



HAL
open science

Three layers of uncertainty

Ilke Aydogan, Loïc Berger, Valentina Bosetti, Ning Liu

► **To cite this version:**

Ilke Aydogan, Loïc Berger, Valentina Bosetti, Ning Liu. Three layers of uncertainty. 2022. hal-03031751v2

HAL Id: hal-03031751

<https://hal.science/hal-03031751v2>

Preprint submitted on 15 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



November 2022

iRisk WORKING PAPER SERIES

2022-iRisk-01

Three layers of uncertainty

Ilke Aydogan

IESEG School of Management, Univ. Lille, CNRS, UMR 9221 - LEM - Lille Economic Management, F-59000 Lille, France; and iRisk Research Center on Risk and Uncertainty (i.aydogan@ieseg.fr).

Loïc Berger

CNRS, Univ. Lille, IESEG School of Management, UMR 9221 - LEM - Lille Economic Management, F-59000 Lille, France; iRisk Research Center on Risk and Uncertainty; RFF-CMCC European Institute on Economics and the Environment (EIEE), and Centro Euro-Mediterraneo sui Cambiamenti Climatici, Italy (l.berger@ieseg.fr).

Valentina Bosetti

Department of Economics and IGIER, Bocconi University, and RFF-CMCC European Institute on Economics and the Environment (EIEE), Centro Euro-Mediterraneo sui Cambiamenti Climatici, Italy (valentina.bosetti@unibocconi.it).

Ning Liu

School of Economics and Management, Beihang University and Laboratory for Low-carbon Intelligent Governance, Beihang University, China (nliu2018@buaa.edu.cn)

IESEG School of Management, Lille Catholic University, 3, rue de la Digue F-59000 Lille Tel: 33(0)3 20 54 58 92
www.ieseg.fr

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of IESEG School of Management or its partner institutions.

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorization of the author(s).

For all questions related to author rights and copyrights, please contact directly the author(s).

THREE LAYERS OF UNCERTAINTY*

Ilke Aydogan[†] Loïc Berger[‡] Valentina Bosetti[§] Ning Liu[¶]**Abstract**

We explore decision-making under uncertainty using a framework that decomposes uncertainty into three distinct layers: (1) *risk*, which entails inherent randomness within a given probability model; (2) *model ambiguity*, which entails uncertainty about the probability model to be used; and (3) *model misspecification*, which entails uncertainty about the presence of the correct probability model among the set of models considered. Using a new experimental design, we isolate and measure attitudes towards each layer separately. We conduct our experiment on three different subject pools and document the existence of a behavioral distinction between the three layers. In addition to providing new insights into the underlying processes behind ambiguity aversion, we provide the first empirical evidence of the role of model misspecification in decision-making under uncertainty.

Keywords: Ambiguity aversion, model uncertainty, model misspecification, non-expected utility, reduction of compound lotteries

JEL Classification: D81

*We thank Mohammed Abdellaoui, Aurélien Baillon, Han Bleichrodt, David Crainich, Philippe Colo, Thomas Epper, Itzhak Gilboa, Lars Peter Hansen, Peter Klibanoff, Olivier L’Haridon, Pietro Ortoleva, Vincent Théroude, Peter Wakker, and especially Massimo Marinacci for helpful comments and suggestions. We also thank seminar and conference participants at PSE, Bocconi University, HEC Paris, LEM, Paris-8, D-TEA, Rislab and FUR for insightful comments and discussions. We thank Diana Valerio for her help with the experiment. This research is supported by a grant from the French Agence Nationale de la Recherche (ANR-17-CE03-0008-01 INDUCED), the Region Hauts-de-France (2021.00865 CLAM), the I-SITE UNLE (project IBEBACC), the European Research Council (ERC-2013-StG 336703-RISICO), and the National Natural Science Foundation of China (Grant 72103015). Logistic supports from the Bocconi Experimental Laboratory in the Social Sciences (BELSS) and from the Anthro- Lab are kindly acknowledged.

[†]IESEG School of Management, Univ. Lille, CNRS, UMR 9221 - LEM - Lille Économie Management, F-59000 Lille, France; and iRisk Research Center on Risk and Uncertainty (i.aydogan@ieseg.fr).

[‡]CNRS, Univ. Lille, IESEG School of Management, UMR 9221 - LEM - Lille Économie Management, F-59000 Lille, France; iRisk Research Center on Risk and Uncertainty; RFF-CMCC European Institute on Economics and the Environment (EIEE), and Centro Euro-Mediterraneo sui Cambiamenti Climatici, Italy (l.berger@ieseg.fr).

[§]Department of Economics and IGIER, Bocconi University, and RFF-CMCC European Institute on Economics and the Environment (EIEE), Centro Euro-Mediterraneo sui Cambiamenti Climatici, Italy (valentina.bosetti@unibocconi.it).

[¶]School of Economics and Management, Beihang University and Laboratory for Low-carbon Intelligent Governance, Beihang University, China (nliu2018@buaa.edu.cn).

1 Introduction

Uncertainty is pervasive and plays a major role in economics. Whether economic agents pursuing individual goals, or policymakers pursuing social objectives, decision-makers (DMs) rarely have complete information about the likelihoods of the relevant states of the world. A valid understanding of individual behavior in the face of uncertainty is therefore of great importance for the construction of realistic economic models capable of making accurate predictions, as well as for prescriptive applications guiding decision-making processes.

Model uncertainty Uncertainty can take many forms. Most economic applications typically focus on the notion of *risk*, in which the DM knows the correct *model* that quantifies the uncertainty about the possible states of the world but not necessarily the correct state. A model is viewed here as a probability distribution over states that governs the exogenous contingencies on which the consequences of a decision may depend. It is typically induced by a mechanism that represents some natural or social phenomenon of interest. For example, a model may be used to predict inflation rates in macroeconomics, the temperature response to increased atmospheric CO₂ concentration in climate science, or the effective reproduction number of a virus in epidemiology.

In most real-life decision problems, including the examples above, the assumption that the DM knows the correct model is hardly satisfied. In particular, DMs typically do not know the exact mechanism at play, thus giving rise to model uncertainty (Marinacci, 2015). For example, there might exist different models trying to describe the same phenomenon or, because each model is by design a simplification, there might be concerns about the specifications of the models themselves and the way they describe the regular features of the phenomenon. How do DMs react to the existence of alternative models? Does the approximate nature of models affect their decisions? The aim of this paper is to explore and measure attitudes towards model uncertainty.

As is the case in many real-life situations, we consider situations in which the DM is able to formulate models but is uncertain about the “correct” one. In line with Hansen and Sargent (2022), these formulated models are referred to as “structured models”. They are usually explicitly featured because they possess a substantive motivation. For example, they may be based on scientific knowledge or well-behaved statistical distributions, relying on empirical

evidence or on theoretical arguments. Yet, as the structured models are possibly misspecified, it might be the case that the DM decides to consider a potentially richer collection of probability distributions to characterize the phenomenon of interest. Such alternative models, which do not possess a formal substantive motivation, are referred to as “unstructured models”. To illustrate, consider the situation of a policymaker having to decide a climate policy based on existing alternative climate models. Each structured model is characterized by a given probability distribution over the long-term temperature response to CO₂ emission pathways. As multiple instances of this distribution exist (depending on the modeling assumptions made by climate scientists or the type of data used to estimate the probabilistic relationship), there are both uncertainty *across* structured models and uncertainty *about* the models themselves, thus leaving the possibility to the policymaker to consider alternative unstructured models.¹ In this paper, we draw on the crucial distinction between *structured* and *unstructured* models to decompose uncertainty into different layers of analysis.

Three layers Building on the early insights of Arrow (1951) and the recent contributions of Hansen (2014), Marinacci (2015), and Hansen and Marinacci (2016), our investigation focuses on a decomposition of uncertainty into three distinct *layers*: (i) *risk*, in which the uncertainty is about the possible states (or outcomes) within a given probability model; (ii) *model ambiguity*, in which the uncertainty is about which alternative probability model one should use among a posited set of structured models; and (iii) *model misspecification*, in which there is uncertainty regarding whether the correct probability model belongs to the set of structured models or not.

More specifically, the first layer of risk characterizes situations in which the consequences of the DM’s actions depend on states of the world over which there is an *objectively* known probability distribution. This uncertainty about states represents the inherent variability *within* a particular probability model and, as such, is considered as being of an aleatory type, analogous to chance mechanisms. The extra layer of model ambiguity arises when the DM is not able to identify the correct model among the set of structured ones. This uncertainty *across* models has an epistemic nature, which may be quantified by means of *subjective* probabilities (Marinacci, 2015). Finally,

¹For a further discussion of model uncertainty in the context of climate change economics, see Berger and Marinacci (2020).

because models are, by design, approximations of more complex phenomena, they are often misspecified. In consequence, the set of structured models under consideration might not include the correct model, thus giving rise to the third layer, or epistemic uncertainty *about* models being correct (Hansen 2014; Hansen and Marinacci 2016, see also White 1982).

Figure 1 illustrates situations with different layers of uncertainty. Situations (a) and (b) encompass the layer of risk only, which is presented in a single stage and in two stages, respectively. Situation (c) encompasses both the layers of model ambiguity and risk. In that case, no objective probabilities can be assigned to the two structured models. Situation (d) is an instance encompassing all the three layers together: in addition to risk and model ambiguity, there is uncertainty about the set of models to be considered.

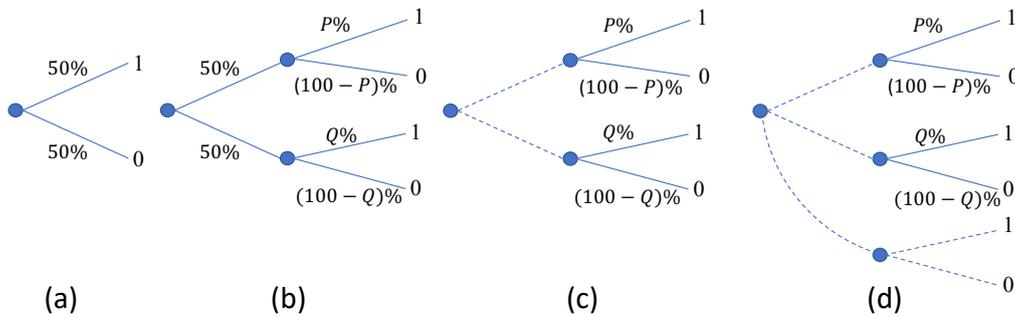


Figure 1: Situations with different layers of uncertainty.

Notes: Branches whose probabilities can be defined with an “objective” measure (i.e., aleatory uncertainty) are represented with solid lines. Branches whose probabilities cannot be defined with an objective measure (i.e., epistemic uncertainty) are represented with dashed lines, without any probability attached.

The three layers of uncertainty are inherent in any decision problem under uncertainty in which a DM adopts probabilistic theories about the outcomes of a phenomenon and forms beliefs about their relevance. Therefore, they provide a useful framework to analyze the vast majority of real-life decision problems under ambiguity.² An example is when different experts provide opinions about the probability of an event (e.g., developing a disease, fire risk in buildings, aircraft accidents, etc.). From the DM’s perspective, each expert’s probability model can be interpreted as an instance of risk. As experts may provide different assessments, the second layer of model ambiguity emerges regarding which expert’s model to rely on. Finally, the third layer

²Since Ellsberg (1961), the term “ambiguity” has emerged in the literature to characterize situations in which probabilities are unknown (Knightian uncertainty). Accordingly, situations encompassing the layers of model ambiguity or model misspecification are often called *ambiguous*.

of model misspecification captures the possibility that none of the experts consulted is correct. Such a decomposition of uncertainty into different layers has been recently used to describe the problem of a policymaker during a pandemic (Berger et al., 2021), and in the context of climate change (Brock and Hansen, 2017; Barnett et al., 2020; Berger and Marinacci, 2020).

This article This paper presents an empirical investigation of attitudes towards each of the three layers of uncertainty. As previous experimental research under the standard Ellsberg (1961) paradigm has so far concerned the layers of risk and model ambiguity exclusively, this paper is also, to our knowledge, the first examination of model misspecification in a laboratory environment.³ Indeed, the ambiguous two-color Ellsberg urn containing N balls provides implicitly $N + 1$ potential compositions, each of which constitutes a risk, whereas the distribution over the compositions, unknown to the DM, relates to model ambiguity. However, as the set of possible distributions is fully specified in an Ellsberg urn, such a setting necessarily leaves out the issue of model misspecification. For that reason, our experiment entails a new, *extended* Ellsberg setting, in which the number of possible compositions is first left unspecified. Then, by changing the information about the possible compositions, we are able to implement different situations of (compound) risk, model ambiguity, and model misspecification.

Our design enables us to isolate the effect of model ambiguity from that of risk by comparing situations characterized by a known probability distribution over possible urn compositions with situations in which this distribution is unknown (e.g., comparing situations (b) and (c) in Figure 1). We further isolate the effect due to model misspecification by considering situations in which alternative compositions, outside the set of structured models, cannot be excluded (e.g., comparing situations (c) and (d)).

Our investigation complements and extends previous empirical research on the multi-stage presentation of uncertainty and its relation to ambiguity (e.g., Halevy, 2007; Abdellaoui et al., 2015; Chew et al., 2017, see also the discussion in Section 7). In what follows, we specifically distinguish the *stages* and *layers* of uncertainty. In particular, whereas our instances of compound risk, model ambiguity and model misspecification are all presented in two *stages*, they

³Note that there also exists a recent, but distinct, experimental literature on learning with misspecified models (e.g., Esponda et al., 2020; Götte et al., 2020). This literature focuses on the behavior of agents who behave optimally but learn with a possibly misspecified model (Berk, 1966; Esponda and Pouzo, 2016).

differ in terms of the *layers* of uncertainty they encompass (see Figure 1). Our conjecture is that there may exist distinct attitudes towards different layers of uncertainty, beyond those towards the multi-stage presentation. If this is the case, taking into account and calibrating these attitudes could reveal essential for analyzing a vast majority of applied decision problems under uncertainty.

We test our conjecture by conducting three experiments. The main experiment is run as a standard laboratory experiment on a pool of university students. The experimental results show that there exists a distinction between both stages and layers of uncertainty. Specifically, while we find that students are typically averse to uncertainty presented in multiple stages (i.e., situations (b), (c), and (d) in Figure 1), they also treat situations with different layers of uncertainty differently. In particular, our subjects are willing to pay on average 8.4% of their expected payoff to avoid being faced with the layer of model ambiguity, and an extra 5.3% to avoid the layer of model misspecification. We explore the robustness of these findings using two follow-up experiments. First, we run the experiment on a pool of risk professionals to understand to what extent our results depend on the specificity of our subject pool and its potential limitations to deal with the relevant complexity of the multi-stage presentation of uncertainty. Second, we run the experiment online using a between-subject (rather than within-subject) design to test the comparative ignorance hypothesis of Fox and Tversky (1995) while evaluating situations with more or less layers of uncertainty. The follow-up experiments support our main conclusions and rule out several alternative explanations. On the one hand, the experiment with risk professionals shows that the role played by the layers in explaining overall uncertainty attitudes is stable when considering a more sophisticated subject pool. On the other hand, the online experiment shows that the results obtained are not due to order effects or contagion across treatments. Overall, the two follow-up experiments thus provide further evidence in favor of the behavioral distinction between the three layers.

Organization The paper is organized as follows. Section 2 presents the main features of our experiment, which is designed to study attitudes towards the three layers of uncertainty. Section 3 introduces the notions of total and differential premia that are used to analyze the attitudes elicited. Section 4 considers different modern theories of choice under uncertainty, discusses how they deal with the three layers, and states our predictions. The results of the

main experiment, run in the lab on students, are presented in Section 5. The results of the two follow-up experiments are summarized in Section 6. We discuss our results in relation to the extant literature and conclude in Section 7.

2 Experiments

We examine choices under different *sources* of uncertainty that potentially encompass different layers.⁴ We run three distinct experiments with the same stimuli. The main experiment took place in a laboratory with university students, whereas the follow-up experiments took place, respectively, in the field with risk professionals, and online. All experiments used real monetary incentives.

2.1 The sources of uncertainty

We consider five sources of uncertainty. These sources are constructed in an *extended* Ellsberg two-color setting using decks, from which a card is randomly drawn. All the decks contain an unspecified number of cards.⁵ The sources are characterized by different deck compositions, defined in terms of their proportion of black cards (and the complementary proportion of red cards).

1. *Simple risk*, denoted SR , entails a deck containing an equal proportion of black and red cards;
2. *Compound risk*, denoted CR , entails a deck that contains either $P\%$ or $Q\%$ black cards, with equal probability;
3. *Model ambiguity*, denoted MA , entails a deck that contains either $P\%$ or $Q\%$ black cards, with unknown probability;
4. *Model misspecification*, denoted MM , entails a deck that is likely (i.e., with at least 50% probability) to contain either $P\%$ or $Q\%$ black cards, but may also contain another (unknown) proportion of black cards;

⁴We refer to sources of uncertainty as “groups of events that are generated by the same mechanism of uncertainty, which implies that they have similar characteristics” (Abdellaoui et al., 2011, p. 696).

⁵For the sake of comprehensiveness and to allow comparisons with previous literature, we also consider the standard two-color Ellsberg ambiguous situation in which the ambiguous deck contains 100 cards. This source was always presented at last to prevent a priming effect about the number of cards in the decks. A discussion of this extra source is presented in the Online Appendix.

5. *Extended Ellsberg*, denoted *EE*, entails a deck that contains an unknown proportion of black cards.

As illustrated in Figure 1, the sources *CR*, *MA*, and *MM* are all presented in two stages, while they differ in terms of the layers of uncertainty they encompass. The source *CR* entails only the layer of risk, although it is presented in a compound way. Under *CR*, the two possible deck compositions, $P\%$ and $Q\%$ black cards, are each unambiguously assigned an objective probability of 50%. Conversely, the source *MA* entails both the layers of model ambiguity and risk. Under *MA*, the two possible deck compositions can only be assigned *subjective* probabilities.⁶ The source *MM* entails all the three layers together, adding the extra layer of model misspecification as follows: Although the probability distribution is likely to be characterized by one of the two compositions as in *MA*, we cannot exclude the possibility of an alternative composition.⁷ Finally, whereas *SR* is an instance of single-stage risk, *EE* corresponds, in spirit, to Ellsberg’s (1961) ambiguous situation in which “numerical probabilities are inapplicable.”

2.2 Procedure in the main experiment

The main experiment was run on computers at the Bocconi Experimental Laboratory for Social Sciences (Italy) using a within-subject design. Subjects were seated in cubicles and could not communicate with one another during the experiment. Each session started with the experimental instructions, examples of the stimuli, and comprehension questions. A typical session lasted approximately one hour, including instructions and payment. Complete instructions are provided in the Online Appendix.

Subjects Five experimental sessions were organized, with a total of 125 university students (average age 20.5 years, 52 women), recruited on a voluntary basis.

⁶Under a symmetry assumption, these subjective probabilities might be assumed to be 50%. Such a symmetry assumption has been supported empirically in various studies (e.g., Abdellaoui et al., 2011; Chew et al., 2017; Epstein and Halevy, 2019), as well as in an additional experiment that we conducted under the same conditions as in our main experiment. Specifically, in an additional experiment, we show that the symmetry assumption holds both at the aggregate and individual level (see Online Appendix S3).

⁷It must be clear that a crucial distinction between *MA* and *MM* is that the set of probability models includes only explicitly featured *structured* models in *MA*, whereas this is not the case in *MM*. Specifically, to explore misspecification, we allow for the possibility to consider a potentially richer collection of probability distributions, which includes *unstructured* models lacking substantive motivation.

Stimuli Subjects faced monetary prospects under the five sources of uncertainty introduced previously. For *CR*, *MA*, and *MM*, we considered two cases with different sets of structured models. In each case, the proportion of black cards is either $P\%$ or $Q\%$, where $Q = 100 - P$. In the first case, $P = 0$ and in the second $P = 25$. For example, *CR* with $P = 0$ entails a deck that contains either 0% black or 100% black cards with equal probability, whereas *CR* with $P = 25$ entails a deck that contains either 25% black or 75% black cards with equal probability. *MA* and *MM* are built analogously. We denote the respective cases as *CR0*, *CR25*, and so forth. Thus, there were eight monetary prospects in total: *SR*, *CR0*, *CR25*, *MA0*, *MA25*, *MM0*, *MM25*, and *EE*. Each prospect gave the subjects either €20 or €0, depending on the color of a card randomly drawn from a deck. In every prospect, the color giving €20 (black or red) was selected by the subjects themselves.

Table 1 summarizes the main characteristics of the eight prospects. They

Table 1: CHARACTERISTICS OF THE PROSPECTS USED IN THE EXPERIMENT

Source	Prospect	# of layers of uncertainty	Set of structured models (M)	Information available (μ) ^a
<i>SR</i>	<i>SR</i>	1	{50%}	$\mu(50\%) = 1$
<i>CR</i>	<i>CR0</i>	1	{0%, 100%}	$\mu(0\%) = \mu(100\%) = 0.5$
	<i>CR25</i>	1	{25%, 75%}	$\mu(25\%) = \mu(75\%) = 0.5$
<i>MA</i>	<i>MA0</i>	2	{0%, 100%}	$\mu(0\%) + \mu(100\%) = 1$
	<i>MA25</i>	2	{25%, 75%}	$\mu(25\%) + \mu(75\%) = 1$
<i>MM</i>	<i>MM0</i>	3	{0%, 100%}	$\mu(0\%) + \mu(100\%) \geq 0.5$
	<i>MM25</i>	3	{25%, 75%}	$\mu(25\%) + \mu(75\%) \geq 0.5$
<i>EE</i>	<i>EE</i>	3	/	/

^a $\mu(m)$ represents the subjective prior probability attached to the structured model $m \in M$. Note that under rational expectations, subjective and ‘true’ probabilities coincide (i.e., for *SR* and *CR*).

were implemented as follows: In *SR*, the subjects were instructed that the deck contained an equal proportion of red and black cards. In the cases of *CR*, *MA*, *MM*, and *EE*, subjects were instructed that the deck would be picked randomly from a pile of decks. In *CR0* (*CR25*), the pile consisted of decks containing 0% (25%) black cards and decks containing 100% (75%) black cards, with an equal proportion of each. In *MA0* (*MA25*), the pile also consisted of decks containing 0% (25%) black cards and decks containing 100% (75%) black cards, but with an unknown proportion of each. In *MM0* (*MM25*), the subjects were instructed that at least half of the pile consisted of

decks containing 0% (25%) black and decks containing 100% (75%) black cards with an unknown proportion of each. Notably, the subjects were also informed that the pile considered may (or may not) contain decks whose composition is different than the ones described. In *EE*, the pile consisted of decks containing red and black cards, each with an unknown composition.

All the decks and piles were constructed in advance by one of the collaborators, who was not present in the room during the experimental sessions.⁸ The subjects were also reminded that they could check the piles and the decks at the end of the experiment to verify the truthfulness of the information provided. During the instructions, subjects were given examples of the different sources of uncertainty that they would face throughout the experiment. We tested their understanding of the differences between the sources through comprehension questions, for which they were given automatic feedback.

We elicited the certainty equivalents (CEs) of the eight prospects using a choice list design.⁹ Specifically, in each prospect, the subjects were asked to make 12 binary choices between the prospect of receiving €20 and receiving a sure monetary amount ranging from €0 to €20. The sure amounts were incremented by €2 between €1 and €19, and the order of the prospects was randomized. After completing the choice lists, the subjects answered a short sociodemographic survey.

Incentives Subjects received a €5 show-up fee. In addition, they received a variable amount depending on one of the choices they made during the experiment (i.e., random incentive). The choice situation on which the payment was based was the same for all the subjects in the same session. In practice, the 12 binary choice questions of the choice lists and the descriptions of the uncertain situations were printed on paper and physically enclosed in sealed envelopes before every session. In each session, a volunteer from the subject pool randomly picked two envelopes before the experiment started: one con-

⁸Thus, no one in the room, including the experimenters, had any additional information about the content of the decks and piles, other than what was described in the experimental instructions. The subjects were informed accordingly to prevent the effects of comparative ignorance, according to which a comparison with a more knowledgeable individual (in this case, the experimenter) may induce ambiguity aversion (Fox and Tversky, 1995).

⁹The validity of CE elicitation has been a topic of debate due to anomalies such as preference reversals (e.g., Grether and Plott, 1979; Hara et al., 2019). The consensus in the experimental literature is that choice-based elicitation of CEs (as in choice list designs) leads to more reliable measurements than asking subjects directly for their CEs by generating fewer such inconsistencies (Bostic et al., 1990; Attema and Brouwer, 2013). For a more detailed discussion, see Online Appendix S5.

taining an uncertain situation and the other containing a question from the choice lists. After the two envelopes had been picked, they were attached, still sealed, to a white board visible to all participants. The subjects were informed that the choice situation that would matter for their payment was contained in the envelopes, which would remain closed and visible until the end of the experiment. When all subjects had completed the questionnaire, the envelopes were opened and their contents were revealed. The draws from the piles and/or from the decks were made according to the uncertain situation contained in the first envelope and the subjects were paid according to their response to the choice question contained in the second envelope.¹⁰

2.3 Procedure in follow-up experiments

Field experiment with actuaries The subject pool in the first follow-up experiment consists of 84 risk professionals (average age 40 years, 37 women), who attended the 31st International Congress of Actuaries in Berlin, Germany. The majority of the subjects were highly educated in the fields of mathematics, statistics, and actuarial science. Subjects had an average of 13 years of work experience in the insurance and finance industries. They faced the same stimuli as in the lab experiment, except that the stakes were multiplied by a factor of 10 (a correct bet yielded €200, instead of €20). To reduce the expected experimental costs and monetary transactions during the conference schedule, only a fraction of the subjects (one out of 10) was paid based on one of their choices. Experimental details are provided in Online Appendix S1.

Online experiment The second follow-up experiment adopts a between-subject (rather than within-subject) design while mostly using the same stimuli as in the main experiment.¹¹ The experiment was conducted in the online platform Prolific. The subject pool consists of the members of the platform from all over the world. To remain as close as possible to the main experiment, we recruited subjects who had declared a student status and were currently

¹⁰The random incentive system (RIS) is one of the most commonly used mechanisms for individual choice experiments in economics. The *prior* incentive system of Johnson et al. (2021) that we use performs the randomization *before*, rather than after, the choices and the resolutions of uncertainty. Hence it aims to reinforce the isolation assumption of Kahneman and Tversky (1979), which is equivalent to monotonicity and therefore sufficient to guarantee the incentive compatibility of the RIS (see Azrieli et al., 2018; Baillon et al., 2022).

¹¹The main difference concerns the implementation of the choice lists. In this follow-up experiment, we use (i) smaller incremental steps for sure amounts on the list (£1 between £0.5 and £19.5), which increase the level of resolution in CE elicitation, and (ii) an automatic filling system.

enrolled in either undergraduate, graduate or doctorate degree. Subjects were randomly assigned to one of the following three treatments: (1) Compound Risk ($N = 277$), (2) Model Ambiguity ($N = 229$), and (3) Model Misspecification ($N = 234$). Each subject faced the two corresponding prospects of their treatment (CR , MA or MM) with $P = 0$ and $P = 25$, as well as the sources R and EE . The stake size of the prospects was £20. Experimental details are given in Online Appendix S2.

3 Measures

In the analysis that follows, we take the midpoint of an indifference interval implied by the switching point in the choice list as a proxy for the CE of the prospect. Switching in the middle of the list implies a CE equal to the expected value (EV) of the prospect, which is €10 in the lab experiment (€100 in the field experiment and £10 in the online experiment).

3.1 Total uncertainty premia

We use the following notions of uncertainty premium to analyze our results.

Definition 1. The *total uncertainty premium* Π_i is defined as

$$\Pi_i \equiv EV_i - CE_i$$

for all $i \in \{SR, CR0, CR25, MA0, MA25, MM0, MM25, EE\}$.

In words, the total premium represents the amount of money that an individual is willing to pay to receive the expected value of the prospect with certainty, rather than facing the uncertainty. The premium is positive (resp. zero, or negative) when a subject is averse (resp. neutral, or loving) to the uncertainty in prospect i . The most well-known total uncertainty premium is the standard risk premium, noted Π_{SR} .

3.2 Differential uncertainty premia

Because we are interested in comparing the way people react to different layers of uncertainty, a more directly relevant measure for our purposes is the notion of differential premium.

Definition 2. The *differential uncertainty premium* $\Pi_{i,j}$ is defined as

$$\Pi_{i,j} \equiv \Pi_j - \Pi_i$$

for all $i, j \in \{SR, CR0, CR25, MA0, MA25, MM0, MM25, EE\}$.

Given that all the prospects in our experiment have the same EV under symmetry, the differential premium may equivalently be expressed as the difference between the CEs of the prospects i and j :

$$\Pi_{i,j} = CE_i - CE_j.$$

The differential premium is positive (resp. zero, or negative) when a subject is more (resp. as much, or less) averse to the uncertainty in prospect j than that in prospect i . The version of this premium that has been typically considered in the literature is the one relative to simple risk, so that $\Pi_{SR,j}$ refers to the so-called *compound risk premium* when $j \in \{CR0, CR25\}$ (see, e.g., Abdellaoui et al., 2015), and to the *ambiguity premium* when $j \in \{MA0, MA25, MM0, MM25, EE\}$ (see, e.g., Maccheroni et al., 2013). In our experiment, two other differential premia will be of special importance for analyzing the results. The first, which is called the *MA (model ambiguity)-premium*, refers to the marginal effect between the first and the second layers of uncertainty. It measures what an individual is ready to pay to avoid being confronted with epistemic uncertainty in the first stage. To isolate the effect of model ambiguity alone (filtering out the effect due to the multi-stage presentation), it is measured in relation to compound risk, by $\Pi_{CR0,MA0}$ and $\Pi_{CR25,MA25}$. The second important premium for our analysis, called the *MM (model misspecification)-premium*, refers to the marginal effect between the second and the third layers of uncertainty. It measures what an individual is ready to pay to be sure that the correct model belongs to the set of structured models. It is measured by $\Pi_{MA0,MM0}$ or $\Pi_{MA25,MM25}$. Overall, the total effect due to model uncertainty can thus be measured by $\Pi_{CR,MM} \equiv \Pi_{CR,MA} + \Pi_{MA,MM}$.

4 Theory and predictions

We can now use these premia to make predictions following some prominent theories of decision-making under uncertainty that accommodate a multi-stage representation. After introducing basic concepts and notation, we use

the sources CR , MA , and MM to draw general predictions under four general hypotheses.¹²

4.1 Preliminaries

Let S denote the set of states of the world and C the set of consequences. Formally, a prospect (or an act) is a function $\mathcal{P} : S \rightarrow C$ mapping states into consequences. We consider a DM who has a complete and transitive preference relation \succsim over prospects. We consider a classic setup in the spirit of Wald (1950), in which M is a set of structured models m representing the probabilistic behavior of a phenomenon of interest. Each model m thus characterizes the inherent randomness governing states' realizations, while the set M is taken as a datum of the decision problem.¹³ In line with Wald (1950), most of the classical frameworks that we present make the assumption that the DM knows that the true model belongs to M , and so face model ambiguity and risk only (for an in-depth review, see Marinacci, 2015). An exception may be found in the recent work of Cerreia-Vioglio et al. (2022), in which model misspecification potentially arises by removing the assumption that the correct model belongs to M .

4.2 Expected Utility hypothesis

Traditionally, economists have dealt with uncertainty by following the subjective expected utility (SEU) approach (Savage 1954). In line with the Bayesian tradition, this approach favors quantifying uncertainty in probabilistic terms and treats any source of uncertainty as *risk*, reducing uncertainty *de facto* to its first layer. A two-layer version of SEU distinguishing risk from model ambiguity has been axiomatized by Cerreia-Vioglio et al. (2013b). In this version, it is assumed that the DM has a subjective prior probability measure $\mu : 2^M \rightarrow [0, 1]$ quantifying the epistemic uncertainty about the models $m \in M$. The subjective probabilities (or priors) reflect the structural information received and some personal information that the DM may have on the models. The SEU of a prospect is

¹²Interested readers will find further theoretical developments and predictions in the Online Appendix S4.

¹³In particular, following Wald (1950) and Marinacci (2015), the physical information M is taken as a primitive of the decision-making problem. Formally, we assume the existence of a measurable space (S, Σ) , where Σ is an algebra of events of S . A model $m : \Sigma \rightarrow [0, 1]$ is thus a probability measure, and the collection M is a finite subset of $\Delta(S)$, the collection of all probability measures.

$$V_{\text{SEU}}(\mathcal{P}) = \mathbb{E}_\mu(\mathbb{E}_m u(\mathcal{P})), \quad (1)$$

where $u : C \rightarrow \mathbb{R}$ is the von Neumann-Morgenstern utility function capturing risk attitude and \mathbb{E} is the expectation operator (taken with respect to the measures μ and m , respectively).¹⁴ Criterion (1) is a two-stage criterion that describes the different layers of uncertainty through standard probability measures. The same attitude is considered towards the layers of risk and model ambiguity, therefore implying compound risk and ambiguity neutrality. The expected utility hypothesis can thus be summarized as follows.

EU hypothesis:

$$\Pi_{SR,CR} = \Pi_{SR,MA} = \Pi_{SR,MM} = 0. \quad (2)$$

In words, distinct *layers* of uncertainty are treated in the same way, and different *stages* are reduced to a single one under the EU hypothesis.¹⁵ In consequence, this hypothesis also predicts zero *MA* and *MM* premia, i.e., $\Pi_{CR,MA} = 0$ and $\Pi_{MA,MM} = 0$.

4.3 Layer hypothesis

Next, we consider families of approaches that allow for different attitudes towards different *layers* of uncertainty. We start by focusing on the theories considering the layers of risk and model ambiguity only and alternatively present the maxmin criteria of Wald (1950) and Gilboa and Schmeidler (1989), and the smooth criterion of Klibanoff, Marinacci, and Mukerji (2005). We then present a general framework that incorporates the fear of model misspecification.

Wald The decision criterion of Wald (1950) considers only the worst possible model among all the *structured* ones in M :

$$V_{\text{Wald}}(\mathcal{P}) = \min_{m \in M} \mathbb{E}_m u(\mathcal{P}). \quad (3)$$

The layer of risk is not affected by the extreme cautiousness entailed by this criterion. For example, it is perfectly conceivable for a DM to be neutral to

¹⁴In other words, expression (1) corresponds to $V_{\text{SEU}}(\mathcal{P}) = \sum_m \mu(m) (\sum_s p(s|m) u(\mathcal{P}(s)))$ in its discrete version, where $p(s|m)$ is the objective probability of state s conditional on model m , and to $V_{\text{SEU}}(\mathcal{P}) = \int_M (\int_S u(\mathcal{P}(s)) dm(s)) d\mu(m)$ in its continuous version.

¹⁵Note that, while there exists no formal three-layer version of the SEU, the last equality in (2) directly follows from the symmetry assumption, which would lead to treat any set of unstructured models symmetrically.

compound risk while being extremely averse to model ambiguity.

Multiple priors The multiple priors (MP) criterion axiomatized by Gilboa and Schmeidler (1989) is less extreme than that of Wald. In this framework, the DM’s information about epistemic uncertainty is quantified by a compact set $\mathcal{C} \subseteq \Delta(M)$ of priors μ over the structured models, and the decision is based on the prior giving the least favorable SEU. The two-layer version of this criterion is written¹⁶

$$V_{\text{MP}}(\mathcal{P}) = \min_{\mu \in \mathcal{C}} \mathbb{E}_{\mu} [\mathbb{E}_m u(\mathcal{P})]. \quad (4)$$

Under the MP model, which is built within the Anscombe and Aumann (1963) framework, the layer of risk is evaluated under the expected utility criterion of von Neumann and Morgenstern (1944).

Smooth ambiguity model While the two previous criteria depart from the Bayesian paradigm, the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005, hereafter KMM) adheres to the Bayesian framework, but generalizes the classical SEU approach by allowing different attitudes towards the layers of risk and model ambiguity. The utility of a prospect \mathcal{P} under this criterion is

$$V_{\text{KMM}}(\mathcal{P}) = \mathbb{E}_{\mu} \phi [\mathbb{E}_m u(\mathcal{P})], \quad (5)$$

where $\phi \equiv v \circ u^{-1}$. The strictly increasing continuous function $v : C \rightarrow \mathbb{R}$ captures the attitude towards the layer of model ambiguity (Marinacci, 2015). Ambiguity aversion results from the concavity of ϕ , which corresponds to a stronger aversion towards the layer of model ambiguity than that of risk (i.e., v being more concave than u).¹⁷ Note that when two stages of risk are involved, each stage is evaluated by the same function u , so that criterion (5) collapses to criterion (1).

¹⁶This version was studied by Cerreia-Vioglio et al. (2013a), while the original version of the MP model proposed by Gilboa and Schmeidler (1989) is recovered when considering the predictive subjective probabilities $\bar{\mu}(s) \equiv \sum_m \mu(m)p(s|m)$ (see Online Appendix S4 for details).

¹⁷When $v = u$ —thus when ϕ is linear— the attitudes towards the two layers of risk and model ambiguity are identical and the smooth ambiguity criterion reduces to the SEU representation (1). See also the model of Nau (2006), which, at least in special cases, takes the same representation as (5) and shares the same interpretation as KMM.

Fear of model misspecification Finally, in a recent contribution, Cerreia-Vioglio et al. (2022, hereafter CHMM) propose a decision-theoretic approach able to deal with the three layers of uncertainty together. Starting from a setup in which the DM is able to posit a set M of structured models m that are motivated by the information available, this approach explicitly removes the assumption that the correct model belongs to M . It then allows the DM to contemplate unstructured models when ranking prospects and to express a fear of model misspecification by following the criterion

$$V_{CHMM}(\mathcal{P}) = \min_{p \in \Delta(S)} [\mathbb{E}_p u(\mathcal{P}) + c_M(p)], \quad (6)$$

where $c_M(p) = \min_{m \in M} c(p, m)$ is a measure of distance between a model p and the posited set M of structured models. Intuitively, the DM here may consider any type of models but penalizes those that are unstructured (and thus lack the substantive motivation of structured ones) by a cost function $c_M(p) \neq 0$ for all $p \notin M$. In particular, the closer a p is to the set M of structured models, the more plausible it is in view of the DM's information and the lower is the adverse impact on the preferences. Representation (6) is special case of the *variational* criterion axiomatized by Maccheroni, Marinacci, and Rustichini (2006). When the cost function has the relative entropy form, criterion (6) coincides with the criterion proposed by Hansen and Sargent (2022).¹⁸ Finally, when the cost function has a particular form (that assigns 0 to models inside M and infinity otherwise), the misspecification fear is absent and criterion (6) coincides with criterion (3).

Overall, if a distinction is made between different *layers* of uncertainty, we expect the *MA* premium and/or the *MM* premium to be non-zero. Because both sources *SR* and *CR* only entail the layer of risk, these theories often assume compound risk reduction (see, e.g., Klibanoff et al. (2005), or Gilboa and Schmeidler (1989)). Hence, the layer hypothesis is summarized as follows.

¹⁸When model ambiguity is absent and the set M is a singleton, it further reduces to the the *multiplier* decision criterion of Hansen and Sargent (2001, 2008). The multiplier decision criterion can also be equivalently written in the smooth ambiguity form when $\phi(x) = -e^{-\lambda x}$ (Hansen and Sargent, 2007; Cerreia-Vioglio et al., 2011). Note that alternative attempts to deal with a general concern about the epistemic uncertainty surrounding the correct probability model also appear in the vast literature on robust control theory (e.g., see Petersen, James, and Dupuis, 2000; Hansen and Sargent, 2001, 2008).

Layer hypothesis:

$$\begin{cases} \Pi_{SR,CR} = 0 \\ \Pi_{CR,MA} \neq 0 \text{ or } \Pi_{MA,MM} \neq 0 \text{ (or both)}. \end{cases} \quad (7)$$

4.4 Stage hypothesis

Finally, there exists a family of theories that models ambiguity as multiple *stages* of uncertainty while not necessarily making a distinction between the layers. For example, Segal’s (1987; 1990) and Seo’s (2009) approaches take any source of ambiguity as compound risk and relax the reduction principle to capture non-neutral attitudes towards ambiguity.

Recursive rank dependent utility model Segal’s (1987; 1990) recursive rank dependent utility (RRDU) approach proposes evaluating the first and second stages of uncertainty by using Quiggin’s (1982) rank dependent utility.¹⁹ In the RDU model of Quiggin, the lottery $x = (x_1, p_1; \dots; x_n, p_n)$, with $x_1 \geq \dots \geq x_n$, is evaluated by

$$V_{\text{RDU}}(x) = u(x_n) + \sum_{s=2}^n [u(x_{s-1}) - u(x_s)] f \left(\sum_{t=1}^{s-1} p_t \right). \quad (8)$$

In this expression, $f : [0, 1] \rightarrow [0, 1]$, with $f(0) = 0$ and $f(1) = 1$, is a strictly increasing transformation function, which is also convex under (global) uncertainty aversion.²⁰ Segal’s RRDU first computes and ranks the certainty equivalents (CEs) derived for each model $m \in M$ using (8) and then applies formulation (8) recursively to the distribution of these CEs induced by the probability measure μ .

Seo’s approach Seo (2009) assumes distinct expected utilities in the different *stages*, using a criterion analogous to that presented in (5). Under this interpretation, u and v each capture the attitude towards one particular stage of uncertainty. As a consequence, ambiguity aversion arises in the same way as nonreduction of compound risk.

¹⁹An alternative approach, using Gul’s (1991) disappointment aversion, is presented in Online Appendix S4.

²⁰Note that the common empirical finding in the literature is local uncertainty seeking for low likelihood events and uncertainty aversion for moderate and high likelihood events, implying an inverse S-shaped-concave and then convex- f function (see Wakker, 2010).

The theories that relate ambiguity attitudes to attitudes towards different *stages* of uncertainty thus only predict violation of the reduction of compound risk axiom (i.e., non-indifference between a compound risk and its reduced –simple risk– form),²¹ and make no distinction between layers:

Stage hypothesis:

$$\begin{cases} \Pi_{SR,CR} \neq 0 \\ \Pi_{CR,MA} = \Pi_{MA,MM} = 0. \end{cases} \quad (9)$$

Finally, note that a generalization of the stage hypothesis that further distinguishes between the different layers can be derived from Ergin and Gul (2009). This generalization of Segal’s approach allows for different transformation functions for the aleatory and epistemic layers of uncertainty present at each stage, thus giving rise to the hybrid:

Stage and layer hypothesis:

$$\begin{cases} \Pi_{SR,CR} \neq 0 \\ \Pi_{CR,MA} \neq 0 \text{ or } \Pi_{MA,MM} \neq 0 \text{ (or both)}. \end{cases} \quad (10)$$

5 Results of the lab experiment

5.1 Data

The data we collected in the lab experiment consist of 124 observations for *MA25* and 125 observations for the rest of the prospects.²² We excluded 36 (3.6% of all) choice lists from 13 different subjects, because they involve multiple-switching, no-switching, or reverse-switching patterns.²³ The proportion of these patterns is significantly lower than the typical 10% observed

²¹Note that Segal’s (1987; 1990) theory assumes a weaker condition than the reduction of compound risk, which is known as *time neutrality* (i.e., indifference between the resolution of uncertainty in the first or in the second stage) and implies $\Pi_{SR,CR0} = 0$ in our experiment. Other models in general also predict violation of time neutrality by assuming distinct attitudes within different stages (e.g., Seo, 2009).

²²One subject omitted answering the choice situation *MA25* by mistake.

²³Such a procedure is standard in the experimental literature (see, e.g., Dean and Ortleva, 2019). It is justified on the grounds that these data are not compatible with standard assumptions on preferences (e.g., monotonicity in money) and that they might be due to a lack of understanding of the choice tasks. Our results are, however, also robust to the inclusion of such data with multiple switching patterns.

in the literature (Yu, Zhang, and Zuo, 2020). We do not observe any order treatment effect.²⁴

5.2 General attitudes towards uncertainty

We start by looking at the general attitudes towards different sources of uncertainty before decomposing them into distinct layers. Unless mentioned otherwise, we report the results with two-sided t -tests. The use of non-parametric tests does not alter our main conclusions.

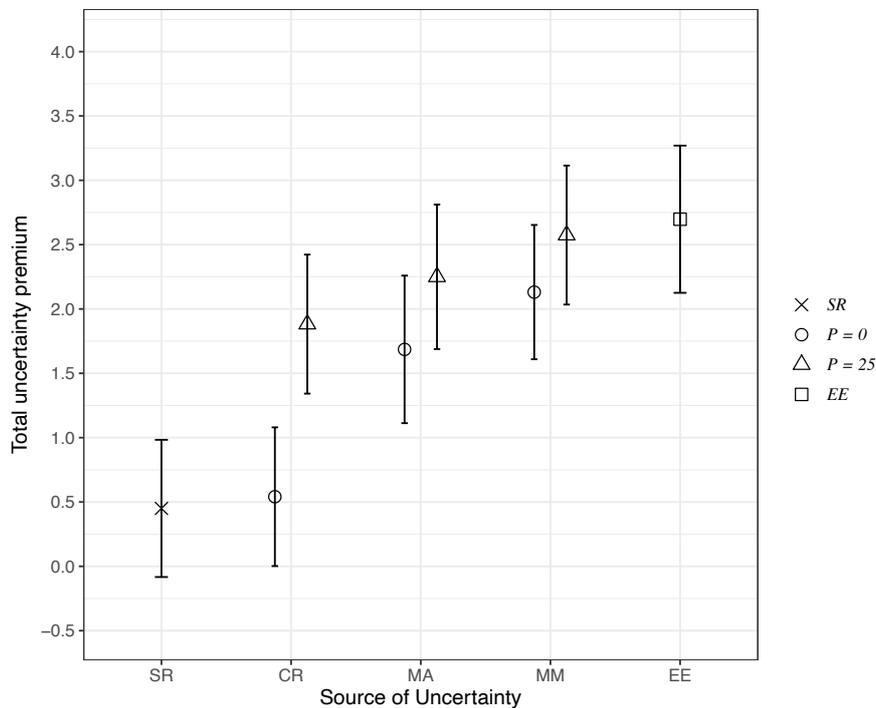


Figure 2: Mean total uncertainty premia and 90% confidence intervals

Figure 2 summarizes the mean *total* uncertainty premia Π_i . We find positive uncertainty premia for all the prospects ($p < 0.001$ for $CR25, MA0, MA25, MM0, MM25, EE$ and $p = 0.099$ for $CR0$) except for SR ($p = 0.16$), which indicates risk neutrality (such a result is reasonable for the small and moderate monetary gains considered in our experiment). Moreover, we observe an increasing trend in total premia from CR to MA and MM (ANOVA with repeated measures, $p < 0.001$ and $p = 0.001$ for treatments with $P = 0$

²⁴The results testing potential order effects are reported in Online Appendix S5. An additional experiment using a between-subject design confirms that the results are not due to contagion (see Online Appendix S2).

and $P = 25$, respectively), suggesting differences between these sources.

Table 2 summarizes the mean *differential* uncertainty premia relative to simple risk. Our data indicate consistent ambiguity aversion, as shown by

Table 2: COMPOUND RISK AND AMBIGUITY PREMIA

	$P = 0$	$P = 25$	Average
$\Pi_{SR,CR}$	0.10 ($N = 118$)	1.40*** ($N = 117$)	0.74*** ($N = 120$)
$\Pi_{SR,MA}$	1.18*** ($N = 119$)	1.81*** ($N = 117$)	1.46*** ($N = 120$)
$\Pi_{SR,MM}$	1.71*** ($N = 119$)	2.10*** ($N = 120$)	1.91*** ($N = 120$)
$\Pi_{SR,EE}$			2.30*** ($N = 116$)

Notes: The number of observations is in parentheses. Average premia are based on all non-missing values. *** p -value<0.01, ** p -value<0.05, * p -value<0.1, based on two-sided t -tests.

the positive ambiguity premia under the sources MA , MM , and EE . The compound risk premium is also positive when $P = 25$ but not when $P = 0$ ($p=0.58$). The proportions of zero, positive, and negative premia are reported in Online Appendix S5.

5.3 Decomposing attitudes towards uncertainty: the model ambiguity and misspecification premia

We now focus on the specific attitudes towards the layers of model ambiguity and model misspecification. To explore the empirical relevance of these attitudes and their respective differential premia, we first conduct a principal component analysis (PCA) of the premia for compound risk ($\Pi_{SR,CR}$) and ambiguity ($\Pi_{SR,MA}$ and $\Pi_{SR,MM}$). The results are reported in Table 3.

Table 3: PRINCIPAL COMPONENT ANALYSIS OF THE RELATIVE COMPOUND RISK AND AMBIGUITY PREMIA

Variable	$P = 0$			$P = 25$		
	1st comp.	2nd comp.	3rd comp.	1st comp.	2nd comp.	3rd comp.
$\Pi_{SR,CR}$	0.54	0.82	0.21	0.57	-0.68	0.45
$\Pi_{SR,MA}$	0.58	-0.54	0.61	0.57	0.73	0.37
$\Pi_{SR,MM}$	0.61	-0.20	-0.77	0.58	-0.04	-0.81
Eigenvalue of the component	1.86	0.68	0.48	2.43	0.31	0.26
Proportion of variance explained	0.62	0.22	0.16	0.81	0.10	0.09

We interpret the first component, having positive and roughly equal loadings on the three premia in both treatments with $P = 0$ and $P = 25$, as capturing a general attitude towards the multi-stage presentation of uncertainty, which is shared by all the three sources. This component explains 62% and 81% of the variance in the treatment $P = 0$ and $P = 25$, respectively.²⁵ In contrast, the second component has loadings of opposite signs on compound risk and model ambiguity, whereas its loading on model misspecification is relatively small. Hence, the second component is interpreted as capturing the attitude towards the layer of model ambiguity. The second component explains 22% and 10% of the variance in the two treatments, respectively. Finally, the third component has a high and negative loading on model misspecification compared to the positive loadings on the other sources. The third component may thus be interpreted as capturing the attitude towards the layer of model misspecification. This last component explains respectively 16% and 9% of the variance in the two treatments.

Next, we confirm these interpretations by looking at the actual differential premia for the layers of model ambiguity ($\Pi_{CR,MA}$) and model misspecification ($\Pi_{MA,MM}$). These premia are indeed highly correlated with, respectively, the second and third components of the PCA ($r = 0.93$ when $P = 0$ and $r = 0.998$ when $P = 25$ for the second component, and $r = 0.96$ when $P = 0$ and $r = 0.82$ when $P = 25$ for the third component). Corroborating the hypothesis of distinct attitudes towards different layers, we observe, in Table 4, that the premia for the specific layers of model ambiguity and model misspecification, as well as that for the general notion of model uncertainty ($\Pi_{CR,MM}$) are all positive. The table, for example, indicates that our subjects are ready to pay

Table 4: MODEL AMBIGUITY AND MODEL MISSPECIFICATION PREMIA

	$P = 0$	$P = 25$	Average
$\Pi_{CR,MA}$	1.11*** ($N = 119$)	0.50*** ($N = 116$)	0.84*** ($N = 120$)
$\Pi_{MA,MM}$	0.50** ($N = 120$)	0.37* ($N = 118$)	0.53*** ($N = 122$)
$\Pi_{CR,MM}$	1.48*** ($N = 120$)	0.67*** ($N = 118$)	1.05*** ($N = 122$)

Notes: The number of observations is in parentheses. Average premia are based on all non-missing values. *** p -value <0.01 , ** p -value <0.05 , * p -value <0.1 , based on two-sided t -tests.

²⁵This result suggests a lower importance for the multi-stage presentation in the treatment with $P = 0$ than that with $P = 25$. This discrepancy can be attributed to the relative simplicity of the treatment with $P = 0$: Despite being technically presented in two stages, this treatment can be seen as featuring only one stage of uncertainty (the second stage being degenerate). We further address the role of complexity in these data in another study (see Aydogan et al., 2019).

on average 8.4% of their expected payoff to avoid facing the layer of model ambiguity, and an extra 5.3% to avoid the layer of model misspecification.

5.4 Explaining the uncertainty premium: the role of layers

Having documented the existence of positive premia for the different layers of uncertainty, we now investigate the relevance of distinguishing between different layers while explaining the overall uncertainty premium. To do so, we run a regression analysis of total premia with subject fixed effects. Our analysis compares two baseline models that do not entail any distinction between the layers with models that capture model ambiguity and model misspecification.

Following the EU hypothesis, the first baseline model assumes the same uncertainty premium across all prospects:

$$\text{Model 1 (EU):} \quad \Pi_i^s = \alpha_0 + \gamma^s + \varepsilon_i^s, \quad (11)$$

where Π_i^s is the total uncertainty premium for prospect $i \in \{SR, CR0, CR25, MA0, MA25, MM0, MM25\}$ for subject s , γ^s is the individual subject fixed effect, and ε_i^s is the error term. The second baseline model is in line with the stage hypothesis, assuming a distinction between simple and multiple-stage presentations of uncertainty, but without distinguishing the layers:

$$\text{Model 2 (Stages):} \quad \Pi_i^s = \beta_0 + \beta_1 TS0_i + \beta_2 TS25_i + \gamma^s + \varepsilon_i^s, \quad (12)$$

where $TS0_i$ and $TS25_i$ are dummy variables for prospects presented in two stages, as opposed to the single-stage simple risk (SR) in the base category. Specifically, $TS0_i$ takes value 1 in the treatment $P = 0$, which presents one degenerate and one non-degenerate stage of uncertainty, whereas $TS25_i$ takes value 1 in the treatment $P = 25$, which entail two non-degenerate stages of uncertainty.

Then, we investigate the potential improvements in the specification when distinguishing the layers of model ambiguity and model misspecification by considering the models

$$\text{Model 1' (Layers): } \Pi_i^s = \alpha_0 + \alpha_1 MA_i + \alpha_2 MM_i + \gamma^s + \varepsilon_i^s \quad (13)$$

and

$$\text{Model 2' (Stages \& layers): } \Pi_i^s = \beta_0 + \beta_1 TS0_i + \beta_2 TS25_i + \beta_3 MA_i + \beta_4 MM_i + \gamma^s + \varepsilon_i^s, \quad (14)$$

where MA_i and MM_i are dummy variables for the prospects entailing respectively the second and third layers.

Table 5 reports the estimation results. The estimations indicate that the introduction of layers or stages always leads to significant increases in the uncertainty premium. Importantly, a model comparison exercise shows that

Table 5: FIXED EFFECT REGRESSIONS OF TOTAL PREMIA

	No distinction between stages		Distinction between stages	
	Model 1 (EU)	Model 1' (Layers)	Model 2 (Stages)	Model 2' (Stages & layers)
<i>MA</i>		1.003*** (0.164)		0.758*** (0.168)
<i>MM</i>		1.413*** (0.171)		1.166*** (0.167)
<i>TS0</i>			1.001*** (0.172)	0.357** (0.170)
<i>TS25</i>			1.760*** (0.217)	1.113*** (0.207)
Constant	1.644*** (0.066)	0.952*** (0.083)	0.462*** (0.158)	0.464*** (0.157)
Observations	845	845	845	845

Notes: Robust standard errors, cluster-corrected at individual level in parentheses,

*** p -value<0.01, ** p -value<0.05, * p -value<0.1

the models incorporating the distinction between the layers systematically outperform the ones that do not. Specifically, we reject the EU and stage hypotheses, which assume the existence of the layer of risk only (F -test, $p < 0.001$ for both tests $\alpha_1 = \alpha_2 = 0$ in model 1' and $\beta_3 = \beta_4 = 0$ in model 2'). Furthermore, we also reject the hypothesis that no distinction exists between the layers of model ambiguity and model misspecification ($p=0.015$ and $p=0.016$ respectively for tests $\alpha_1 = \alpha_2$ in model 1' and $\beta_3 = \beta_4$ in model 2'). Finally, we observe that the models making a distinction between single and two-stage

presentations of uncertainty also outperform the models not making such a distinction ($p < 0.001$ for the test $\beta_1 = \beta_2 = 0$ in both models 2 and 2').

5.5 Individual-level analysis

We now present an individual-level analysis classifying subjects' preferences according to the theoretical predictions. Specifically, we test the compatibility of each individual's preferences over R , CR , MA , and MM with our four hypotheses in both treatments $P = 0$ and $P = 25$.²⁶ Table 6 summarizes the results.

Table 6: CLASSIFICATION OF INDIVIDUALS

EU hyp.	Layer hyp.	Stage hyp.	Stage and layer hyp.
19.7%	31.2%	8.1%	41.0%

As can be observed, the most common preference pattern (41 % of the subjects) is consistent with the hybrid stage and layer hypotheses. In that case, the subject exhibits non-reduction of compound risk as well as distinct attitudes towards layers. The second most common pattern (31%) is in line with the layer hypothesis, in which the subject reduces compound risk while at the same time exhibits non-neutrality towards ambiguity. Next, 20% of the subjects made choices consistent with the EU hypothesis, thus exhibiting neutrality towards compound risk and ambiguity. Finally, the pure stage hypothesis holds for 8% of the subjects.

The consistency of this classification analysis is confirmed by examining the behavior under the source EE . First, we remark that 83% of the EU subjects are also EE -ambiguity neutral (i.e. $\Pi_{R,EE} = 0$), whereas 77% of all non-EU subjects are non-neutral towards EE (Fisher's exact test, $p < 0.001$). Second, the proportions of EE -ambiguity non-neutrality is consistently high within each type of non-EU individuals (75% under the layer hypothesis, 79%

²⁶For subjects whose preferences support a different hypothesis in each treatment, the classification score is split between the two hypotheses. For example, a subject whose preferences conforms to the SEU hypothesis in treatment with $P = 0$ and the layer hypotheses in treatment with $P = 25$ is classified as following 50% SEU and 50% layer hypotheses. The complete classification under the two treatments is reported in Online Appendix S5.

under the stage hypothesis, and 78.5% under the stage & layer hypothesis), suggesting that the source of EE may have been considered as a source encompassing both stages and layers of uncertainty.

Altogether, the results obtained in our experiment thus confirm that each layer of uncertainty plays a distinct and significant role while explaining uncertainty attitudes. Consistent with the previous literature, our findings also indicate the importance of the multi-stage presentation of uncertainty in explaining ambiguity attitudes. In the next section, we summarize the main findings of the two follow-up experiments testing the robustness of these findings.

6 Robustness: Results of the follow-up experiments

Our follow-up experiments address two potential concerns about our results. The first concern is that our results may depend on the specificity of our subject pool. It may, for example, be the case that our findings are artifacts of the potential limitations of the subject pool to deal with the relevant complexity of our sources, which can result in an aversion towards sources with several stages and layers of uncertainty. To evaluate this interpretation, we ran our experiment on the field with a pool of risk professionals. These subjects are expected to be more (quantitatively) sophisticated and thus better able to deal with the complexity inherent to a multi-stage presentation of uncertainty.

The second concern is that our results may be a consequence of order effects or contagion between sources encompassing distinct layers, which results from our within-subjects design. Accordingly, a potential factor behind the distinction between the sources CR , MA , and MM can be the successive evaluation of the sources by the same subject, causing her to be on average more averse to additional layers of uncertainty (comparative ignorance hypothesis of Fox and Tversky, 1995). To remove this concern, we run the experiment online using a between-subject design, in which each source is evaluated separately. Additional details on these two experiments are provided in Online Appendices S1 and S2.

Figure 3 presents the total premia in the follow-up experiments. As can be observed, we replicate the positive trend in $CR - MA - MM$ in both experiments, suggesting different attitudes towards those sources (ANOVA

with repeated measures, $p < 0.001$ for both $P = 0$ and $P = 25$ in the field experiment with actuaries; and ANOVA, $p = 0.004$ and $p = 0.024$ respectively for $P = 0$ and $P = 25$ in the online experiment). Interestingly, we also find

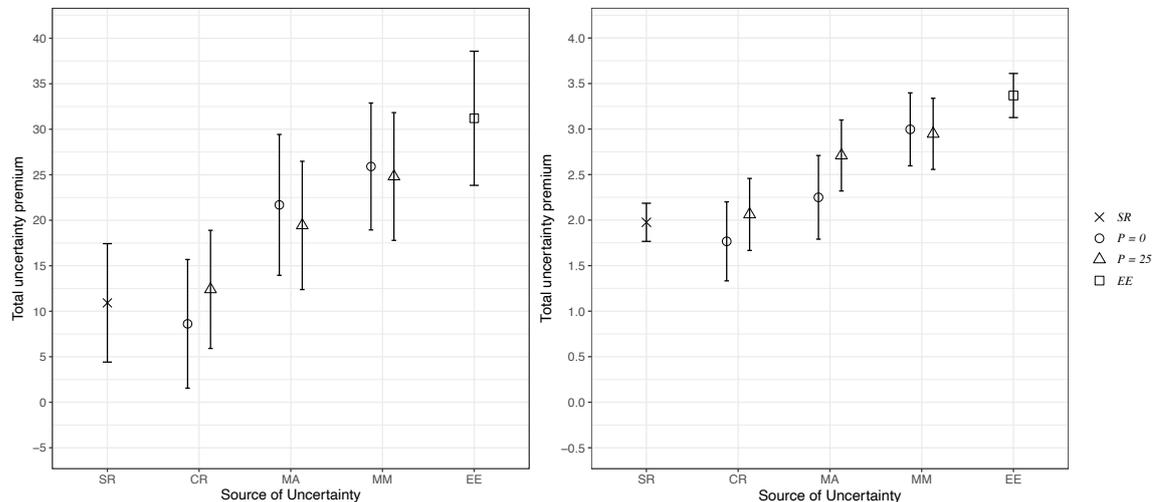


Figure 3: Left: Mean total premia in the field experiment with actuaries (left) and in the online experiment (right)

that the subjects in these experiments were better able to reduce compound risk, as indicated by the average compound risk premium that does not differ from zero ($p > 0.10$ for average $\Pi_{SR,CR}$ in both experiments). The attitudes towards the layers of uncertainty are nevertheless robust and consistent with those in our main experiment. In particular, the average model uncertainty premium amounts to 13.13% of the expected payoff in the experiment with risk professionals and to 11.2% in the online experiment ($p < 0.001$ for both), which is comparable to the 10.5% observed in the main experiment.

The role played by the layers in explaining the total uncertainty premia is also confirmed in replications of the regression analysis. In both follow-up experiments, we find that the presence of different layers increases the uncertainty premia and that incorporating them into the analysis improves the estimation results of the premia. Yet, as also suggested by the weaker compound risk premia, we find less evidence on specific attitudes toward the multi-stage presentation of uncertainty in these experiments than in the main one (see Online Appendix for details).

The stability of attitudes towards layers (rather than stages) is further observed in a replication of the individual level analysis with the pool of risk professionals. Specifically, the proportion of risk professionals following the

layer hypothesis is 28%, which does not differ from the 31.2% observed in our main experiment (two-sample Z -test of proportions, $p = 0.59$), whereas the proportion of them exhibiting attitudes towards a multiple-stage presentation is significantly lower than among students (49% vs. 17%, $p < 0.001$). Overall, these additional results suggest that our main findings on the layers of uncertainty are neither an artifact of a specific subject pool, nor due to any order effect or contagion.

7 Concluding remarks and related literature

In this paper, we use a simple experimental environment to demonstrate the existence of an empirical distinction between attitudes towards the three layers of uncertainty. In what follows, we clarify the contribution of our study in relation with the existing literature.

Characterization of ambiguity. Since Ellsberg (1961), ambiguity aversion has been one of the most intensively investigated phenomena in experimental economics (see Trautmann and van de Kuilen, 2015). Among the different theoretical approaches that have been proposed to explain Ellsberg-type behaviors, the multi-stage representation of ambiguity, which assumes a set of possible probability models (or distributions), has received much attention.²⁷ The theories adopting a multi-stage approach are reviewed in Marinacci (2015), and recent empirical applications include Halevy (2007), Baillon and Bleichrodt (2015), Chew et al. (2017), and Cubitt et al. (2019). Among other findings, a tight association between ambiguity and compound risk attitudes, as for example highlighted in Halevy (2007), has been replicated (with varying degrees of success) in the empirical literature (see Abdellaoui et al., 2015; Berger and Bosetti, 2020; Dean and Ortoleva, 2019; Gillen et al., 2019; Chapman et al., 2018, and further results in the Online Appendix). Such an association supports a characterization of ambiguity by means of compound risk, in line with the theoretical claims of Segal (1987; 1990).

Following the recent discussions in Hansen (2014) and Hansen and Marinacci (2016), what we argue in this paper is that a decomposition into three distinct layers of analysis –which arise naturally in many real life decision prob-

²⁷Another important approach that assumes non-additive beliefs is due to Schmeidler (1989) and the cumulative prospect theory of Tversky and Kahneman (1992). Note that there is also a close link between these models and the multiple priors approach assuming sets of probabilities.

lems where a set of possible probabilities can be posited— can provide a better characterization of ambiguity in applications. Going beyond the preceding contributions in the literature, our experiments aim to quantify the attitudes towards each of the three layers and test their respective role in attitudes towards ambiguity. Our results demonstrate that, although a multi-stage presentation of uncertainty indeed plays a role in ambiguity attitudes, ambiguity cannot simply be reduced to the layer of risk. In particular, we show that (1) individuals are ready to pay positive premia to avoid being confronted to the the specific layers of model ambiguity and model misspecification; (2) statistical models that assume only a distinction between stages underperform when estimating ambiguity premia, compared to models that distinguish different layers; and (3) individual behaviors can be characterized by three main types, namely (i) those who distinguish between both the layers and stages of uncertainty, (ii) those who distinguish between the layers of uncertainty only, and (iii) those following the expected utility hypothesis, whereas a fourth type associated with a distinction between stages only is marginal. Overall, our results call for further empirical and theoretical developments accommodating the three layers of uncertainty.

More on model misspecification In this paper, we follow the decomposition of uncertainty into three layers proposed by Hansen (2014) and Hansen and Marinacci (2016). These layers are distinguished based on the knowledge of the DM. The most challenging layer to accommodate is probably model misspecification, in which uncertainty is induced by the approximate nature of the models considered. How to deal with model misspecification in a fruitful way is a concern that has occupied statisticians, econometricians, and control theorists for a long time. For example, it has long been a challenge for statisticians, whose objective is to find the correct statistical model but who have been trained with the idea that “essentially, all models are wrong [i.e., misspecified], but some are useful” (Box, 1976; Watson and Holmes, 2016),²⁸ or for econometricians, who stand outside an economic model and are asked to

²⁸Box (1976, p. 792) wrote “Since all models are wrong the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad”, and Cox (1995): “Finally it does not seem helpful just to say that all models are wrong. The very word model implies simplification and idealization. The idea that complex physical, biological or sociological systems can be exactly described by a few formulae is patently absurd. The construction of idealized representations that capture important stable aspects of such systems is, however, a vital part of general scientific analysis and statistical models, especially substantive ones, do not seem essentially different from other kinds of model.”

choose among different parameters characterizing a family of possible models to best explain real-world data and to test the implications of these models (White, 1982). Considering model misspecification under this perspective thus relates to the process of examining models to identify their flaws and potential improvements. While this uncertainty about the correct specification of a model is certainly relevant for a statistician or econometrician outside the model, it is also potentially important for agents inside an economic model, be it individuals pursuing individual goals, or policymakers pursuing social objectives, who confront uncertainty as they make decisions.

Yet, surprisingly enough, the concern of how to account for model misspecification in a way that guides the use of purposefully simplified models in a sensible way has been largely absent (with the exception of Cerreia-Vioglio et al., 2022) in decision theory, whose objective is precisely to describe how a person should behave in an uncertain environment. One reason for this shortcoming is that the distinction between model ambiguity and potential model misspecification can arguably be fuzzy. Typical approaches in practice blur the distinction mathematically by simply embedding the existing substantive models in a much bigger space of potential probability models and treating them symmetrically within a bigger universe that could capture the misspecification. However, in reality, the fear of a potential model misspecification is conceptually more complicated to address than the already challenging problem of how much credibility we should give to the different models we consider. In a variety of policy applications, including public health, climate change, and macroeconomics, multiple models are on the table with potentially unknown (or in principle unknowable) parameters. At the same time, we know that all those models under consideration are misspecified, in some ways by design. Consequently, the issue of how much weight to give to the different models is different from how we should confront potential model misspecification. The latter leads to a concern that our simplified models could promote misleading policy conclusions if taken too literally.

This paper highlights empirically the importance of considering the fear of model misspecification as a distinct layer from model ambiguity. While trying to capture such a distinction experimentally with a simple urn necessarily renders some aspects of decision-making under uncertainty simplistic, our results suggest that artificially enlarging the spread between potential probability models and treating them symmetrically does not fully capture concerns for potential misspecification. Overall, our treatment of model misspecification

should thus be viewed as a good proxy to provide insights into the relevance of the third layer in a controlled environment. Extrapolating our experimental findings to real-life situations, in which model misspecification arises naturally, we conjecture that the role of model misspecification is potentially more important than what we capture in our experiment.

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., P. Klibanoff, and L. Placido (2015). Experiments on compound risk in relation to simple risk and to ambiguity. *Management Science* 61(6), 1306–1322.
- Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *The annals of mathematical statistics* 34(1), 199–205.
- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica* 19, 404–437.
- Attema, A. E. and W. B. Brouwer (2013). In search of a preferred preference elicitation method: A test of the internal consistency of choice and matching tasks. *Journal of Economic Psychology* 39, 126–140.
- Aydogan, I., L. Berger, and V. Bosetti (2019). Economic rationality: Investigating the links between uncertainty, complexity, and sophistication. IGIER Working Paper 653, Bocconi University.
- Azrieli, Y., C. P. Chambers, and P. J. Healy (2018). Incentives in experiments: A theoretical analysis. *Journal of Political Economy* 126(4), 1472–1503.
- Baillon, A. and H. Bleichrodt (2015). Testing ambiguity models through the measurement of probabilities for gains and losses. *American Economic Journal: Microeconomics* 7(2), 77–100.
- Baillon, A., Y. Halevy, and C. Li (2022). Experimental elicitation of ambiguity attitude using the random incentive system. *Experimental Economics*, 1–22.
- Barnett, M., W. Brock, and L. P. Hansen (2020). Pricing uncertainty induced by climate change. *The Review of Financial Studies* 33(3), 1024–1066.
- Berger, L., N. Berger, V. Bosetti, I. Gilboa, L. P. Hansen, C. Jarvis, M. Marinacci, and R. D. Smith (2021). Rational policymaking during a pandemic. *Proceedings of the National Academy of Science* 118(4), e2012704118.
- Berger, L. and V. Bosetti (2020). Are policymakers ambiguity averse? *The Economic Journal* 130, 331–355.
- Berger, L. and M. Marinacci (2020). Model uncertainty in climate change economics: A review and proposed framework for future research. *Environmental and Resource Economics* 77, 475–501.
- Berk, R. H. (1966). Limiting behavior of posterior distributions when the model is incorrect. *The Annals of Mathematical Statistics*, 51–58.
- Bostic, R., R. J. Herrnstein, and R. D. Luce (1990). The effect on the preference-reversal phenomenon of using choice indifference. *Journal of Economic Behavior & Organization* 13(2), 193–212.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association* 71(356), 791–799.
- Brock, W. A. and L. P. Hansen (2017). Wrestling with uncertainty in climate economic models. mimeo.
- Cerreia-Vioglio, S., P. Hansen, Lars, F. Maccheroni, and M. Marinacci (2022). Making decisions under model misspecification. mimeo.

- 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.
- Chapman, J., M. Dean, P. Ortoleva, E. Snowberg, and C. Camerer (2018). Econographics. Technical report, National Bureau of Economic Research.
- Chew, S. H., B. Miao, and S. Zhong (2017). Partial ambiguity. *Econometrica* 85(4), 1239–1260.
- Cox, D. R. (1995). Comment on "model uncertainty, data mining and statistical inference". *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 158, 455–456.
- Cubitt, R., G. van de Kuilen, and S. Mukerji (2019, 04). Discriminating Between Models of Ambiguity Attitude: a Qualitative Test. *Journal of the European Economic Association*.
- Dean, M. and P. Ortoleva (2019). The empirical relationship between nonstandard economic behaviors. *Proceedings of the National Academy of Sciences* 116(33), 16262–16267.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics* 75, 643–669.
- Epstein, L. G. and Y. Halevy (2019). Ambiguous Correlation. *The Review of Economic Studies* 86(2), 668–693.
- Ergin, H. and F. Gul (2009). A theory of subjective compound lotteries. *Journal of Economic Theory* 144(3), 899–929.
- Esponda, I. and D. Pouzo (2016). Berk–nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica* 84(3), 1093–1130.
- Esponda, I., E. Vespa, and S. Yuksel (2020). Mental models and learning: The case of base-rate neglect. Technical report, mimeo.
- Fox, C. R. and A. Tversky (1995). Ambiguity aversion and comparative ignorance. *The quarterly journal of economics* 110(3), 585–603.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18(2), 141–154.
- Gillen, B., E. Snowberg, and L. Yariv (2019). Experimenting with measurement error: Techniques with applications to the caltech cohort study. *Journal of Political Economy* 127(4), 1826–1863.
- Götte, L., M. Kozakiewicz, et al. (2020). Experimental evidence on misguided learning. Technical report, mimeo.
- Grether, D. M. and C. R. Plott (1979). Economic theory of choice and the preference reversal phenomenon. *The American Economic Review* 69(4), 623–638.
- Gul, F. (1991). A theory of disappointment aversion. *Econometrica* 59(3), 667–686.
- Halevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica* 75(2), 503–536.
- Hansen, L. and T. Sargent (2008). *Robustness*. Princeton University Press.
- Hansen, L. P. (2014). Nobel lecture: Uncertainty outside and inside economic

- models. *Journal of Political Economy* 122(5), 945–987.
- Hansen, L. P. and M. Marinacci (2016). Ambiguity aversion and model misspecification: An economic perspective. *Statistical Science* 31, 511–515.
- Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. *American Economic Review* 91(2), 60–66.
- Hansen, L. P. and T. J. Sargent (2007). Recursive robust estimation and control without commitment. *Journal of Economic Theory* 136(1), 1–27.
- Hansen, L. P. and T. J. Sargent (2022). Structured ambiguity and model misspecification. *Journal of Economic Theory* 199, 105165.
- Hara, K., E. A. Ok, and G. Riella (2019). Coalitional expected multi-utility theory. *Econometrica* 87(3), 933–980.
- Johnson, C., A. Baillon, H. Bleichrodt, Z. Li, D. Van Dolder, and P. P. Wakker (2021). Prince: An improved method for measuring incentivized preferences. *Journal of Risk and Uncertainty* 62(1), 1–28.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 363–391.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica* 73, 1849–1892.
- Maccheroni, F., M. Marinacci, and D. Ruffino (2013). Alpha as ambiguity: Robust mean-variance portfolio analysis. *Econometrica* 81(3), 1075–1113.
- 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), 1000–1076.
- Nau, R. F. (2006). Uncertainty aversion with second-order utilities and probabilities. *Management Science* 52(1), 136–145.
- Petersen, I. R., M. R. James, and P. Dupuis (2000). Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Transactions on Automatic Control* 45(3), 398–412.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization* 3(4), 323–343.
- 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.
- Segal, U. (1987). The Ellsberg paradox and risk aversion: An anticipated utility approach. *International Economic Review*, 175–202.
- Segal, U. (1990). Two-stage lotteries without the reduction axiom. *Econometrica*, 349–377.
- Seo, K. (2009). Ambiguity and second-order belief. *Econometrica* 77(5), 1575–1605.
- Trautmann, S. and G. van de Kuilen (2015). Ambiguity attitudes. In G. Keren and G. Wu (Eds.), *Blackwell Handbook of Judgment and Decision Making*. Blackwell.
- Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5(4), 297–323.
- von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic*

Behavior. Princeton University Press.

Wakker, P. P. (2010). *Prospect theory: For risk and ambiguity*. Cambridge university press.

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

Watson, J. and C. Holmes (2016). Approximate models and robust decisions. *Statistical Science* 31(4), 465–489.

White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica* 50(1), 1–25.

Yu, C. W., Y. J. Zhang, and S. X. Zuo (2020). Multiple switching and data quality in the multiple price list. *The Review of Economics and Statistics* (forthcoming), 1–45.