



The uncertainty triangle – Uncovering heterogeneity in attitudes towards uncertainty

Daniel R. Burghart¹ · Thomas Epper^{2,3} · Ernst Fehr³

Published online: 13 August 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

This paper develops a graphical tool – the uncertainty triangle – that allows for testing whether choices under uncertainty obey the generalized axiom of revealed preference (GARP). We find that more than 95% of subjects made choices that can be rationalized by the maximization of a well-behaved utility function. The uncertainty triangle also makes it straightforward to characterize heterogeneity in attitudes towards uncertainty. To accomplish this we propose a one-parameter extension of Expected Utility in which uncertainty attitude is everywhere constant in the triangle. Experimental data indicate that about 60% of participants made choices consistent with the model and, within this group, 48% were uncertainty averse, 22% uncertainty seeking, and 30% uncertainty neutral. The remaining 40% of participants appear to hold variable uncertainty attitudes. A model that can accommodate this variability is proposed and calibrated.

Keywords Uncertainty · Risk · Expected utility · Revealed preferences · Decision-making

JEL Classifications D81 · C91

1 Introduction

The analysis of decisions under risk involves prospects that have fully known probabilities. Real-world settings, however, rarely involve full knowledge about prospects. When individuals do not possess full knowledge about prospects they are said to

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s11166-020-09331-8>) contains supplementary material, which is available to authorized users.

✉ Daniel R. Burghart
dan.burghart@gmail.com

Extended author information available on the last page of the article.

make choices in the face of Knightian uncertainty (Knight 1921). While individual variability in risk attitudes is a well-established finding in that domain, heterogeneity of uncertainty attitudes is not as widely studied. One reason is that studying uncertainty attitudes poses additional challenges. The goal of this paper is to tackle some of these challenges and provide a characterization of individual variability in attitudes towards uncertainty. The first step we take is to assume that the objects of choice are two-outcome *lower envelope lotteries*.¹ Lower envelope lotteries specify lower bounds on probabilities, $\{\underline{p}_h, \underline{p}_l\}$, for a high and low outcome, z_h and z_l with $z_h > z_l$, and the amount of unassigned probability mass, $y = 1 - (\underline{p}_h + \underline{p}_l)$, an objective quantity henceforth called *uncertainty*.² Formally, we denote a lower envelope lottery as $L = (\underline{p}_h, \underline{p}_l, y)$. Analogous to the probability distributions studied in risk, the entries in a lower envelope lottery must be non-negative and sum to one.

To parsimoniously model choices in this setting we propose a one parameter extension of Expected Utility. We call this model Partial Ignorance Expected Utility (PEU) and it evaluates a lower envelope lottery (L) as³

$$\text{PEU}[L] = \text{EU}\left[\alpha \underline{L} + (1 - \alpha) \bar{L}\right], \quad \alpha \in [0, 1]. \quad (1)$$

In terms of notation, \underline{L} is a transformation of the lower envelope lottery L into a risky lottery with all of the uncertainty in L added to the minimum probability of the worst outcome: $\underline{L} = (\underline{p}_h, \underline{p}_l + y)$. In contrast, \bar{L} is a transformation of L into a risky lottery with all of the uncertainty added to the minimum probability of the best outcome: $\bar{L} = (\underline{p}_h + y, \underline{p}_l)$. The parameter α controls a mixture of these two lotteries and is interpreted as a choosers attitude towards uncertainty. Values of $\alpha > 1/2$ place more weight on \underline{L} with, correspondingly, less weight on \bar{L} . This indicates a pessimistic attitude towards the uncertainty in L and we label choosers with $\alpha > 1/2$ as *uncertainty averse*. Values of $\alpha < 1/2$ put more weight on \bar{L} and less weight on \underline{L} . This indicates a more optimistic attitude towards uncertainty and we label choosers with $\alpha < 1/2$ as *uncertainty seeking*. Choosers with $\alpha = 1/2$ weight equally between \underline{L} and \bar{L} and are called *uncertainty neutral*. Whatever a chooser's α , the mixture

¹The *lower envelope* nomenclature is borrowed from the literature on imprecise probabilities. See, for example, Dempster (1967) and Shafer (1976).

²This paper uses the term ‘uncertainty’ as a concept that is related-to but distinct-from ‘ambiguity.’ Here, uncertainty is an objective quantity corresponding to the amount of unassigned probability mass in a lower envelope lottery. Ambiguity, in contrast, is “a quality depending on the amount, type, reliability, and ‘unanimity’ of information” (Ellsberg 1961, pg. 657). Thus, ambiguity is a more general concept that can include subjective components. The distinction between uncertainty and ambiguity is important because our approach permits an examination of uncertainty attitudes in the context of lower envelope lotteries. Our approach does not, however, permit an examination of ambiguity, or ambiguity attitudes, in the manner of Ghirardato et al. (2004).

³The name ‘Partial Ignorance Expected Utility’ was inspired by Chapter 13.5 in Luce and Raiffa (1958).

$\alpha \underline{L} + (1-\alpha) \overline{L}$ is assumed to be evaluated with a von Neumann-Morgenstern Expected Utility, EU $[\cdot]$.⁴

To help illustrate how α captures uncertainty attitude consider Fig. 1a. We call this the *uncertainty triangle* and it is a graphical depiction of all two outcome lower envelope lotteries.⁵ The vertices of the uncertainty triangle represent three special cases: (i) certainty of receiving the good outcome ($\underline{p}_h = 1$), (ii) certainty of receiving the bad outcome ($\underline{p}_l = 1$), and (iii) a situation in which nothing is known about the probability distribution over the two outcomes ($y = 1$). Lower envelope lotteries on the hypotenuse of the triangle are analogous to fully specified lotteries because they have no uncertainty ($y = 0$). Horizontal movements in the triangle result in a one-to-one tradeoff between uncertainty (y) and the minimum probability for the bad outcome (\underline{p}_l) while holding the minimum probability for the good outcome (\underline{p}_h) constant. Vertical movements trade off y and \underline{p}_h holding \underline{p}_l constant while movements parallel to the hypotenuse hold y fixed while trading off \underline{p}_h and \underline{p}_l .

The Partial Ignorance Expected Utility model can be written as:

$$\text{PEU}(\underline{p}_h, \underline{p}_l, y) = \alpha [\underline{p}_h u_h + (\underline{p}_l + y) u_l] + (1 - \alpha) [(\underline{p}_h + y) u_h + \underline{p}_l u_l]. \quad (2)$$

Normalizing utilities such that $u_h = 1$ and $u_l = 0$ gives:

$$\text{PEU}(\underline{p}_h, \underline{p}_l, y) = \underline{p}_h + (1 - \alpha)y. \quad (3)$$

This means PEU produces indifference curves that are linear and parallel in the uncertainty triangle. The slope of these indifference curves is determined solely by the parameter $\alpha \in [0, 1]$. To see this consider a choice between a lower envelope lottery with 100% uncertainty, $(0, 0, 1)$, and a lower envelope lottery analogous to a 50/50 lottery $(1/2, 1/2, 0)$. Figure 1b plots these two alternatives in the uncertainty triangle. The 100% uncertain alternative is at the lower-right vertex and the 50/50 lottery is halfway between the endpoints of the hypotenuse. The PEU for the 100% uncertain option is $\text{PEU}(0, 0, 1) = 1 - \alpha$. The PEU for the 50/50 lottery is $\text{PEU}(1/2, 1/2, 0) = 1/2$. A PEU maximizer is indifferent between these two options if and only if they are uncertainty neutral (i.e. $\alpha = 1/2$). Indifference curves for an uncertainty neutral chooser are depicted in green in Fig. 1b. An uncertainty averse

⁴The PEU model has been axiomatized in a more general ‘sets of lotteries’ setting in Olszewski (2007). Also related to our approach Gul (2015) introduces a source-dependent theory they coin Hurwicz Expected Utility. The Hurwicz Expected Utility Theory is itself a special case of a more general theory (Expected Uncertain Utility Theory) introduced in Gul and Pesendorfer (2014). The model in Gul (2015) is a sub-class of α -MaxMin Expected Utility (Marinacci 2002; Ghirardato et al. 2004) as it imposes further restrictions on priors. In this sense, the PEU model can also be interpreted as a variant of the α -MaxMin model that is adapted to the setting of lower envelope lotteries (Marinacci 2002; Ghirardato et al. 2004). PEU differs, however, by means of the objects of choice and its interpretation. First, PEU preferences are defined over lower envelope lotteries whereas α -MaxMin preferences are defined over acts (bets). Second, in the PEU model the best and the worst possibilities (\overline{L} and \underline{L}) are objectively defined entities. For the α -MaxMin model, however, the best and worst possibilities are determined by an individual’s subjective beliefs. Section 4 discusses some implications of these distinctions.

⁵Burghart (2018) utilizes a triangular figure to explore a model in the setting of “upper envelope lotteries.” These objects of choice are distinct from lower envelope lotteries in that they list the maximum probability for each outcome and a term called “information” which captures how much is known about the probabilities at the time of choice.

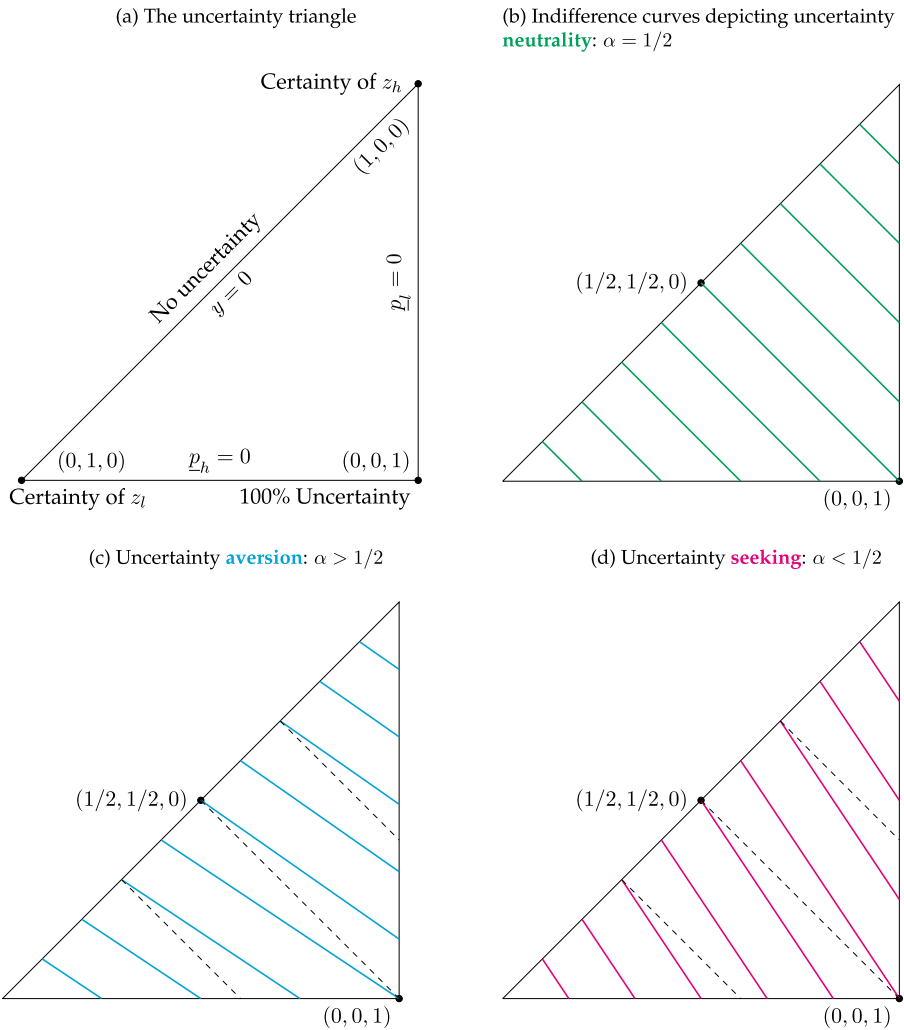


Fig. 1 A triangular diagram can be used to plot two outcome lower envelope lotteries and indifference curves for various uncertainty attitudes

chooser (i.e. $\alpha > 1/2$) would prefer the 50/50 lottery to the 100% uncertain option. Indifference curves for an uncertainty averse chooser are depicted in **blue** in Fig. 1c. Someone who is uncertainty seeking ($\alpha < 1/2$) would prefer the 100% uncertain option to the 50/50 lottery. The indifference curves for an uncertainty seeking chooser are shown in **red** in Fig. 1d.

The PEU model places significant structure on preferences. To explore whether actual behavior is consistent with this structure we designed and conducted an experiment in which the objects of choice were two-outcome lower envelope lotteries. Our experiment systematically varied the tradeoffs between minimum probabilities and uncertainty. This design permits a non-parametric assessment of the PEU model.

This assessment indicates that about 60% of participants made choices consistent with PEU maximization. For these PEU maximizers we empirically characterize the heterogeneity of their uncertainty attitudes with finite mixture models. Finite mixture methods provide an endogenous classification of individuals to different preference types and estimate precise preferences ($\hat{\alpha}$) for each type. Our estimation results show that there is substantial heterogeneity of uncertainty attitudes amongst the PEU maximizers. This heterogeneity is best characterized by the existence of three distinct types that, coincidentally, correspond to uncertainty aversion, seeking, and, approximately speaking, neutrality. We also obtain estimated proportions of participants that belong to each of these distinct types. One type is comprised of 30% of the PEU maximizers and corresponds to (near) uncertainty neutrality ($\hat{\alpha} = 0.517$). Roughly 48% of PEU maximizers exhibited uncertainty aversion ($\hat{\alpha} = 0.583$). The third type, comprising roughly 22% of PEU maximizers, display uncertainty seeking behavior ($\hat{\alpha} = 0.395$). Importantly, the finite-mixture approach does not assume the ex-ante existence of these three types. Instead, these three types emerge endogenously from the finite mixture methodology. In addition, almost all of the PEU maximizers are cleanly assigned to one distinct type or another (i.e. the subjects' posterior probabilities of belonging to, for example, uncertainty aversion, is almost exclusively close to one).

About 40% of experimental participants cannot be considered PEU maximizers because they exhibited non-constant uncertainty attitudes. That is, in the uncertainty triangle, their indifference curves are linear but non-parallel.⁶ To parametrically explore preferences for this group we introduce the β -PEU model. The β -PEU model has one more parameter than PEU. It retains the linear indifference curves of PEU but allows indifference curves to 'fan-in' ($\beta > 0$) or 'fan-out' ($\beta < 0$) across the uncertainty triangle. Indifference curves that 'fan-in' imply increasing aversion to uncertainty when moving northeast in the triangle (i.e. as the minimum probability of the good outcome increases). Indifference curves that 'fan-out' imply the opposite – decreasing aversion to uncertainty with northeast movements in the triangle. Again using finite mixture models we show that all of the β -PEU maximizers in our sample have indifference curves that 'fan-in' (i.e. increasing aversion to uncertainty, or $\beta > 0$). But again there is substantial heterogeneity mainly with regard to their baseline uncertainty aversion: About 62% of β -PEU maximizers exhibit indifference curves that fan in yet are, on average, not significantly different from uncertainty neutrality. The remaining 38% exhibit very similar fanning behavior but, their average uncertainty attitudes are best characterized as uncertainty seeking.

The next section of the paper lays out our experimental design and methods in more detail. Section 3 presents our empirical tests and results. Section 4 discusses the relationship between (i) the PEU model and other theoretical approaches towards uncertainty and (ii) our empirical results and other empirical studies of uncertainty and ambiguity attitudes. This section also discusses recent empirical studies of uncertainty/ambiguity. Section 5 concludes and provides directions for future work.

⁶It is tempting to draw parallels between non-constant uncertainty attitudes and non-constant risk attitudes as manifested by the Allais paradox (Allais 1953). Indeed, the β -PEU model shares similarities with risky choice models that can accommodate variable risk attitudes such as weighted utility (Chew and MacCrimmon 1979, 1983, 1989; Fishburn 1983). We further explore these similarities in Section 5.

2 Experimental design and methods


To study preferences for lower envelope lotteries we collected choices made in an experiment. The experiment contained 25 choice situations. Each choice situation contained six lower envelope lotteries. Each of the six lower envelope lotteries had only two possible outcomes: 60 or 20 Swiss Francs (CHF).⁷ In addition, one of the six lower envelope lotteries always had zero uncertainty, making it analogous to a risky lottery. The lower envelope lotteries in a choice situation were arranged in such a way that there was a constant tradeoff between minimum probabilities and uncertainty. This design can be interpreted as a traditional demand elicitation mechanism, or ‘budget experiment’, for lower envelope lotteries. The design varies (i) the relative tradeoff between minimum probabilities and uncertainty, a value we interpret as a price (i.e. the slope of a linear budget), and (ii) the probability assigned to the CHF 60 outcome for the lower envelope lottery with zero uncertainty.

To better illustrate how the design can be viewed as a traditional budget experiment, consider the three example budgets in Fig. 2a. The thin, solid line segment represents a choice situation in which the tradeoff between uncertainty and the minimum probabilities of receiving each outcome is one-to-one. Specifically, if uncertainty is reduced by two units this results in a one unit increase in each of the minimum probabilities. This is easy to see when moving from the end of the budget at the lower-right vertex (100% uncertainty), to the end of the budget at the midpoint on the hypotenuse (the 50/50 lottery). This movement completely reduces the 100% uncertainty into an equal division of probabilities. Next, consider the thick budget in Fig. 2a that connects the points (0.4, 0, 0.6) on the vertical leg of the triangle and the 70/30 lottery on the hypotenuse. This line retains the tradeoff between uncertainty and minimum probabilities, but increases the probability for the CHF 60 outcome in the zero uncertainty option. We call this type of movement a ‘tradeoff-constant increase in likelihood,’ or just ‘likelihood increase.’ This is analogous to an increase in wealth while holding price constant. Finally, consider the dashed line in Fig. 2a. The dashed line and thick line have the same probability for the CHF 60 outcome in the zero uncertainty option (i.e. both budgets intersect the hypotenuse of the uncertainty triangle at the same point). The dashed line, however, changes the tradeoff between minimum probabilities and uncertainty such that a two unit reduction in uncertainty results in a more-than-one unit increase in the minimum probability of CHF 60, and a smaller-than-one-unit increase in the minimum probability of CHF 20. Put simply, for the dashed line segment in Fig. 2a, the tradeoffs are such that the minimum probability for CHF 60 is cheap.

The 25 budgets used in our experiment are plotted in Fig. 2b. The black lines depict the budgets. The dots on each line depict the alternatives available on that budget. This wide range of variation in tradeoffs (prices) and likelihoods (wealth) permits a thorough examination of preferences towards lower envelope lotteries. Table 3 in Section C of the [Electronic Supplementary Material](#) provides a detailed summary of each budget. Also in Section C of the [Electronic Supplementary Material](#), Table 4

⁷ At the time of the experiment, one CHF was worth approximately \$1.10.

(a) Three budgets illustrating (i) a tradeoff-constant increase in likelihood (thin line to thick line) and (ii) a likelihood-constant tradeoff change (thick line to dashed line)



(b) The 25 budgets used in our experiment, each featuring six alternatives




Fig. 2 The Experimental Design

provides a comprehensive list of the lower envelope lotteries available on each of the 25 budgets.

2.1 Elicitation software

Participant choices were elicited using software programmed with the Psychtoolbox Matlab libraries freely available at <http://www.psychtoolbox.org>. Figure 3 shows screenshots from our elicitation software for one choice situation (budget). The row of boxes at the bottom of each screenshot, each with a thumbnail image, provide a visual depiction of all the lower envelope lotteries available from that budget. When participants moved their mouse over one of these boxes, the corresponding thumbnail image was drawn as a large image in the center of the screen. The large image provided detailed outcome, minimum probability, and uncertainty information for the alternative in the highlighted thumbnail.

Figures 3a, b, and c depict alternatives available from one budget in the experiment. More specifically, these figures depict alternatives available from the budget depicted by the thin solid line segment in Fig. 2a. In the row of boxes at the bottom of each screenshot in Fig. 3 the left-most alternative is the 50/50 lottery (i.e. the lower envelope lottery $(0.5, 0.5, 0)$ which is the midpoint on the hypotenuse in the uncertainty triangle). The right-most alternative is the 100% uncertain alternative (i.e. the lower envelope lottery $(0, 0, 1)$ which is represented as the lower-right vertex in the uncertainty triangle). Figure 3a illustrates when the mouse was moved over the thumbnail depicting the 100% uncertain alternative – a green outline surrounds the right-most thumbnail and the large image shows the outcome information for this alternative and that no minimum probability information is provided. Figure 3b

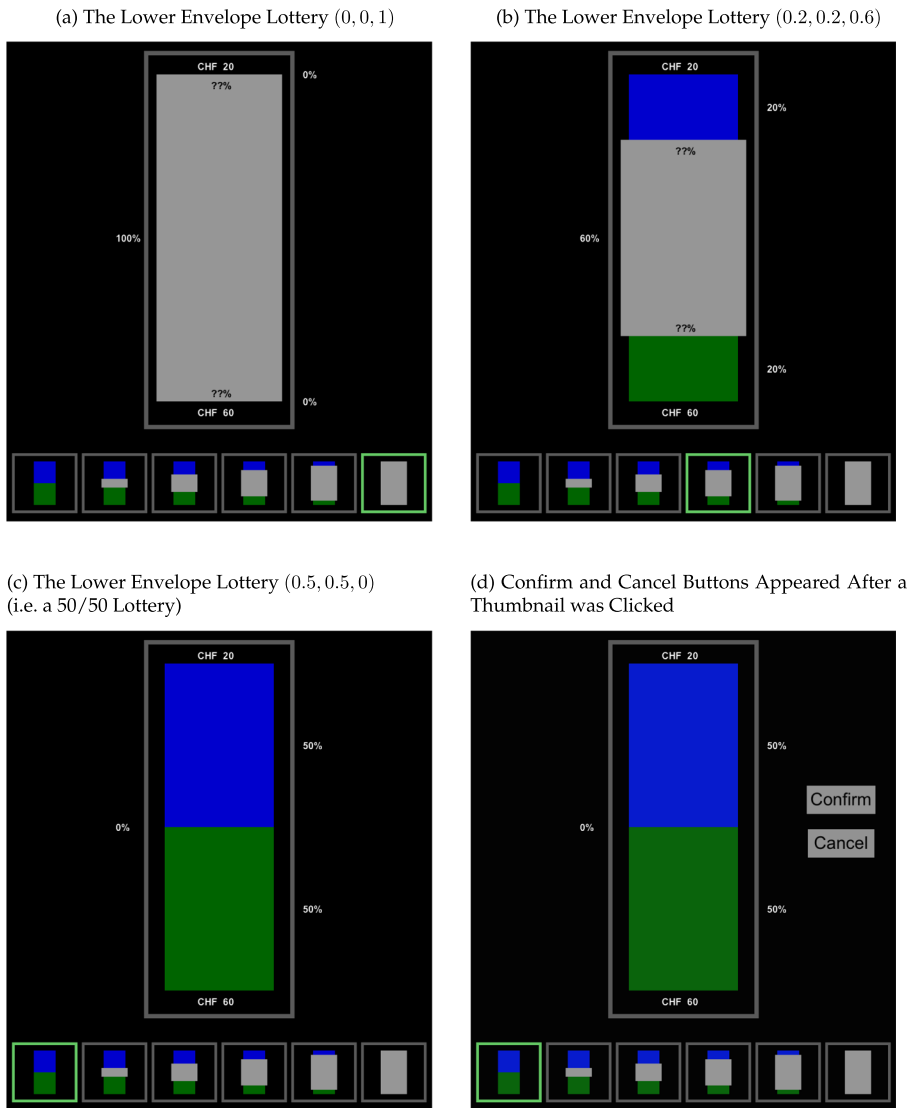


Fig. 3 Screenshots from our Elicitation Software. Panels (a)–(c) depict different alternatives from the same budget that corresponds to the thin solid line-segment in Fig. 2a (i.e. the budget with endpoints (0, 0, 1) and (0.5, 0.5, 0). Panel (d) depicts “Confirm” and “Cancel” buttons that appeared after participants clicked one of the thumbnails at the bottom of the screen


depicts when the mouse highlighted an alternative from the interior of the budget – a green outline surrounds the thumbnail and the specific minimum probability and uncertainty information are shown in the large image. Figure 3c depicts when the mouse highlighted the lower envelope lottery with no uncertainty that was

available from this budget – a green outline surrounds the left-most thumbnail and the probability information is shown in the large image.


Participants were instructed to select the alternative they wanted most in each choice-situation. Participants indicated their choice by moving the mouse over the relevant thumbnail, such that it was depicted as the large image in the middle of the screen, and then clicking that thumbnail. After clicking the thumbnail the green outline was locked and “Confirm” and “Cancel” buttons appeared (Fig. 3d shows these buttons when the zero uncertainty lower envelope lottery was selected). If participants clicked the “Cancel” button, the “Confirm” and “Cancel” buttons vanished, and the green outline was unlocked so as to allow a different alternative to be selected in that choice situation. If participants clicked the “Confirm” button the software recorded the participant’s choice and moved on to the next choice-situation.

The row of thumbnail images at the bottom of the screen provided an easy-to-see representation of the tradeoff (price) and likelihood (wealth) in a given budget which helped to minimize participant confusion. To see how this works consider Figs. 4a, b, and c. Figure 4a is a closeup of the thumbnails in the screen shots in Fig. 3. It shows a one-to-one tradeoff between uncertainty and minimum probabilities. Figure 4b illustrates how the thumbnails made it easy to see a tradeoff constant likelihood shift – these thumbnails correspond to alternatives available from the thick budget in Fig. 2b. Figure 4c illustrates how the thumbnails made it easy to see a likelihood constant tradeoff shift – these thumbnails correspond to alternatives available from the dashed budget in Fig. 2a.

(a) Thumbnails representing the alternatives available from the thin line segment in Figure 2a



(b) A Tradeoff Constant Likelihood Increase – Thumbnails representing the alternatives available from the thick line segment in Figure 2a



(c) A Likelihood Constant Tradeoff Change – Thumbnails representing the alternatives available from the dashed line segment in Figure 2a




Fig. 4 A Zoomed in View of the Thumbnails Depicting the Alternatives Available from the Three Budgets in Fig. 2a

2.2 Resolution of uncertainty

The choice-situations in our experiment made available many distinct lower envelope lotteries. To resolve the uncertainty for each of these lower envelope lotteries, with a single source, and in one stage, we constructed an urn with 100 balls in it. To illustrate how the urn permitted single source and single stage uncertainty resolution consider the lower envelope lottery (0.4, 0.4, 0.2). Figure 5 illustrates how this lower envelope lottery appeared in the experiment while Figure 6 illustrates how the urn would be used to resolve uncertainty from this lower envelope lottery. In the urn there were exactly 40 balls numbered 1 to 40 and exactly 40 balls numbered 61 to 100 (illustrated in blue and green, respectively, in Fig. 6). There were also exactly 20 balls in the interval 41 to 60. It was unknown, however, how many of these 20 balls had an odd-number on them and how many had an even-number on them. This uncertainty in odds and evens is what allowed the urn to resolve the uncertain component of the lower envelope lotteries in our experiment. Specific to the lower envelope lottery (0.4, 0.4, 0.2) shown in Fig. 5: If the ball drawn from the urn had any number 1 to 40 on it, the lottery paid out CHF 20. If the ball drawn from the urn had any number 61 to 100 on it, the lottery paid out CHF 60. If the ball was between 41 and 60 and odd numbered, the lottery paid out CHF 20. And if the ball was between 41 and 60 and even-numbered the lottery paid out CHF 60. Of course the ratio of odd- and even-numbered balls was unknown – in the extreme they could have been all odd or all even. Thus, the key was that odd and even balls in the range of uncertainty were used to resolve the unknown portion of the lower envelope lottery. In this way our urn allowed us to resolve the uncertainty in one stage for any of the lower envelope lotteries in our experiment.⁸

2.3 Making draws from the urn

To make draws from our urn we used a commercially available bingo blower. As can be seen in Fig. 7 we occluded the windows on our bingo blower with purple masking tape so that balls could be seen mixing inside the blower but the numbers on the balls could not be seen. Before participants entered the lab they were told that a bingo blower would be used as a randomization device to determine their payment in the experiment. To familiarize participants with the bingo blower it was turned on and placed in the corridor where participants entered the laboratory.

The bingo-blower used can create a queue of balls in the drawing tube. To avoid ball-queuing in the drawing tube, we inserted a stopper. As can be seen in Fig. 7,

⁸Following Hey et al. (2010), and the re-assessment of those data in Kothiyal et al. (2014), we wanted to insure that participants were not concerned about experimenters ‘stacking the urn’ against the participants. Thus, experimenters were blind to the exact composition of the urn. This was accomplished by giving various students, faculty, and staff in the authors’ home department envelopes with two balls, a pen, and instructions telling them to, for example, write 7 on both balls, write 8 on both balls, or 7 on one ball and 8 on the other ball, and then to seal the envelope. The balls were placed in the urn, and quality control for readability and consistency with the instructions, was done by a person not involved with the experiment. Experimental participants were told how the urn was constructed during the instructions, and informed that experimenters were blinded to the exact composition of the urn.

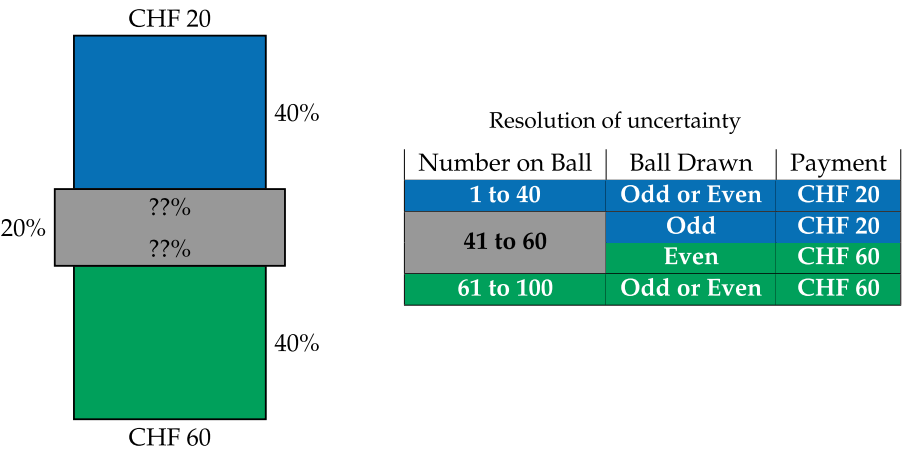


Fig. 5 How the Urn Resolves Uncertainty For the Lower Envelope Lottery (0.4, 0.4, 0.2)

the stopper was a wooden dowel, cut to the same length as the drawing tube, with a large wooden ball attached to the top. When participants were at the cashier’s desk, they removed the stopper so that one ball was “drawn” by the bingo blower. After the participant left the cashier’s desk, the drawn-ball and stopper were replaced in the bingo blower by the experimenter.

3 Analyzing choices

In this section we analyze choices made by 203 experimental participants. All experimental sessions were conducted during one week in July, 2014 and all experimental





Fig. 6 How The Urn Resolves Uncertainty For the Lower Envelope Lottery (0.4, 0.4, 0.2)

Fig. 7 The Bingo Blower and Stopper Used to Make Draws from the Urn



procedures were consistent with a protocol approved by a human subjects ethics committee.

Our analysis proceeds in three steps. The first step assesses whether participants' choices adhere to the generalized axiom of revealed preference (GARP). Choices that adhere to GARP can be rationalized with a preference relation that is complete, transitive, continuous, and monotone. These are all foundational assumptions which underlie the PEU model along with many other theories of choice under uncertainty.

The second analytical step uses a non-parametric test to assess whether choices can be assumed to arise through PEU maximization. Using this non-parametric test, we identify participants whose choices can be assumed to have arisen from PEU maximization. For these PEU-maximizers we calibrate uncertainty attitudes, as captured by α , using parametric methods. We start by assuming homogeneous α 's and, one-by-one, increase the number of possible preference types.

This two-step analytical procedure, which combines non-parametric and parametric assessments, is distinct from the purely parametric approach used by Stahl (2014) and Hey and Pace (2014). For example, Hey and Pace (2014) used experimental data in an assessment of five single-stage models of decision-making under uncertainty in which beliefs were taken to be endogenous. Their experimental design had participants betting on which color of ball (with three possible colors) would be drawn from an urn that was mixed by a bingo blower. Participants were shown the actively mixing bingo blower that contained the urn in a transparent plexiglass box.⁹ Hey and Pace (2014) then used parametric estimation techniques to describe and predict choices. They found support for α -MaxMin Expected Utility (Marinacci 2002) as a

⁹The transparency of the bingo blower is an important consideration when interpreting the results in Hey and Pace (2014). Transparency means that some information regarding the composition of the urn would be available to participants through visual inspection. For example, the extreme possibility of the urn being comprised of just one color of ball could be easily ruled out.

representation for behavior in their experimental data, in addition to Vector Expected Utility (Siniscalchi 2009).

The third and final step in our analysis examines choices made by the subset of participants who were not PEU maximizers. We closely examine the choice data to understand why PEU failed as a parsimonious model. We propose an extension, β -Partial Ignorance Expected Utility, that provides a good fit for this sub-sample of participants. The β -PEU model retains the linear indifference curves of PEU but, its additional parameter, β , allows for fanning-in or fanning-out across the uncertainty triangle.

3.1 Assessing the generalized axiom of revealed preference (GARP)


Choices from a set of linear budgets (with strictly positive prices) are consistent with the maximization of a preference relation that is complete, transitive, continuous, and monotone if and only if they adhere to the generalized axiom of revealed preference (Afriat 1967; Varian 1982). In an experimental setting such as ours, the generalized axiom is an easy-to-test condition and Varian (1982) provides an easy-to-implement algorithm. Technical details are in Section A.1 of the [Electronic Supplementary Material](#).

Adherence to an axiomatic criterion is binary: Either choices conform perfectly to GARP or they do not. While taking a binary perspective on GARP-compliance can be informative, a widely accepted approach is to determine, conditional on observing a violation, just “how badly” choices departed from perfect GARP-compliance. A common measure used in this regard is Houtman-Maks (HM). Houtman-Maks is the largest subset of choices which are GARP-compliant (Houtman and Maks 1985). So, for example, if we see that a full set of 25 choices violates GARP but, by removing one offending choice, the remaining choices are GARP-compliant, this would represent a Houtman-Maks of 24.

To determine whether a participant’s choices are GARP-compliant, we compare their Houtman-Maks to a critical value derived from a Monte Carlo simulation. Our simulation uses choices for 5,000 synthetic experimental participants, each with a uniform random choice rule (see e.g. Bronars 1987). We calculated Houtman-Maks for each synthetic. We interpret the distribution of all 5,000 synthetics’ Houtman-Maks as a sampling distribution and use it to determine critical values. This sampling distribution is illustrated with the blue bars in Fig. 8a, and exact values are reported in Fig. 8b. The critical value for the 95% confidence level implied by this sampling distribution is 21. Relative to this critical value, 198 of our 203 experimental participants (98%) had a Houtman-Maks that met or exceeded this threshold (i.e. $HM \geq 21$). If the cutoff for GARP-compliance is raised to the critical value for the 99% confidence level of 22, 187 participants (92%) had a Houtman-Maks that met or exceeded this threshold (i.e. $HM \geq 22$). Generally speaking, we find very high rates of GARP-compliance.¹⁰

¹⁰ The procedure introduced by Beatty and Crawford (2011) yields the same conclusion: Given our set of budgets it is very unlikely to pass GARP with uniform random choice. Put differently, our GARP test has very high power.

(a) A High Rate of GARP-compliance Illustrated by the Distribution of Houtman-Maks for Experimental Participants (red) and a Sampling Distribution Implied by 5,000 Synthetics With a Uniform Random Choice Rule (Blue)



(b) Supporting Data

Houtman-Maks	Humans			Synthetics		
	Freq.	%	Cum.	Freq.	%	Cum.
12	0	0	0	2	0.04	0.04
13	0	0	0	10	0.2	0.24
14	0	0	0	86	1.72	1.96
15	0	0	0	309	6.18	8.14
16	0	0	0	728	14.56	22.7
17	0	0	0	1,093	21.86	44.56
18	0	0	0	1,158	23.16	67.72
19	1	0.49	0.49	885	17.7	85.42
20	4	1.97	2.46	489	9.78	95.2
21	11 [†]	5.42	7.88	188	3.76	98.96
22	21 [†]	10.34	18.23	43	0.86	99.82
23	39 [†]	19.21	37.44	7	0.14	99.96
24	45 [†]	22.17	59.61	1	0.02	99.98
25	82 [†]	40.39	100	1	0.02	100
Total	203	100	-	5,000	100	-


[†] – Exceeds the critical value for the 95% confidence level established by 5,000 synthetics with a uniform random choice rule.

Fig. 8 Houtman-Maks for the Generalized Axiom of Revealed Preference (GARP)

3.2 Assessing the Partial Ignorance Expected Utility maximization hypothesis

In the uncertainty triangle, PEU implies indifference curves that are linear and parallel. This geometry necessitates a particular pattern of choices in an experiment like ours. Consider, for example, the three budgets, each with a distinct slope, depicted in Fig. 9. Suppose that the 50/50 lottery was chosen from the middle budget (i.e.

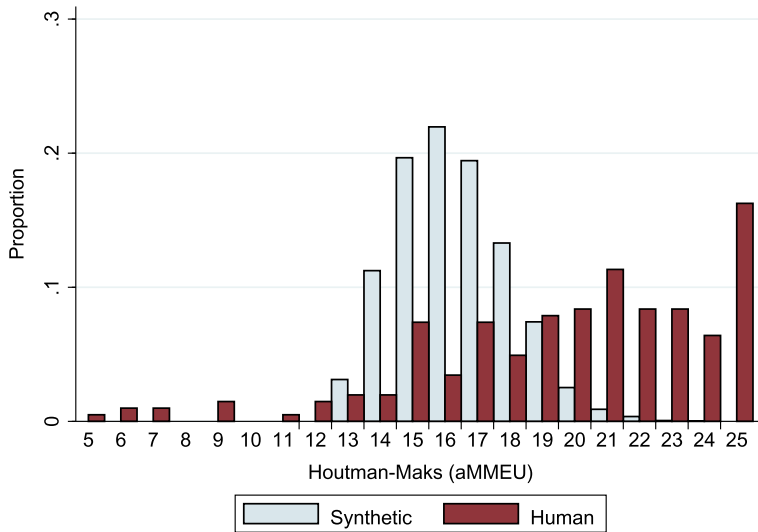
Fig. 9 Testing for Linear and Parallel Indifference Curves As Required Under Partial Ignorance Expected Utility



the budget connecting the 100% uncertain alternative to the 50/50 lottery). Any linear indifference curve that rationalizes this choice must lie in the shaded area. And because PEU indifference curves must be parallel, the choice from any steeper budget, like the upper-most budget in Fig. 9, must also be the option on the hypotenuse. When a budget is flatter than the middle budget, however, no such prediction can be placed on choices. If, however, the original choice from the middle budget was the 100% uncertain alternative then the reverse structure would be required. Flatter budgets would require the most uncertain alternative available to be selected, while no such requirement could be placed on choices from steeper budgets. If the original choice was an alternative from the strict interior of the middle budget, any PEU indifference curve that rationalizes this choice must lie on top of that budget line. So, for any steeper budget the lottery must be selected while for any flatter budget the most uncertain alternative must be selected. The comprehensive technical details for testing PEU maximization in this way, and algorithmic details, can be found in Section A.2 of the [Electronic Supplementary Material](#).

Paralleling the logic of assessing “how badly” choices can depart from the generalized axiom, we adapt the logic of Houtman-Maks (HM) to our test of PEU maximization. We define “Houtman-Maks PEU” (HM-PEU) as the largest subset of choices that conform to our test. Figure 10a is a histogram of HM-PEU for the 203 participants in our experiment (red bars). The supporting data are in Fig. 10b. The blue bars in Fig. 10a show the sampling distribution for HM-PEU implied by 5,000 synthetic experimental participants who had a 50/50 choice rule for just the endpoints of the budgets – these 5,000 synthetics’ choices were restricted to be either the fully-specified lottery or the most uncertain alternative. The critical value for the 95% confidence level for this sampling distribution was 20. Of the 203 participants in our experiment, 120 (59.1%) had an HM-PEU that met or exceeded this critical value. Considering the critical value for the 99% confidence level of 21, 103 participants (50.7%) had an HM-PEU that met or exceeded this value. Generally speaking, the assumption of PEU maximization is appropriate for the majority of our experimental participants.

(a) Many Subjects Made Choices Consistent With PEU Maximization



(b) Supporting Data

Houtman-Maks	Humans			Synthetics		
	Freq.	%	Cum.	Freq.	%	Cum.
5	1	0.49	0.49	0	0	0
6	2	0.99	1.48	0	0	0
7	2	0.99	2.46	0	0	0
9	3	1.48	3.94	0	0	0
10	0	0	3.94	0	0	0
11	1	0.49	4.43	0	0	0
12	3	1.48	5.91	0	0	0
13	4	1.97	7.88	156	3.12	3.12
14	4	1.97	9.85	562	11.24	14.36
15	15	7.39	17.24	983	19.66	34.02
16	7	3.45	20.69	1,098	21.96	55.98
17	15	7.39	28.08	972	19.44	75.42
18	10	4.93	33	665	13.3	88.72
19	16	7.88	40.89	371	7.42	96.14
20	17 [‡]	8.37	49.26	126	2.52	98.66
21	23 [‡]	11.33	60.59	45	0.9	99.56
22	17 [‡]	8.37	68.97	18	0.36	99.92
23	17 [‡]	8.37	77.34	3	0.06	99.98
24	13 [‡]	6.4	83.74	1	0.02	100
25	33 [‡]	16.26	100	0	0	100
Total	203	100	-	5,000	100	-

‡ – Exceeds the critical value for the 95% confidence level established by 5,000 synthetics with a 50/50 choice rule for only the end-points of a budget.

Fig. 10 Houtman-Maks for the Partial Ignorance Expected Utility (PEU) Representation

3.2.1 Estimating Partial Ignorance Expected Utility parameters

This section describes the procedures we used to estimate PEU preference parameters. We assume random utility (RU) with additively separable choice noise (McFadden 1981). Denoting L^{jk} as the k^{th} lower envelope lottery available in the j^{th} choice-situation, the random utility from choosing L^{jk} is given by

$$RU(L^{jk}) = \text{PEU}(L^{jk}) + \varepsilon^{jk}, \quad (4)$$

where ε^{jk} is an independent and identically distributed random variable. The systematic, PEU-component of $RU(L^{jk})$, for the CHF 60 and CHF 20 outcomes available in our experiment, can be written as

$$\begin{aligned} \text{PEU}(\underline{p}_{60}, \underline{p}_{20}, y) = & \alpha \left[\underline{p}_{60} u_{60} + (\underline{p}_{20} + y) u_{20} \right] \\ & + (1 - \alpha) \left[(\underline{p}_{60} + y) u_{60} + \underline{p}_{20} u_{20} \right]. \end{aligned} \quad (5)$$

Normalizing outcome utilities such that $u_{20} = 0$ and $u_{60} = 1$, Expression (5) simplifies to

$$\text{PEU}(\underline{p}_{60}, \underline{p}_{20}, y) = \underline{p}_{60} + (1 - \alpha)y, \quad (6)$$

an expression with just one parameter (α) that captures an individual's everywhere-constant uncertainty attitude. A participant chooses the k^{th} lower envelope lottery (i.e. $L^{j*} = L^{jk}$) if and only if $RU(L^{j*}) \geq RU(L^{jk})$. Thus, the probability that the subject chooses the k^{th} lower envelope lottery in choice situation j is given by

$$\text{Prob}[L^{j*} = L^{jk}] = \text{Prob}[\text{PEU}(L^{j*}) - \text{PEU}(L^{jk}) \geq \varepsilon^{jk} - \varepsilon^{j*}]. \quad (7)$$

We assume that ε^{jk} follows a type I extreme value distribution with scale parameter σ . So the choice probabilities take the standard form:

$$\text{Prob}[L^{j*} = L^{jk}] = \frac{\exp\left(\frac{\text{PEU}(L^{jk})}{\sigma}\right)}{\sum_{k=1}^6 \exp\left(\frac{\text{PEU}(L^{jk})}{\sigma}\right)}. \quad (8)$$

From these choice probabilities we obtain individual i 's contribution to the likelihood function

$$f(\mathbb{L}; \theta) = \prod_{j=1}^{25} \prod_{k=1}^6 \text{Prob}[L^{jk}]^{\mathbb{1}[L_i^{j*} = L_i^{jk}]}, \quad (9)$$

where \mathbb{L} is the set of all lower envelope lotteries in the choice experiment and the function $\mathbb{1}$ is an indicator that returns one if the condition in the bracket is true and zero otherwise.

Under the assumption that all participants have the same α (i.e. homogeneous preferences), the log-likelihood function is

$$\mathcal{L}_{\text{PEU}}^1 = \sum_{i=1}^n \ln f(\mathbb{L}; \alpha, \sigma), \quad (10)$$

where the subscript i denotes the i^{th} of the $n = 120$ individuals with a PEU representation. The PEU model involves estimation of two parameters, the index of uncertainty aversion α , and the dispersion parameter σ .

We then relax the assumption of homogenous preferences and explore the possibility that our data were generated by multiple and distinct preference types. To do so, we employ a finite mixture approach.¹¹ The principal idea of such models is assigning each subject to one of C different preference types. Each type is endogenously characterized by a distinct vector of parameters, (α_c, σ_c) , with $c \in \{1, \dots, C\}$. In addition, the estimation procedure yields estimates of the proportion of the sample that belongs to each type, π_c . Summing over all C behavioral types yields the complete log-likelihood function

$$\mathcal{L}_{\text{PEU}}^C = \sum_{i=1}^n \ln \sum_{c=1}^C \pi_c f(\mathbb{I}; \alpha_c, \sigma_c). \quad (11)$$

In finite mixture models the number of behavioral types must be fixed prior to estimation. There is a large literature discussing the optimal number of types but there is no common agreement on which measure is optimal (for a summary see McLachlan and Peel 2000). Here, we follow a similar route as Bruhin et al. (2010) and stop incrementing the number of types as soon as adding more types does not generate novel qualitative insights. This means, for example, that if a model with only one type indicates that this type is uncertainty averse while a model with two types shows that there is also a substantial minority of uncertainty seeking types, the characterization of the population with two types is preferable because otherwise we would overlook the substantial minority of subjects whose behavior is qualitatively different (i.e., uncertainty seeking).

To estimate parameters we use the iterative expectation maximization (EM) algorithm. Details can be found in Dempster et al. (1977) and the Supplementary Material of Bruhin et al. (2010). Reported standard errors are obtained by bootstrapping with 1,000 replications (Efron and Tibshirani 1993). Resampling is done at the level of participants. We prevent the order of types in the parameter vector from changing during resampling (i.e. label switching) by first calculating the Euclidean distance between the replication types' parameter vectors and the original types' parameter vectors, and then reordering the full parameter vector accordingly.

3.2.2 Estimation results: Examining uncertainty attitude types

Estimation results are reported in Table 1. Model 1 assumes homogeneous preferences across the sub-sample of 120 participants that are PEU representable. The parameter estimates and bootstrapped 95% confidence intervals for the α 's are plotted in blue in the top portion of Fig. 11. That the confidence interval lies above 0.5 indicates that, on average, participants exhibited uncertainty aversion.

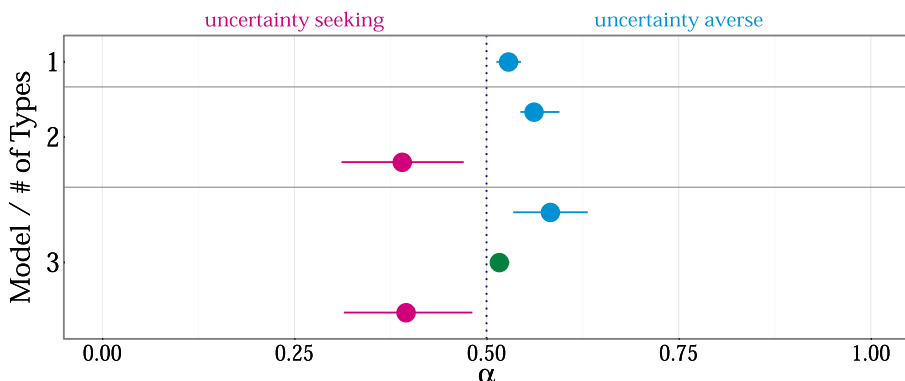
¹¹ Applications of finite mixture models are discussed in El Gamal and Grether (1995), Stahl and Wilson (1995), Houser et al. (2004), Bruhin et al. (2010), Fehr-Duda et al. (2010), and Fehr-Duda and Epper (2012).

Table 1 Model estimates for the 120 participants representable with Partial Ignorance Expected Utility

		Model 1 (One Type)		Model 2 (Two Types)		Model 3 (Three Types)	
Type 1	π_1 (s.e.)	1.000	(-)	0.795*	(0.096)	0.483*	(0.081)
	α_1 (s.e.)	0.529 ^{*1/2}	(0.008)	0.562 ^{*1/2}	(0.017)	0.583 ^{*1/2}	(0.024)
	σ_1 (s.e.)	0.025*	(0.002)	0.020*	(0.002)	0.026*	(0.004)
Type 2	π_2 (s.e.)	-		0.205	(-)	0.302*	(0.070)
	α_2 (s.e.)	-		0.390 ^{*1/2}	(0.042)	0.517 ^{*1/2}	(0.008)
	σ_2 (s.e.)	-		0.029*	(0.005)	0.005*	(0.004)
Type 3	π_3	-		-		0.215	(-)
	α_3 (s.e.)	-		-		0.395 ^{*1/2}	(0.053)
	σ_3 (s.e.)	-		-		0.029*	(0.004)
Individuals		120		120		120	
Observations		3,000		3,000		3,000	
\mathcal{L}		-3,190		-2,882		-2,681	
BIC		6,396		5,804		5,425	

Remarks ^{*1/2} denotes that α_j is different from 1/2 (uncertainty neutrality) at the 95% confidence level. * denotes difference from 0 at the 95% confidence level. Standard errors are calculated using 1,000 bootstrap replications. For further details see Section B of the [Electronic Supplementary Material](#)

Model 2 relaxes the restriction of homogeneous preferences and allows for two preference types. Fitted uncertainty parameters, α_1 and α_2 , and their standard errors (s.e.) are shown in the middle column of Table 1. The middle portion of Fig. 11 shows the bootstrapped 95% confidence intervals for these two preference types. The confidence interval for Type 1 preferences, colored in blue, is greater than 0.5, consistent with uncertainty aversion. The confidence interval for Type 2 preferences, colored in red, is less than 0.5, consistent with uncertainty seeking. The finite mixture model

**Fig. 11** Bootstrapped 95% Confidence Intervals For The Uncertainty Parameters (α) in Table 1

also yields an estimate for the parameter π_1 , the proportion of the sample that is associated with the Type 1 preferences. This estimate, shown in the first row for the Type 1 preferences in Table 1, indicates that about 80% of participants are uncertainty averse. The remaining 20% of participants belong to the Type 2 preference type, or uncertainty seeking.

The estimates for Model 3, which has three preference types, are in the right-most column of Table 1 and the bootstrapped 95% confidence intervals for the three preference types are shown in the lower portion of Fig. 11. As with Model 2, the **blue** and **red** confidence intervals correspond to uncertainty aversion and neutrality. The **green** confidence interval, from a strictly statistical perspective, also corresponds to uncertainty aversion. However, the difference between the fitted parameter and 0.5 is negligible – it is only 3% larger than 0.5. We therefore interpret this type as exhibiting (near) uncertainty neutrality. In terms of the estimated proportions for each of these three preference types, **48%** are uncertainty averse, **30%** are nearly uncertainty neutral, and **22%** are uncertainty seeking. If we further increase the number of types to four or five, no new qualitative insights emerge. There is always a type best characterized by near uncertainty neutrality and comparatively little choice noise and the overall share of uncertainty averse types remains relatively stable. Thus, the main consequence of increasing the number of types to four or five is the emergence of subdivisions among the uncertainty averse and the uncertainty seeking types.¹²

3.3 Examining why PEU maximization failed

Section 3.2 reveals that a significant proportion of participants made choices that cannot be modeled with PEU maximization (41%, or 83 of 203 participants; see Figure 10a for details). Two explanations for why PEU failed as a parsimonious model seem plausible. First, it could be that the everywhere-constant uncertainty attitude embedded in the PEU representation was violated, while the linearity of indifference curves was retained. That indifference curves would still be linear means that choices would be predominantly at the endpoints (corners) of the budgets in our experiment. In this case, the fraction of choices between the fully-specified lottery endpoint and the ‘most uncertain’ endpoint should vary across the uncertainty triangle. A second explanation for why PEU failed as a parsimonious model would be that indifference curves in the uncertainty triangle were non-linear. A pattern of choices consistent with this explanation would mean that choices were, to a large extent, on the interior of the budgets in our experiment.

For this sub-sample of 83 participants we observe very low proportions of choices on the interior of the budgets. And the bubble plot in Fig. 12 is illustrative of this general pattern. It displays choices made from five budgets, each with the same tradeoffs but differing elevations in the triangle. As can be seen in the figure, only a small proportion of choices are on the interior of these budgets. Put another way, there are a large proportion of choices at the endpoints of the budgets. The ratio of choices between the fully-specified lottery endpoint and the most uncertain endpoint,

¹²More details on the quality of classification can be found in Section A.2 of the [Electronic Supplementary Material](#).




Fig. 12 A Bubble Plot of the Choices Made by 83 Participants on Five Budgets With a Constant Price Illustrating Non-constant Uncertainty Attitudes Across the Uncertainty Triangle

however, varies across the uncertainty triangle. This pattern of choices is roughly consistent with the first explanation above. We therefore consider a generalization of PEU that relaxes the everywhere constant uncertainty attitude only.

3.3.1 The β -Partial Ignorance Expected Utility model

As an empirical expedient for accommodating a non-constant uncertainty attitude we introduce a generalization of PEU. We call this generalization the β -Partial Ignorance Expected Utility (β -PEU) model. The β -PEU model weakens the everywhere-constant uncertainty attitude embedded in PEU by allowing indifference curves to fan-in or fan-out across the uncertainty triangle. The additional parameter, β , controls the extent to which indifference curves fan-in or fan-out across the uncertainty triangle. Formally, the β -PEU of a lower envelope lottery in our two-outcome experiment is given by

$$\beta\text{-PEU}(\underline{p}_{60}, \underline{p}_{20}, y) = \tilde{\alpha} \left[(\underline{p}_{20} + y) u_{20} + \underline{p}_{60} u_{60} \right] + (1 - \tilde{\alpha}_i) \left[\underline{p}_{20} u_{20} + (\underline{p}_{60} + y) u_{60} \right], \quad (12)$$

where $\tilde{\alpha}$ is a parameter that systematically varies based on the lower envelope lottery being evaluated. Specifically, for a lower envelope lottery $(\underline{p}_{60}, \underline{p}_{20}, y)$ in the experiment,

$$\tilde{\alpha} = \alpha + \beta \left(\underline{p}_{60} - \frac{1}{2} + \frac{y}{\sqrt{2}} \right). \quad (13)$$


Here, the parameter α is “baseline” uncertainty aversion while β captures whether uncertainty attitude is increasing ($\beta > 0$) or decreasing ($\beta < 0$) while moving north-east in the triangle. The term $p_{60} + y/\sqrt{2}$ is the probability of the high outcome for the fully specified lottery that an uncertainty neutral person would be indifferent to. This provides a value for the “elevation” of the lower envelope lottery. Geometrically, y is the length of a line segment orthogonal to the hypotenuse and the lower envelope lottery (i.e. the Euclidean distance from the hypotenuse to the lower envelope lottery’s location in the triangle). The fraction $y/\sqrt{2}$ gives the elevation, or distance, along the hypotenuse, of that orthogonal connector. Subtracting one half centers this value on the line segment connecting the lower right vertex and the midpoint on the hypotenuse. The interpretation is that the β -PEU model has a systematically varying uncertainty attitude, where the parameter β determines whether uncertainty attitude is increasing or decreasing while moving northeast in the triangle.

Figure 13 shows example indifference curves for two parameterizations of β -PEU. Both examples have $\alpha = \frac{1}{2}$ so that the indifference curve intersecting the uncertainty triangle at the 50/50 lottery (i.e. the midpoint of the hypotenuse) has a slope equivalent to uncertainty neutrality. Figure 13a shows a positive value of β which produces indifference curves that ‘fan-in’ across the uncertainty triangle. Put another way, positive values of β mean that aversion to uncertainty is increasing when moving northeast in the triangle. Negative values of β , like those depicted in Fig. 13b, have the opposite pattern – indifference curves ‘fan-out’ when moving northeast in the triangle.

3.3.2 Estimating β -Partial Ignorance Expected Utility parameters

In general we follow the same procedures as those in Section 3.2.1. The only differences being that we now consider the $n = 83$ participants whose choices could

(a) Increasing aversion to uncertainty:
 $\alpha = 1/2, \beta > 0$ (‘Fanning In’)



(b) Decreasing aversion to uncertainty:
 $\alpha = 1/2, \beta < 0$ (‘Fanning Out’)




Fig. 13 Indifference Curves for the β -Partial Ignorance Expected Utility Model

not be rationalized with a PEU representation. Accordingly, we replace the systematic component of random utility with β -PEU and estimate three parameters per type: $(\alpha_c, \beta_c, \sigma_c)$. Normalizing the outcome utilities for the systematic β -PEU component of random utility, so that $u_{20} = 0$ and $u_{60} = 1$, we have

$$\beta\text{-PEU}(\underline{p}_{60}, \underline{p}_{20}, y) = \underline{p}_{60} + (1 - \tilde{\alpha}) y. \quad (14)$$

Again, the number of behavioral types (C) has to be fixed prior to estimation. We follow the same procedure as in Section 3.2.1 by incrementing the number of types until additional types provide no new qualitative insights. Here, we end up with $C = 2$ types, both showing considerable fanning across the simplex. Interestingly, one type exhibits uncertainty seeking at low elevations in the triangle and uncertainty aversion at high elevations. Perforce, these preferences appear uncertainty neutral near the ‘middle’ of the uncertainty triangle. The other type reveals an average behavior which is everywhere uncertainty seeking, but decreasingly so at higher elevations in the triangle.¹³

3.3.3 Estimation results: Examining β -Partial Ignorance Expected Utility preference types

Estimation results are reported in Table 2. Model 1 assumes homogeneous preferences across the sub-sample of 83 participants. The parameter estimates and a bootstrapped 95% confidence region for preference parameters α and β are plotted in red in Fig. 14a. The confidence intervals indicate that α is less than 0.5 and β is positive. In general, these homogeneous parameter estimates exhibit uncertainty seeking in the lower portions of the triangle, neutrality through the middle, and slight to moderate aversion in the upper elevations. This is easy to see in Fig. 15a which plots indifference curves for Model 1 as red lines in the uncertainty triangle.

Model 2 estimates two types of preferences. Fitted preference parameters, α_1 , β_1 , α_2 , and β_2 , and their standard errors (s.e.), are shown in the right column of Table 2. Bootstrapped 95% confidence regions for both preference types are shown in Fig. 14b. The confidence region for Type 1 preferences is colored green while the confidence region for Type 2 preferences is colored red. Indifference curves for these two types are plotted in Fig. 15b. Type 1 preferences exhibit a pattern of behavior similar to the homogeneous preferences in Model 1: uncertainty seeking in the lower portion of the triangle, neutrality through the middle, and uncertainty aversion in the upper elevations. Type 2 preferences are everywhere uncertainty seeking through the lower and middle portions of the triangle and uncertainty neutral through the upper elevations. Approximately 62% of this sample is estimated to be associated with the Type 1 preferences (see π_1 for Model 2 in Table 2). Broadly speaking, the parameter

¹³We also estimated models with $C = 3$ types. The additional type can be broadly described as a subgroup of the ‘everywhere uncertainty averse’ type when $C = 2$. We also validated the β -PEU model by estimating it with the sub-sample of 120 ‘vanilla’ PEU types from Section 3.2 above. As expected for these validations, all estimated β ’s are not statistically different than zero.

Table 2 Estimates for the 83 participants modeled with β -Partial Ignorance Expected Utility

		Model 1 (One Type)		Model 2 (Two Types)	
Type 1	π_1 (s.e.)	1.000	(-)	0.618*	(0.093)
	α_1 (s.e.)	0.458 ^{*1/2}	(0.014)	0.506	(0.021)
	β_1 (s.e.)	0.464 *	(0.021)	0.439 *	(0.028)
	σ_1 (s.e.)	0.042*	(0.003)	0.031*	(0.003)
Type 2	π_2 (s.e.)	-		0.382	(-)
	α_2 (s.e.)	-		0.359 ^{*1/2}	(0.036)
	β_2 (s.e.)	-		0.528 *	(0.058)
	σ_2 (s.e.)	-		0.045*	(0.010)
Individuals		83		83	
Observations		2,075		2,075	
\mathcal{L}		-2,868		-2,747	
BIC		5,760		5,548	

Remarks: ^{*1/2} denotes that α_j is different from 1/2 at the 95% confidence level. * denotes difference from 0 at the 95% confidence level. Standard errors are calculated using 1,000 bootstrap replications. For further details see Section B of the [Electronic Supplementary Material](#)

estimates for Model 2 are consistent with the bubble plots discussed above – aversion to uncertainty increases when moving northeast in the triangle.¹⁴

4 Related literature

We do not attempt to review the large literature on ambiguity and uncertainty in this section – Camerer and Weber (1992), Camerer (1995), Etner et al. (2010), and Machina and Viscusi (2014b) already provide excellent reviews. Instead, we limit ourselves to a discussion of recent empirical research and refer to relevant theories where appropriate. Also, in contrast to the body of this paper where uncertainty and ambiguity are clearly distinguished, we soften this distinction here to facilitate discussion. While we want to emphasize that the findings in this paper are appropriate only for the two outcome lower envelope setting studied here, we take liberties in this section to make connections with existing papers that examine ambiguity preferences.

Motivated by the so-called ‘Ellsberg paradox,’ there has been extensive theoretical work focused on models that can accommodate ambiguity aversion and, as a result, canonical Ellsberg behavior (Ellsberg 1961). For example, Schmeidler (1989) put forward Choquet Expected Utility. This model assumes unique beliefs that can

¹⁴More details on the quality of classification can be found in Section B of the [Electronic Supplementary Material](#).

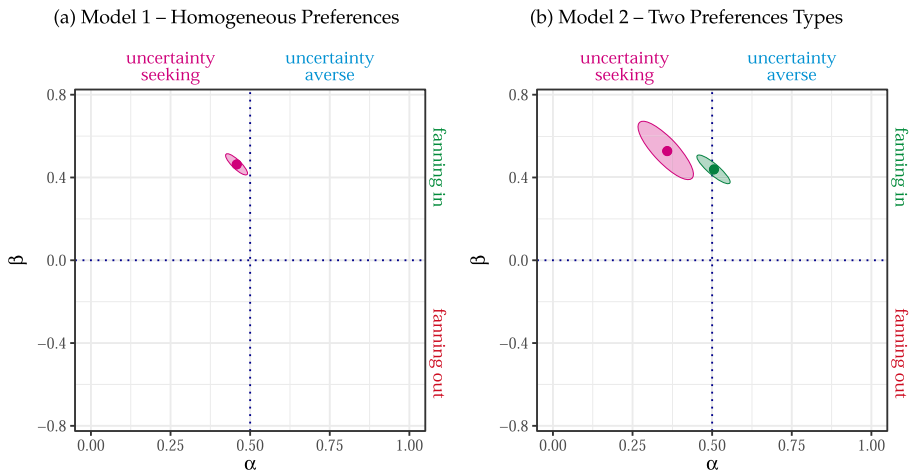


Fig. 14 Bootstrapped 95% Confidence Regions Based on 1,000 Replications for the β -Partial Ignorance Expected Utility Preference Parameters Reported in Table 2




Fig. 15 Indifference Curves for the β -Partial Ignorance Expected Utility Preference Parameters Reported in Table 2. The dashed lines are for reference and represent uncertainty neutrality. Note that we only observe choices in a sub-domain of the uncertainty triangle – Indifference curves in the upper-right and lower-left corner are therefore extrapolations

be represented with *capacities*. Unlike probabilistic beliefs, capacities need not be additive which makes it possible to rationalize Ellsberg-type behavior.

In contrast to the unique beliefs assumed in the Subjective and Choquet Expected Utility models, Gilboa and Schmeidler (1989) proposed a model in the multiple priors class. The so-called MaxMin Expected Utility model assumes that bets induce sets of probabilistic beliefs. A bet is then evaluated by the belief that has the lowest Expected Utility. A natural counterpart to the MaxMin Expected Utility model is the MaxMax Expected Utility model – a bet is evaluated by the belief that has the highest Expected Utility. Marinacci (2002) and Ghirardato et al. (2004) proposed an amalgam of the MaxMin and MaxMax Expected Utility models. The so-called α -MaxMin Expected Utility model represents a mixture of MaxMin and MaxMax evaluations of bets.¹⁵ In this way α -MaxMin Expected Utility is based on both beliefs and ambiguity attitudes or tastes (as captured by α). From an empirical perspective, this can present challenges when jointly estimating beliefs and preference parameters (Hey et al. 2010; Kothiyal et al. 2014).

At first blush, α -MaxMin and PEU appear quite similar. They are distinguished, however, in two important ways. First, PEU preferences are defined over sets of lotteries whereas α -MaxMin preferences are defined over acts (bets). Second, in the PEU model the best and the worst possibilities (\bar{L} and \underline{L}) are objectively defined entities. For the α -MaxMin model, however, the worst and the best probability distribution over outcomes are determined by an individual's subjective beliefs.¹⁶ As mentioned before, empirical identification of these subjective components is very challenging and requires comprehensive data or additional, potentially critical, assumptions. Thus, from an empirical perspective, the PEU approach can be advantageous because the slope of the indifference curve captures all the relevant information about uncertainty attitudes. Moreover, this slope can be easily measured with budgets in the space of lower envelope lotteries.

Our use of lower envelope lotteries to study uncertainty attitudes contributes to the small, but growing, literature that explores preferences over ‘sets of lotteries.’ While lower envelope lotteries are simpler than the more general ‘sets of lotteries’ setting investigated theoretically by Olszewski (2007) and Ahn (2008), this simplicity can be advantageous for the empirical study of uncertainty attitudes. It allows, for example, applications of non-parametric revealed preference assessments of model classes. More specifically, we can use the simplicity of lower envelope lotteries to, roughly speaking, assign a price to uncertainty. This, in turn, permits non-parametric tests of (i) the generalized axiom of revealed preference and (ii) whether preferences can be characterized by linear and parallel indifference curves. The parsimony of lower envelope lotteries also permits a tractable and intuitive way to describe the pattern of

¹⁵The theory introduced in Ghirardato et al. (2004) considers beliefs as a component of the theory (i.e. beliefs are endogenous). An α -MaxMin Expected Utility model without endogenous beliefs was axiomatized by Jaffray (1994).

¹⁶Putting aside the difference in domain (i.e. lower envelope lotteries vs. acts), these two models can make identical behavioral predictions if an α -MaxMin Expected Utility maximizer has priors over the full range of probabilities.

choices that results from a violation of constant uncertainty attitudes such as when indifference curves “fan-in” or “fan-out”.

Our paper is also related to a growing literature that examines ambiguity attitudes in experimental settings (e.g. Abdellaoui et al. 2011; Ahn et al. 2014; Halevy 2007; Hey et al. 2010, 2014; Stahl 2014; Baillon and Placido 2019). For example, Abdellaoui et al. (2011) shows that ambiguity attitudes can vary significantly depending on the source of ambiguity. The goal of that study was, however, different from ours as it did not examine the GARP-compliance of subjects nor the heterogeneity in preferences in those data. Instead, Abdellaoui et al. (2011) developed a new tool (source functions) for translating events (i.e. collections of states) into a willingness to bet. By comparing source functions elicited for urns with known and unknown compositions, the paper provides a rich characterization of attitudes towards ambiguity. Although Abdellaoui et al. (2011) focused on different questions, those data also suggest that varying attitudes towards ambiguity/uncertainty are a prominent feature of behavior in these types of settings. Specifically, that paper found that for large probabilities ($p > 0.5$) source functions were significantly higher for an urn with a known composition (urn K) than for an urn with a completely unknown composition (urn U). In contrast, when probabilities were small ($p < 0.5$) there was no meaningful difference in source functions between the urns. Because the difference between the source functions for K and U provides a quantitative measure of ambiguity aversion, this finding suggests that ambiguity aversion increases with the likelihood of the good outcome. Our non-parametric test for constant uncertainty aversion suggests, however, that increasing uncertainty aversion is not a property of all subjects; 60% of our subjects make choices consistent with constant uncertainty attitudes. And within this group we find individuals who exhibit uncertainty aversion, uncertainty seeking and uncertainty neutrality. But we also find that a substantial share of subjects (40%) exhibit increasing uncertainty aversion. This property of our data can be parsimoniously captured using the β -PEU model with indifference curves that ‘fan-in’ across the uncertainty triangle.

Ahn et al. (2014) examined ambiguity preferences by asking subjects to allocate a budget between three Arrow-securities. Each security paid unity in one of three mutually exclusive states, and zero otherwise. For one of the states the probability that it would be realized was known to be $\frac{1}{3}$. For the other two states the exact probabilities were unknown. The approach employed in Ahn et al. (2014) highlights a natural complementarity between examining ambiguity preferences with Arrow-securities and the lower envelope lottery approach employed here. Ahn et al. (2014) relied on varying monetary outcomes under a fixed set of probability distributions. In contrast, the lower envelope lottery approach used here varies minimum probabilities and uncertainty (i.e. the set of lotteries) while holding outcomes fixed. Our paper also differs from Ahn et al. (2014) because we test, non-parametrically, for constant uncertainty attitudes and show that roughly 40% of our subjects display variable uncertainty attitudes. Our findings suggest that the assumption of the everywhere-constant ambiguity attitude made in Ahn et al. (2014) may not be appropriate for all subjects. In addition, we find a much higher percentage of subjects that display uncertainty aversion: Ahn et al. (2014) found that only 10% of their subjects exhibited ambiguity aversion. In our data 29% of subjects exhibited constant uncertainty

aversion and 23% displayed varying uncertainty aversion that increased from moderate to high levels as the likelihood for the good outcome increases. Future research should clarify whether these differences are methodological or due to differences in sampled populations.

In a setup distinct from ours, Halevy (2007) examined different explanations for non-neutrality towards ambiguity. The experiments in Halevy (2007) distinguished between theories that model ambiguity aversion as violations of probabilistic sophistication and theories that model choice under uncertainty using two-stage prospects. He found substantial support for the latter explanation. While our paper asks different questions, by focusing on GARP compliance and parsimonious empirical characterizations of heterogeneity in uncertainty attitudes, we believe that Halevy (2007)'s finding suggests a promising extension of our present study. In particular, in our setup uncertainty is objectively resolved in one stage. So, violations of the reduction of compound lottery axiom are unlikely to play a role in our observed non-neutral uncertainty attitudes. We could, however, easily introduce a two-stage resolution of uncertainty in our setup. By comparing behavior under one-stage and two-stage resolution of uncertainty it would be straightforward to measure the change in non-neutral uncertainty attitudes that is due to reducing compound lotteries.

Hey et al. (2010, 2014) examine data from experiments involving bets on balls mixing in a transparent bingo blower. Both of these papers employed parametric estimation techniques to examine the differential fits for a range of models at the individual level. As demonstrated in these ambitious studies (see also Kothiyal et al. 2014), joint identification of preference parameters and beliefs using, for example, an α -MaxMin-Expected Utility specification, can be challenging. Due to the large number of parameters that have to be estimated, identification requires both a large number of observations and strong auxiliary parametric assumptions. The examination of indifference curves in the uncertainty triangle, however, obviates these challenges. In our setting uncertainty is an objective quantity and there are not beliefs. So, the slope of indifference curves in the triangle captures uncertainty attitudes under PEU (or simple extensions of PEU in the case of varying uncertainty attitudes).

Stahl (2014) examined an experiment in which participants made a series of choices in settings similar to Ellsberg's two- and three-color configurations. The primary focus of Stahl (2014) was to characterize heterogeneity in ambiguity attitudes. To accomplish this the paper provides estimates of (i) individual-level models, (ii) a representative agent model, and (iii) finite mixture models for a representation which characterized agents/types by two parameters: an index of ambiguity aversion and an error dispersion parameter. These estimates do not permit an examination of varying ambiguity attitudes. In terms of results, Stahl (2014) found a comparatively small proportion of subjects that exhibited ambiguity aversion ($\approx 12\%$) and a rather large group of subjects ($> 60\%$) that revealed a behavior "... suggest[ing] that the majority of participants found the choice tasks profoundly confusing" (Stahl 2014, p.617). We believe that participants in our experiment were not confused because (i) our choice situations were very easy to understand (see Section 2), (ii) the rate of GARP-compliance was extremely high, (iii) we found no types with a random choice rule, and (iv) the assignment of individuals to different preference classes is very clean.

5 Discussion and conclusion

This paper introduced two-outcome lower envelope lotteries as a simple framework for examining choice in the face of uncertainty. This simplicity has advantages. For example, because the slope of indifference curves in the uncertainty triangle captures uncertainty attitudes this constitutes a useful tool for examining basic, yet important, empirical patterns of behavior under uncertainty. It enables the examination of GARP-compliance, allows for a non-parametric assessment of whether uncertainty attitudes are constant or vary across the triangle and, in combination with finite mixture methods, characterizations of the heterogeneity of uncertainty attitudes.

For future research the setting of lower envelope lotteries can be used as an input toward a richer characterization of choices made in the face of uncertainty. One potential direction could be an exploration of whether individuals that exhibit increasing uncertainty aversion also commit Allais-type violations of Expected Utility (Allais 1953; Burghart 2020). For Allais-type violations it is widely taken that risk aversion is increasing as the likelihood for the good outcome increases (see Machina (1982), Hypothesis II and Figure 5b). Here, we show that many people exhibit increasing uncertainty aversion as likelihood increases. This raises the question of whether individuals who exhibit increasing uncertainty aversion will also exhibit increasing risk aversion.

Another direction that seems promising is using the lower envelope lottery setting to explore uncertainty preferences when there are three or more outcomes. Machina (2014a) examines such a setting and highlights how many models of choice under uncertainty include an *informational symmetry* assumption. In a three outcome lower envelope setting this would require, for example, $(0, 0, 0, 1) \sim (1/3, 1/3, 1/3, 0)$. Generalizing PEU to three (or more) outcomes, however, would imply $(0, 0, 0, 1) \sim (1/2, 0, 1/2, 0)$. This represents a situation where empirical interrogation of these contrasting predictions could prove fruitful.¹⁷

Yet another direction that seems promising is using lower envelope lotteries to examine attitudes towards uncertainty when a risky alternative is not available in the choice set. If behavior differs between environments with and without a risky alternative it could challenge existing theories of choice under uncertainty. Further, in a setting where lower envelope lotteries are defined on complementary events our setup could shed more light on violations of probabilistic sophistication (Baillon et al. 2018). For example, Baillon and Bleichrodt (2015) demonstrate that the additivity of disjoint events, as required by probabilistic sophistication, fails for naturally occurring uncertainties. In addition, Baillon and Placido (2019) highlight how decreasing ambiguity aversion is important both theoretically and empirically. The lower envelope setting could shed additional light on this important feature of behavior.

Exploring whether one- or two-stage resolution of lower envelope lotteries has an effect on behavior would be another area worth exploring. Such an extension seems natural given that Halevy (2007) documents a relationship between canonical

¹⁷The observations in this paragraph were generously provided by a referee.

ambiguity aversion and violations of the reduction of compound lotteries. Lower envelope lotteries provide a simple and intuitive setting for deeper explorations of these questions, and others, in the domain of uncertainty.

Acknowledgments Special thanks go to Aurélien Baillon and Peter Wakker as they gave us detailed feedback on an early draft of this paper, in addition to some key references. All authors gratefully acknowledge valuable input from Jim Cox, Charles Efferson, Tony Williams, anonymous referees, participants in the seminars at Princeton University, University of Zurich, Virginia Tech, and Wesleyan University, participants at the 2014 Economic Science Association North American meeting held in Fort Lauderdale, Florida, participants at RUD 2016 (Paris-Dauphine), and participants at the 2016 European Meeting of the Econometric Society in Geneva. Sonja Vogt was instrumental in preparing the urn used to resolve uncertainty. And we also thank the students, faculty and staff in the Department of Economics at UZH for their assistance in composing the urn. Stefan Wehrli, Oliver Brägger and Nadja Jehli provided technical and practical assistance during data collection. Christina Stoddard provided excellent editing and typesetting advice. This work was supported by a grant from the European Research Council on the Foundations of Economic Preferences (#295642).

References

- Abdellaoui, M., Baillon, A., Placido, L., Wakker, P.P. (2011). The rich domain of uncertainty: Source functions and their experimental implementation. *The American Economic Review*, 101, 695–723.
- Afriat, S.N. (1967). The construction of utility functions from expenditure data. *International Economic Review*, 8, 67–77.
- Ahn, D.S. (2008). Ambiguity without a state space. *The Review of Economic Studies*, 75, 3–28.
- Ahn, D., Choi, S., Gale, D., Kariv, S. (2014). Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics*, 5, 195–223.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21, 503–546.
- Baillon, A., & Bleichrodt, H. (2015). Testing ambiguity models through the measurement of probabilities for gains and losses. *American Economic Journal: Microeconomics*, 7, 77–100.
- Baillon, A., Huang, Z., Selim, A., Wakker, P.P. (2018). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, 86, 1839–1858.
- Baillon, A., & Placido, L. (2019). Testing constant absolute and relative ambiguity aversion. *Journal of Economic Theory*, 181, 309–332.
- Beatty, T.K.M., & Crawford, I.A. (2011). How demanding is the revealed preference approach to demand? *The American Economic Review*, 101, 2782–2795.
- Bronars, S.G. (1987). The power of nonparametric tests of preference maximization. *Econometrica*, 55, 693–698.
- Bruhin, A., Fehr-Duda, H., Epper, T. (2010). Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78, 1375–1412.
- Burghart, D.R. (2018). Maximum probabilities, information, and choice under uncertainty. *Economics Letters*, 167, 43–47.
- Burghart, D.R. (2020). The two faces of independence: Betweenness and homotheticity. *Theory and Decision*, 88, 567–593.
- Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5, 325–370.
- Camerer, C. (1995). Individual decision making. In Kagel, A., & Roth, J. (Eds.) *The handbook of experimental economics* (pp. 587–683). Princeton: Princeton University Press.
- Chew, S., & MacCrimmon, K. (1979). Alpha utility theory, lottery compositions and the Allais paradox. Unpublished manuscript, University of British Columbia.
- Chew, S. (1989). Axiomatic utility theories with the betweenness property. *Annals of Operations Research*, 19, 273–298.

- Chew, S.H. (1983). A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica*, 51, 1065–1092.
- Dempster, A.P. (1967). Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38, 325–339.
- Dempster, A., Laird, N., Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B*, 39, 1–38.
- Efron, B., & Tibshirani, R.J. (1993). *An introduction to the bootstrap*. New York: Chapman & Hall.
- El Gamal, M.A., & Grether, D.M. (1995). Are people Bayesian? Uncovering behavioral strategies. *Journal of the American Statistical Association*, 90, 1137–1145.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Etner, J., Jeleva, M., Tallon, J.-M. (2010). Decision theory under ambiguity. *Journal of Economic Surveys*, 26, 234–270.
- Fehr-Duda, H., Bruhin, A., Epper, T., Schubert, R. (2010). Rationality on the rise: Why relative risk aversion increases with stake size. *Journal of Risk and Uncertainty*, 40, 147–180.
- Fehr-Duda, H., & Epper, T. (2012). Probability and risk: Foundations and economic implications of probability-dependent risk preferences. *Annual Review of Economics*, 4, 567–593.
- Fishburn, P.C. (1983). Transitive measurable utility. *Journal of Economic Theory*, 31, 293–317.
- Ghirardato, P., Maccheroni, F., Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118, 133–173.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18, 141–153.
- Gul, F., & Pesendorfer, W. (2014). Expected uncertain utility theory. *Econometrica*, 82, 1–39.
- Gul, F. (2015). Hurwicz expected utility and subjective sources. *Journal of Economic Theory*, 159, 465–488.
- Halevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica*, 75, 503–536.
- Hey, J.D., Lotito, G., Maffioletti, A. (2010). The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity. *Journal of Risk and Uncertainty*, 41, 81–111.
- Hey, J.D., & Pace, N. (2014). The explanatory and predictive power of non two-stage-probability theories of decision making under ambiguity. *Journal of Risk and Uncertainty*, 49, 1–29.
- Houser, D., Keane, M., McCabe, K. (2004). Behavior in a dynamic decision problem: an analysis of experimental evidence using a Bayesian type classification algorithm. *Econometrica*, 72, 781–822.
- Houtman, M., & Moks, J.A.H. (1985). Determining all maximal data subsets consistent with revealed preference. *Kwantitatieve Methoden*, 19, 89–104.
- Jaffray, J.-Y. (1994). Dynamic decision making with belief functions. In Yager, R.R., Fedrizzi, M., Kacprzyk, J. (Eds.) *Advances in the Dempster-Shafer theory of evidence* (pp. 331–352). New York: Wiley.
- Knight, F.H. (1921). *Risk, uncertainty and profit*. Boston: Hart, Schaffner & Marx; Houghton Mifflin Co.
- Kothiyal, A., Spinu, V., Wakker, P.P. (2014). An experimental test of prospect theory for predicting choice under ambiguity. *Journal of Risk and Uncertainty*, 48, 1–17.
- Luce, R.D., & Raiffa, H. (1958). *Games and decisions*, 3rd edn. New York: Wiley.
- Machina, M. (1982). Expected utility: Analysis without the Independence Axiom. *Econometrica*, 50, 277–323.
- Machina, M.J. (2014a). Ambiguity aversion with three or more outcomes. *The American Economic Review*, 104, 3814–3840.
- Machina, M., & Viscusi, W.K. (2014b). *Ambiguity and Ambiguity Aversion*, 1st ed, North-Holland.
- Marinacci, M. (2002). Probabilistic sophistication and multiple priors. *Econometrica*, 70, 755–764.
- McFadden, D. (1981). In Manski, C. (Ed.) *Econometric models of probabilistic choice*. Cambridge: MIT Press.
- McLachlan, G., & Peel, D. (2000). *Finite mixture models*. Wiley series in probabilities and statistics. New York: Wiley.
- Olszewski, W. (2007). Preferences over sets of lotteries. *The Review of Economic Studies*, 74, 567–595.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57, 571–587.
- Shafer, G. (1976). *A mathematical theory of evidence*. Princeton: Princeton University Press.
- Siniscalchi, M. (2009). Vector expected utility and attitudes toward variation. *Econometrica*, 77, 801–855.

- Stahl, D.O., & Wilson, P.W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10, 218–254.
- Stahl, D.O. (2014). Heterogeneity of ambiguity preferences. *The Review of Economics and Statistics*, 96, 609–617.
- Varian, H.R. (1982). The nonparametric approach to demand analysis. *Econometrica*, 50, 945–973.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Daniel R. Burghart¹ · Thomas Epper^{2,3} · Ernst Fehr³

Thomas Epper
thomas.epper@unisg.ch

Ernst Fehr
ernst.fehr@econ.uzh.ch

- ¹ Department of Economics, California State University Sacramento, 6000 J Street, Sacramento, CA 95757, USA
- ² School of Economics and Political Science, University of St. Gallen, Varnbühlstrasse 19, 9000 St. Gallen, Switzerland
- ³ Department of Economics, University of Zürich, Blümlisalpstrasse 10, 8006 Zürich, Switzerland