Eliciting Multiple Prior Beliefs

Mohammed Abdellaoui* Philippe Colo† Brian Hill‡

Abstract

Despite the increasing importance of multiple priors in various domains of economics, choice-based incentive-compatible multiple-prior elicitation remains an open problem. This paper develops a solution, comprising a preference-based identification of a subject's probability interval for an event, and a method for eliciting it. The method applies under weak decision-theoretic assumptions, with no need for probabilistic sophistication. To demonstrate its feasibility, we implement it in three incentivized experiments on artificial and natural sources of uncertainty. Intervals elicited by our method are sensitive to the direction and amount of information, and are typically consistent with 'objective' probabilities where available. We find a predominance of non-degenerate probability intervals, with intervals being wider when there is less information or predictability. The probability intervals elicited with our method are similar to those stated by subjects on aggregate, suggesting that the method can provide behavioral foundations for the use of stated probability-interval techniques in the field.

Key words: Multiple Priors, Probability Intervals, Belief Measurement, α-maxmin EU, Imprecise Probability.

JEL Codes: D9, D81

*HEC Paris & CNRS, 1 rue de la Libération, 78351 Jouy-en-Josas, France. E-mail: abdellaoui@hec.fr.

†ETH Zurich & Universität Bern, Länggassstrasse 49a, 3012 Bern, Switzerland. E-mail: colo.philippe@gmail.com.

‡HEC Paris & CNRS, 1 rue de la Libération, 78351 Jouy-en-Josas, France. E-mail: hill@hec.fr.

1. Introduction

The standard Bayesian model of decision under uncertainty stipulates that a decision maker's beliefs are fully captured by a single probability measure (Savage, 1954; Anscombe and Aumann, 1963). Empirical applications often call for elicitation of subjective beliefs (Manski, 2004), and a wide array of probability elicitation methods have been proposed including, beyond *stated* probabilities, scoring rules and matching-probability based approaches. Importantly, the latter are choice based and incentive compatible, and hence can be used to evaluate simpler methods and ground their use in the field. For instance, studies showing that stated probability elicitation methods often lead to limited performance loss compared to choice-based approaches provide a principled foundation for their use in large-scale field studies (Trautmann and Kuilen, 2015).

Elicitation of subjective probabilities plays a significant role in areas such as macroeconomics—with interest in beliefs concerning future demand or inflation (Guiso and Parigi, 1999; Engelberg et al., 2011)—as well as development and agricultural economics—where important factors include agents' beliefs about future weather, market factors or outcomes of crop, technological or entrepreneurial choices (Delavande et al., 2011; Cerroni, 2020). Such future events often involve significant uncertainties, especially in times of crisis, change or innovation. Uncertainties of this scale, and behavioural evidence concerning them, have motivated the development of multiple prior decision models (Gilboa and Schmeidler, 1989; Ghirardato et al., 2004), which replace the Bayesian single-prior representation of beliefs by a set of priors, generating a probability interval for each event. A rich theoretical literature has documented characteristic differences in insurance and investment decisions taken by multiple prior agents as compared to Bayesian ones (Dow and da Costa Werlang, 1992), with qualitatively distinct consequences in macroeconomics (Ilut and Schneider, 2014), asset pricing (Garlappi et al., 2007; Epstein and Schneider, 2010), mechanism design (Bose and Renou, 2014), health economics (Giustinelli et al., 2022) and climate economics (Hill, 2024). However, despite this evidence that multiple-prior-generated imprecision is a potential driver of various economic phenomena—and indeed, despite its use for communication by several institutions, e.g. the Intergovernmental Panel on Climate Change and central banks (Mastrandrea et al., 2010; Carney et al., 2019)—probability elicitation methods rule it out. They thus cannot allow us to properly ascertain its role and leverage its potential. Elicitation of multiple prior beliefs is needed.

The situation concerning multiple prior *elicitation* is markedly different from that for Bayesian probability, with almost all attempts to date focusing on subjects' *stated* probability intervals (Giustinelli et al., 2022; Kriegler et al., 2009). In particular, the absence of *bona fide* theoretically well-founded, choice-based and incentive-compatible multiple-prior elicitation methods deprives stated intervals of a firm grounding, hence raising potential questions about the findings based on them. This paper proposes such an elicitation method for multiple-prior probability intervals, implements it in a series of laboratory experiments and compares it to stated intervals.

An impossibility result from the statistics literature (Seidenfeld et al., 2012, Prop 5) provides a flavour of the challenge posed by incentive-compatible elicitation of multiple priors: there exist no real-valued continuous strictly proper scoring rule for multiple-prior probability intervals. Related issues affect the matching probability method (Borel, 1939; Anscombe and Aumann, 1963). It elicits the matching probability (MP)—that is, the proportion of red balls in an unambiguous red-and-blue-balled urn at which the subject is indifferent between betting on red from the urn and betting on a target event E. Under Subjective Expected Utility (SEU), the MP of E coincides with the subject's probability of it. For multiple-prior preferences, however, this is no longer the case. For instance, under the popular (Hurwicz) α-maxmin EU model, the MP reflects the bounds of the subject's probability interval for the event, but also her *attitude to uncertainty* or *ambiguity*. Indeed, even eliciting the MPs of E and its complement E^c (which, beyond SEU, need not add to one) does not allow identification of the subject's probability interval in general, due to the confounding ambiguity attitude factor.[1] Existing theoretical and experimental approaches to this well-known issue (e.g. Ghirardato et al. 2004; Eichberger et al. 2011; Section 5) assume that the subject's set of priors is generated by precise probabilistic beliefs, i.e. preferences are probabilistically sophisticated (Chateauneuf et al. 2007; Baillon et al. 2018b, 2021; Gul and Pesendorfer 2015; Section 5). However, such assumptions are least warranted in situations where multiple priors are most relevant—and hence undermine the suitability of such *precision-laden* methods for multiple-prior probability-interval elicitation. Indeed, to meet the challenge of multiple-prior elicitation, an incentive-compatible, fully general and hence *precision-free* method is required.

[1]For instance, under the Hurwicz α-maxmin EU model, the MP of an event E, $MP(E)$, satisfies:

$$MP(E) = \alpha \underline{p}(E) + (1-\alpha)\overline{p}(E)$$

where the subject's probability interval for E is $[\underline{p}(E), \overline{p}(E)]$ and α is typically interpreted as a reflection of the subject's ambiguity attitude (Section 2.3). Since the probability interval for the complement event E^c is $[1-\overline{p}(E), 1-\underline{p}(E)]$, eliciting the MPs for the event and its complement yields two equations in three unknowns—and hence does not allow identification of the subject's probability interval. See Section 2.4.

Drawing on theoretical results that provide a solution to the identification problem for α-maxmin EU and a wide range of generalisations (Hill, 2023), we develop an MP-like elicitation method that uses extraneous random devices with *interval-valued* rather than precise probabilities. To illustrate, consider an urn containing only red and blue balls, where all that is known is that at least proportion r of the balls in the urn are red, and at least proportion b are blue (with $r+b \leq 1$). Here, the probabilities of getting red or blue on the next draw from the urn are summarized by the intervals $[r, 1-b]$ and $[b, 1-r]$, respectively. To identify the bounds of the subject's probability interval for E, it suffices to find such an urn where the subject is indifferent between betting on E and betting on red from the urn, and between betting against E and betting against red (i.e. on blue). As we show in Section 2.4, under α-maxmin EU and a range of generalisations, the subject's probability interval is given by the interval $[r, 1-b]$ corresponding to this urn. Moreover, this identification holds independently of the subject's ambiguity attitudes.[2] This thus yields a choice-based association of an 'interval-valued' urn to each event, which identifies the subject's probability interval for it. This *matching probability interval* (MPI) notion resolves the problem of choice-based incentive-compatible probability-interval elicitation in theory.

Our approach resolves the aforementioned foundational challenges. First of all, it is theoretically robust, insofar as it operates under Hurwicz α-maxmin expected utility as well as an array of generalisations—and hence without assumptions on subjects' ambiguity attitudes. Moreover, it is precision free, requiring no assumption of precise probabilities underpinning subjects' probability intervals. As discussed in Section 5, beyond distinguishing our approach from those mentioned above, this also differentiates it from scoring rules for most-likely intervals for the value of an unknown parameter (Winkler and Murphy, 1979; Schlag, 2015).

To operationalize elicitation of matching probability intervals, we develop an MPI version of the two-step MP elicitation method adopted by Abdellaoui et al. (2021, 2023). Under their method, a subject undergoes a 'bisection' binary-choice procedure followed by a 'confirmation' choice list; we develop an analogue binary choice procedure and 'two-dimensional' choice list, tailored for eliciting (two-dimensional) probability intervals instead of (one-dimensional) probability values. We implement our method in three laboratory studies. EXP A involves an artificial source of uncertainty—the colour of the next chip drawn from a bag—where prior information was provided through sampling. This controlled environment allows validation testing of the method, via the observed relationship between the elicited intervals and the exogenous

[2]Technically, under α-maxmin EU, these indifferences yield a pair of equations where the ambiguity attitude factor α cancels out, hence leading to a unique solution for the subject's probability interval; see Section 2.4.

information. Moreover, by eliciting stated probability intervals as well, it permits a comparison of the two elicitation approaches. EXP N1 and EXP N2 focus on natural sources of uncertainty, based on continuous variables. There, the method is used to elicit the interval-valued cumulative distribution functions (CDFs) generated by subjects' multiple priors.[3] Interval-valued CDFs are commonly used in applications to go beyond the assumption of precise subjective probabilities (Karanki et al., 2009); our elicitation of CDFs provides a test of our approach, showing that it can operate in such contexts.

Our method passes the validation tests in EXP A, providing intervals that are sensitive to both the direction (e.g. sample frequency) and quantity (e.g. sample size) of information, and that are typically consistent with 'objective' probabilities. On natural sources (EXP N1 and EXP N2), it elicits, for the vast majority of subjects, non-degenerate interval-valued CDFs. All experiments suggest that imprecise beliefs—i.e. intervals of non-zero width—are widespread, providing a choice-based confirmation of the finding of Giustinelli et al. (2022) using stated probability intervals. We also find that the width of elicited intervals decreases when there is more information, familiarity or predictability—a correlation that could be taken to corroborate the solidity of our method. On aggregate, the intervals elicited by our incentive-compatible method in EXP A are generally similar to stated intervals, suggesting that our method provides foundations for some uses of the latter methods in large-scale field studies. Some interesting differences do however emerge, with stated intervals tending to be larger than choice-based ones in information-rich contexts.

The paper is structured as follows. Section 2 sets out the theoretical background and presents the central planks of our approach (the 'matching probability interval' notion and the elicitation method), with the relevant theoretical results. Section 3 describes our experimental implementations, in the form of three studies. Section 4 contains our results and supporting analyses, whereas in Section 5 we discuss connected issues, related literature and future directions. Proofs, further details, data analyses and experimental details are contained in the Appendices.

[3]Many elicitation applications in economics and beyond require subjects' probability distributions or CDFs over a continuous variable of interest (e.g. US inflation in 2025, Eurozone GDP in 2024, average global temperature in 2030).

2. Theoretical Background

In this section, we first set out the general setup, the objects of elicitation and the underlying decision model (Sections 2.1–2.3). Then we present the elements of our method. First, we propose an analogue of MPs for probability intervals and show that they are sufficient to yield the subject's probability interval for an event, in theory (Section 2.4). Then we turn to implementation, presenting, in Section 2.5, an MPI analogue of the two-step of the MP elicitation method developed by Abdellaoui et al. (2021, 2023).

2.1. Bets on events and interval-valued urns

We consider decision-making situations where the objects of choice are two-outcome prospects that pay a fixed monetary outcome z if an event occurs, and nothing otherwise. Prospects with general winning event E and winning amount z are denoted $(z,E,0)$ and called *bets*. The *complementary* bet, which pays out when the event E does not occur, is denoted $(0,E,z)$. Prospects where the probability of winning is exogenously provided in the form of an interval $[\underline{p},\overline{p}]$ are denoted $(z,[\underline{p},\overline{p}],0)$, and are called *interval lotteries* (IL).[4] As for bets, the complementary IL, where the probability of losing is an objectively given interval $[\underline{p},\overline{p}]$, is denoted $(0,[\underline{p},\overline{p}],z)$.

As mentioned previously, interval lotteries are operationalized by urns containing red and blue balls with partial information about the composition. For instance, consider a red-and-blue-balled urn with at least a proportion r of red balls, at least a proportion b of blue balls (with $r+b \leq 1$), but where there is no information about the colour composition of the remaining balls. For such an urn, the information only allows assignment of the interval $[r,1-b]$ for the probability of the next ball drawn from the urn being red; similarly, there is the interval $[b,1-r]$ for the next ball being blue. For the sake of simplicity, we denote the urn with at least proportion r of red balls and at least proportion b of blue balls by $[r,1-b]$. We refer to the set of such *interval-valued urns* by $\mathscr{I}$.[5]

Each urn $[r,1-b]$ in $\mathscr{I}$ can be related to two (sorts of) prospects. One is the prospect that pays z if the next ball drawn from the urn is red, and nothing otherwise. For such a prospect, the probability of winning is characterized by the interval $[r,1-b]$; this thus realises

[4]Our notion of interval lottery is distinct from that used by Gul and Pesendorfer (2014). They use 'interval lottery' to denote (precise) probability measures over the set of intervals of (monetary) prizes; here, 'interval lottery' denotes assignments of probability intervals to (fully determined, precise) outcomes. In particular, the interval lotteries $(z,[r,1-b],0)$ used here clearly do not belong to the concept used by Gul and Pesendorfer (zero probability is assigned to each outcome in the interior of the interval $[0,z]$).

[5]Formally: $\mathscr{I} = \{[x,y] : (x,y) \in [0,1]^2,\ 0 \leq x \leq y \leq 1\}$.

the interval lottery $(z,[r,1-b],0)$. The other prospect involves the complementary bet on this urn—that is, the bet on the next ball drawn from it being blue. Note that the probability of *losing* here is characterised by the interval $[r,1-b]$, so the probability of winning is given by $[b,1-r]$. Hence this prospect realises the complementary IL $(0,[r,1-b],z)$ or equivalently, $(z,[b,1-r],0)$. Standard lotteries correspond to the special case where the composition of the urn is fully known—i.e. $r = 1-b$. So, for instance, the *matching probability* (MP) of an event E can be defined in this setup as the r such that $(z,[r,r],0) \sim (z,E,0)$.

2.2. Probability intervals and interval-valued CDFs

Multiple prior belief representations involve a convex, closed set $\mathscr{C}$ of probability measures. For each event E, the set of priors generates a *probability interval* $\{p(E) : p \in \mathscr{C}\} = [\underline{p}(E), \overline{p}(E)]$, where $\underline{p}(E) = \min\{p(E) : p \in \mathscr{C}\}$ and $\overline{p}(E) = \max\{p(E) : p \in \mathscr{C}\}$ are the *lower* and *upper probabilities for E* respectively. In our experiment on artificial sources of uncertainty, the aim is to elicit probability intervals of the relevant events.

The natural sources of uncertainty in our other experiments are real-valued variables, e.g. the daily minimum temperature in Paris between November and March. In the precise probability case, elicitation aims at revealing the subjective probability over the variable, which can be represented as a subjective cumulative distribution function (CDF). One common way of doing so, for a variable taking values in a real interval T, is by eliciting subjective probabilities of events of the form $E_t = \{t' \in T : t' \leq t\}$, i.e. corresponding to the variable lying below certain fixed values. Indeed, for a probability measure $p \in \Delta(T)$, the CDF is defined as $F_p(t) = p(E_t)$. Analogously, a set of priors $\mathscr{C} \subseteq \Delta(T)$ generates the *interval-valued CDF* $F_{\mathscr{C}}(t) = \{p(E_t) : p \in \mathscr{C}\}$, which takes the probability interval corresponding to E_t as value, for each t. This can be visually represented in terms of two (real-valued) functions: the *lower CDF*, $\underline{F_{\mathscr{C}}}(t) = \min\{p(E_t) : p \in \mathscr{C}\} = \underline{p}(E_t)$, and the *upper CDF*, $\overline{F_{\mathscr{C}}}(t) = \max\{p(E_t) : p \in \mathscr{C}\} = \overline{p}(E_t)$. In these experiments, the aim is to elicit subjects' interval-valued CDFs. Although probability intervals and interval-valued CDFs involve an information loss as compared to sets of priors, they are often sufficient for applications, and sometimes preferable insofar as they are easier to communicate. Indeed, interval-valued CDFs are widely used for representing, communicating and studying sets of priors over continuous variables, where they often go under the name of distribution bands or p-boxes (Berger et al., 2000; Karanki et al., 2009).

2.3. Decision model

We only assume that subjects have preferences over bets and interval lotteries. Large parts of our method hold under the representation where a bet $(z,E,0)$ or interval lottery $(z,[r,1-b],0)$ is evaluated according to:

$$W([\underline{p},\overline{p}])u(z) \tag{1}$$

where $[\underline{p},\overline{p}] = [\underline{p}(E),\overline{p}(E)]$ (the probability interval for E generated by the subjects' set of priors; Section 2.2) in the case of the bet, and $[\underline{p},\overline{p}] = [r,1-b]$ in the case of the IL. In (1), u is a utility function normalized so that $u(0) = 0$, and W is a (real-valued) 'willingness-to-bet function' that is continuous and increasing in both bounds, normalised (i.e. $W([x,x]) = x$ for all x) and strictly increasing in the lower bound. For presentation purposes, we will focus on the special case where W is linear, i.e. where (1) reduces to the Hurwicz α-maxmin EU evaluation of bets and ILs according to:

$$\alpha\underline{p}u(z)+(1-\alpha)\overline{p}u(z) \tag{2}$$

with $\underline{p},\overline{p}$ and u as above. The mixture coefficient $0<\alpha\leq 1$ reflects ambiguity attitude in this model, with higher values being associated with more aversion.[6] For instance, $\alpha = 1$ yields the 'maximally ambiguity averse' Gilboa-Schmeidler (1989) maxmin-EU model. For $1>\alpha>\frac{1}{2}$, (1) accommodates the standard Ellsberg (1961) ambiguity averse preference for a bet on the color of a ball drawn from an urn of known 50-50 composition over a bet on the color of a ball drawn from a 2-color urn of unknown composition, as well as ambiguity seeking behavior at low probabilities.[7] By contrast, such behavior cannot be accommodated when $\alpha<\frac{1}{2}$. Since typical findings suggest some ambiguity seeking behavior at low probabilities, but ambiguity aversion at larger ones (Abdellaoui et al., 2011; Kocher et al., 2018), we take $\alpha>\frac{1}{2}$ to be typical; at a certain point in the presentation, we shall assume that preferences are represented according to (2) with $\alpha>\frac{1}{2}$ (see Sections 2.5 and 5).

Note that the general form (1), which underpins most of the method developed here, can accommodate non-linear, Prospect-Theory-style weighting of the lower and upper probabilities, for instance taking $W([\underline{p},\overline{p}]) = \alpha w(\underline{p})+(1-\alpha)w(\overline{p})$, where w is a weighting function and α

[6] The assumption that W is strictly increasing in the lower bound—i.e. decision makers are sensitive to the lower winning probability—rules out the $\alpha = 0$ case of this model, maxmax-EU. However, there is basically no evidence for such preferences in the population.

[7] For instance, when the probability of red from the 2-color red-and-blue unknown urn is characterized by the interval $[0,1]$, a bet on red from this urn is evaluated as $(1-\alpha)u(z)$ under (2), which is less than the evaluation of a bet on red from the known urn, $\frac{1}{2}u(z)$, when $\alpha>\frac{1}{2}$. However, the evaluation of a bet on the color of a ball drawn from a 10-color urn of unknown composition, $(1-\alpha)u(z)$, is higher than that of a bet on the color of a ball drawn from a 10-equiprobable-color known urn, $0.1u(z)$, whenever $\alpha<0.9$.

is as in the α-maxmin EU model (2). It can also accommodate transformations of the probability interval $[\underline{p}, \overline{p}]$, taking $W([\underline{p}, \overline{p}]) = \alpha \underline{\varphi([\underline{p}, \overline{p}])} + (1-\alpha)\overline{\varphi([\underline{p}, \overline{p}])}$ for α as above and some transformation φ taking probability intervals to probability intervals. Hence it covers cases where the subject's 'real' probability interval is transformed, for instance to incorporate certain ambiguity attitudes, before being used for decision, as in Gajdos et al. (2008). As discussed in detail in Section 5 and Appendix B, the heart of the method applies canonically under such weightings or transformations.

Hill (2023) sets out a formal framework that allows for axiomatic foundations for (2), as well as a range of generalizations.[8] In particular, he shows that introducing ILs allows one to overcome the well-known problem of separating the α factor from the set of priors under α-maxmin EU (Ghirardato et al., 2004; Eichberger et al., 2011), and obtain complete identification of the model. As discussed in more detail there (see notably Hill, 2023, Section 3.3), to the extent that the mixture coefficient reflects a taste (for ambiguity), the use of a single α (or W under (1)) in the evaluation of bets and ILs is consistent with the common practice of using a single utility function for the evaluation of both risky and uncertain prospects, or with the insistence in some parts of the ambiguity literature on the 'portability' of the parameters representing ambiguity attitudes across decision situations (see e.g. Marinacci, 2015, p 1051).

2.4. Matching Probability Intervals

To illustrate our approach, take, as in our EXP A, a bag containing 100 green and yellow chips, where the only information available about its composition comes in the form of four prior draws with replacement, one of which was green. Consider the event E: "the next randomly drawn chip will be yellow". Concerning this event, a multiple-prior decision maker will form a probability interval $[\underline{p}(E), \overline{p}(E)]$ on the basis of the information provided and her beliefs about the proportion of yellow chips in the bag; for instance, it might be $[0.5, 0.9]$. The corresponding interval for the complementary event—which tracks the proportion of green chips in the bag—is $[1-\overline{p}(E), 1-\underline{p}(E)]$, i.e. $[0.1, 0.5]$ in this example. Our aim is to elicit the interval $[\underline{p}(E), \overline{p}(E)]$ for event E.

Standard matching probabilities do not suffice to reveal this subjective probability interval. Indeed, under α-maxmin EU, eliciting matching probabilities for the bets on yellow (E) and green (E^C), $MP(E)$ and $MP(E^c)$, results in the following two equations

[8] See Grant et al. (2019) for an axiomatisation of a special case of (1).

$$
\begin{aligned}
MP(E) &= \alpha\underline{p}(E)+(1-\alpha)\overline{p}(E), \\
MP(E^c) &= \alpha(1-\overline{p}(E))+(1-\alpha)(1-\underline{p}(E)).
\end{aligned} \tag{3}
$$

in three unknowns. Suppose, for instance, that the elicited matching probabilities for yellow and green were 0.66 and 0.26 respectively. Simple calculation reveals that this is consistent with the DM's actual interval for yellow (E) being $[0.5, 0.9]$ if $\alpha = 0.6$, but it would yield the interval $[0.65, 0.75]$ under $\alpha = 0.9$. Since α is unknown, the probability interval for the target event is not uniquely determined by matching probabilities. To solve this identification problem, Hill (2023) supplements the formal setup of the original α-maxmin EU model with the possibility to calibrate subjective probability intervals against objective interval lotteries. As noted previously, in the current paper, the latter are operationalized through red-and-blue-balled urns with partially known composition.

The decision maker's probability interval for E can be mapped to a unique (objective) interval by finding the interval-valued urn $[r, 1-b]$ such that she is indifferent between betting on the yellow chip from the bag (E) and the red ball from the urn, and between betting on the green chip from the bag and the blue ball from the urn. Formally, this yields:

$$
(z, [r, 1-b], 0) \sim (z, E, 0), \tag{4}
$$

$$
(0, [r, 1-b], z) \sim (0, E, z). \tag{5}
$$

We call $[r, 1-b] \in \mathscr{I}$ such that indifferences (4) and (5) hold the *matching probability interval* (MPI) of the event E.

Plugging these indifferences into (1) yields a pair of equations that are clearly satisfied by $r = \underline{p}(E)$, $1-b = \overline{p}(E)$. Under the α-maxmin EU model (2) with $\alpha \neq \frac{1}{2}$, this is the unique solution (Proposition A.2, Appendix A): hence there is a unique MPI, which identifies the subjective probability interval $[\underline{p}(E), \overline{p}(E)]$.[9] As noted in Appendix B, under generic cases of the weighting or probability-interval transformation generalizations of α-maxmin EU discussed in Section 2.3, the MPI is also unique. So to elicit a subject's probability interval for the event

[9] More precisely, under (2) with $\alpha \neq \frac{1}{2}$, the indifferences yield the equations:

$$
\begin{aligned}
\alpha r + (1-\alpha)(1-b) &= \alpha\underline{p}(E)+(1-\alpha)\overline{p}(E), \\
\alpha(1-(1-b)) + (1-\alpha)(1-r) &= \alpha(1-\overline{p}(E))+(1-\alpha)(1-\underline{p}(E)).
\end{aligned} \tag{6}
$$

from which α drops out, yielding a unique solution for $\underline{p}, \overline{p}$.

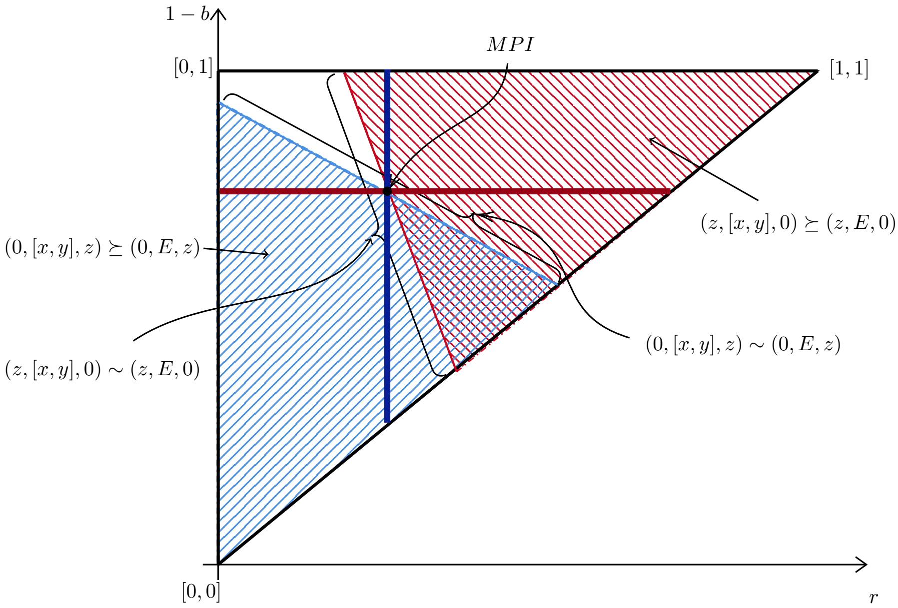

Figure 1: Matching Probability Interval in the space $\mathscr{I}$ of interval-valued urns, for an event E.

E under the main cases of (1), it suffices to find the MPI of E.

The MPI can be conceptually illustrated on Figure 1. A point in the black-edged triangle, (x,y), represents the urn $[x,y]$—i.e. with at least proportion x of red balls and at least proportion $1-y$ of blue ones. As such, it represents two interval lotteries: $(z,[x,y],0)$, the bet on red from the urn, and $(0,[x,y],z)$, the bet on blue. The red hatched area represents the upper contour set (under (2)) of the bet $(z,E,0)$ in the space of interval lotteries corresponding to bets on red: i.e., the set of (x,y) such that $(z,[x,y],0) \succeq (z,E,0)$. The blue hatched area is the upper contour set of the complementary bet $(0,E,z)$ in the space of complementary ILs (corresponding to bets on blue): it is the set of (x,y) such that $(0,[x,y],z) \succeq (0,E,z)$. The boundaries of these sets (the diagonal red and blue lines respectively) represent the indifference curves of $(z,E,0)$ (resp. $(0,E,z)$), in the space of 'red' (resp. 'blue') ILs. The matching probability interval corresponds to the black point at the intersection of these two lines.

Note that standard lotteries and urns with fully known composition correspond to the points on the diagonal ($x=y$) in Figure 1. So the MP of the bet on E is given by the point where the red indifference curve meets the diagonal; and similarly for the MP of E^c and the blue curve. It clear from the Figure that one cannot derive the subject's probability interval from these MPs without knowing the slope of the indifference curves, and this is determined by the ambiguity attitude coefficient α in (2). This is a graphical representation of the previously discussed identification difficulty with MPs. By contrast, the MPI will coincide with the subjective probability interval, independently of the coefficient α.

Name	Preferences	Colour (in Figure 1)
R-B	$(z,[x,y],0) \succeq (z,E,0)$ & $(0,[x,y],z) \succeq (0,E,z)$	Red & Blue
Wh	$(z,[x,y],0) \preceq (z,E,0)$ & $(0,[x,y],z) \preceq (0,E,z)$	White (neither Red nor Blue)
R	$(z,[x,y],0) \succeq (z,E,0)$ & $(0,[x,y],z) \preceq (0,E,z)$	Red
B	$(z,[x,y],0) \preceq (z,E,0)$ & $(0,[x,y],z) \succeq (0,E,z)$	Blue

Table 1: Preference-based division of $\mathscr{I}$

2.5. Elicitation of Matching Probability Intervals

Our strategy for eliciting MPIs is based on an extension of the two-step MP method adopted by Abdellaoui et al. (2021, 2023), where a subject undergoes a 'bisection' binary-choice procedure followed by a 'confirm-or-correct' choice list. Whilst subjects' payments depend solely on the choice list, the binary choice part serves as an aid to filling it in. Here, we develop an analogous two-step procedure that consists in a sequence of binary-choice questions followed by a 'two-dimensional' choice list, tailored for eliciting (two-dimensional) probability intervals instead of (one-dimensional) probability values.

Binary-Choice Procedure. For each event, our subjects first undertake a chained sequence of binary-choice tasks (Section 3.2). Here we set out the general principles of this procedure, leaving full details for Appendix D.1. The logic can be illustrated on Figure 1, notably by dividing the space of interval-valued urns into four preference-defined areas, summarised in Table 1. The procedure is based on the following observation.

Proposition 1. *Suppose preferences are represented according to (2) with $\alpha > \frac{1}{2}$, and let E be an event.*

a. For any $[x,y]$ in the R-B region (i.e. such that the corresponding preferences in Table 1 hold, for E), $\underline{p}(E) \leq x$ and $\overline{p}(E) \geq y$. Moreover, for any $[x,y]$ in the Wh region, $\underline{p}(E) \geq x$ and $\overline{p}(E) \leq y$.

b. For any $[x,y]$ in the R region (i.e. such that the corresponding preferences in Table 1 hold, for E), every $[x',y']$ with $x' \geq x$ and $y' \geq y$ is also in R. Moreover, for any $[x,y]$ in the B region, every $[x',y']$ with $x' \leq x$ and $y' \leq y$ is also in B.

It follows from part a. that if the experimenter has found an interval-valued urn $[x_{RB}, y_{RB}]$ in the R-B region (i.e. with the preference pattern in Table 1, row 1), and a $[x_{Wh}, y_{Wh}]$ in the Wh region, then the MPI is contained in the 'box generated' by these points, i.e. it is in the set $\{[x,y] : x_{Wh} \leq x \leq x_{RB}, y_{RB} \leq y \leq y_{Wh}\}$. The procedure works by searching the smallest such generated box for further points in R-B or Wh, in order to 'reduce' the size of the boxes and hence 'home into' the MPI. In this sense, it is analogous to the bisection procedure for MPs, where preferences indicate that the MP is in a particular interval, and the procedure searches to reduce the width of that interval.

Note that a similar result to Proposition 1 a. does not hold for the R and B regions. However, by part b. it can be concluded, for any interval-valued urn $[x,y]$ in R, that every point North-East of $[x,y]$ is also in R; and similarly for an interval-valued urn in B. So, if the experimenter has just discovered an urn in R (i.e. the elicited preferences for that urn are as specified in Table 1, row 3), then, to seek a point in R-B or Wh, she need not look North-East of this point; and analogously for urns in B. The procedure works, after eliciting preferences for an urn in R and B, by performing a bisection along one-dimensional cuts of the space $\mathscr{I}$ guided by this observation, until an urn in R-B or Wh is found, whence the procedure in the previous paragraph applies again. Details are provided in Appendix D.1. In particular, as shown there (Proposition **??**), the procedure canonically converges to the subject's probability interval for the event, not only under the α-maxmin EU model (2) but also under the generalizations discussed above (see also Section 5 and Appendices A and B.2).

Note finally that the procedure used has an in-built 'precision bias'. Whenever no urn in the R-B or Wh regions has been found, the procedure deliberately moves closer to the space of precise urns (see Appendix D.1). In this way, if there is any misclassification of subjects due to no urns being found in the R-B or Wh regions, the tendency would be for the procedure to represent them as more precise than they actually are.

Two-dimensional Choice List Procedure. For each event, after the binary-choice questions, subjects face a two-dimensional choice list, already filled in according to their responses on the previous procedure. They may modify the preferences encoded on this choice list before confirming. Only the confirmed preferences qualify for payment: so it is essential that the mechanism realized by such choice lists is incentive compatible. We now set out the theory underlying the two-dimensional choice lists. It relies on the following Proposition.

Proposition 2. *Suppose preferences are represented according to* (1). *For any event E, and any* $[r, 1-b] \in \mathscr{I}$, $[r, 1-b]$ *is a matching probability interval of E if and only if*

$$\begin{aligned} (z,[q,1-b],0) &\succ (z,E,0) \quad \textit{for all } q > r, \\ (z,[q,1-b],0) &\prec (z,E,0) \quad \textit{for all } q < r, \end{aligned} \tag{7}$$

and

$$\begin{aligned} (0,[r,q],z) &\prec (0,E,z) \quad \textit{for all } q > 1-b, \\ (0,[r,q],z) &\succ (0,E,z) \quad \textit{for all } q < 1-b. \end{aligned} \tag{8}$$

To illustrate, consider any MPI $[r,1-b]$ of an event E, so that the indifference (4) is satisfied. Clearly, under (1), it follows that (7) holds. On Figure 1, this determines the preferences involving the 'red' ILs corresponding to the bold red (horizontal) line. To the left of the MPI, the bet on E is preferred to the IL corresponding to the bet on red from the urn $[q,1-b]$ (i.e. with probability $[q,1-b]$ of winning); to the right of the MPI, the IL is preferred to the bet; and at the MPI, the two are indifferent. Likewise, the indifference (5) concerning the complementary bet determines preferences involving the 'blue' ILs corresponding to the bold blue (vertical) line, with the MPI being the point on that line where preferences 'switch' from the bet concerning E to the IL. The red (horizontal) and blue (vertical) bold lines in Figure 1 are thus analogous to a pair of choice lists, and the MPI is the switching point on each of them. We henceforth refer to the combination of the two as a *2D choice list*. By Proposition 2, we know that any urn that is a switching point on both branches of a 2D choice list is an MPI.

Inspired by this observation, consider an incentivization mechanism in which a subject who confirms the interval-valued urn $[r,1-b]$ for event E is remunerated as follows. Each urn $[x,y]$ in the 2D choice list—i.e. each $[x,y] \in \{[x,y] \in \mathscr{I} : y = 1-b\} \cup \{[x,y] \in \mathscr{I} : x = r\}$—determines a binary choice $\Phi_{[r,1-b],E}([x,y])$, defined as follows:

$$\Phi_{[r,1-b],E}([x,y]) = \begin{cases} \{(z,E,0),(z,[x,y],0)\} & \text{if } y = 1-b \\ \{(0,E,z),(0,[x,y],z)\} & \text{if } x = r, y \neq 1-b \end{cases} \tag{9}$$

In terms of Figure 1, if the urn is on the horizontal line going through $[r,1-b]$ (e.g. the bold red horizontal line in the Figure, if $[r,1-b]$ is the MPI), the choice is between the bet on the event and the bet on red from the urn; if it is on the vertical line going through $[r,1-b]$, the choice is between the complementary bet and the complementary IL. The incentive scheme selects an option for each of these possible choices: for the choice corresponding to urn $[x,y]$ it selects $\phi_{[r,1-b],E}([x,y])$, defined by:

$$\phi_{[r,1-b],E}([x,y]) = \begin{cases} (z,E,0) & \text{if } y = 1-b,\ x < r \\ (z,[x,y],0) & \text{if } y = 1-b,\ x \geq r \\ (0,[x,y],z) & \text{if } x = r, y < 1-b \\ (0,E,z) & \text{if } x = r, y > 1-b \end{cases} \tag{10}$$

I.e. if the urn $[x,y]$ has $y = 1-b$, $x < r$, then this selects the option $(z,E,0)$—the subject 'plays' the bet on E—and similarly for the other cases. The incentive mechanism first draws an urn $[x,y]$ at random from $\{[x,y] \in \mathscr{I} : y = 1-b\} \cup \{[x,y] \in \mathscr{I} : x = r\}$, and hence the choice $\Phi_{[r,1-b],E}([x,y])$; it then pays the subject according to the outcome of the selected bet or IL, $\phi_{[r,1-b],E}([x,y])$. It follows immediately from Proposition 2 that this mechanism is incentive compatible in the sense of weak dominance.

Corollary 1. *Suppose preferences are represented according to* (1)*, and let* $[r,1-b]$ *be such that, for every urn* $[x,y] \in \{[x,y] \in \mathscr{I} : y = 1-b\} \cup \{[x,y] \in \mathscr{I} : x = r\}$*,* $\phi_{[r,1-b],E}([x,y])$ *is a weakly dominant option in* $\Phi_{[r,1-b],E}([x,y])$*. Then* $[r,1-b]$ *is a matching probability interval of* E*.*

In other words, among all probability intervals that the subject could report, only matching probability intervals are such that the option selected by the mechanism is (weakly) preferred, no matter the choice in the 2D choice list that is played 'for real'. Hence implementing this incentive scheme on a subject's confirmed 2D choice list incentivizes reporting her MPI for the event. Since precise probabilities (and SEU) are a special case of multiple priors (respectively, Eq. (1)), this mechanism functions equally for Bayesian decision makers, who are incentivized to report their precise probabilities. We set out the experimental implementation of 2D choice lists in Section 3.2.

Note finally the depth of the analogy with choice lists for MPs. There, MPs are determined by the switching point, i.e. the maximum probability for which the subject prefers the bet on the target event over the lottery with that probability of winning. Similarly, the proposed probability-interval incentive mechanism elicits a single point, which is the switching point on each branch of the 2D choice list. Moreover, in standard MP choice lists, the switching point determines the preferences in the rest of the choice list by stochastic dominance. Similarly,

	(Frequency, Sample size)					
	Group A			Group B		
Choice-based MPI	(0.50, 4)	(0.25, 20)	(0.50, 100)	(0.25, 4)	(0.50, 20)	(0.25, 100)
Stated probability interval	(0.25, 4)	(0.50, 20)	(0.25, 100)	(0.50, 4)	(0.25, 20)	(0.50, 100)

Table 2: Prior information faced by subjects in EXP A: frequency of green chips in the previous sample, and sample size.

Proposition 2 guarantees that the elicited point determines the other preferences in the 2D choice list according to a probability-interval analogue of stochastic dominance, which states that, between ILs $(z,[r,1-b],0)$ and $(z,[r',1-b],0)$, decision makers prefer the prospect with higher lower probability.[10] Finally, in standard MP choice lists, this property underlies the incentive compatibility: it ensures that only the MP is such that, no matter the choice played 'for real' from the choice list, the selected option is preferred by the subject. Corollary 1 establishes an analogous result for MPIs and the proposed incentive mechanism.

3. Experimental Methods

We applied our probability-interval elicitation method in three experiments. One, EXP A, evaluated the effectiveness of our method in a simple setup using artificial sources of uncertainty. The target events were the color of the next chip randomly drawn from bags filled with 100 yellow and green chips, where the only information about each bag's content came from earlier draws conducted with replacement. The other two experiments, EXP N1 and EXP N2, involved uncertain events stemming from natural sources: the minimum winter temperatures in Paris and Sydney in EXP N1, and the test scores for two admission pathways at a French Business School in EXP N2.

3.1. Subjects

233 students completed the experiment: 101 from the INSEAD-Sorbonne Behavioral Lab (Paris, France) for experiment EXP A, 80 from university of Paris 1 for EXP N1 and 52 from HEC Paris Business School for EXP N2. Subjects' choices were collected through computer-based individual interviews that lasted about one hour in each study. Each individual interview started with a video presentation of the experimental instructions, followed by comprehension

[10]This 'Lower Stochastic Dominance' property, which is equivalent to the assumption (Section 2.3) that W is strictly increasing in the lower bound, is behind the preference patterns in Proposition 2; see Appendix A.

questions and one training MPI elicitation task (on an event not involved in the ensuing experiment). Appendix D.2 contains screenshots and a link to the video instructions for EXP A.[11] In all experiments, subjects were told that there were no right or wrong answers, and that they could ask any question regarding the experiment. Differences in experimental instructions between the experiments are explained in the sequel.

3.2. Artificial sources of uncertainty

Sources and choice tasks. The sources of uncertainty in EXP A were physical, opaque, labeled bags containing 100 green or yellow chips, with prior information about the composition coming in the form of prior draws with replacement (Appendix D.2). Different bags corresponded to different prior information, i.e., different sample size and frequency of green in the preceding draws.

Subjects were randomly allocated to one of two groups. The tasks for each group are specified in Table 2. Each group carried out two blocks of tasks. Each block involved three different bags. For a given group of subjects and a given bag in the *choice-based elicitation* block, subjects' probability intervals for the event that the next chip drawn from the bag was green were elicited using the proposed method. Then the subjects were asked to state their (precise) probability that the next chip was green (on a one-cursor slider). In the *stated* block, for each bag, subject were asked to state their probability interval for the next chip being green (on a standard, two-cursor slider), and then, as for the choice-based task, give their (precise) probability. The order of the bags within blocks was randomized, as was the order of the blocks in each group. As is clear from Table 2, all subjects provided probability intervals for each bag: for one group, these were elicited using our method; for the other they were stated directly.

Elicitation procedure. As noted in Section 2.5, our choice-based method for eliciting probability intervals follows the general two-step structure adopted by Abdellaoui et al. (2021, 2023) for MP elicitation: a binary-choice procedure is first used to aid subjects to fill in responses on a choice list, which they then confirm or modify.

More specifically, for each event E (e.g. drawing green from a specific bag), we first applied the binary-choice procedure set out in Section 2.5 and Appendix D.1. Each stage of the procedure consisted of two binary choices involving bets concerning E and bets on the color

[11] Beyond those who completed the experiment, 19 subjects in EXP A, 12 in EXP N1 and 0 in EXP N2 did not pass the comprehension check. They received a flat payment but were not given the possibility to continue the experiment.

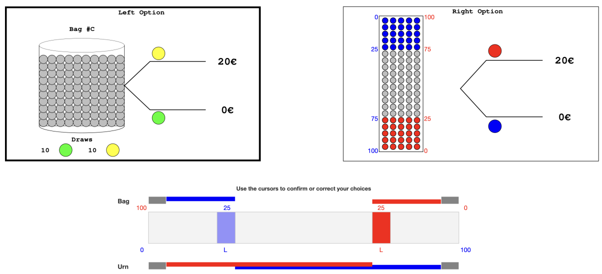

(a) Initial setting

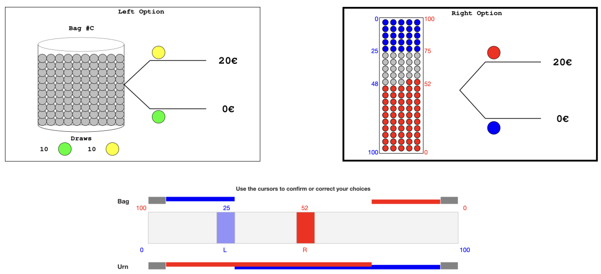

(b) After having moved red cursor

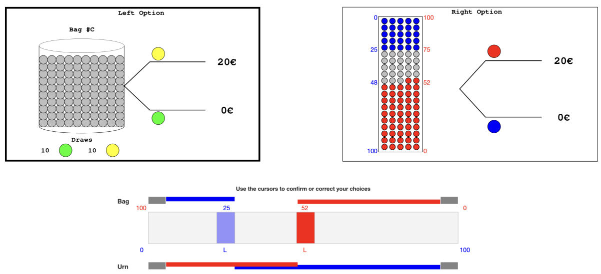

(c) After having clicked on left option

Figure 2: 2D confirmation choice list: displays.

of a ball drawn from a partially known 100-ball urn with a specified minimum proportions b and $r \leq 1-b$ of blue and red balls, respectively. All bets involved the same winning and losing outcomes. For each event and urn used, we collected the subject's choice in the decision between the bet on the event E and the bet on the next ball drawn from the urn being blue; in the subsequent choice question, we elicited their choice in the decision between the bet on E^c (or against E) and the bet on the next ball drawn from the same urn being red. (See Figure D.6, Appendix D.2 for illustrative examples of binary choices in our experiments.) These elicitations situated the urn in one of the areas in Table 1. The urn proposed in the next stage depended on the preferences elicited in the previous choices according to the binary-choice procedure (Section 2.5 and Appendix D.1). The subjective probability interval for E elicited at the end of the procedure is deduced from the preferences over such bets, as specified in the cited sections. The procedure continued until the interval was estimated to a precision of 0.15 if it was not degenerate, 0.05 if it was degenerate (i.e. corresponded to a precise probability), or up to 12 stages, whichever came first. The probability interval produced was fed into the next, two-dimensional choice-list 'confirmation' step of the elicitation procedure.

To illustrate the 'confirmation' procedure, Figure 2a shows the screen that a subject would see after the binary-choice procedure returns an interval $[0.25, 0.75]$ for the draw of a green ball from the specified bag. The corresponding 2D choice list is materialized by means of a two-cursor scrollbar. The red and blue cursors in the scrollbar determine minimum number of red

	Sources	Events $E_{t_i} = \{t' \in T : t' \leq t_i\}$ for t_i:
EXP N1	Paris	$-2, 2, 5, 8$
	Sydney	$15, 18, 20, 22$
EXP N2	Maths	$7, 10, 12, 15, 17$
	Contraction	$7, 10, 12, 15, 17$

Table 3: Natural sources of uncertainty and events in EXP N1 and EXP N2

and blue balls respectively, hence specifying the urn on the right. The chosen option between the bet on the bag (on the left) and the bet on the specified urn is highlighted. By moving the cursors, the subject can scan the choices associated to different urns. In particular, when moving the red cursor, the blue cursor remains fixed at the pre-specified value: so the subject scans all the urns with the same minimum number of blue balls but differing minimum numbers of red balls. In terms of Figure 1, this corresponds to the choices represented by the horizontal line through $[0.25, 0.75]$. When the red cursor is set far to the left, the minimum number of winning red balls in the urn is low, and the bet on the urn is less attractive (this corresponds to urns on the left of the horizontal line). As the red cursor is shifted further to the right, the minimum number of winning red balls increases, and the bet on the urn becomes more attractive: as indicated on Figure 1, the preference switches in favour of the bet on the urn at some point. Figure 2b illustrates the state of the scrollbar when the red cursor is moved to 52. The subject can modify her choices, for any setting of the red cursor, by clicking on the preferred bet. As illustrated in Figure 2c, which shows the result of clicking on the bag for the previous cursor setting, this updates the slider. Note in particular that the horizontal lines above and below the slider, which indicate the preferred options for each choice, are updated. Similarly, the subject can scan and modify choices involving urns with different minimum number of blue balls (for a fixed minimum number of red balls) by moving the blue cursor. See Appendix D.2 for further details, as well as a link to an online version of EXP A.

After any modifications, subjects had to reconfirm all of the associated choices, by moving one cursor then the other, before continuing on to the next phase of the experiment. The precision of the scrollbar, and hence subject responses, was to the nearest 0.01 (to the precise minimum number of red and blue balls out of 100 respectively).

3.3. Natural sources of uncertainty

Sources. Each of EXP N1 and EXP N2 involved two comparable natural sources of uncertainty. The type of source in EXP N1 was the minimum daily temperature over the previous

November–March period; the sources differed in the city whose temperature was of interest—Paris, where the experiment was carried out, and Sydney. The typical winning event E_{t_i} in this case was of the form: "the minimum temperature on day D in Paris (or Sydney) was less than or equal to t_i", where D was a randomly chosen day in the specified period (see Section 3.4). For each source in EXP N1, we chose temperature value t_i's close to the 10%, 33%, 66% and 90% percentiles of the true distribution (Table 3).

EXP N2 involved marks in two of the previous year's entrance exams for admission at undergraduate level to a prominent French business school, HEC Paris.[12] The subjects in the experiment had sat these exams either in the previous Spring or in the one before. The sources differed in the exam considered: a Maths exam, which is generally considered to be 'objectively marked', and the 'Contraction' exam—a summary of a philosophical or literary text—whose marking is considered more 'unpredictable' by candidates and students. Indeed, the marks in the latter exam have higher variance.[13] The typical winning event E_{t_i} here corresponds to: "candidate C obtained a mark less than or equal to t_i in the Maths (or Contraction) exam" for a randomly drawn candidate C. We used the same values for both sources (Maths and Contraction), picked so they would seem to reasonably scan the range and correspond to comparable points in the true distribution over Contraction scores, where they were at the 3%, 15%, 33%, 68% and 86% percentiles (Table 3).[14]

Choice tasks in EXP N1. Each subject undertook three blocks of tasks. Each of the first two blocks concerned a single source (Paris or Sydney), and involved the elicitation of the probability intervals for each of the events in the source (Table 3). The order of these two blocks was randomized. In each block, the subject first declared, in an non-incentivized manner and using a scrollbar, her estimated maximum and minimum values for the minimum temperature on the unidentified day D (see Section 3.4). This is standard procedure in expert elicitation for unbounded sources, aimed at combating anchoring bias (Morgan, 2014), and played no role in our elicitation. Then the elicitation procedure set out in Section 2.5 and implemented as in EXP A (Section 3.2) was applied for each event in the source. Within each block, the two extreme events (i.e. lowest and highest temperature points) were asked first, in a random order, followed by the other two events, in a random order.

[12]All candidates to this school at undergraduate level must apply in an entrance stream, each of which involves a different set of exams. The exams whose marks were involved in this experiment were sat by all candidates in both the 'ECS' (scientific) and 'ECE' (economics) streams. All subjects in this experiment were students admitted to the school through one of these streams.

[13]The variance of marks for Maths is 3.77, where it is 9.92 for Contraction.

[14]They were at the 0%, 0%, 2%, 21% and 60% of the true distribution of Maths scores.

The final block involved the elicitation of MPs for the events in Paris treatment, using the two-step bisection-then-choice-list procedure from Abdellaoui et al. (2021) (see Appendix D.2 for details). MPs were elicited for each event E_{t_i} in this source and its complement $E_{t_i}^c$ (Table 3). The order of elicitations was randomized in this block.

Choice tasks in EXP N2. Each subject undertook two blocks of tasks. Each of the blocks concerned a single source (Maths or Contraction), and involved the elicitation of the probability intervals for each of the events in the source (Table 3). The order of the blocks was randomized, as was the order of the events in each block. In each block, the elicitation procedure set out in Section 2.5 and implemented as in EXPs A and N1 was applied for each event in the source. Each block ended with an omnibus confirmation screen, in which the interval-valued urns elicited for each of the events in the source were displayed in graphical form (see Appendix D.2 for details). The subject was given the opportunity to select and modify any of her responses for the events in the source. This screen, the sources and the larger number of events elicited per source (see Table 3) were the central differences with respect to EXP N1.

3.4. Incentivizing Subjects

Participants in all studies received a flat payment of €10. Additionally, a random incentive system was implemented, which was entirely analogous to those standardly used to implement elicitation of MPs. In EXPs N1 and N2, after the presentation of the instructions and before the beginning of the experiment, the subject chose a number from a given range, which identified an individual case of the variable of interest (the day D, if the source was minimum temperature; the candidate C, if the source was the mark). The exact case identified was specified according to a spreadsheet that would only be revealed to the subject at the end of the experiment. At the end of each of the three experiments, a choice list (a 2D choice list in EXPs A and N2; a 2D choice list or MP-choice list in EXP N1) and choice on it were selected at random by the computer.[15] The subject was then paid according to the decision she had made on that choice. If she had chosen, say, the bet on the event that the minimum temperature in Paris is less than or equal to 2°C, then the day which she chose was revealed, as well as daily temperature data for the November–March period, and she won if the minimum temperature on that day was indeed 2°C or less; if not, she lost. Or, in EXP A, if she had chosen the bet that the next

[15] More precisely, for the selected choice list, a color—red or blue—was selected at random, and then an urn on the branch of the 2D choice list corresponding to that color was selected at random.

chip in a certain bag was green, then a chip was drawn randomly from that bag and she won according to its color. If she had chosen the urn, then she composed the appropriate urn—she counted the specified minimum numbers of red and blue balls, with the remaining balls coming from pre-constructed Ellsberg urns (of unknown composition). Then a ball was drawn from the constructed urn, and she was paid according to whether she bet on the color of that ball or not. All bets yielded €20 if won, and nothing otherwise.

4. Results

4.1. Performance and Validation

Figure 3 plots the 25%, 50% and 75% quantiles (Interquartile Ranges, i.e. IQRs) of the upper and lower probabilities and CDFs for all events elicited used our method and all experiments (see Table C.1 in Appendix C.1 for basic descriptive statistics). This Figure already gives some early indications about our results, and the performance of our elicitation method.

EXP A. First of all, Panel (a) of Figure 3 shows how the 'balance' of evidence, as represented by the observed frequency of green chips in EXP A, affects the position of the elicited probability intervals: they are higher when the observed frequency is larger. For each sample size, unpaired t-tests and Mann-Whitney tests reject the null hypothesis of equal midpoints of the elicited interval for different observed frequencies in the previous draws ($p < 0.001$ in all cases), with the means being higher for higher frequencies. Since one would expect such sensitivity of posterior beliefs, the method passes this first 'validation' check.

Another possible test for the method, that can be applied under the artificial source of uncertainty, is the comparison with the posterior probabilities of a Bayesian who updates a uniform prior with the same information by Bayes rule. For 80% of events, across all subjects, this 'objective Bayesian' probability was contained in the elicited probability interval, suggesting that in the vast majority of cases, subjects did not rule out the Bayesian probability in forming their posterior probability intervals. Moreover, even in the cases where the Bayesian probability was not in the interval, it was not far, with the average minimal distance to the interval among instances where the Bayesian probability was not contained in it being less than 0.06. Unsurprisingly, the midpoints of the elicited intervals were substantially correlated with the Bayesian probability: the Spearman correlation was 0.65.

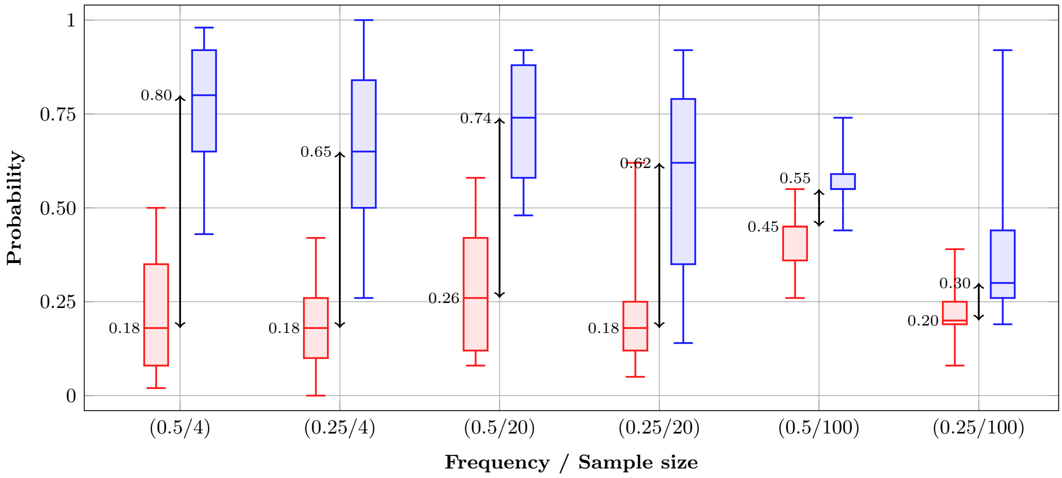

(a) EXP A

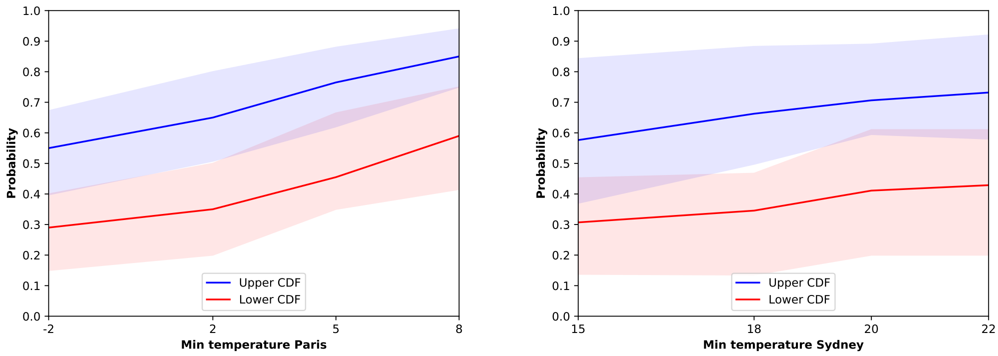

(b) EXP N1

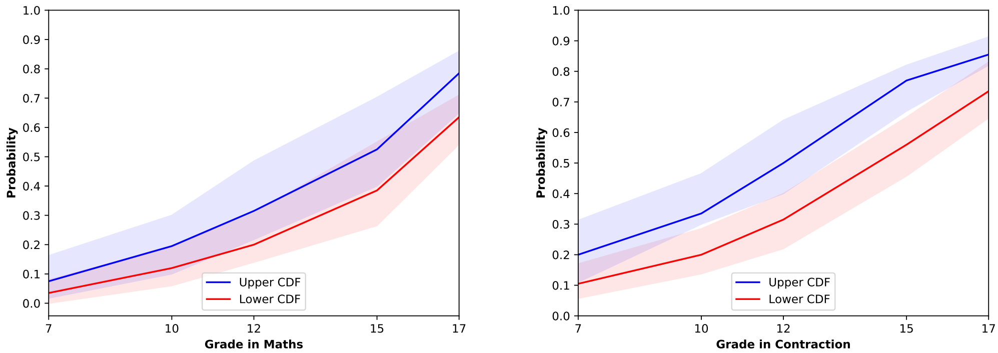

(c) EXP N2

Figure 3: IQRs of upper and lower probabilities and CDFs

EXPs N1 & N2. The experiments eliciting interval-valued CDFs for natural sources of uncertainty suggest that the general message of validity extends to such contexts. Echoing the sensitivity to frequencies found in EXP A, the upper and lower CDFs in the other experiments differ across subjects and events—thus suggesting the consistency of the method. A validity

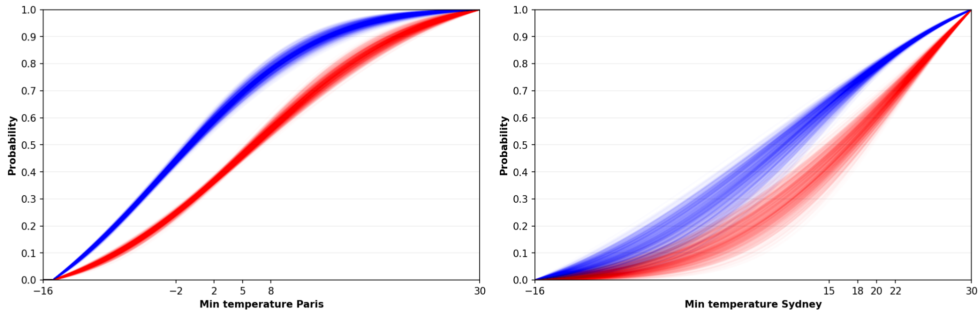

(a) EXP N1 (Truncated Normal distribution: Min temperature)

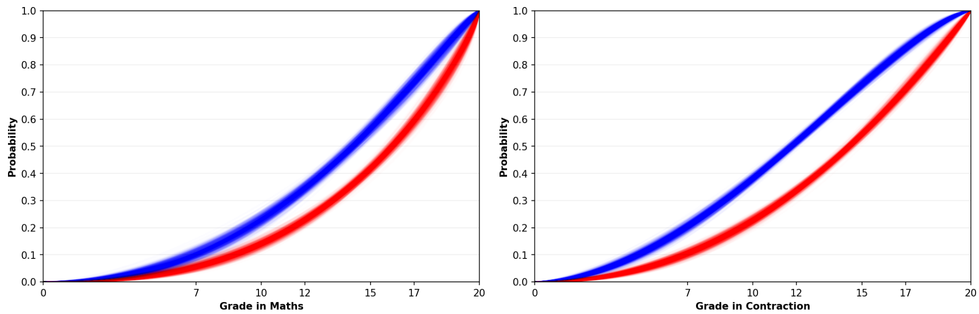

(b) EXP N2 (Beta distribution: Grade)

Figure 4: Bayesian estimation of lower and upper CDFs: plots of 1000 samples from MCMC.

test in this context would examine whether the upper and lower CDFs are increasing—an issue that can be investigated using the Kendall rank.[16] As reported in Table C.5, Appendix C.1, the median Kendall τ_b is far greater than 0 for all sources, pointing to increasing upper and lower CDFs. In EXP N2, where subjects were given the opportunity to confirm their replies on all the 2D choice lists for a source (Section 3.3), CDFs were strictly increasing (Kendall rank of 1) for the vast majority of subjects. In EXP N1, there were more violations of monotonicity; however, the median Kendall ranks for upper and lower CDFs in the Paris treatment were similar those obtained under the more standard MP elicitation (Table C.5). This suggests that violations were not unique to the probability-interval elicitation method proposed here.

Finally, we re-analyse the data from these experiments under a standard Bayesian approach, estimating posterior distributions over upper and lower CDFs using a MCMC procedure. Figure 4 plots 1000 MCMC samples for each of the upper and lower distributions, for each source, under the parametric distributions for upper and lower CDFs that offer the best fit (see Tables C.15, C.16 and C.17; Appendix C.2). They suggest that the proposed elicitation technique supports parametric estimation of subjective probability intervals in the population, insofar as

[16]The Kendall τ_b is an indicator of ordinal association: the value 1 indicates that the CDFs or MPs are strictly increasing; 0 suggests that there is no association between the elicited probability and the size of the event; -1 indicates a strictly decreasing relationship between the two.

they chime with expectations given the nature of the events. For instance, they suggest that the dispersion of subjective upper and lower probabilities is larger for the temperature source (EXP N1) than the grade source (EXP N2), which could be related to the fact that all subjects in EXP N2 had sat both exams, and were very interested in the marking, several months before. Also, within EXP N1, there is more dispersion in the estimated distributions for Sydney than for Paris, as would be expected given the less familiar nature of the former source for Paris subjects.[17]

4.2. Imprecision

Overall Imprecision. Our raw data (Figure 3) suggest that subjects' beliefs are often *imprecise*: i.e. there is a gap between their upper and lower probabilities, as indicated in Figure 3a by the arrows connecting the median upper and lower probabilities for each bag. For further analysis, we define a subject's *Imprecision concerning an event E* to be the width of her elicited probability interval for E, i.e. $\overline{p}(E) - \underline{p}(E)$. A subject's *Average Imprecision* across all elicited events in EXP A, or across all elicited events in a source in EXPs N1 and N2, gives an indication of how imprecise the subject's beliefs are, on average, across the relevant events. Naturally, an SEU decision maker will assign precise probabilities to all events, and hence have imprecision 0 (for all events and sources).

Figure 5 displays the 25%, 50% and 75% quantiles, and max and min subject-level Average Imprecision across all sources in all experiments (see also Table C.7, Appendix C.1). It clearly suggests a prevalence of imprecision, with mean and median Average Imprecision greater than 0.1 for most sources and experiments. Binomial tests reject the hypothesis of equal probability for the Average Imprecision to be equal to vs. greater than 0 for each source ($p < 0.001$ in all cases), with a clear majority of subjects—99 out of 101 in EXP A, 79 out of 80 in EXP N1, and 52 out of 52 in EXP N2—having strictly positive Imprecision on average. The data on the number of precise events—events for which the subject's probability interval has zero width—tells a similar story, with not more than around 5% of subjects giving precise probabilities for all events in a single source (Table C.8, Appendix C.1).[18]

[17] More precisely, it is clear from Tables C.17a and C.17b that the standard deviations of the parameters for the Paris source are typically lower than for Sydney.

[18] Further analysis, reported in Appendix C.1, confirms that the observed imprecision in elicited probability intervals cannot be explained by imprecision in the elicitation procedure.

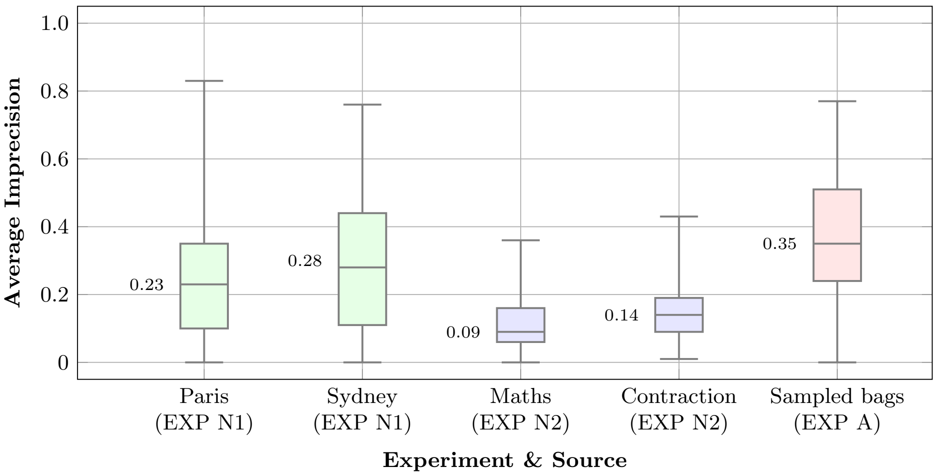

Figure 5: IQRs of Average Imprecision for all events in EXP A and across sources in EXPs N1 and N2

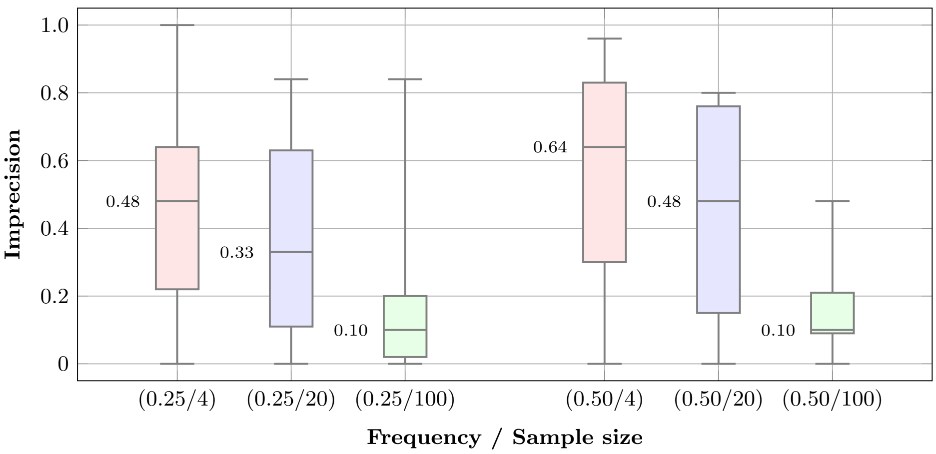

Figure 6: IQRs of Imprecision at varying frequencies and sample sizes, EXP A

Information and familiarity. A reasonable hypothesis is that, *ceteris paribus,* subjects' beliefs are more imprecise concerning events with which they are less familiar, or about which they feel as if they have less knowledge or information. In terms of multiple priors models, this corresponds to wider probability intervals for events for which there is less information. Given the explicit control on the information available via the observed sample size, EXP A allows for a particularly clear examination of the effect of information on imprecision.

As is clear from Figure 6, which displays the 25%, 50% and 75% quantiles, and max and min Imprecision across subjects for each frequency and sample size observed, Imprecision decreases with sample size. Recall (Table 2, Section 3.2) that every subject's interval was elicited, for a given frequency, at sample sizes 4 and 100, so the relationship between Imprecision and

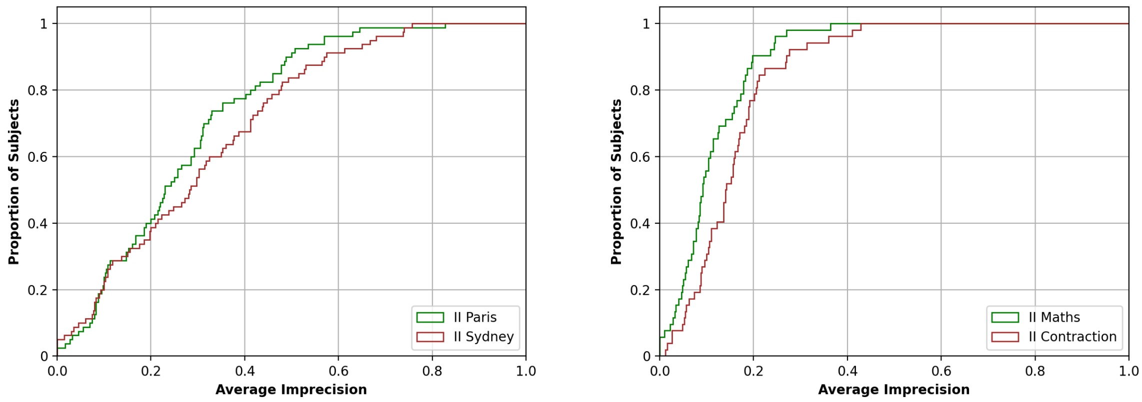

(a) EXP N1 (b) EXP N2

Figure 7: CDFs of Average Imprecision per source across subjects

information can be tested at a within-subject level. Paired t-tests reject the null hypothesis of equal Imprecision across these sample sizes, for both frequencies explored ($p < 0.001$ in both cases). Binomial tests of the null hypothesis of equal chance of one Imprecision being higher come to the same conclusion ($p < 0.001$ in both cases), with 49 out of 51 subjects (respectively 45 out of 50 subjects) having a more imprecise interval for the smaller sample size at observed frequency 0.5 (resp. 0.25). That the proposed elicitation method captures an expected relationship between imprecision and information in the carefully controlled environment of EXP A further bolsters its credentials.

Whilst allowing for less control, the natural sources of uncertainty used in EXPs N1 and N2 also support differences in perceived information, with (Paris-based) subjects likely to be less familiar with the weather in Sydney than that in Paris, and the Contraction exam in EXP N2 generally being considered to be 'less predictable' than the Maths one (Section 3.2). The CDFs of the Average Imprecision per source across subjects, plotted in Figure 7, suggest a relationship between imprecision and familiarity or predictability of the source. The CDF for Contraction—known as the less predictable exam—is entirely to the right of that for Math, indicating a larger Average Imprecision at the subject level. Similarly, the CDFs for Sydney—the less familiar source—is to the right of that for Paris for a large range of values, suggesting that more imprecision for this source. A paired t-test barely fails to reject the null hypothesis of identical Average Imprecision across the sources in EXP N1 ($p = 0.0895$), whilst it rejects it for EXP N2 ($p = 0.0016$). A Binomial test comes to similar conclusions ($p = 0.576$ for EXP N1; $p = 0.017$ for EXP N2), with 45 out of 80 (resp. 35 out of 52) subjects having a more imprecise interval on average for Sydney in EXP N1 (resp. Contraction in EXP N2).

Event-level Imprecision. We also investigate imprecision at the event level within sources in EXPs N1 and N2. One-way ANOVAs of the Imprecision (dependent variable) against the event (factor) reject the null hypothesis of identical imprecision across all events for the sources in EXP N2 ($p < 0.001$ for Maths; $p = 0.003$ for Contraction), whilst failing to reject it for the sources in EXP N1 (Table C.9, Appendix C.1). These conclusions are also illustrated in CDFs of the Imprecision for each elicited event in each source, across subjects (Figure C.1, Appendix C.1). This suggests not only that imprecision is widespread, but that imprecision may be event dependent within sources, as one would expect if some events are intuitively more uncertain than others. For instance, the least imprecise event in EXP N2 involves, for both sources, the lowest grade, where many subjects are presumably more sure of their judgements.

In summary, our method reveals that, when beliefs are elicited with a method allowing for (non-degenerate) probability intervals, imprecision is widespread, at least for the events considered here. Crucially, we recover an expected relationship between imprecision and perceived information in both controlled artificial sources and natural ones. This can be seen as providing further indirect evidence for the solidity of the proposed elicitation method. Finally, at least within some sources, the extent of imprecision may depend on the event.

4.3. Matching versus Stated Probability Intervals

Recall that in EXP A, the same events (concerning bags as characterised by frequencies and sample sizes) were elicited across different subjects using different methods: some subjects underwent the proposed incentive-compatible method, while others were asked for stated probability intervals, as in Giustinelli et al. (2022). This permits between-subject comparison of the results of interval elicitation under the two methods.

Figure 8 displays the 25%, 50% and 75% quantiles, and min and max of the distribution of the interval midpoints and Imprecision across subjects, for each event (concerning a bag characterised by frequency and sample size) and elicitation method. In the aggregate, the stated intervals are roughly comparable to those elicited under the proposed incentive-compatible method for most events, though the dispersion across subjects may differ at some points. Given the theoretical well-foundedness of the method developed here, this could be understood as providing validation for the use of stated intervals in large-scale field studies aimed at eliciting aggregate characteristics.

The Figure does however suggest some interesting differences between the intervals elicited by our method and stated intervals. For one, there is a greater dispersion in the *position* of inter-

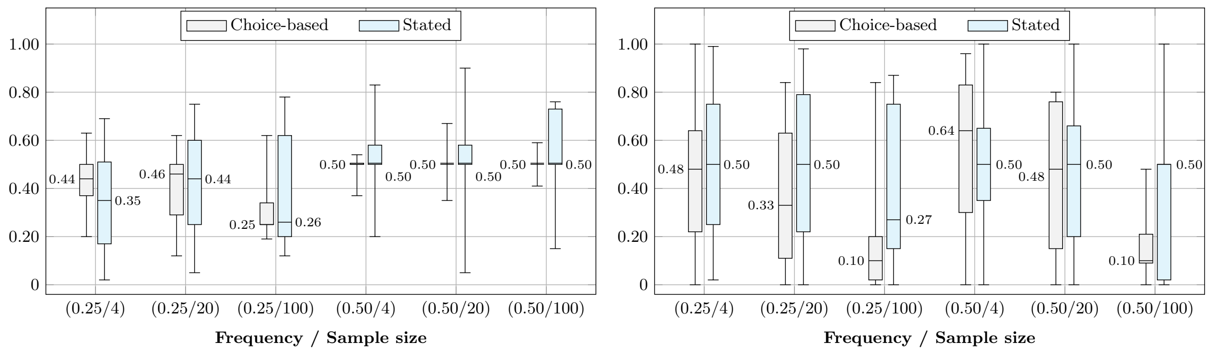

(a) Midpoints (b) Imprecision

Figure 8: IQRs of probability interval midpoints and Imprecision across frequencies and sample sizes in EXP A, for the proposed choice-based elicitation method and stated intervals.

vals, as is clearly visible on the plots of the midpoints in Figure 8a. This is particularly notable for the 'symmetric' case of frequency 0.5, which are more tightly centred on the midpoint of 0.5 under the choice-based elicitation method. More importantly, the imprecision tends to be *lower* under incentive-compatible elicitation as compared to stated intervals, especially when there is lots of information (Figure 8b). For instance, both unpaired t-tests and Mann-Whitney tests reject the null hypothesis of equal imprecision for both frequencies under sample size 100 ($p < 0.001$ in all cases), though this is not the case for smaller sample sizes. This suggests that the expected link between information and imprecision is tighter for intervals elicited under our method. Moreover, it hints that the use of stated interval elicitation may tend to overestimate imprecision, especially in information-rich environments.

Moreover, we can undertake analysis at the individual subject level, using the Average Imprecision, defined in Section 4.2, for each subject and each elicitation method. The Spearman correlation between subjects' Average Imprecision under our elicitation method vs stated probabilities intervals is 0.23 ($p < 0.05$), suggesting that the stated method does fairly well at identifying subjects whose intervals are wider on average.

4.4. Stated probabilities, matching probabilities and the α-maxmin EU mixture coefficient

Recall that, in EXP A, subjects provided their stated (precise) probabilities for the events; comparison with the elicited probability intervals may provide another 'sanity check' for the method. For 77% of events, across all subjects, the stated probabilities were contained in the intervals elicited by our method; this figure rose to 81% if one removes stated probabilities suggesting limited effort on the part of subjects.[19] Across such responses, stated probabilities were strongly correlated with the midpoints of the elicited intervals (Spearman correlation of 0.49).[20]

Further insight can be gleaned from EXP N1, which contained a choice-based elicitation of MPs for the events in the Paris source (Section 3.3). As noted in the Introduction and Section 2.4, under the α-maxmin EU model (2), MPs for an event and its complement generate a pair of equations (Eq. (3)) that cannot be solved for α and the upper and lower probabilities of the event in general. However, drawing on the elicited MPs and our elicitations of upper and lower probabilities, they can be used to elicit the mixture coefficient α. Under analysis using the raw data, the median α across subjects is 0.80 (Table C.19, Appendix C.3); a Bayesian estimation of the α in tandem with the lower and upper CDFs (see Appendix C.2, Table C.17a) yields mean value 0.81. As discussed at more length in Section 5, this is, to our knowledge, the first direct choice-based elicitation of the α in the α-maxmin EU model that fully controls for the set of priors by eliciting the relevant information about them without invoking supplementary assumptions. Moreover, it is consistent with the findings for stated probabilities in EXP A: for $\alpha > 0.5$, the MP of an event is below the midpoint of the corresponding probability interval.

5. Discussion

Several general conclusions emerge from our implementation of the proposed multiple-prior elicitation method over a range of different sources of uncertainty. The first concerns its feasibility and validity. The method produces probability intervals that are reasonable: they are sensitive to the aspects of available information (in experimental contexts where that is controlled), and in general consistent with 'objective' (Bayesian update) probabilities where available. The second concerns the extent and determinants of imprecision. Although precise probabilities are permitted by our procedure, many subjects' elicited probability intervals are imprecise – they do not reduce to a single probability value – for the events considered here. Crucially, we recover an expected relationship between imprecision and perceived information in both controlled artificial sources and natural ones. This can be seen as providing further indirect

[19] More precisely, recalling that the stated probability task was not incentivised, and that it was completed on a one-cursor slider (Section 3.2), in just under 10% of cases, subjects provided stated probabilities that were very close to the default slider setting of 0. The reported proportion removes all responses below 0.1 (noting that the lowest observed frequency is 0.25).

[20] Moreover, stated probabilities were typically close to, yet below, the midpoints of the corresponding elicited intervals: in the mean, stated probabilities were 0.05 less than the midpoints.

evidence for the solidity of the proposed elicitation method.

Our third general conclusion concerns the comparison of our method with probability intervals stated directly by subjects. Several studies have compared different methods for precise probability elicitation in laboratory settings (e.g. Trautmann and Kuilen, 2015; Hollard et al., 2016). These could be used to ground and improve belief elicitation beyond the lab. After all, if a simpler method yields similar results to a more complex one with better properties – in terms of incentive compatibility, for instance – then such comparisons can bolster confidence in the use of the former method in the field. Our between-subject comparison in EXP A between the proposed elicitation method and stated probability intervals is the first such exercise for multiple prior or imprecise probability beliefs, to our knowledge (see also the related literature discussion below). It suggests that, in aggregate, stated probability-interval methods give a fairly good approximation to the intervals provided by our choice-based method, though they may overestimate the extent of imprecision in information-rich environments. Whilst this is a first study, and others are doubtless required, this bodes well for the use of stated probability-interval methods in the field (e.g. Giustinelli et al., 2022), as well as for the solidity of their conclusions.

Finally, as concerns the well-known identification problem for the α-maxmin EU model, we draw on our probability-interval elicitation to perform the first elicitation of the model's mixture coefficient that is fully general and controls for beliefs (see discussion below).

We now discuss the robustness of our procedure, related literature, and some directions for future development.

Robustness. Although Hurwicz α-maxmin EU is one of most general decision models in the literature taking as belief component a set of priors, the core of our method applies beyond this model. As set out in Sections 2.3–2.5, it operates on a general model (Eq. (1)) that can incorporate probability weighting in the style of Prospect Theory (Wakker, 2010) or transformations of probability intervals in the style of Gajdos et al. (2008). More precisely, the notion of MPI remains well-defined for all such extensions, and the decision maker's subjective probability interval is always a MPI (Section 2.4). Moreover, MPIs are essentially unique for generic cases of such extensions (Appendix B). The incentivization mechanism, which is implemented in the 2D choice list part of the method, is incentive compatible under the most general form of model (1), as evidenced by Proposition 2 and the accompanying discussion in Section 2.5. So it applies not only under α-maxmin EU, but also under the probability weighting or probability-interval transformation generalizations just mentioned. As concerns the first, binary-choice part of our method, although Proposition 1 underlying it is stated under α-maxmin EU with $\alpha > \frac{1}{2}$, it is actually a corollary of a stronger result (Proposition A.1, Appendix A). This result covers the aforementioned generalizations under conditions analogous to $\alpha > \frac{1}{2}$ (Appendix B). Moreover, as discussed in Appendix B, there is independent evidence that such conditions hold for most of our subjects. Note that the generality of the second, 2D choice list part of the procedure is more important, for this is the part that counts for incentivizing subjects' responses (Sections 2.5 and 3.4). Finally, the method also applies under decision models that do not belong to the family (1), such as the multiple-prior minimax expected regret model (Appendix B, footnote 1).

The 2D choice list mechanism is incentive compatible in the sense of weak dominance (Corollary 1, Section 2.5). This demands, for the reported interval-valued urn $[r, 1-b]$, that, for every choice between the bet on red from an urn $[x, 1-b]$ and the bet on the event E, the subject prefers the option selected by the incentive mechanism; and similarly for choices between bets on blue from urns $[r, y]$, for varying y, and the bet against E. The set of choices between the bet on the event and bets on urns with varying minimal numbers of red balls forms a branch of the 2D choice list, and can itself be thought of as a (standard one-dimensional) choice list. The weak-dominance notion of incentive compatibility focuses on the choices in this list in isolation from the way the number of blue balls b is set; and similarly for the other branch. Our implementation was designed to favor such isolation, notably via the realization of 2D choice lists by a single scrollbar with two cursors (Figure 2, Section 3.2 and Appendix D.2). Visually very different from Figure 1, this presentation is less suggestive of opportunities for strategically reporting the interval to influence the set of choices used for remuneration. Notwithstanding this, the extent to which such strategic reasoning has been employed by the subjects in our experiments is ultimately an empirical question, and we treat it as such. On this front, our elicitation method has the advantage that such reasoning leads to easily recognizable choice patterns. As discussed in Appendix B.2, for a subject represented by (2) with $\alpha \in (0,1)$ and any set of priors, her optimal response to the 2D choice list task when reasoning strategically[21] is one of the intervals $[0,0], [0,1], [1,1]$. However, no subjects gave such responses for all events elicited, with only one subject across all three experiments giving such an interval for more than half of the elicited events (Table C.6, Appendix C.1; see Appendix B.2 for further details). This suggests that strategic reasoning is extremely infrequent among our subjects.

[21] As set out in the cited Appendix, under strategic reasoning, the subject considers the choice of MPI as a choice of a (second-order) lottery over ILs and particular bets for or against E.

Related literature. Our elicitation method relates to existing experimental and theoretical literature on multiple prior models, and the α-maxmin EU model in particular. Part of this literature is concerned with testing or comparing such models (e.g. Hey et al., 2010; Baillon and Bleichrodt, 2015); by contrast, the aim here is to elicit probability intervals in the context of a general multiple prior model. Likewise, there is a literature using matching probabilities or certainty equivalents to study willingness to bet on objectively-given probability intervals (e.g. Baillon et al., 2012; Chew et al., 2017) . The present paper, by contrast, uses such interval-valued urns with the distinct aim of eliciting subjective probability intervals.

Multiple prior beliefs are most relevant in situations where agents typically do not hold precise probability distributions determining preferences, so to be generally applicable, a multiple-prior elicitation method should avoid assuming underlying precise probabilities. The assumption that subjects have precise probabilistic beliefs which completely determine the contributions of events to their (potentially non-expected utility) preferences is called *probabilistic sophistication* (Machina and Schmeidler, 1992; Chew and Sagi, 2006). As emphasized in the Introduction, our method avoids all assumptions of this sort. This arguably sets it apart from much of the theoretical literature and virtually all of the experimental literature on multiple prior models.

On the theory side, the challenge of incentive-compatible elicitation under α-maxmin EU (2) is compounded by identification issues, arising from the fact that different pairs of mixture coefficient α and sets of priors can represent the same preferences (see Introduction and Section 2.4). Proposed approaches include pinning down the set of priors using 'unambiguous preferences' (Ghirardato et al., 2004), though this has problems in finite state spaces (Eichberger et al., 2011), or enrichening the state space to include an infinite product structure and invoking symmetry axioms (Klibanoff et al., 2021). Another line of attack concentrates on special cases where the set of priors is generated by a precise probability distribution. For instance, Gul and Pesendorfer (2014, 2015) and Chateauneuf et al. (2007) obtain a unique identification of α and the set of priors: the former when the set is generated as extensions of a precise probability measure on a subalgebra of events; the latter when it is generated from a precise probability measure via ε-contamination, i.e. the mixture with the set of all probability measures.[22] Since in both cases, preferences and sets of priors are generated from precise probabilities, they assume some form of probabilistic sophistication. Our approach, by contrast, deliberately eschews such assumptions as inadmissible in many situations of interest. Rather, it follows

[22]Formally, the assumption is that the set of priors $\mathscr{C} = \{(1-\varepsilon)p + \varepsilon q : q \in \Delta\}$, where Δ is the space of all probability measures, p is an element of Δ and $\varepsilon \in [0,1]$.

the theoretical approach developed by Hill (2023), who resolves the identification issue for α-maxmin EU and a range of extensions by using interval lotteries, with no need for specific richness assumptions on the state space, probabilistic sophistication, or any other non-standard assumptions on the set of priors.

On the experimental front, there is a small literature dedicated to incentive-compatible elicitation of multiple priors. One family of approaches purports to elicit them as the support of second-order beliefs, represented as a probability measure over the space of probability measures. Beyond the assumption of second-order probabilities, which is foreign to the original multiple prior models (Gilboa and Schmeidler, 1989; Bewley, 2002; Ghirardato et al., 2004), and the fact that they import an assumption of probabilistic sophistication, albeit at the second-order level, these often make further assumptions about the role of these second-order beliefs in choice. For instance, Qiu and Weitzel (2016) propose a method that relies on the assumption that a subject's opinions about others subjects' matching probabilities coincides with the uncertainty surrounding her own assessment.

Another family of approaches draws on the probabilistically-sophisticated special case of α-maxmin EU studied by Chateauneuf et al. (2007), where the subject's set of priors is generated as the ε-contamination of a single probability measure. Dimmock et al. (2015); Baillon et al. (2018b,a) use elicitation of MPs or certainty equivalents to estimate 'ambiguity indices', which they claim can be used to back out the mixture coefficient α and the set of priors. However, as shown by Baillon et al. (2021, Theorem 16 & Section 7.3, Eq. (20)), these indices are only guaranteed to yield the subject's set of priors if they are generated from a precise probability measure by ε-contamination, in which case preferences are represented by the Chateauneuf et al. (2007) model. So, though unwarranted in situations where multiple prior decision models come to the fore, this elicitation technique assumes probabilistic sophistication. In fact, our data provides empirical insight into the relevance of their assumption. Whilst Chateauneuf et al. (2007) implies that the imprecision (in the sense of Section 4.2) is the same for all events,[23] our observations reject this equality for the sources in EXP N2 (Section 4.2; see also Table C.9 and Figure C.1, Appendix C.1): these are thus sources for which their method's underlying assumption does not hold. Of course, this does not bode well for the general applicability of their method. It does not follow that it is never viable; indeed, our data indicate that the Paris source in EXP N1 may satisfy their assumptions. Moreover, we can estimate the ambiguity indices used in the aforementioned papers on the basis of the data from our study (EXP N1,

[23] If the set of priors is as defined in footnote 22, then, for any E, $(1-\varepsilon)p(E) \in [0,1-\varepsilon]$, so the probability interval for event E is $[(1-\varepsilon)p(E),(1-\varepsilon)p(E)+\varepsilon]$, and hence the event has imprecision ε.

Paris treatment) and under their assumption about the set of priors;[24] doing so, we find, for instance, that they yield the value 0.82 for the mixture coefficient α—which, reassuringly, is close to the Bayesian and raw estimates reported in Section 4.4. So our elicitation method is not only more general and robust, insofar as it applies in situations where the assumptions underlying their approach do not hold; moreover, it can evaluate precisely in which cases they do hold. In those cases, their approach, implemented on our data, gives the same result as our 'direct' elicitation.

Another related branch of literature focuses on scoring rules. Hossain and Okui (2013) provides a scoring rule in the absence of expected utility preferences: since it elicits precise probabilities under probabilistic sophistication, it does not tackle the issue of multiple-prior elicitation. Scoring rules have also been proposed for most-likely intervals for the value of an unknown parameter (Winkler and Murphy, 1979; Schlag, 2015). Typically, they are incentive compatible under the assumption that the subject is a Subjective Expected Utility maximiser with a precise probability distribution (Schlag, 2015, Section 5), and hence under an assumption stronger than probabilistic sophistication. By applying them where the unknown parameter at issue is itself a probability, such scoring rules could conceivably be repurposed to elicit probability intervals. However, given that they rely on the assumption of expected utility—here at the second-order level—they would need to suppose precise probabilities in order to elicit imprecise ones; as noted, this seems inappropriate for situations where multiple priors are relevant. As mentioned in the Introduction, the difficulty of developing scoring rules that avoid such probabilistic assumptions is further underlined by an impossibility result showing that there are no real-valued continuous strictly proper scoring rules for multiple-prior probability intervals (Seidenfeld et al., 2012, Prop 5).

Going beyond the lab, there is a large and growing literature on elicitation of multiple priors or imprecise probabilities in a range of disciplines, from economics to climate science. All such elicitation exercises of which we are aware use stated probability intervals, and as such are not incentive compatible. For instance, Giustinelli et al. (2022) elicit beliefs on dementia and long-term care decisions in a large-scale representative survey (over 1000 subjects), allowing stated probabilities to be interval-valued. Consistently with our results (Section 4.2), they find

[24]Specifically, Baillon et al. (2018b) propose the average of $1-MP(E)-MP(E^c)$ over a selection of events as their measure of the 'ambiguity aversion index' b. The average for the events elicited here can be deduced directly from Table C.20 (Appendix C.4), as around 0.16. On the other hand, under (2) with the specified form for the set of priors (see footnote 22), their 'a-insensitivity index' $a=\varepsilon$. Under such sets of priors, as noted in footnote 23, every E has imprecision ε. The Average Imprecision measured by our method (Section 4.2 and Table C.7) thus gives an estimate of their a: it is around 0.25. The mixture coefficient α is related to these indices by $\alpha=\frac{1}{2}\left(\frac{b}{a}+1\right)$ (Baillon et al., 2021), yielding the value in the text.

widespread imprecision. They argue forcefully for the importance of probability-interval elicitation for reducing survey bias and understanding attitudes to and behavior in the face of high-uncertainty events, such as whether one will develop dementia and whether to insure against it. In another approach, in different domain, Kriegler et al. (2009) elicit beliefs of selected scientists (around 50 subjects) concerning climate tipping points, allowing participants to state probability intervals for these (notoriously uncertain) events. Such expert elicitations, which involve often time-consuming and individualised sessions with selected experts, have emerged as a central tool for managing complex uncertainties (Morgan, 2014). Though they have traditionally aimed at eliciting precise probabilities, Kriegler et al. (2009) shows that imprecision is widespread for some events, which once again argues for the relevance of probability-interval elicitation.

Future Directions Our method can shed some much-needed light on the criticisms of stated approaches centred on their lack of incentive compatibility and theoretical grounding. The preliminary comparison from EXP A shows that, in the aggregate, the stated approach yields similar results to our incentive-compatible decision-theoretically-well-founded method. As reported in Section 4.3, beyond this general match, there is a significant correlation in the Average Imprecision between the two methods across subjects. This suggests that, roughly, subjects with larger intervals as elicited by our method will tend to provide larger intervals in the stated task. Our comparison thus arguably provides justification for certain uses of stated elicitation: results found using stated methods that bear on mean imprecision or on tendencies across subjects promise to hold up under our more theoretically rigorous method. Other results concerning the comparison—for instance, the fact that stated intervals are considerably wider than those elicited by our method in information-rich situations (Section 4.3)—flag potential limits. If the aim is to study absolute amounts of imprecision in contexts where there is plenty of information, perhaps stated probability intervals are not a sufficiently robust tool.

This suggests one direction for future research. As noted, stated probability intervals are typically used in large-scale surveys (such as Giustinelli et al. 2022). The sorts of comparisons conducted in lab settings in EXP A shed light on their performance, and in particular the performance loss with respect to incentive-compatible methods for specific research questions. As such, they provide indications of expected performance in the field. Further research can expand our comparison, by identifying more precisely the sorts of characteristics of intervals where stated methods fair well, by extending the comparison to natural (as opposed to artificial) sources of uncertainty, or by mapping the performance of different refinements of the stated approach. A properly grounded probability-interval elicitation method, of the sort developed in this paper, can serve as a tool for designing and evaluating simpler methods for use in large-scale studies.

Moreover, although our method was developed with the aim of demonstrating the possibility and feasibility of choice-based incentive-compatible probability-interval elicitation and investigating some basic characteristics of subjective probability intervals, future research could operationalise simpler, parametrised versions, with fewer choice questions. Such versions could be more implementable, for instance in field studies. Some large-scale surveys use choice tasks without necessarily incentivising them (e.g. Falk et al., 2018), and questions formulated in terms of bets may trigger different cognitive mechanisms to those formulated in terms of stated judgements. Our method could thus lay the foundations of a bet-based approach to add to the arsenal of probability-interval elicitation procedures used in practice.

Finally, analogous possibilities exist for expert elicitation exercises (Kriegler et al., 2009; Morgan, 2014). Compared to survey studies, these typically involve fewer subjects (experts), with each spending more time; accordingly, more precision is desired of the elicitation at the individual level. Aggregate-level performance of an elicitation method—of the sort suggested for stated methods by the results in Section 4.3—is less relevant for such exercises. Our experiments suggest the promise of our method to provide individual-level probability-interval elicitation with theoretically well-founded incentive-compatibility properties. Probability elicitation exercises in decision analysis often use bet-based choice tasks without necessarily incentivising them (e.g. Clemen and Reilly, 2013); again, our method, applied in this context, complements existing stated approaches to eliciting probability intervals.

6. Conclusion

This paper proposes and implements a solution to the open problem of choice-based incentive-compatible elicitation of multiple prior beliefs. It comprises a new preference-based notion—Matching Probability Intervals—and a probability-interval analogue of a state-of-the-art elicitation procedure for matching probabilities. Our elicitation operates under the Hurwicz α-maxmin EU model as well as a range of generalizations, and in the absence of strong assumptions about subjects' sets of priors, most notably any form of probabilistic sophistication.

Our implementation of the elicitation method, in three experiments to elicit subjective probability intervals and upper and lower CDFs over artificial and natural sources of uncertainty, testifies to its validity and feasibility. It finds a predominance of imprecision—intervals of non-zero width—across our subjects, for all explored sources, showing it to be related to information, familiarity or predictability. It also compares our choice-based elicitation with stated probability-interval methods, showing that they yield similar results in aggregate. Our method also allows us to perform what, to our knowledge, is the first elicitation of the mixture coefficient in the α-maxmin EU model that fully controls for beliefs.

Acknowledgements

We would like to thank Itzhak Gilboa, Simon Grant, Charles Manski and Peter Wakker, as well as audiences at D-TEA, RUD and SAET, for stimulating discussion and excellent feedback on early versions of this paper. We thank the Université Paris 1 Experimental Economics lab, the INSEAD-Sorbonne Behavioral Lab, Raffael Grasberger, Pierre-Antoine Amiand-Leroy, Maxime Frolov, Nawel Ayadi and Esso-Wazam Tchalanga for research assistance.

Funding

This research is supported by grants of the French National Research Agency (ANR), "Investissements d'Avenir" (LabEx Ecodec/ANR-11-LABX-0047) (ANR) and the ANR project DUSUCA (ANR-14-CE29-0003-01). We gratefully acknowledge the support of Hi! Paris.

Data Availability Section

The data and code underlying this research is available on Zenodo at https://doi.org/10.5281/zenodo.14652288.

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, 695–723.

ABDELLAOUI, M., H. BLEICHRODT, AND C. GUTIERREZ (2023): "Unpacking Overconfident Behavior When Betting on Oneself," *Management Science*.

ABDELLAOUI, M., E. DIECIDUE, E. KEMEL, AND A. ONCULER (2021): "Temporal Risk Resolution: Utility versus Probability Weighting Approaches," *Management Science*.

ANSCOMBE, F. J. AND R. J. AUMANN (1963): "A Definition of Subjective Probability," *The Annals of Mathematical Statistics*, 34, 199–205.

BAILLON, A. AND H. BLEICHRODT (2015): "Testing ambiguity models through the measurement of probabilities for gains and losses," *American Economic Journal: Microeconomics*, 7, 77–100.

BAILLON, A., H. BLEICHRODT, U. KESKIN, O. LÂHARIDON, AND C. LI (2018a): "The Effect of Learning on Ambiguity Attitudes," *Management Science*, 64, 2181–2198.

BAILLON, A., H. BLEICHRODT, C. LI, AND P. P. WAKKER (2021): "Belief hedges: Measuring ambiguity for all events and all models," *Journal of Economic Theory*, 198, 105353.

BAILLON, A., L. CABANTOUS, AND P. P. WAKKER (2012): "Aggregating imprecise or conflicting beliefs: An experimental investigation using modern ambiguity theories," *Journal of Risk and Uncertainty*, 44, 115–147.

BAILLON, A., Z. HUANG, A. SELIM, AND P. P. WAKKER (2018b): "Measuring ambiguity attitudes for all (natural) events," *Econometrica*, 86, 1839–1858.

BERGER, J. O., D. R. INSUA, AND F. RUGGERI (2000): "Bayesian robustness," *Robust Bayesian Analysis*, 1–32.

BEWLEY, T. F. (2002): "Knightian decision theory. Part I," *Decisions in Economics and Finance*, 25, 79–110.

BOREL, E. (1939): "Valeur pratique et philosophie des probabilitÃ©s.[Practical and philosophical value of probability]," *TraitÃ© du calcul des probabilitÃ©s et de ses applications [Treatise of probability calculus and its applications]*, 4.

BOSE, S. AND L. RENOU (2014): "Mechanism design with ambiguous communication devices," *Econometrica*, 82, 1853–1872.

CARNEY, M., JON CUNLIFFE, DAVE RAMSDEN, ANDREW HALDANE, JONATHAN HASKEL, MICHAEL SAUNDERS, SILVANA TENREYRO, AND GERTJAN VLIEGHE (2019): "Inflation Report," *Bank of England*.

CERRONI, S. (2020): "Eliciting farmersâ subjective probabilities, risk, and uncertainty preferences using contextualized field experiments," *Agricultural Economics*, 51, 707–724, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/agec.12587.

CHATEAUNEUF, A., J. EICHBERGER, AND S. GRANT (2007): "Choice under uncertainty with the best and worst in mind: Neo-additive capacities," *Journal of Economic Theory*, 137, 538–567.

CHEW AND J. S. SAGI (2006): "Event exchangeability: Probabilistic sophistication without continuity or monotonicity," *Econometrica*, 74, 771–786.

CHEW, S. H., B. MIAO, AND S. ZHONG (2017): "Partial ambiguity," *Econometrica*, 85, 1239–1260, publisher: Wiley Online Library.

CLEMEN, R. T. AND T. REILLY (2013): *Making Hard Decisions with Decision Tools*, Cengage Learning.

DELAVANDE, A., X. GINÃ©, AND D. MCKENZIE (2011): "Measuring subjective expectations in developing countries: A critical review and new evidence," *Journal of development economics*, 94, 151–163, publisher: Elsevier.

DIMMOCK, S. G., R. KOUWENBERG, AND P. P. WAKKER (2015): "Ambiguity attitudes in a large representative sample," *Management Science*.

DOW, J. AND S. R. DA COSTA WERLANG (1992): "Uncertainty aversion, risk aversion, and the optimal choice of portfolio," *Econometrica: Journal of the Econometric Society*, 197–204, publisher: JSTOR.

EICHBERGER, J., S. GRANT, D. KELSEY, AND G. A. KOSHEVOY (2011): "The alpha-MEU model: A comment," *Journal of Economic Theory*, 146, 1684–1698.

ELLSBERG, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *Quart. J. Econ.*, 75, 643–669.

ENGELBERG, J., C. F. MANSKI, AND J. WILLIAMS (2011): "Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition," *Journal of Applied Econometrics*, 26, 1059–1078, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.1206.

EPSTEIN, L. G. AND M. SCHNEIDER (2010): "Ambiguity and Asset Markets," *Annual Review of Financial Economics*, 2, 315–346.

FALK, A., A. BECKER, T. DOHMEN, B. ENKE, D. HUFFMAN, AND U. SUNDE (2018): "Global evidence on economic preferences," *The Quarterly Journal of Economics*, 133, 1645–1692.

GAJDOS, T., T. HAYASHI, J.-M. TALLON, AND J.-C. VERGNAUD (2008): "Attitude toward imprecise information," *J. Econ. Theory*, 140, 27 – 65.

GARLAPPI, L., R. UPPAL, AND T. WANG (2007): "Portfolio selection with parameter and model uncertainty: A multi-prior approach," *Review of Financial Studies*, 20, 41–81.

GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): "Differentiating ambiguity and ambiguity attitude," *J. Econ. Theory*, 118, 133–173.

GILBOA, I. AND D. SCHMEIDLER (1989): "Maxmin expected utility with non-unique prior," *J. Math. Econ.*, 18, 141–153.

GIUSTINELLI, P., C. F. MANSKI, AND F. MOLINARI (2022): "Precise or Imprecise Probabilities? Evidence from Survey Response on Late-onset Dementia," *Journal of the European Economic Association*, 20.

GRANT, S., P. RICH, AND J. D. STECHER (2019): "An Ordinal Theory of Worst- and Best-Case Expected Utility," SSRN Scholarly Paper ID 3369078, Social Science Research Network, Rochester, NY.

GUISO, L. AND G. PARIGI (1999): "Investment and demand uncertainty," *The Quarterly Journal of Economics*, 114, 185–227, publisher: MIT Press.

GUL, F. AND W. PESENDORFER (2014): "Expected uncertain utility theory," *Econometrica*, 82, 1–39.

——— (2015): "Hurwicz expected utility and subjective sources," *Journal of Economic Theory*, 159, 465–488.

HEY, J., G. LOTITO, AND A. MAFFIOLETTI (2010): "The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity," *Journal of Risk and Uncertainty*, 41, 81–111.

HILL, B. (2023): "Beyond uncertainty aversion," *Games and Economic Behavior*, 141, 196–222.

——— (2024): "Prices vs Quantities under Severe Uncertainty," SSRN 4826405, HEC Paris.

HOLLARD, G., S. MASSONI, AND J.-C. VERGNAUD (2016): "In search of good probability assessors: an experimental comparison of elicitation rules for confidence judgments," *Theory and Decision*, 80, 363–387, publisher: Springer.

HOSSAIN, T. AND R. OKUI (2013): "The binarized scoring rule," *Review of Economic Studies*, 80, 984–1001.

ILUT, C. L. AND M. SCHNEIDER (2014): "Ambiguous business cycles," *American Economic Review*, 104, 2368–99.

KARANKI, D. R., H. S. KUSHWAHA, A. K. VERMA, AND S. AJIT (2009): "Uncertainty Analysis Based on Probability Bounds (P-Box) Approach in Probabilistic Safety Assessment," *Risk Analysis*, 29, 662–675.

KLIBANOFF, P., S. MUKERJI, K. SEO, AND L. STANCA (2021): "Foundations of ambiguity models under symmetry: alpha-MEU and smooth ambiguity," *Journal of Economic Theory*, 105202.

KOCHER, M. G., A. M. LAHNO, AND S. T. TRAUTMANN (2018): "Ambiguity aversion is not universal," *European Economic Review*, 101, 268–283.

KRIEGLER, E., J. W. HALL, H. HELD, R. DAWSON, AND H. J. SCHELLNHUBER (2009): "Imprecise probability assessment of tipping points in the climate system," *Proceedings of the national Academy of Sciences*, 106, 5041–5046.

MACHINA, M. J. AND D. SCHMEIDLER (1992): "A More Robust Definition of Subjective Probability," *Econometrica*, 60, 745–780.

MANSKI, C. F. (2004): "Measuring expectations," *Econometrica*, 72, 1329–1376, publisher: Wiley Online Library.

MARINACCI, M. (2015): "Model uncertainty," *Journal of the European Economic Association*, 13, 1022–1100.

MASTRANDREA, M. D., C. B. FIELD, T. F. STOCKER, O. EDENHOFER, K. L. EBI, D. J. FRAME, H. HELD, E. KRIEGLER, K. J. MACH, AND P. R. MATSCHOSS (2010): "Guidance note for lead authors of the IPCC fifth assessment report on consistent treatment of uncertainties," *Intergovernmental Panel on Climate Change (IPCC)*.

MORGAN, M. G. (2014): "Use (and abuse) of expert elicitation in support of decision making for public policy," *Proceedings of the National Academy of Sciences*, 111, 7176–7184.

QIU, J. AND U. WEITZEL (2016): "Experimental evidence on valuation with multiple priors," *Journal of risk and uncertainty*, 53, 55–74.

SAVAGE, L. J. (1954): *The Foundations of Statistics*, New York: Dover, 2nd edn 1971.

SCHLAG, K. H. (2015): "A method to elicit beliefs as most likely intervals," *Judgment and Decision Making*, 10, 456.

SEIDENFELD, T., M. J. SCHERVISH, AND J. B. KADANE (2012): "Forecasting with imprecise probabilities," *International Journal of Approximate Reasoning*, 53, 1248–1261.

TRAUTMANN, S. T. AND G. KUILEN (2015): "Belief elicitation: A horse race among truth serums," *The Economic Journal*, 125, 2116–2135.

WAKKER, P. P. (2010): *Prospect Theory For Risk and Ambiguity*, Cambridge: Cambridge University Press.

WINKLER, R. L. AND A. H. MURPHY (1979): "Use of Probabilities in Forecasts of Maximum and Minimum Temperatures," *Meteorological Magazine*, 108, 317–329.

A. Proofs

For differentiable W as in representation (1), let $\partial_1 W([x,y])$ denote the partial derivative of W with respect to the first coordinate, x, at $[x,y]$, and similarly for $\partial_2 W([x,y])$ and the second coordinate.

Proposition 1 is a corollary of the following Proposition, the uniqueness of the MPI (Proposition A.2 below), and the fact that (2) corresponds to a case of (1) where W is differentiable, $\partial_1 W([x,y]) = \alpha$ and $\partial_2 W([x,y]) = 1-\alpha$.

Proposition A.1. *Let E be an event, and suppose preferences are represented according to (1) with a unique MPI for E and W differentiable with $\partial_1 W([\underline{p}(E),\overline{p}(E)]) > \partial_2 W([\underline{p}(E),\overline{p}(E)])$.*

a. For any $[x,y]$ in the R-B region (i.e. such that the corresponding preferences in Table 1 hold, for E), $\underline{p}(E) \leq x$ and $\overline{p}(E) \geq y$. Moreover, for any $[x,y]$ in the Wh region, $\underline{p}(E) \geq x$ and $\overline{p}(E) \leq y$.

b. *For any $[x,y]$ in the R region (i.e. such that the corresponding preferences in Table 1 hold, for E), every $[x',y']$ with $x' \geq x$ and $y' \geq y$ is also in R. Moreover, for any $[x,y]$ in the B region, every $[x',y']$ with $x' \leq x$ and $y' \leq y$ is also in B.*

Proof. Part b. follows directly from the fact that, under (1), given that W is increasing in both bounds, whenever $x \leq x'$ and $y \leq y'$, then $(z,[x,y],0) \preceq (z,[x',y'],0)$ and $(0,[x,y],z) \succeq (0,[x',y'],z)$.

As concerns part a., since W is increasing in both bounds, if $[x,y]$ is such that $x < \underline{p}(E)$ and $y < \overline{p}(E)$, then $(z,[x,y],0) \preceq (z,E,0)$ and $(0,[x,y],z) \succeq (0,E,z)$, so $[x,y]$ is in the B region. Similarly, for $[x,y]$ such that $x > \underline{p}(E)$ and $y > \overline{p}(E)$, $[x,y]$ is in the R region.

For each x such that there exists y with $(0,[x,y],z) \sim (0,E,z)$, let $J(x)$ be y such that this indifference holds. By Lower Stochastic Dominance, i.e. the fact that W is strictly increasing in the first coordinate (footnote 10), there is a unique $J(x)$ for all such x. By construction $\{[x,y] : (0,[x,y],z) \sim (0,E,z)\} = \{[x,J(x)]\}$. This set, which we call Ind_B, is the indifference curve for $[\underline{p}(E),\overline{p}(E)]$ in the space of bets on blue. Note that, since W is differentiable, so is J with $\frac{dJ}{dx}(x) = -\frac{\partial_2 W([x,J(x)])}{\partial_1 W([x,J(x)])}$ on its domain.

For each y such that there exists x with $(z,[x,y],0) \sim (z,E,0)$, let $I(y)$ be x such that this indifference holds. By Lower Stochastic Dominance, i.e. the fact that W is strictly increasing in the first coordinate (footnote 10), there is a unique $I(y)$ for all such y. By construction $\{[x,y] : (z,[x,y],0) \sim (z,E,0)\} = \{[I(y),y]\}$. This set, which we call Ind_R, is the indifference curve for $[\underline{p}(E),\overline{p}(E)]$ in the space of bets on red. Note that, since W is differentiable, so is I with $\frac{dI}{dy}(y) = -\frac{\partial_2 W([I(y),y])}{\partial_1 W([I(y),y])}$ on its domain.

Since there is a unique MPI, there exists a unique $[x,y]$ at the intersection of the two indifference curves; i.e. a unique $[x,y]$ with $y = J(x)$ and $x = I(J(x))$. For any sufficiently small $dx > 0$, $[x+dx, J(x)+J'(x)dx]$ belongs to the blue indifference curve Ind_B. Similarly, $[x+I'(J(x))J'(x)dx, J(x)+J'(x)dx]$ belongs to the red indifference curve Ind_R. Hence, for $x+dx$ with small $dx > 0$, the blue indifference curve Ind_B is 'above' the red indifference curve Ind_R (as in Figure 1) if and only if $I'(J(x))J'(x) < 1$. Substituting in the derivatives of I and J, this holds if and only if $\partial_1 W([x,y]) > \partial_2 W([x,y])$. Since the MPI is unique, it follows that the blue indifference curve is 'above' the red indifference curve for all $x' > x$. By similar reasoning, the blue indifference curve is 'below' the red one for all $x' < x$. The result follows from the fact that $[\underline{p}(E),\overline{p}(E)]$ is the MPI, the previously noted fact about points to the South-West and North-East of $[\underline{p}(E),\overline{p}(E)]$, and the definition of the R-B region (respectively Wh region) in Table 1 as those urns 'below' the blue indifference curve and 'above' the red (resp. 'above' the blue one and 'below' the red one).

□

Proof of Proposition 2. Under (1), it follows from the first preference pattern in Proposition 2 that $W([q,1-b]) > W([\underline{p}(E),\overline{p}(E)])$ for all $q > r$ and $W([q,1-b]) < W([\underline{p}(E),\overline{p}(E)])$ for all $q < r$, and similarly for the others. By the continuity of W, it thus follows from the first two preferences that $W([r, 1-b) = W([\underline{p}(E), \overline{p}(E)])$, and from the second pair of preferences that $W([b, 1-r]) = W([1-\overline{p}(E), 1-\underline{p}(E)])$. It thus follows that $(z, [r, 1-b], 0) \sim (z, E, 0)$ and $(0, [r, 1-b], z) \sim (0, E, z)$, so $[r, 1-b]$ is a MPI for E, as required. The converse direction is an immediate consequence of the fact that W is strictly increasing in the lower bound. □

Finally, we state for completeness the result on the uniqueness of the MPI.

Proposition A.2. *For any decision maker represented according to (2) with $\alpha \neq \frac{1}{2}$, and for any event E, there is a unique MPI for E.*

Proof. Existence is immediate from Eqs. (4) and (5). Uniqueness is immediate from the linearity of the indifference curves in $\mathscr{I}$-space (see Figure 1). □