are reducible : people are indifferent between risk and compound risk when the expected values of the lotteries are identical. This rational behavior of subjects has however been seriously challenged in the literature. An early example dates back to Yates and Zukowski (1976), while more recent ones comprise of Chow and Sarin (2002); Halevy (2007). In particular, Halevy (2007) reports the results of an experiment suggesting that people are generally compound risk averse, and that attitudes towards compound risk and ambiguity are tightly associated. However, he also shows that people, on average, prefer compound risk situations to ambiguous ones. Qualitatively similar results were obtained by Dean and Ortoleva (2015) and by Armantier and Treich (2015), who show that not only was attitude to compound risk tightly associated to attitude towards ambiguity, but so was attitude towards complex risk.[15] Abdellaoui et al. (2015) also find, in a setup close to Halevy's, an association between compound risk reduction and ambiguity neutrality. The association they find is however weaker than in Halevy's data. In particular, these authors show that compound and simple risks are valued differently, but also find pronounced differences between compound risk and ambiguity attitudes. Interestingly, they show that, for mathematically more sophisticated subjects, compound risk reduction is compatible with ambiguity non-neutrality, suggesting that failure to reduce compound risk and ambiguity non-neutrality do not necessarily share the same behavioral grounds. Relatedly, Prokosheva (2016) obtains a significant relationship between arithmetic test scores and compound risk reduction in an experiment with adolescents subjects (interestingly no such relationship is found between ambiguity neutrality and these test scores). Finally, in a recent study Harrison et al. (2015) specifically test the reduction of

[15] A complex risk in Armantier and Treich's (2015) design refers to a situation in which the probabilities associated with the different events are non-trivial to compute. For example if a pair of colored balls draw simultaneously from two transparent urns (one from each urn).

compound lotteries with objective probabilities both in a setup with multiple choices associated with a random incentive system, and in a setup with a unique choice. They find evidence of violation of reduction of compound lotteries in the first case, but not in the second.

While evidence is far from definitive, overall the existing literature suggests that subjects in general manifest aversion towards compound risks. This behavior presents similarities to that reported concerning ambiguity aversion. The tight association between ambiguity and compound lottery attitude lends itself to two possible hypotheses. First, ambiguous situations may be perceived by individuals as compound risks and individuals fail to reduce them (in violation of the reduction –or independence– axiom). Alternatively, compound risks may be perceived as ambiguous situations to which individuals feel an aversion (in violation with the independence axiom). Although we do not explicitly test any specific theory that might explain why compound risk may be associated with ambiguity, we try to shed light on this issue through an experimental setup where the compound lottery is extremely simple. If cognitive inability is at the basis of failures to reduce compound probabilities (Abdellaoui et al., 2015; Harrison et al., 2015; Prokosheva, 2016), and aversion to compound lottery reflects a deficiency of the 'human intuitive statistician' (Budescu and Fischer, 2001), then by designing a compound risk situation that is very easily reducible, we partly rule out instances based on limited cognitive ability, and we expect subjects to effectively reduce compound risk if the probabilities of the two layers of uncertainty are objectively given.

The situation is different when the probability assessments are described by experts. In this case the second layer of uncertainty is no longer objective, and the two situations are expected to be evaluated differently. In particular, we expect to observe $C2 \geq C3$ when the only difference between the uncertain situations is whether the probabilities of being confronted with a given risk are subjectively determined or given by a known random device. In the extreme case with dogmatic experts (e.g. in RLP and CE tasks), the situation is analogous to Schmeidler's (1989) two coins examples.[16] While a Savagian expected utility maximizer would be indifferent between the two uncertain situations, we predict that most subjects in our experiment will not evaluate the two distributions in the same way. In the case of the risky urn (or fair coin), the distribution is based on objective information that supports a symmetric assessment while in the case of conflicting dogmatic experts (or unknown coin), the same estimates are subjective and rely on symmetry in the absence of information. From the DM's point of view, this distinction is essential and we expect to observe a majority of subjects opting for the risk rather than the model uncertainty situation, revealing in this way higher aversion towards subjective uncertainty. By transitivity, we also expect any uncertain situation being proposed with given probabilities to be preferred to a comparable situation with similar expected value and expected probabilities, but where these probabilities come from different expert's assessments (i.e. $C1 \geq C3$).

[16]In this example a subject is given the choice between betting on the result of a known fair coin coming up *Heads* or *Tails*, and a coin that has never been tested and is absolutely unknown. In this case, since no information is available about the probability of each side coming up, it is a symmetry argument which suggests the probabilities 50% to be considered. Notice that the Bayesian approach does not permit any distinction between the 50%-50% distribution based on information (from experts for example) and the one based on lack of information. See Gilboa et al. (2012) for a discussion on this subject.

3.3 Predictions based on a specific model of choice under uncertainty

The last two hypotheses we are testing are based on the smooth representation of preferences under model uncertainty (Marinacci, 2015). Under this representation, model uncertainty is addressed by considering a single prior probability measure over the set M of models. Specifically, it is assumed that the different possible models are indexed by a parameter θ and fully characterized by $P_\theta(r)$, the probability of drawing a red ball conditional on θ. This probability is uncertain and takes value $P_\theta(r)$ with probability q_θ, reflecting the DM's belief about the accuracy of the particular model θ. Under the smooth model uncertainty framework, the decision maker chooses the act that maximizes her utility given by:

$$U(f_i) = \mathrm{E}_\theta(v \circ u^{-1}) \left(\sum_{c \in \{r_i, b_i\}} \tilde{P}_\theta(c) u\left(f_i(c)\right) \right), \tag{6}$$

where E_θ is the expectation operator taken over the prior distribution indexed by θ,[17] u is the standard von Neumann and Morgenstern (1944) utility index capturing the DM's attitude towards risk (i.e. over objective probabilities), and v captures the attitude towards model uncertainty (i.e. over subjective probabilities). These functions are assumed to be strictly increasing and continuous. They are both cardinally unique. The sum within brackets is nothing but the expected utility of an act for a given objective model. It is then expressed in monetary terms by considering the certainty equivalent for each model P_θ. Since P_θ is itself uncertain, the different certainty equivalents are then evaluated by considering the expected utility using function v. In the case where both attitudes towards the different types of uncertainty are identical (whenever v is equal, up to an affine transformation, to u), we recover the classical subjective expected utility model of Cerreia-Vioglio et al. (2013b). This model therefore encompasses both the Savagian subjective expected utility[18] and the classical von Neumann-Morgernstern (when M is a singleton) representations. When attitudes towards risk and model uncertainty are different, representation (6) corresponds to Marinacci's (2015) setup, which consists of an enriched version of Klibanoff et al. (2005) in the presence of Waldean information (the smooth ambiguity function is recovered by setting $\phi \equiv v \circ u^{-1}$). In that sense, the DM is ambiguity averse if she is more averse to model uncertainty than to risk. In Appendix B, we characterize the different urns considered in the RLP task under this specific model of choice under uncertainty.

The last two hypotheses concern particular properties of risk and model uncertainty aversion (functions u and v) that are investigated in the second part of our experiment.

Hypothesis 3. Analogously to what is widely accepted in the risk theory literature and given the similarity of our procedure with the one used by Holt and Laury (2002, 2005), we expect to observe decreasing absolute risk aversion (DARA) and increasing relative risk aversion (IRRA)

[17] In particular, $\mathrm{E}_\theta \tilde{X}_\theta = \sum_\theta q_\theta X_\theta$ if the set M is finite, and $\mathrm{E}_\theta \tilde{X}_\theta = \int_M X(\theta) dq(\theta)$ if M is infinite.

[18] Remark that for each prior distribution q, there exists a distribution $\bar{P}(c)$ such that $U(f_i) = \sum_c \bar{P}(c) u(f_i(c))$, as in the original Savagian SEU representation.

for both functions u and v.[19] Specifically, by changing the values of the gains proposed and the probabilities that are associated to these gains, we expect to observe:

$$\frac{\partial}{\partial w_0}\left[-\frac{u''(w_0)}{u'(w_0)}\right] \leq 0 \quad \text{and} \quad \frac{\partial}{\partial w_0}\left[-\frac{u''(w_0)}{u'(w_0)}w_0\right] \geq 0, \tag{7}$$

where w_0 denotes the individual's wealth level, which is composed of the individual's background wealth ω, and the expected gain in each lottery $\bar{o}_p\underline{o}$: $w_0 = \omega + p\bar{o} + (1-p)\underline{o}$. Similarly, we are interested in the DARA and IRRA properties of function v, and therefore test whether we observe:

$$\frac{\partial}{\partial w_1}\left[-\frac{v''(w_1)}{v'(w_1)}\right] \leq 0 \quad \text{and} \quad \frac{\partial}{\partial w_1}\left[-\frac{v''(w_1)}{v'(w_1)}w_1\right] \geq 0, \tag{8}$$

where the individual's wealth level w_1, in situations of model uncertainty, is an average of certainty equivalent wealth levels under the two expert's models.

Hypothesis 4. Since ambiguity aversion in this setup results from the combination of attitudes towards both risk and model uncertainty, we should be able to indirectly characterize the properties of the ambiguity function. We are thus interested in knowing whether the absolute ambiguity aversion is constant or whether it is increasing or decreasing, in the sense that agents are willing to pay more or less to remove all source of uncertainty as their level of expected utility increases. Constant absolute ambiguity aversion (CAAA), as argued by Grant and Polak (2013), is an implicit characteristic of many of the ambiguity models proposed in the theoretical literature. It is for example implicitly assumed in the models by Gilboa and Schmeidler (1989); Hansen and Sargent (2001); Maccheroni et al. (2006). On the contrary, decreasing absolute ambiguity aversion (DAAA) is a condition that has been shown to play an important role in the determination of the precautionary saving motive under ambiguity (Gierlinger and Gollier, 2008; Berger, 2014), in the chances of survival of ambiguity averse investors (Guerdjikova and Sciubba, 2015), or in the choice of optimal abatement policies under scientific model uncertainty (Berger et al., 2016). The general representation (6) we describe in this section does not assume any particular type of preference functional. As previously mentioned, it could be seen as an enriched version of the model developed by Klibanoff et al. (2005). The domain of the ambiguity function $\phi = v \circ u^{-1}$ is however different than that of u and v, which are defined over monetary sets. Indeed, ϕ takes arguments that belong to a set of expected utilities. Considering this difference, we are interested in knowing whether the ambiguity function exhibits CAAA or instead if, as is the case for risk aversion, individuals are less ambiguity averse at higher levels of expected utility and therefore exhibit DAAA. Formally, what we are testing is the sign of $\frac{\partial}{\partial U}\left[-\frac{\phi''(U)}{\phi'(U)}\right]$, where U is the individual's expected utility level when the probabilities given by experts are averaged.[20]

[19] To be completely precise, we should talk about "decreasing absolute model uncertainty aversion" and "increasing relative model uncertainty aversion" in the case of function v, but for the sake of simplicity we prefer to refer to the widely used acronyms DARA and IRRA for the v function as well. While the DARA property seems well accepted in the literature, note however that the IRRA property is subject to debate when investigated outside of the lab environment (Harrison et al., 2007; Brunnermeier and Nagel, 2008; Chiappori and Paiella, 2011).

[20] Note that in a recent contribution, Baillon and Placido (2016) also tested the CAAA and DAAA hypotheses using a framework different from ours and found evidence of DAAA under Ellsberg's type of uncertainty.

4 General results

Before reporting the results obtained from the direct comparisons between the urns (RLP task) and from the certainty equivalent tasks (CE task), we discuss two important preliminary remarks. The first is related to the fact that the experts in these tasks are dogmatic. Given that the design of these tasks (unlike the PL task) does not allow us to trace subjects' attention to and incorporation of experts' information, we cannot guarantee that subjects did not have any additional and alternative models in mind when making decisions (i.e. misspecification issues). The presence of dogmatic experts ensures that the results we obtained concerning model uncertainty aversion are at worst underestimated. Indeed, if subjects were to consider any other probabilities between 0% and 100% as a possibility, the preferences associated with the observed choices would have to reflect a higher aversion to model uncertainty than what is presented in our results.[21] Second, in the following analysis we assume that subjects assign symmetric probabilities to the two experts. This assumption relies on a symmetry of information argument: since the information about the experts is perfectly symmetric, there is a priori no reason to believe that one of them may deserve more weight than the other. The prior distribution over the models should in consequence reflect this symmetry. This argument mirrors what Schmeidler (1989) calls an "unwritten rule saying that symmetric information with respect to the occurrence of events results in equal probabilities", or more generally to the "Principle of Insufficient Reason" or "Principle of Indifference" (Bernoulli, 1713; Laplace, 1814). In practice, it might however well be the case that some of the subjects unequally weighted the two experts, for example by overweighting the most optimistic or pessimistic expert. If this were indeed to be the case, we would expect subjects' choices to be biased in favor of model uncertainty situations, given that subjects are given the choice of the color on which to bet. The results concerning the degree of model uncertainty aversion we found and discuss in this section therefore only represent a lower bound of what might be individuals' real preferences.

4.1 Ambiguity neutrality, reduction and the nature of probabilities

Table 2 reports the results of the RLP task that was used to test our predictions and confront the expected utility theory with the model uncertainty framework. Specifically, it presents the results of the pairwise comparisons between Urn 1 and Urns 2, 3 and 4 respectively. In the spirit of Ellsberg, subjects are called ambiguity neutral if they express indifference between Urn 1 and Urn 4 (i.e. $C1 = C4$). A non-neutral attitude may either express ambiguity aversion or ambiguity seeking. Analogously, subjects may either be compound risk neutral (i.e. $C1 = C2$) or not, and exhibit the same attitude towards risk and model uncertainty (i.e. $C1 = C3$) or not. The results reveal the anticipated pattern: 70.9% of our subjects reduce compound lottery, 68.8% express a different attitude towards objective probabilities (risk) than towards subjective ones (model uncertainty), and 79.4% are non-neutral towards ambiguity.[22] As in Halevy (2007),

[21] Given that any other symmetric distribution in the space of expected utilities would consist in a mean preserving contraction of the dogmatic experts' distribution.

[22] In particular, 62.4% of our subjects are more model uncertainty averse than risk averse, and 70.4% are ambiguity averse. The detailed contingency table is provided in Appendix S1 (Table S1).

Table 2: Association between ambiguity neutrality, ROCL and ROCU

Ambiguity Neutrality		**ROCL** No	**ROCL** Yes	**ROCU** No	**ROCU** Yes	**Total**
No	Count	48 (25.4%)	102 (53.97%)	118 (62.43%)	32 (16.93%)	**150** (79.37%)
	Expected	43.7	106.3	103.2	46.8	
Yes	Count	7 (3.7%)	32 (16.93%)	12 (6.35%)	27 (14.29%)	**39** (20.63%)
	Expected	11.3	27.7	26.8	12.2	
Total		**55** (29.1%)	**134** (70.9%)	**130** (68.78%)	**59** (31.22%)	**189** (100%)

Notes: Relative frequencies in parentheses. Chi-square test: 0.085 (ROCL). Chi-square test: 8.9e-09 (ROCU).

we recover the association between ambiguity neutrality and reduction of compound objective risks. The association we found is however weaker than the one found by Halevy, with 82% of the ambiguity neutral subjects reducing compound risks (32 out of 39 subjects), but only 24% of the compound risk neutral subjects being also ambiguity neutral (32 out of 134). In comparison with the expected frequency under a null hypothesis of independence, the observed number of subjects indifferent between Urns 1, 2 and 4 is off by 16%. On the contrary, Table 2 reveals a stronger association between attitudes towards model uncertainty and ambiguity. Out of the 59 subjects who expressed the same attitude towards risk and model uncertainty, 46% of them (27 subjects) also expressed ambiguity neutrality, representing 69% of the 39 ambiguity neutral subjects. The observed frequency of subjects implicitly expressing $C1 = C3 = C4$ is therefore 2.2 times more than the expected frequency under the null hypothesis of independence. Similarly, out of the 130 subjects who did not reduce the two layers of uncertainty when being confronted with subjective probabilities, only 9% of them (12 subjects, which represents less than half of the expected frequency under the hypothesis of independence) were also ambiguity neutral. In comparison, the numbers obtained with two layers of objective probabilities suggest a weaker association, with 7 subjects being ambiguity neutral out of the 55 subjects who did not reduce the compound risk. Finally, out of the 150 subjects who did not express ambiguity neutrality, 48 did not reduce the compound uncertainty with objective probabilities (9.8% more than under the independence hypothesis), while 118 did not reduce the compound uncertainty in the presence of subjective probabilities (14.3% more than under the independence hypothesis). From Table 2, we conclude that the association observed between ambiguity neutrality and reduction of compound uncertainty is stronger when probabilities are subjective than when they are objectively given. This result is confirmed by a Chi-square test rejecting (p-value=8.9e-09) the independence hypothesis between ambiguity neutrality and reduction of compound uncertainty in the presence of subjective probabilities, and the Chi-square test that does not reject the one (p-value=0.085) between ambiguity neutrality and reduction of compound objective risks. A theory accounting for ambiguity non-neutrality should therefore be able to make a distinction between situations in which probabilities are objectively known and situations in which probabilities are subjective.

We then perform a couple of logistic regressions, reported in Table 3, where the probability of ambiguity neutrality is explained by attitudes towards compound risk and model uncertainty. Exhibiting the same attitude towards objective and subjective probabilities predicts ambiguity

Table 3: Characteristics of Ambiguity Neutrality: Logistic Regressions

	Odds Ratio	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
ROCL	2.151	0.973	0.886	5.222
	(2.138)	(1.045)	(0.820)	(5.573)
ROCU	8.297***	3.320	3.787	18.176
	(8.282***)	(3.343)	(3.754)	(18.270)

Notes: Logistic regressions. Adjusted results in parentheses. Dependent variable: Ambiguity neutrality. 189 observations. * $p-value < 0.05$, ** $p-value < 0.01$, *** $p-value < 0.001$

neutrality (odds of being ambiguity neutral are 8.3 times higher, p-value = 1.23e-7), while reduction of compound risks does not (p-value = 0.09). The probability of being ambiguity neutral is, in our sample, 20.6% and increases to 46% when the individual exhibits the same attitude towards objective as towards subjective probabilities, while it diminishes to 9% otherwise. This means that the predicted probability of ambiguity neutrality is 37 percentage points greater (p-value =1.56e-7) for an individual exhibiting the same attitude towards risk and model uncertainty than for one who does not. In comparison, compound risk neutrality only increases the predicted probability of ambiguity neutrality by 11 percentage points (from 13% to 24%, p-value =0.055).

Robustness round. The original experiment presented in this section is not free from being subject to biases and imperfections: the €15 payoff may seem too small or the length of the whole experiment too long for subjects to carefully pay attention to the properties of the problems they are facing; the recruitment of the subjects is based on an online sign-up form and the sample consists only of university students; the experiment takes place in a laboratory environment and the questionnaire is fully computerized. While the procedure and organization we follow is standard in experimental economics, they might have introduced different biases into our results. To assess the robustness of the results we conduct a second experiment with slightly modified conditions as described in Appendix A. The purpose of this robustness round is to confirm that objective and subjective probabilities are evaluated differently, and to study the link between attitudes towards risks (simple and compound), model uncertainty and ambiguity. The robustness round is an artefactual field experiment differing from the original one in several dimensions: the experiment took place during the 21st session of the Conference of the Parties (COP21) to the United Nations Framework Convention on Climate Change (UNFCCC) held in Paris in December, 2015; the subject pool consisting of policy makers or actors in the decision making process (mainly climate negotiators and NGO representatives) originated from 52 countries; the experiment was conducted with pen and paper during individual in-person interviews lasting about 15 minutes; and the prize was scaled from €15 to €50. An interesting characteristic of this sample is that subjects involved are, a priori, used to being confronted with model uncertainty in their professional activities. Climate negotiators and other participants are indeed aware different climate and economic models exist, with each of them giving different predictions regarding the evolution of the climate system and its economic consequences. Results of the robustness round are described in Appendix A. In general, the robustness round reinforces the results we previously found, with 69% of our subjects revealing a non-neutral attitude towards ambiguity, 70% exhibiting a different attitude towards objective and subjective probabilities, and 52% reducing compound risks. In total, the share of non ambiguity neutral subjects who did not reduce compound uncertainty with subjective probabilities is 87%, while the share of non-ambiguity neutral subjects who did not reduce compound risk is 60% (see Table A.1). Although the independence hypothesis between compound risk and ambiguity *is* rejected in this round, the results of the logistic regressions confirm the much stronger association between ambiguity neutrality and similar attitudes towards risk and model uncertainty, than between ambiguity neutrality and compound risk reduction (Table A.2).

Finally, in online Appendix S1 we go beyond the dichotomous analysis of neutral/non-neutral attitudes, and separate preferences into "averse", "neutral" and "loving" behaviors for the different types of uncertain situations we propose. We show that the results obtained for both the main experiment and the robustness round are maintained and are even strengthened in this case. All in all, model uncertainty attitudes seem to be the driving force in determining the attitudes subjects exhibit towards ambiguity.

4.2 Risk vs. Model uncertainty

If a non-neutral attitude towards ambiguity is tightly associated with the difference in the attitudes towards risk and model uncertainty, it is interesting to investigate in more detail the extent to which individuals place different value on these two types of uncertain situations. The two CE tasks enable us to obtain a direct measure of the strength of model uncertainty aversion relative to risk aversion. It is achieved using pairwise comparisons where the individual is offered a series of choices between a sure amount, and a risky or model uncertainty outcome respectively. As previously mentioned, we chose to discard inconsistent data from this analysis. The consistent sub-sample is made of 169 subjects. Table 4 reports the descriptive statistics of the intervals to which the certainty equivalents of Urns 1 and 3 (i.e. uncertain situations $25_{.5}4$ and $25_{1\,0}4$, respectively) belong. Each interval defines the highest outcome for which the uncertain situation is preferred and the lowest outcome which is preferred to the uncertain situation. The results

Table 4: Descriptive Statistics

	Mean	**Median**	**Mode**	**SD**	**Min**	**Max**	**Obs**
$C1$	[11.92; 13.22]	[13 ; 14]	[14 ; 15]	[2.35 ; 1.98][a]	[6 ; 8]	[15 ; 18]	169
$C3$	[10.25; 11.80]	[10 ; 12]	[8 ; 10]	[3.16 ; 2.82][a]	< 4	> 25	169

[a]The first (second) number corresponds to the standard deviation of the lower (upper) bound.

from Table 4 confirm our predictions: subjects are on average ready to pay a higher premium –measured as the difference between the certainty equivalent and the expected gain– to avoid a situation where probabilities are subjective than to avoid a situation in which probabilities are objective. In particular, for an expected outcome of €14.50, the mean amount that our subjects deem equivalent to the risky situation is between €11.92 and €13.22, while under model uncertainty the mean lies in the interval €10.25-11.80. The distribution of certainty equivalents of the risky situation ($C1$) second order stochastically dominates the distribution of the certainty equivalents of the model uncertainty situation ($C3$). While $C1$ does not first order stochastically dominate (FOSD) $C3$, this is only because there is one subject who systematically preferred the model uncertainty situation to the sure outcome (even when the choice was made between €25 for sure, and a situation in which one expert expressed a 100% probability the gain would be €25, while the other expert expressed a 100% probability the gain would be €4). We can only speculate what the preferences of this subject are. He/she could be a very optimistic subject who always trusts the expert predicting the highest outcome. In that sense, his/her first choice would express indifference between two situations yielding €25. Once we remove this subject from the sample, we recover the result that $C1$ FOSD $C3$. This result is illustrated in Figure 1, which displays the proportions of safe choices –expressed by preference for the sure amount– for each of the the ten decisions between Urn 1 or Urn 3. The list of sure amounts is written in descending order. The dashed line represents the prediction under the assumption of either risk or model

Figure 1: Proportion of safe choices in the CE tasks and predictions

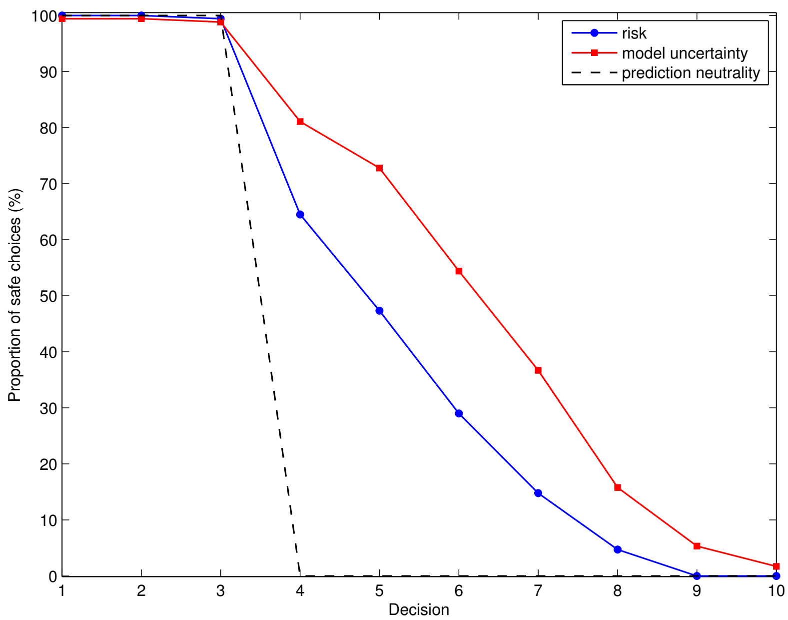

uncertainty neutrality. In this case, the certainty equivalents of both uncertain situations are the same, and the probability that the sure outcome is chosen is 1 for the first three decisions, and then 0 for the remaining ones. The blue line presents the observed choice frequency of the sure outcome option for each of the decisions in the risky situation. As can be observed, it is at the right of the risk neutral prediction, indicating a tendency for risk aversion among subjects

($C1 \leq C_0$). The red line represents the observed frequency in the case of model uncertainty. It lies to the right of the blue line, suggesting our subjects manifest a stronger aversion to model uncertainty than to risk ($C3 \leq C1 \leq C_0$). As predicted, we finally note that our subjects did not express infinite model uncertainty aversion (which would have consisted of a proportion of 100% for each decision). In fact, only two subjects (representing a proportion of 1.2%) expressed an extreme form of pessimism by systematically selecting the certain outcome when confronted to model uncertainty.[23] We can therefore confidently reject the maxmin expected utility hypothesis (Wald, 1950) in which subjects only consider the worst possible existing model in their decision making process. A Wilcoxon signed-rank test statistically confirms (p-value=3.2e-12) that the risky alternative is valued differently than the corresponding model uncertainty situation.[24]

5 Characterizing preferences under uncertainty

We use the choices made in the various PL tasks to further characterize preferences under uncertainty in the context of the smooth model. In particular, we use the 80 binary choices each subject typically provided to infer attitudes towards risk and model uncertainty, and use this information to quantify the degree of ambiguity aversion. In total about 14% of choices in the eight PL tasks were deemed inconsistent (reverse choices or multiple switching points) and were discarded from the analysis. This number is in line with what is found in other laboratory experiments (e.g., Holt and Laury, 2002). We feel confident that subjects who are left in the sample understood the instructions correctly and were revealing their true preferences.

5.1 Preliminary remark

Since the model of decision under uncertainty we study involves two distinct behavioral characteristics of the decision maker, the experimental procedure has to be designed such that it generates data that are rich enough to disentangle the different components of the subjects' attitude towards ambiguity. The double PL procedure, which presents choices in the presence of both objective and subjective probabilities, is designed for this purpose. It enables us to jointly elicit risk and model uncertainty attitudes. To see the importance of using a joint procedure, consider the identification of risk and model uncertainty under the general model uncertainty theory presented in expression (6). Assuming this expression correctly describes choices made by our subjects over uncertain alternatives, a subject would be indifferent between two options $\bar{o}_p\underline{o}$ and $\bar{o}_{\hat{p}_1\hat{p}_2}\underline{o}$ if and only if:

$$u^{-1}\Big(pu(\omega+\bar{o})+(1-p)u(\omega+\underline{o})\Big) = v^{-1}\bigg(\frac{1}{2}\left(v\circ u^{-1}\right)\Big(\hat{p}_1 u(\omega+\bar{o})+(1-\hat{p}_1)u(\omega+\underline{o})\Big) + \frac{1}{2}\left(v\circ u^{-1}\right)\Big(\hat{p}_2 u(\omega+\bar{o})+(1-\hat{p}_2)u(\omega+\underline{o})\Big)\bigg), \tag{9}$$

[23]Note that relaxing the definition of expressing an extreme form of pessimism by considering those subjects who expressed nine safe choices before switching to the model uncertainty situation as indifferent between the two options, the number of extreme model uncertainty averse individuals increases to 9 out of the 169 subjects (5.3%).

[24]The significance of the one-sided test, where the alternative hypothesis is that the median of the switching point in the model uncertainty task is greater than in the risk aversion task, is 4.2e-12.

where ω represents background wealth. When considered in terms of attitude towards ambiguity, the identity $\phi = v \circ u^{-1}$ enables us to rewrite (9) as

$$pu(\omega+\bar{o})+(1-p)u(\omega+\underline{o}) = \phi^{-1}\biggl(\frac{1}{2}\phi\Bigl(\hat{p}_1 u(\omega+\bar{o})+(1-\hat{p}_1)u(\omega+\underline{o})\Bigr) + \frac{1}{2}\phi\Bigl(\hat{p}_2 u(\omega+\bar{o})+(1-\hat{p}_2)u(\omega+\underline{o})\Bigr)\biggr). \tag{10}$$

From (9) and (10), it is clear that estimating model uncertainty aversion or ambiguity aversion under the assumption of risk neutrality yields exactly the same results. Ambiguity aversion is therefore significantly overestimated when risk neutrality is assumed.[25] If we relax the assumption of risk neutrality and let risk aversion $-u''/u'$ be positive, it becomes clear from the relationship $-\phi''/\phi' = (-v''/v' + u''/u')/u'$ that the implied degree of absolute ambiguity aversion is lower. One can therefore not capture the distinction between model uncertainty and ambiguity aversion without estimating the level of risk aversion, for which separated risky tasks also need to be performed.

5.2 The double PL tasks

In the double PL tasks, we exploit comparisons between urns of type 2 and 3, with experts who are no longer dogmatic. The risky tasks only deal with urns of type 2. They are based on Holt and Laury's (2002) mechanism, which has become a standard for elicitation of risk aversion. The model uncertain tasks are constructed analogously. Table 5 illustrates the type of choices our subjects were confronted with in this part of the experiment. In this example, Option A offers either €35 or €1 with equal probability, while Option B offers the same outcomes with unknown probabilities (although respondents are also given additional information in the form of the two experts' assessments). In the first decision for example, Expert 1 assesses the probability of obtaining €35 to be 50%, while Expert 2 is 100% sure the outcome will be €1. The expected value of Option A ($\mathrm{EV^A}$), the expected value of Option B if either Expert 1 or Expert 2 is correct (respectively $\mathrm{EV_1^B}$ and $\mathrm{EV_2^B}$), the average expected value of Option B ($\mathrm{EV^B}$) under the assumption of equal weights attached to each expert, and its standard deviation ($\mathrm{SD^B}$), are also provided in Table 5, but were not given to subjects during the experiment. While the expected value of Option A is kept constant throughout the various choices, the expected value of Option B is increasing as one proceeds down the table. The standard deviation, on the other hand, is decreasing (except between the first and second decision). Overall, the decision table is constructed in such a way that, for any increasing utility function, Option B always stochastically dominates (in the first or second order sense) the previous decision as one proceeds down the table.[26] Following the theoretical model presented in Section 3.1, this feature should induce subject to switch only once, from Option A to Option B, while progressing down the table. Our subjects went through four tasks similar to the one illustrated in Table 5, which vary in the

[25] We discuss the estimation results of this particular case in Appendix S5 provided online.

[26] It is for example easy to see that Option B in the second decision first order stochastically dominates Option B in the first decision, and that Option B in the fourth decision second order stochastically dominates Option B in the third decision.

Table 5: Payoff table in the model uncertainty aversion tasks

Option A			Option B				EV^{A}	EV_1^{B}	EV_2^{B}	EV^{B}	SD^{B}
$\bar{o}$	p	$\underline{o}$	$\bar{o}$	$\hat{p}_1$	$\hat{p}_2$	$\underline{o}$	(€)	(€)	(€)	(€)	(€)
35	0.5	1	35	0.5	0	1	18	18	1.0	9.5	8.5
35	0.5	1	35	0.9	0	1	18	31.6	1.0	16.3	15.3
35	0.5	1	35	0.9	0.09	1	18	31.6	4.1	17.8	13.8
35	0.5	1	35	0.8	0.19	1	18	28.2	7.5	17.8	10.4
35	0.5	1	35	0.8	0.21	1	18	28.2	8.1	18.2	10.0
35	0.5	1	35	0.7	0.31	1	18	24.8	11.5	18.2	6.6
35	0.5	1	35	0.6	0.41	1	18	21.4	14.9	18.2	3.2
35	0.5	1	35	0.55	0.46	1	18	19.7	16.6	18.2	1.5
35	0.5	1	35	0.51	0.50	1	18	18.3	18.0	18.2	0.2
35	0.5	1	35	0.61	0.60	1	18	21.7	21.4	21.6	0.2

Notes: Probabilities always refer to the outcome $\bar{o} \geq \underline{o} \in O$. $EV^{B} = \frac{1}{2}EV_1^{B} + \frac{1}{2}EV_2^{B}$; $SD^{B} = \left(\frac{1}{2}(EV_1^{B} - EV^{B})^2 + \frac{1}{2}(EV_2^{B} - EV^{B})^2\right)^{0.5}$

proposed payoffs and probabilities. The set of payoffs and probabilities is designed in a way that the final payoffs span the range of income over which we are estimating model uncertainty aversion, which is the same as the one over which risk aversion is estimated.[27]

5.3 Eliciting risk and model uncertainty attitudes

We use each of the subjects' binary choices to estimate the parameters of two latent utility functions that explain these choices. We allow for a stochastic error structure, as opposed to a strictly deterministic structural estimation procedure, as we want to allow for subjects to make some errors and, at the same time, to account for the panel structure of the data. Given the important support for the CRRA hypothesis in the empirical literature on risk aversion (Harrison et al., 2007; Brunnermeier and Nagel, 2008; Chiappori and Paiella, 2011), but at the same time the experimental evidence found in favor of increasing relative risk aversion (IRRA) (Holt and Laury, 2002), we maintain a generic parametric structure for the identification problem. We let both utility functions representing risk and model uncertainty attitudes be of the expo-power (EP) form (Saha, 1993). In the case of risk, this means that the utility function takes the following form:

$$u(x) = \frac{1 - \exp\left(-a_u(\omega + x)^{1-r_u}\right)}{a_u}. \tag{11}$$

This representation includes CRRA and constant absolute risk aversion (CARA) as special cases, and exhibits the desirable properties of decreasing absolute risk aversion and increasing relative risk aversion for positive values of the parameters a_u and r_u (Abdellaoui et al., 2007).[28] Note the presence of ω, representing background wealth in expression (11). As is generally the case in the experimental literature, we assume $\omega = 0$. It should however be clear that in a situation in

[27] In particular, Option A in the three other uncertain tasks takes values as follows: $3.5_{0.5}0.1$, $35_{0.1}1$, $35_{0.9}1$, and Option B covers the space around these values analogously to what is presented in Table 5. For a detailed description of our experimental design, please refer to the Online supplemental Appendix.

[28] As is well known, the Arrow-Pratt index of relative risk aversion of the EP function is $-u''(x)(\omega+x)/u'(x) = r_u + a_u(1-r_u)(\omega+x)^{1-r_u}$. It is then easy to see that this function exhibits CRRA of value r_u when $a_u = 0$, and CARA of a_u when $r_u = 0$.

which $\omega > 0$, the same observed choices would imply higher risk aversion. Using the procedure proposed by Andersen et al. (2008), we then construct the expected utility of the two options comprising each decision by using candidate values of parameters a_u and r_u, and a linking index in order to infer the likelihood of the observed choice. The parameters of the latent utility function (11) are then chosen in order to maximizes the likelihood of getting the observed ranking of the different options, taking into account a Luce (1959) error specification with a structural noise parameter.[29] The first part of Table 6 presents the estimates obtained from the risky tasks. Given the prominent position CRRA has achieved in the theoretical and empirical literature, we provide both of the estimates for the cases in which u is of the CRRA and EP type. The estimate

Table 6: Estimates of risk, model uncertainty and ambiguity preferences

	u		v		ϕ	
	CRRA	EP	CRRA	EP	CRAA	EP
a		0.0294*** (0.00215)		0.152*** (0.0542)		-1.802 (0.9655)
r	0.279*** (0.0119)	0.135*** (0.0193)	0.738*** (0.0210)	0.467*** (0.0542)	0.534*** (0.0261)	0.86*** (0.0452)
noise parameter	0.103*** (0.00327)	0.105*** (0.00330)	0.0358*** (0.00237)	0.0534*** (0.00343)	0.0476*** (0.00213)	0.0363*** (0.00184)
Observations	5320	5320	7570	7570	7570	7570
Loglikelihood	-1550.3	-1516.8	-3682.5	-3682.1	-3680.6	-3675.1

Notes: Luce error specification is used in the estimation. Standard errors in parentheses. The EP risk specification is used to estimate v and ϕ. * $p-value < 0.05$, ** $p-value < 0.01$, *** $p-value < 0.001$

of the CRRA parameter we obtain is 0.28, which is lower than the one we found using the CE task only. When the EP specification is considered, we estimate $r_u = 0.135$ and $a_u = 0.029$, which implies IRRA. While the focus of our analysis is on comparing these estimates with the ones obtained for the model uncertainty function v, we note that their absolute magnitudes are consistent with the results obtained by Holt and Laury (2002); Andersen et al. (2008). We however recognize that the estimates we obtain only hold locally over the domain of stakes offered in our experiment. The last two rows of Table 6 present information about the data used (30 risk aversion choices for each of the 189 subjects, minus the inconsistent choices that are discarded) and the resulting log likelihood values. As can be seen, the log likelihood of the EP specification is slightly better than the CRRA one, but this should not be surprising given that the estimates are all significant, and the hypothesis $a_u = 0$ is therefore rejected. Given the superiority of the EP specification in explaining the observed choices in the risky tasks, this is the specification we consider in the remaining part of the estimation procedure. We then estimate the model uncertainty aversion function v, which takes the general EP form:

$$v(CE) = \frac{1 - \exp\left(-a_v (CE)^{1-r_v}\right)}{a_v}, \tag{12}$$

[29]The statistical specification we use allows for taking into account the correlation between responses given by the same subject. Robust estimates considering clustering corrections are provided in the Supplemental Material provided online (Appendix S3). There is essentially no difference in the significance of our estimates in this case.

Figure 2: Absolute (left) and relative (right) risk and model uncertainty aversion using EP estimates (95% confidence in grey).

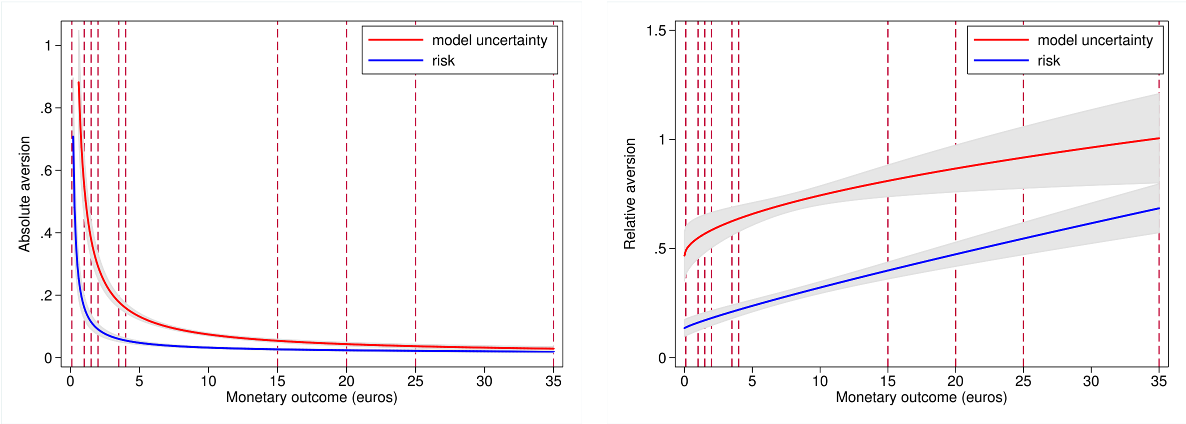

where CE represents the certainty equivalent wealth for a given model θ: $CE \equiv u^{-1}\Big(\hat{p}_\theta u(\bar{o}) + (1-\hat{p}_\theta)u(\underline{o})\Big)$. The second part of Table 6 presents the estimates obtained from our five uncertain PL tasks. Estimates for the special cases of v being of the CRRA type ($a_v = 0$) are also provided for indicative purposes. In that case, the coefficient estimated is significantly higher than the one obtained in the case of risk. It should however be noted that this specification leads to a smaller loglikelihood value than the general expo-power formulation (12). Focusing on the EP specification, we remark that the estimates we obtain ($a_v = 0.152$ and $r_v = 0.467$) with the joint identification procedure are both significantly positive. This implies our subjects exhibit both decreasing absolute model uncertainty aversion and increasing relative model uncertainty aversion. Our interest lies in comparing the estimates obtained for model uncertainty with the ones obtained for risk aversion. In Figure 2, we provide the paths of estimated absolute and relative aversion indexes for both risk and model uncertainty over the experimental prize domain. As predicted, we observe that the indexes are both decreasing in the monetary outcome when considered in absolute terms and increasing in relative terms (DARA and IRRA). Interestingly, we also directly observe from Figure 2 that the degree of model uncertainty aversion is significantly higher (in both absolute and relative terms) than the one of risk aversion. This result confirms our main hypothesis that subjects are more averse to subjective probabilities than to objective ones. Specifically, while the index of relative risk aversion is respectively 0.32 and 0.62 when the monetary outcome considered is either $x = 10$ or $x = 30$, the index of relative model uncertainty aversion takes values of 0.74 and 0.96 for the corresponding outcomes. Note that in the special case where both u and v are of the CRRA type, the indexes of relative aversion to risk and model uncertainty are $r_u = 0.28$ and $r_v = 0.73$ when jointly estimated.[30] These differences observed between the attitudes towards objective and subjective probabilities now enable us to quantify the attitude subjects manifest towards ambiguity.

[30] In order to assess the sensitivity of the model uncertainty aversion index to variations in relative risk aversion, we also used the maximum likelihood procedure to estimate r_v using different (exogenously given) values of r_u. These additional results are presented in Appendix S5 available online.

5.4 The implications for ambiguity attitude

The joint characterization of functions u and v representing the subjects' attitudes towards two different types of uncertainty (risk and model uncertainty) has an important direct implication for the characterization of ambiguity aversion. Indeed using the identity $\phi \equiv v \circ u^{-1}$ and the results obtained in the previous section, we are now able to characterize directly the attitude subjects manifest towards ambiguity, and to compute the indexes of absolute and relative ambiguity aversion (see Online Appendix S2 for the detailed analytical computations under the double EP specification). These indexes are represented in Figure 3. While we observe a

Figure 3: Absolute (left) and relative (right) ambiguity aversion obtained with EP function estimates

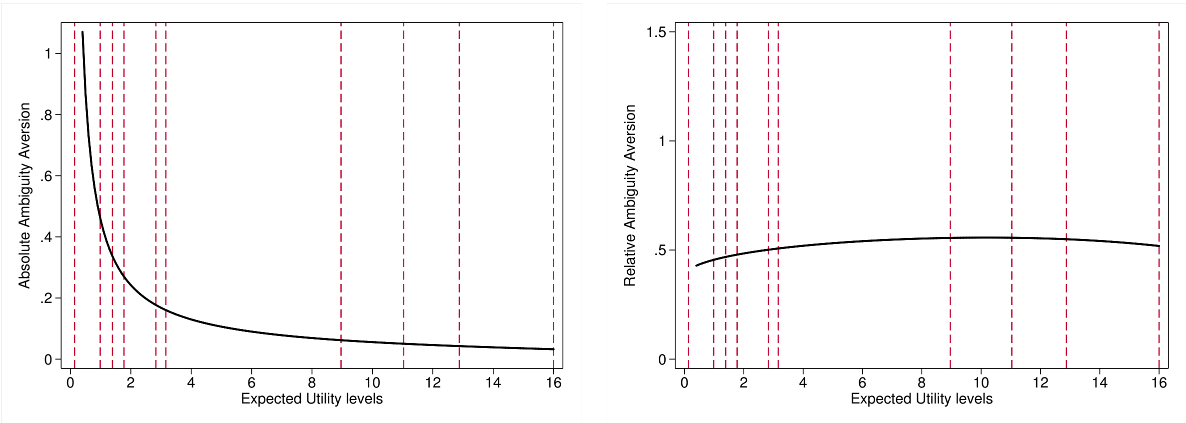

clear decreasing trend in the degree of absolute ambiguity, we remark that the degree of relative ambiguity aversion seems to be fairly constant over the domain considered. It should be noted that the domain of the ambiguity function ϕ is not the same monetary outcome domain used previously in the study of u and v. Instead, ϕ is defined over expected utility levels U. In this sense, the vertical dashed red lines in Figure 3 represent the levels of utility obtained for the corresponding monetary outcomes in Figure 2, when the utility function u is of the EP type and coefficients are as estimated in Table 6.

To assess the robustness of the constant relative ambiguity aversion (CRAA) result presented in Figure 3, we apply the joint estimation procedure directly to u and ϕ (both of which are of the EP type). In particular, this means that we let the ambiguity aversion function be:

$$\phi(U) = \frac{1 - \exp\left(-a_\phi (U)^{1-r_\phi}\right)}{a_\phi}, \tag{13}$$

where U represents the expected utility obtained under a given model θ: $U \equiv \hat{p}_\theta u(\bar{o}) + (1 - \hat{p}_\theta)u(\underline{o})$, and u is defined as in equation (11). The estimated results are provided in the last two columns of Table 6. In this case, the coefficient a_ϕ of the EP formulation is not significant at the 5% level ($p - value = 0.062$). The function describing preferences towards ambiguity should therefore be of the CRAA type instead. Under this particular specification, the constant relative ambiguity aversion index is estimated to be 0.53. It does not correspond exactly to the value observed in Figure 3, but this should not be surprising given that the ambiguity functions do not share the same specification in the two cases. If we instead consider the case of u being CRRA, we also obtain a non significant coefficient a_ϕ ($p - value = 0.093$) under the EP specification, and estimate the coefficient $r_\phi = 0.62$ under constant relative ambiguity aversion.[31]

6 Conclusion

Uncertainty is crucial in collective as well as in individual decision making. During the past few years, a vast literature aiming at better formalizing the decision process in the face of objective and subjective uncertainty has been growing and encompassing multiple academic fields. This body of research investigates how individuals integrate available information in the process of decision making through the development of theoretical frameworks and experimental analyses. In particular, multiple decision models have been developed to account for attitudes towards ambiguity. These models have been adopted to explain individuals' behavior in multiple contexts and are increasingly applied to prescribe optimal strategies in the face of uncertainty. The growing application of ambiguity aversion models calls for the development of experimental efforts enabling both a better understanding of the underlying mechanisms at play, and the quantification of ambiguity preferences, similar to what has been done in the study of risk. In this paper, we provide new experimental evidence on behavior towards compound risk and model uncertainty in relation to simple risk and ambiguity. Our design enables us to disentangle the role played by objective and subjective probabilities in determining individuals' ambiguity attitudes, and to quantify, through a joint elicitation procedure, the extent to which ambiguity aversion exists as well as the properties of the ambiguity aversion function. We conducted both a laboratory experiment with students, and a field experiment with policy makers, and use non-parametric statistics as well as structural econometrics to analyze choice patterns.

There are three main findings emerging from our analysis. First, we confirm that attitudes towards ambiguity and uncertainty presented in a compound way are associated. This association is however much stronger when the second layer of uncertainty is subjective than when it is objective. Provided that the compound probabilities are simple enough, we find that most subjects reduce compound risks but do not reduce compound uncertainty when different models are considered and the probability of each of them being correct is unknown. Second, we show that subjects tend to be both risk and model uncertainty averse, but exhibit stronger aversion to model uncertainty than to risk. Following a generic model of choice under uncertainty (Klibanoff et al., 2005; Marinacci, 2015), we interpret this behavioral characteristic as evidence of ambiguity aversion. Using a joint estimation procedure, we elicit the degree of ambiguity aversion, which we estimate to be around 0.5 when considered in relative terms. Third, investigating in more detail attitude towards model uncertainty, we find that model uncertainty aversion is decreasing in wealth when considered in absolute terms, and increasing when considered in relative terms. In regards to ambiguity attitude, we find evidence of decreasing absolute ambiguity aversion (DAAA) and constant relative ambiguity aversion (CRAA). The results we obtain reveal incon-

[31] In this case, the result could have been obtained directly from the twofold CRRA estimation results provided in online Appendix S2, given that ϕ is of the CRAA type with $r_\phi = \frac{r_v - r_u}{1 - r_u}$, when both u and v are CRRA (Berger et al., 2016).

sistencies with the SEU model (Savage, 1954; Cerreia-Vioglio et al., 2013b): most of our subjects reduce compound objective risk and the majority of them (70%) are ambiguity non-neutral or do not reduce subjective probabilities (62%). The results also enable us to reject the maxmin model: only between 1 and 5% of our subjects' choices are compatible with a decision based only on the most pessimistic model, and the coefficients of both model uncertainty and ambiguity aversion that we estimate are finite. Our findings are however consistent with an interpretation of the two-stage model (Klibanoff et al., 2005; Nau, 2006; Marinacci, 2015) in which ambiguity non-neutrality stems from the non-reduction of objective and subjective uncertainty, rather than from an inability to reduce compound objective risks (Seo, 2009). Finally, our results caution against modeling ambiguity attitude by means of exponential functions (Hansen and Sargent, 2001, 2008), though further research is warranted to make bolder quantitative statements.[32] Overall, the results in these experiments call for a new reading of some important findings previously obtained in the literature in trying to explain the behavioral mechanisms underneath individuals' attitudes towards ambiguity.

[32]The robust control model developed by Hansen and Sargent (2001, 2008) may be seen as a special case of an REU model in which the ambiguity function is of the exponential form (i.e. constant absolute ambiguity aversion) (Cerreia-Vioglio et al., 2011).

Appendix

A Robustness round

We conducted a second experiment (robustness round) at the COP21 to the UNFCCC, held in Paris in December, 2015. The 91 subjects who participated in this robustness round originated from 52 different countries: 46% of them were climate negotiators, 21% represented NGOs, and the remaining ones were either researchers/academics (11%), journalists (5%), representatives of the private sector (4%) or self-identified with a different category (12%). In individual in-person interviews, we prompted respondents who volunteered for the study with a few questions framed in the context of climate change,[33] before giving them the RLP task. Additional procedural differences were: the experiment was conducted with pen and paper, the payoff reached €50 if the bet was correct, and subjects did not have access to calculators. Table A.1 summarizes the association we found in the robustness round between ambiguity neutrality, reduction of compound lotteries with objective probabilities (ROCL) and reduction of compound uncertainty with subjective probabilities (ROCU). As can be observed, most of our subjects (69%) reveal

Table A.1: Association between ambiguity neutrality, ROCL and ROCU (robustness round)

		ROCL		**ROCU**		
Ambiguity neutral		No	Yes	No	Yes	**Total**
No	Count	38	25	55	8	**63**
	Expected	*30.5*	*32.5*	*44.3*	*18.7*	
		(41.76%)	(27.47%)	(60.44%)	(8.79%)	(69.23%)
Yes	Count	6	22	9	19	**28**
	Expected	*13.5*	*14.5*	*19.7*	*8.3*	
		(6.59%)	(24.18%)	(9.89%)	(20.88%)	(30.77%)
Total		**44**	**47**	**64**	**27**	**91**
		(48.35%)	(51.65%)	(70.33%)	(29.67%)	(100%)

Notes: Relative frequencies in parentheses. Chi-square test (ROCL): 6.1e-4. Chi-square test (ROCU): 1.1e-7

a non-neutral attitude towards ambiguity, 48% do not reduce compound risk, and 70% do not reduce compound uncertainty when different probabilistic models exist and the probabilities associated to each of them are subjective. Similarly to Halevy's (2007) results, we observe a relatively strong association between ambiguity neutrality and reduction of compound objective risks. In particular, the observed frequency of subjects implicitly revealing $C1 = C2 = C4$ is 52% higher than the expected frequency under a null hypothesis of independence. However, as observed in the original experiment, the frequency of subjects implicitly expressing $C1 = C3 = C4$ is also more than twice (exactly 2.3 times) the expected frequency under the null hypothesis of independence. We observe a stronger association between ambiguity neutrality and reduction of compound uncertainty when probabilities are subjective than when they are objective, but contrary to the main experiment, we reject the independence hypothesis in both situations (the

[33] Specifically we asked them their assessed probability distribution over 2100 temperature increases based on current "Nationally Determined Contributions".

p-value of the Chi-square test of independence between ambiguity neutrality and reduction of compound uncertainty when probabilities are objective is 6.1e-4, and 1.1e-7 when probabilities are subjective).

The results of the logistic regressions for the probability of being ambiguity neutral are reported in Table A.2. As in the original experiment, ambiguity non-neutral subjects are overrepresented in the sample. The results are very similar to those presented in Table 3. In particular, an identical attitude towards objective and subjective probabilities enables us to predict ambiguity neutrality with statistical significance. The odds of being ambiguity neutral when expressing ROCU corresponds to 14.5 times the odds when it is not the case (p-value = 1.37e-6). Reduction of compound risk alone also now has a significant impact on ambiguity neutrality (p-value = 0.0011). However, we can see from Table A.2 that the odds ratio is much lower than the one corresponding to the attitude towards model uncertainty. Moreover, as we adjust the logistic regression to account for the two effects simultaneously, the effect of compound risk attitude becomes non-significant (p-value = 0.37). The probability of being an ambiguity neutral subject

Table A.2: Characteristics of Ambiguity Neutrality: Logistic Regressions (robustness round)

	Odds Ratio	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
ROCL	5.573**	2.941	1.981	15.677
	(1.793)	(1.174)	(0.497)	(6.472)
ROCU	14.514***	8.041	4.900	42.989
	(10.671***)	(6.790)	(3.066)	(37.138)

Notes: Logistic regressions. Adjusted results in parentheses. Dependent variable: Ambiguity neutrality. 91 observations. * $p-value < 0.05$, ** $p-value < 0.01$, *** $p-value < 0.001$

in this sample is 30.7%. It increases to 70% when the individual exhibits the same attitude towards objective as towards subjective probabilities (and drops to 14% when this is not the case). This means that the change in probability increases by 56 percentage points and is significant (p-value =9.26e-9) when attitude towards model uncertainty goes from 'the same attitude as the one towards risk' to 'a different attitude than the one towards risk'. In comparison, compound risk neutrality only increases the predicted probability of ambiguity neutrality by 33 percentage points when considered in isolation (from 14% to 47%, p-value =0.0002).

B Application to the smooth model

B.1 The choices situations

In this section, we characterize the different situations presented in the experiment under the smooth model. In that case, the different uncertain situations presented in the RLP task are evaluated as follows:

- *Risk (Urn 1)*: In the first risky urn, the set of possible models is a singleton $M = \{P(r) =$

1/2}. There is no model uncertainty and representation (6) collapses to the standard von Neumann-Morgenstern (vNM) expected utility representation:

$$U(f_1) = \sum_{c\in\{r_1,b_1\}} P(c)\, u\left(f_1(c)\right). \tag{B.1}$$

- *Compound risk (Urn 2)*: The case of Urn 2 is very similar since it only deals with risk. A bet on this urn is therefore evaluated using function u only. Two objective models are considered here: $M = \{P(r) = 1, P(r) = 0\}$, and the probability of each model being the correct one is the objective probability associated with the coin toss. The subjective prior beliefs over models therefore coincide with the objective probabilities. This representation of the two layers of risk is mathematically equivalent to a situation in which a single model $\bar{P}(r)$ exists. This therefore means that if the decision maker is able to reduce compound risks, the situation is evaluated exactly the same way as the first one:

$$U(f_2) = \mathrm{E}_\theta(u \circ u^{-1}) \sum_{c\in\{r_2,b_2\}} \tilde{P}_\theta(c)\, u\left(f_2(c)\right) = \sum_{c\in\{r_2,b_2\}} \bar{P}(c)\, u\left(f_2(c)\right), \tag{B.2}$$

  where $\bar{P}(c) := \mathrm{E}_\theta \tilde{P}_\theta(c) = 1/2$ are the reduced probabilities of the events $c \in \{r_2, b_2\}$.

- *Model uncertainty (Urn 3)*: In the case of Urn 3, the DM is not given any information about the total number of balls in the urn to make it difficult for him to construct any possible *objective* urn model. However, direct information about the composition of the urn is given by two "experts" that only differentiate themselves by their names ("Expert 1" and "Expert 2") and their assessments of the urn composition. This information is assumed to be taken as a datum of the problem, and as such considered as objective by the DM. In particular, this information enables the DM to define the set $M = \{P_\theta(r)\}$, where $\theta = \{1, 2\}$ refers to the experts. In the RLP task, the two possible models described by the experts are $M = \{P(r) = 1, P(r) = 0\}$. The probability q_θ of each model being perceived as correct is subjective and the second layer of uncertainty is evaluated using function v. In principle, this urn is therefore evaluated using a combination of both functions u and v. However since the two experts are dogmatic in our RLP task, the risk is degenerate. The evaluation of any act in this special case is then realized using function v only:

$$U^d(f_3) = \sum_{c\in\{r_3,b_3\}} q_\theta v\big(f_3(c)\big), \tag{B.3}$$

  where the superscript d refers to the case where experts are *dogmatic*.

- *Ambiguity à la Ellsberg (Urn 4)*: In the case of Urn 4, the proportion of red and black balls is unknown, but the total number of balls N present in the urn is given. The set of possible models may therefore be restricted to $M = \{P_\theta(r) : P(r) = \frac{\theta-1}{N} \text{ for } \theta = \{1, \ldots, N+1\}\}$, and the DM then subjectively determines to which model he assigns a positive probability.

Act f_4 is then evaluated as follows:

$$U(f_4) = \sum_{\theta=1}^{N+1} q_\theta (v \circ u^{-1}) \left(\sum_{c \in \{r_4, b_4\}} P_\theta(c)\, u\,(f_4(c)) \right). \tag{B.4}$$

B.2 Predictions

Our first two predictions in the context of the smooth model are expressed as follows.

Hypothesis 1. When functions u and v represent attitudes towards risk and model uncertainty respectively, the first hypothesis simply becomes $u'' \leq 0$ and $v'' \leq 0$.

Hypothesis 2. Under the smooth framework, the first part of second hypothesis simply turns out to be v being more concave than u, in the sense that $-\frac{v''}{v'} \geq -\frac{u''}{u'}$. Furthermore, the association between (4) and (5) turns out to be perfect by construction ($\phi = v \circ u^{-1}$). Condition (4) is therefore both necessary and sufficient for (5). The condition $C3 > \underline{C}$ in turn translates to $-\frac{v''}{v'} < \infty$, or equivalently $v^{-1}\left(\mathrm{E}_\theta(v \circ u^{-1})\left(\sum_{c \in \{r_i, b_i\}} \tilde{P}_\theta(c) u\,(f_i(c))\right)\right) > u^{-1}\left(\min_\theta \sum_{c \in \{r_i, b_i\}} P_\theta(c) u\,(f_i(c))\right)$. The first equality of statement (4) is then trivial and comes from the fact that criterion (6) collapses to the standard vNM expected utility model when all sources of uncertainty are objective.

B.3 Risk vs. Model uncertainty in the CE tasks

Choice frequencies for each number of safe choices in the CE tasks and the implied interval for the risk or model uncertainty aversion parameter, in the special case of the CRRA function,[34] are reported in Table B.1. As can be observed, more than 84% of subjects made between 3 and

Table B.1: Classification of uncertain choices (CE tasks)

Number of safe choices	Range of relative risk or model uncertainty aversion: $u(x)$ or $v(x) = x^{1-r}/1-r$	Proportion of choices	
		Risk	Model Uncertainty
0-1	$r < -1.04$	0.00 %	0.59 %
2	$-1.04 < r < -0.12$	0.56 %	0.59 %
3	$-0.12 < r < 0.12$	34.91 %	17.75 %
4	$0.12 < r < 0.34$	17.16 %	8.28 %
5	$0.34 < r < 0.55$	18.34 %	15.38 %
6	$0.55 < r < 1$	14.20 %	20.71 %
7	$1 < r < 1.55$	10.06 %	21.89 %
8	$1.55 < r < 2.58$	4.73 %	9.47 %
9-10	$2.58 < r$	0.00 %	5.33 %

[34] A utility function has the CRRA property if it takes the form $u(x) = \frac{x^{1-r}}{1-r}$, where r is the coefficient of relative risk aversion (when $r = 1$, this collapses to $u(x) = \ln x$).

6 safe choices in the risky task, while in the model uncertainty task the proportion of subjects making choices in this interval is only 62%. When we consider the proportion of subjects making between 2 and 7 safe choices, these numbers increase respectively to 95% and 85%. Finally, using an estimation procedure that will be described in the next section, we also found the best estimates for the coefficients of relative risk and model uncertainty aversion when both u and v are of CRRA type to be respectively $r_u = 0.42$ and $r_v = 0.83$ for the two CE tasks.

References

Abdellaoui, M., A. Baillon, L. Placido, and P. P. Wakker (2011). The rich domain of uncertainty: Source functions and their experimental implementation. *The American Economic Review 101*(2), 695–723.

Abdellaoui, M., C. Barrios, and P. P. Wakker (2007). Reconciling introspective utility with revealed preference: Experimental arguments based on prospect theory. *Journal of Econometrics 138*(1), 356–378.

Abdellaoui, M., P. Klibanoff, and L. Placido (2015). Experiments on compound risk in relation to simple risk and to ambiguity. *Management Science 61*(6), 1306–1322.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2006). Elicitation using multiple price list formats. *Experimental Economics 9*(4), 383–405.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2008). Eliciting risk and time preferences. *Econometrica 76*(3), 583–618.

Armantier, O. and N. Treich (2015). The rich domain of risk. *Management Science*.

Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica 19*, 404–437.

Baillon, A. and L. Placido (2016). Testing constant absolute and relative ambiguity aversion. *Mimeo*.

Berger, L. (2014). Precautionary saving and the notion of ambiguity prudence. *Economics Letters 123*(2), 248 – 251.

Berger, L., J. Emmerling, and M. Tavoni (2016). Managing catastrophic climate risks under model uncertainty aversion. *Management Science*.

Bernoulli, J. (1713). *Ars conjectandi*. Impensis Thurnisiorum, fratrum. [English edition, 2006, The Art of Conjecturing, Together with Letter to a Friend on Sets in Court Tennis, JHU Press].

Brunnermeier, M. K. and S. Nagel (2008). Do wealth fluctuations generate time-varying risk aversion? micro-evidence on individuals? asset allocation (digest summary). *American Economic Review 98*(3), 713–736.

Budescu, D. V. and I. Fischer (2001). The same but different: an empirical investigation of the reducibility principle. *Journal of Behavioral Decision Making 14*(3), 187–206.

Budescu, D. V., A. K. Rantilla, H.-T. Yu, and T. M. Karelitz (2003). The effects of asymmetry among advisors on the aggregation of their opinions. *Organizational Behavior and Human Decision Processes 90*(1), 178–194.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011). Uncertainty averse preferences. *Journal of Economic Theory 146*(4), 1275–1330.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2013a). Ambiguity and robust statistics. *Journal of Economic Theory 148*, 974–1049.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2013b). Classical subjective expected utility. *Proceedings of the National Academy of Sciences 110*(17), 6754–6759.

Charness, G., U. Gneezy, and A. Imas (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization 87*, 43–51.

Chiappori, P.-A. and M. Paiella (2011). Relative risk aversion is constant: Evidence from panel data. *Journal of the European Economic Association 9*(6), 1021–1052.

Chow, C. C. and R. K. Sarin (2002). Known, unknown, and unknowable uncertainties. *Theory and Decision 52*(2), 127–138.

Dean, M. and P. Ortoleva (2015). Is it all connected? a testing ground for unified theories of behavioral economics phenomena. *Unpublished manuscript*.

Dimmock, S. G., R. Kouwenberg, and P. P. Wakker (2015). Ambiguity attitudes in a large representative sample. *Management Science*, null.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Eco-*

nomics 75, 643–669.

Ergin, H. and F. Gul (2009). A theory of subjective compound lotteries. *Journal of Economic Theory 144*(3), 899–929.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics 10*(2), 171–178.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory 118*(2), 133–173.

Gierlinger, J. and C. Gollier (2008, November). Socially efficient discounting under ambiguity aversion. Technical Report 561, IDEI, Toulouse, Toulouse.

Gilboa, I., A. Postlewaite, and D. Schmeidler (2012). Rationality of belief or: why savage's axioms are neither necessary nor sufficient for rationality. *Synthese 187*(1), 11–31.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics 18*(2), 141–154.

Grant, S. and B. Polak (2013). Mean-dispersion preferences and constant absolute uncertainty aversion. *Journal of Economic Theory 148*(4), 1361–1398.

Guerdjikova, A. and E. Sciubba (2015). Survival with ambiguity. *Journal of Economic Theory 155*, 50–94.

Halevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica 75*(2), 503–536.

Hansen, L. and T. Sargent (2008). *Robustness*. Princeton university press Princeton, NJ.

Hansen, L. P. (2014). Nobel lecture: Uncertainty outside and inside economic models. *Journal of Political Economy 122*(5), 945–987.

Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. *American Economic Review 91*(2), 60–66.

Hardy, G. H., E. M. Wright, D. R. Heath-Brown, and J. H. Silverman (1979). *An introduction to the theory of numbers*, Volume 4. Clarendon press Oxford.

Harrison, G. W., M. I. Lau, and E. E. Rutström (2007). Estimating risk attitudes in Denmark: A field experiment. *The Scandinavian Journal of Economics 109*(2), 341–368.

Harrison, G. W. and J. A. List (2004). Field experiments. *Journal of Economic literature 42*(4), 1009–1055.

Harrison, G. W., J. Martínez-Correa, and J. T. Swarthout (2015). Reduction of compound lotteries with objective probabilities: Theory and evidence. *Journal of Economic Behavior & Organization 119*, 32–55.

Harrison, G. W. and E. E. Rutström (2008). Risk aversion in the laboratory. In J. C. Cox and G. W. Harrison (Eds.), *Risk Aversion in Experiments (Research in Experimental Economics, Volume 12)*, Chapter 6, pp. 41–196. Emerald Group Publishing Limited.

Holt, C. A. and S. K. Laury (2002). Risk aversion and incentive effects. *American economic review 92*(5), 1644–1655.

Holt, C. A. and S. K. Laury (2005). Risk aversion and incentive effects: New data without order effects. *American Economic Review*, 902–904.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*, 1849–1892.

Knight, F. (1921). Risk, Uncertainty, and Profit, Boston, MA: Hart, Schaffner & Marx.

Laplace, P. S. (1814). *Théorie analytique des probabilités*. Courcier.

Luce, R. D. (1959). *Individual choice behavior*. New York: Wiley.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association 13*(6), 1022–1100.

Nau, R. F. (2006). Uncertainty aversion with second-order utilities and probabilities. *Management Science 52*(1), 136–145.

Prokosheva, S. (2016). Comparing decisions under compound risk and ambiguity: The importance of cognitive skills. *Journal of Behavioral and Experimental Economics*.

Saha, A. (1993). Expo-power utility: A "flexible" form for absolute and relative risk aversion. *American Journal of Agricultural Economics 75*(4), 905–913.

Savage, L. (1954). *The Foundations of Statistics.* New York: J. Wiley. second revised edition, 1972.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica 57*(3), 571–587.

Seo, K. (2009). Ambiguity and second-order belief. *Econometrica 77*(5), 1575–1605.

Trautmann, S. and G. van de Kuilen (2014). Ambiguity attitudes. In G. Keren and G. Wu (Eds.), *Blackwell Handbook of Judgment and Decision Making.* Blackwell.

von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior.* Princeton University Press.

Wald, A. (1950). *Statistical decision functions.* New York: John Wiley & Sons.

Yates, J. F. and L. G. Zukowski (1976). Characterization of ambiguity in decision making. *Behavioral Science*.

Supplemental material (online only)

S1 Beyond neutrality vs. non-neutrality

We here present the results from the RLP tasks when the distinction is made between three different attitudes towards the type of uncertainty i: aversion ($C1 > Ci$), neutrality ($C1 = Ci$), and loving ($C1 < Ci$), where $i = \{2, 3, 4\}$ represents compound risk, model uncertainty and ambiguity respectively.

Table S1: Association between attitudes towards ambiguity, compound risk and model uncertainty

		Compound risk			**Model uncertainty**			
Ambiguity		$C1 > C2$	$C1 = C2$	$C1 < C2$	$C1 > C3$	$C1 = C3$	$C1 < C3$	**Total**
		Main experiment (N=189)						
$C1 > C4$	Count	23	93	17	98	26	9	**133**
	Expected	*23.2*	*94.3*	*15.5*	*83*	*41.5*	*8.4*	
		(12.17%)	(49.21%)	(8.99%)	(51.85%)	(13.76%)	(4.76%)	(70.37%)
$C1 = C4$	Count	4	32	3	10	27	2	**39**
	Expected	*6.8*	*27.7*	*4.5*	*24.3*	*12.2*	*2.5*	
		(2.12%)	(16.93%)	(1.59%)	(5.29%)	(14.29%)	(1.06%)	(20.63%)
$C1 < C4$	Count	6	9	2	10	6	1	**17**
	Expected	*3*	*12.1*	*2*	*10.6*	*5.3*	*1.1*	
		(3.17%)	(4.76%)	(1.06%)	(5.29%)	(3.17%)	(0.53%)	(8.99%)
Total		**33**	**134**	**22**	**118**	**59**	**12**	**189**
		(17.46%)	(70.9%)	(11.64%)	(62.43%)	(31.22%)	(6.35%)	(100%)
		Fisher's exact test: 0.188			Fisher's exact test: 3.35e-7			
		Robustness round (N=91)						
$C1 > C4$	Count	19	22	7	35	7	6	**48**
	Expected	*13.2*	*24.8*	*10*	*22.7*	*14.2*	*11.1*	
		(20.88%)	(24.18%)	(7.69%)	(38.46%)	(7.69%)	(6.59%)	(52.75%)
$C1 = C4$	Count	3	22	3	5	19	4	**28**
	Expected	*7.7*	*14.5*	*5.8*	*13.2*	*8.3*	*6.5*	
		(3.30%)	(24.18%)	(3.30%)	(5.49%)	(20.88%)	(4.40%)	(30.77%)
$C1 < C4$	Count	3	3	9	3	1	11	**15**
	Expected	*4.1*	*7.7*	*3.1*	*7.1*	*4.5*	*3.5*	
		(3.30%)	(3.30%)	(9.89%)	(3.30%)	(1.10%)	(12.09%)	(16.48%)
Total		**25**	**47**	**19**	**43**	**27**	**21**	**91**
		(27.47%)	(51.65%)	(20.88%)	(47.25%)	(29.67%)	(23.08%)	(100%)
Relative frequencies in parentheses.		Fisher's exact test: 8.31e-5			Fisher's exact test: 4.161e-10			

Contingency. As can be observed from the upper panel of Table S1, among the 189 subjects participating in the main experiment, 133 (70.4%) are ambiguity averse, 134 (70.9%) are compound risk neutral, and 118 (62.4%) are more model uncertainty averse than risk averse. Among those individuals manifesting stronger aversion to model uncertainty than to risk, 83% (98 out of 118 subjects) also exhibit ambiguity aversion. Looking at compound risk attitude, we remark that 69.4% of our subjects (93 out of 134) who reduce compound risks are also ambiguity averse, suggesting separate attitudes towards these two types of uncertain situations. Comparing the observed frequencies with the expected ones under the null hypothesis of independence with respect to ambiguity attitude, we do not observe significant differences in the case of compound risk, but do observe differences in the case of model uncertainty. More specifically, we observe that the number of subjects exhibiting both ambiguity aversion and being more averse to model uncertainty than to risk increases by 18% compared to the null hypothesis of independence, while the number of subjects who are ambiguity neutral and manifest the same attitude towards model uncertainty and towards risk is more than twice the number under the null hypothesis. Interestingly, we do not observe any kind of pattern between ambiguity loving and either compound risk loving or having less aversion to model uncertainty than to risk. To test statistically the relationship between the attitudes towards the different types of uncertainty, we conducted two-by-two independence tests. Fisher exact tests (2-sided) confirm (p-value=3.3e-7) our predictions that attitude towards model uncertainty and towards ambiguity are tightly associated, while we cannot reject the independence hypothesis between the attitudes towards compound lottery and ambiguity (p-value=0.19).

Turning to the pool of policy makers (robustness round), we see from the lower panel of Table S1 that most of our subjects exhibit ambiguity aversion (52.7%), reduce compound risk (51.6%), and are more model uncertainty averse than risk averse (47.3%). As already observed in Appendix A, there is an association between ambiguity attitude and compound risk, but this association seems weaker than the one between ambiguity and model uncertainty. Comparing the observed frequencies with the ones obtained under the null hypothesis of independence, we observe that the number of ambiguity averse subjects that are also compound risk averse increases by 43.9%, while this number rises to 54.2% when ambiguity aversion is considered together with having a stronger aversion to model uncertainty than to risk. The same happens when considering neutrality: the observed frequency of both ambiguity and compound risk neutral subjects is increased by 51.7% with respect to the expected frequency under the null hypothesis of independence, while it is more than doubled (+128.9%) when considering ambiguity neutral subjects expressing the same aversion to risk and model uncertainty. Contrary to the results obtained in the main experiment with students, this association is also extended to those with uncertainty loving attitudes. Specifically, 9 subjects are observed to be both ambiguity and compound risk loving (compared to an expected count of 3.1 under the independence hypothesis), and 11 are observed to be both ambiguity loving and less averse to model uncertainty than to risk (3.5 under the independence hypothesis). The associations we found between the different attitudes towards different types of uncertainty are confirmed by Fisher's exact tests, which enable us to statistically reject the independence hypotheses between ambiguity and both compound risk and model uncertainty respectively.

Multinomial logistic analysis. To further investigate the association between the attitudes towards the different types of uncertainty, we run a couple of multinomial logistic regressions. The results are summarized in Table S1. We present, for both the main experiment and the robustness round, the relative risk ratios for the multinomial logit model, the standard errors and the bounds of the 95% confidence interval. The reference group is ambiguity neutrality ($C1 = C4$). The relative risk ratio in the first row for example compares compound risk aversion

Table S1: Characteristics of Ambiguity Attitude: Multinomial Logistic Regressions

		Relative Risk Ratio	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
		Main experiment (N=189)			
$C1 > C4$					
	$C1 > C2$	1.978	1.146	0.636	6.157
	$C1 < C2$	1.950	1.285	0.536	7.093
	$C1 > C3$	10.177***	4.386	4.373	23.682
	$C1 < C3$	4.673	3.872	0.921	23.710
$C1 < C4$					
	$C1 > C2$	5.333*	3.988	1.232	23.090
	$C1 < C2$	2.370	2.341	0.342	16.429
	$C1 > C3$	4.5*	2.859	1.295	15.633
	$C1 < C3$	2.25	2.937	0.174	29.055
		Robustness round (N=91)			
$C1 > C4$					
	$C1 > C2$	6.333***	4.374	1.636	24.515
	$C1 < C2$	2.333	1.757	0.533	10.209
	$C1 > C3$	19***	12.373	5.302	68.086
	$C1 < C3$	4.071	3.186	0.879	18.868
$C1 < C4$					
	$C1 > C2$	7.333	7.498	0.988	54.404
	$C1 < C2$	22***	19.961	3.716	130.238
	$C1 > C3$	11.4	14.357	0.966	134.545
	$C1 < C3$	52.25***	61.680	5.167	528.342

Notes: Multinomial logistic regressions. Dependent variable: Ambiguity attitude. Coefficients show effects relative to the excluded category "neutrality". * $p - value < 0.05$, ** $p - value < 0.01$

($C1 > C2$) to compound risk neutrality ($C1 = C2$) for being ambiguity averse ($C1 > C4$) relative to being ambiguity neutral. In this case, the relative risk ratio is not statistically different from 1. On the contrary, for being more averse to model uncertainty than to risk ($C1 > C3$) relative to expressing the same attitude towards the two types of uncertainty ($C1 = C3$), the relative risk for being ambiguity averse relative to being ambiguity neutral is expected to increase by a factor of 10.17. This ratio is furthermore higher than 1 and statistically significant, at a confidence level of 0.1%.

Altogether, the results presented in Table S1 reinforce the results obtained in the dichotomous case, in which we compare the cases of neutrality to non-neutrality. In the main experiment, only having a stronger aversion to model uncertainty than to risk explains ambiguity aversion relative to ambiguity neutrality with statistical significance, while a significant part of the association we found between compound risk and ambiguity comes from the ambiguity loving attitude relative to ambiguity neutrality, being explained by compound risk aversion relative to compound risk neutrality (relative risk ratio of 5.3). In the robustness round, the association between attitudes towards ambiguity and towards compound risk is significant, but weaker than the one between ambiguity and model uncertainty. In particular, for being ambiguity averse relative to being ambiguity neutral, the relative risk ratio switching from compound risk neutral to averse is 6.3, and is 19 for switching from having the same aversion towards model uncertainty as to risk, to having a higher aversion towards model uncertainty than to risk. Finally, as already observed in the contingency analysis, we also observe a significant association between the loving attitudes towards the different types of uncertainty.

In Figure S1, we present the predicted probabilities of exhibiting each type of attitude towards ambiguity (aversion, neutrality, loving in the column dimension), at each corresponding attitude towards compound risk (in blue) and model uncertainty (in red). To ease comparisons, we also provide the predicted probabilities of ambiguity attitudes irrespective of the attitudes towards compound risk and model uncertainty (dashed black lines).[35] The pattern we already

Figure S1: Adjusted predictions of model uncertainty and compound risk attitudes

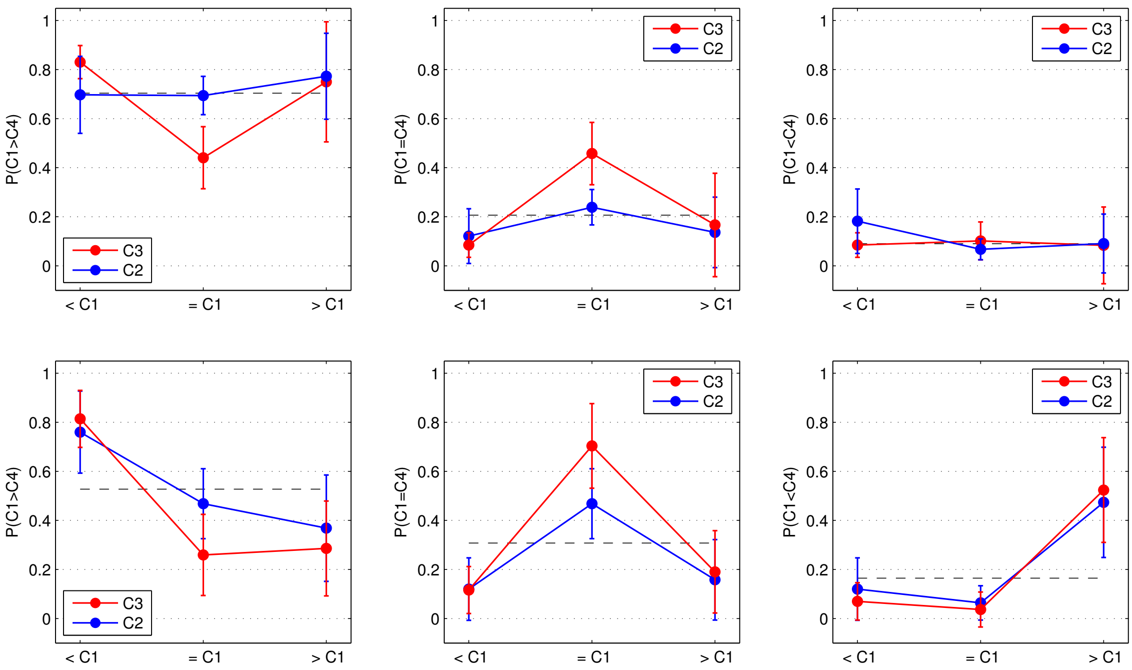

Notes: First row represents results of the main experiment (N=189), second row represents results of robustness round (N=91). Bars represent 95% confidence levels

observed from the analysis of the relative risk ratios is here made clearer. In the main experiment (first row), exhibiting a stronger aversion towards model uncertainty than towards risk increases the probability of being ambiguity averse to 83%, while the probability decreases to 44% if the subject exhibits the same attitude towards risk and model uncertainty. These numbers are significantly different from the 70.4% dashed line, which represents the probability of being ambiguity averse for the whole sample. Similarly we observe that having equal aversion to risk and model uncertainty increases the probability of being ambiguity neutral significantly (from 20.6% in the whole sample, to 45.8%). Compound risk attitude on the contrary does not seem to significantly effect the probabilities associated with the different ambiguity attitudes

[35]Remark that these probabilities exactly correspond to the total proportions of ambiguity averse, neutral and loving subjects provided in the last columns of Table S1.

(except maybe compound risk aversion which increases the probability of being ambiguity loving to 18.2%, as opposed to 9% for the whole sample). Turning to the robustness round, we observe a tight association between ambiguity attitudes and both compound risk attitudes and model uncertainty attitudes. However, the association is stronger when attitude towards model uncertainty is considered. The probability of exhibiting ambiguity aversion is for example 81.4% for a more model uncertainty averse than risk averse individual, as opposed to 52.8% for the whole sample. Similarly, the predictive probability of ambiguity neutrality goes from 30.8% for the whole sample to 70.4% once the subject exhibits equal attitude towards model uncertainty as towards risk. Finally, the association extends to the case of ambiguity loving, which goes from 16.5% to 52.4% once the subject becomes more risk averse than model uncertainty averse.

S2 Absolute and relative ambiguity aversion

When the functions characterizing risk and model uncertainty preferences are both of the expo-power type, and are respectively defined as $u(x) = \frac{1-\exp(-a_u x^{1-r_u})}{a_u}$ and $v(x) = \frac{1-\exp(-a_v x^{1-r_v})}{a_v}$, the ambiguity function $\phi \equiv v \circ u^{-1}$ may be written as:

$$\phi(U) = \frac{1 - \exp\Big(- a_v \big(- \frac{\ln(1-a_u U)}{a_u} \big)^{\frac{1-r_v}{1-r_u}} \Big)}{a_v},$$

where U belongs to the space of expected utilities. In that case, it may be shown that the Arrow-Pratt absolute ambiguity aversion index is:

$$-\frac{\phi''(U)}{\phi'(U)} = \frac{a_v \big(\frac{1-r_v}{1-r_u}\big) \Big(\frac{-\ln(1-a_u U)}{a_u}\Big)^{\frac{r_u - r_v}{1-r_u}} + \big(\frac{a_u}{\ln(1-a_u U)}\big) \frac{r_u - r_v}{1-r_u} - a_u}{1 - a_u U}. \tag{S2.1}$$

It is then easy to see that in the special case in which u and v are both of the CARA type (i.e. when $r_u = r_v = 0$), this index collapses to:

$$-\frac{\phi''(U)}{\phi'(U)} = \frac{a_v - a_u}{1 - a_u U}. \tag{S2.2}$$

This index is positive whenever $a_v > a_u$, so that absolute ambiguity aversion results from higher absolute model uncertainty aversion than absolute risk aversion. Similarly, in the special case in which u and v are both of the CRRA type (i.e. when $a_u = a_v = 0$), the absolute ambiguity aversion becomes:

$$-\frac{\phi''(U)}{\phi'(U)} = \frac{r_v - r_u}{(1 - r_u)U}, \tag{S2.3}$$

and is positive whenever $r_v > r_u$.

S3 Robust estimates of risk and model uncertainty preferences

Since each of our subjects provided multiple choices in the experiment, we may want to correct for the possible correlation of errors associated with a given subject (which may for example be due to unobserved individual effects). In this case, the residuals from the same subject are treated as potentially correlated, and the correction is made when calculating standard errors of estimates. As argued in Andersen et al. (2008), this procedure allows heteroskedasticity between and within clusters, as well as autocorrelation within clusters, and generalizes the "robust standard errors" approach popular in econometrics. The estimates of risk and model uncertainty aversion with robust standard errors are presented in Table S3.

Table S3: Robust estimates of risk, model uncertainty and ambiguity preferences

	u		v	
	CRRA	EP	CRRA	EP
a		0.0294^{***} (0.00398)		0.152^{***} (0.0252)
r	0.279^{***} (0.0281)	0.135^{***} (0.0278)	0.738^{***} (0.0444)	0.467^{***} (0.0842)
noise parameter	0.103^{***} (0.00378)	0.105^{***} (0.00384)	0.0358^{***} (0.00492)	0.0534^{***} (0.00643)
Observations	5320	5320	7570	7570 0
Loglikelihood	-1550.3	-1516.8	-3682.5	-3682.1

Notes: Estimation uses Luce error specification. Standard errors in parentheses. The EP risk specification is used to estimate v. * $p - value < 0.05$, ** $p - value < 0.01$, *** $p - value < 0.001$

S4 Direct comparison of objective and subjective probabilities

In practice ambiguity aversion is almost exclusively measured in Ellsberg like experiments, in which a situation with unknown probabilities is compared to a simple risk situation with known probabilities (Armantier and Treich, 2015). In this sense, our design does not constitute an exception to the common practice, since we in turn consider the pairwise comparisons of the different types of uncertain situations, holding constant the simple risk situation. To get a sense of how the results would be altered if the comparison was realized in a more direct way, we also asked our subjects to express their preferences directly between the compound risk and the model uncertainty situations. In practice, this means that we include a fourth RLP choice in which Urns 2 and 3 are directly compared (see Section 2.1 and Table **??** in the main body of the paper). The results are presented in association with the ambiguity attitude in Table S4 below. It should be mentioned that the association with ambiguity attitude should be treated with caution, given that the reference points are not the same in this case (Urns 1 and 4 are used to measure ambiguity attitude, Urns 2 and 3 are used to measure attitudes towards objective or subjective probabilities). However, the results are still informative. In particular, we see that about sixty percent (59.8% in the main experiment, and 62.6% in the robustness round) of our subjects prefer the situation in which the urn composition is determined by flipping a coin in the air, than the situation in which the urn composition is unknown but information concerning its composition is given by two experts. As expected in our theoretical predictions and suggested by the theoretical model we presented, the association between ambiguity attitude, and a different

Table S4: Association between ambiguity neutrality and difference between Urns 2 and 3

		Main experiment			**Robustness round**		
		$C2 = C3$			$C2 = C3$		
$C1 = C4$		No	Yes	**Total**	No	Yes	**Total**
No	Count	95	55	**150**	48	15	**63**
	Expected	*89.7*	*60.3*		*39.5*	*23.5*	
		(50.26%)	(29.1%)	(79.37%)	(52.75%)	(16.48%)	(69.23%)
Yes	Count	18	21	**39**	9	19	**28**
	Expected	*23.3*	*15.7*		*17.5*	*10.5*	
		(9.52%)	(11.11%)	(20.63%)	(9.89%)	(20.88%)	(30.77%)
Total		**113**	**76**	**189**	**57**	**34**	**91**
		(59.79%)	(40.21%)	(100%)	(62.64%)	(37.36%)	(100%)

Notes: Relative frequencies in parentheses. Chi-square test: 0.051 Chi-square test: 6.1e-5

attitude towards objective than subjective probabilities, seems to be tightly associated. This is particularly the case when analyzing the behavior of policy makers in the robustness round. In this case, 76% of the subjects (48 out of 63 subjects) non-neutral towards ambiguity are not neutral with regards to the type of probabilities (objective vs. subjective) they face in the second layer of uncertainty. In comparison with the expected frequency under a null hypothesis of independence, the observed number of subjects not indifferent between Urns 1-4 and Urns 2-3 increases by 22%. Similarly, out of the 28 subjects who expressed ambiguity neutrality, 19 of them (68%) were also indifferent between the coin and the expert determining the composition of the urn. This represents an increase of 81% in comparison with the expected frequency under a null hypothesis of independence. This result is confirmed by a Chi-square test rejecting (p-value=6.1e-5) the independence hypothesis between ambiguity neutrality and equal attitude towards objective and subjective probabilities ($C2 = C3$). In the case of the lab experiment, results are less clear cut but go in the same direction.

S5 Sensitivity of model uncertainty aversion

In order to get further insight regarding the sensitivity of model uncertainty aversion to risk aversion, we used the results from our PL tasks to estimate the degree of relative model uncertainty aversion alone, for values of relative risk aversion ranging from 0 to 0.75. It should be clear that in this situation the estimation procedure is not anymore jointly realized (given the exogenously chosen values of risk aversion considered). In Figure S2, we present the results of these estimations in the particular case where both u and v are of the CRRA type. For any value of relative risk aversion, we remark that the relative model uncertainty parameter estimated (represented in red) is significantly higher than the 45 degree line (in blue) representing relative risk aversion. In particular, it goes from $r_v = 0.66$ when $r_u = 0$ to $r_v = 0.96$ when $r_u = 0.75$. We also present in black the degree of relative ambiguity aversion implied by the twofold CRRA specification. Following Berger et al. (2016), the index of relative ambiguity aversion in this situation is given by $r_\phi = \frac{r_v - r_u}{1 - r_u}$. As can be observed, r_v and r_ϕ both coincide when risk

Figure S2: Relative model uncertainty aversion for different values of relative risk aversion when both u and v are of the CRRA type

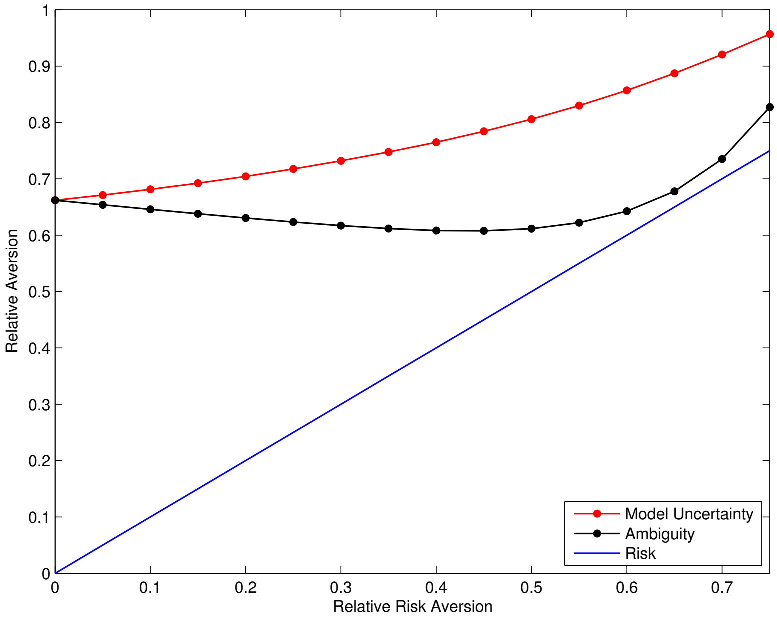

neutrality is considered (i.e. $r_u = 0$). Following the discussion provided in the main body of the paper, we claim that the level of ambiguity aversion is overestimated when measured in situations of risk neutrality, and remark that for positive values of r_u the index of relative ambiguity aversion is always lower than the one of relative model uncertainty aversion.