

ISSN 1745-8587



Department of Economics, Mathematics and Statistics

BWPEF 1710

Eliciting Second-Order Beliefs

Subir Bose

University of Leicester

Arup Daripa

Birkbeck, University of London

December 2017

Eliciting Second-Order Beliefs

Subir Bose

Arup Daripa

University of Leicester
sb345@leicester.ac.uk

Birkbeck, University of London
a.daripa@bbk.ac.uk

December 10, 2017

Abstract

We study elicitation of subjective beliefs of an agent facing ambiguity (model uncertainty): the agent has a non-singleton set of (first-order) priors on an event and a second-order prior on these first-order belief-states. Such a two-stage decomposition of uncertainty and non-reduction of compound lotteries resulting from non-neutrality to the second-order distribution plays an important role in resolving the Ellsberg Paradox. The problem of eliciting beliefs on unobservable belief-states with ambiguity-sensitive agents is novel, and we introduce new elicitation techniques using report-dependent prize variations. We construct a direct revelation mechanism that induces truthful reporting of the first-order belief states as well as the second-order distribution on the belief-states as the unique best response. The technique requires knowledge of the sensitivity function to second-order distribution (capturing ambiguity attitude) and the vN-M utility function, which we also elicit.

JEL CLASSIFICATION: D81, D82

KEYWORDS: Ambiguity, second-order beliefs, elicitation of second-order beliefs (support and distribution), Klibanoff-Marinacci-Mukerji (2005) representation, elicitation of second-order-distribution-sensitivity function, elicitation of the vN-M utility function.

1 Introduction

Second-order probabilities and attitude towards such probabilities play an important role in explaining choice that appear anomalous under expected utility. When the data generating process is known, an agent faces only aleatory uncertainty, captured by a (subjective) probability model. However, the data generating process itself might be uncertain, in which case the agent faces a second layer of uncertainty over possible probability models. Such “model uncertainty” or “ambiguity” is then captured by subjective beliefs—referred to as second-order beliefs—over possible probability models. Non-reduction of consequent two-stage compounding of lotteries is critical for explaining choice behavior categorised as the Ellsberg paradox (Ellsberg, 1961).

Introduced by Segal (1987,1990), such non-reduction is key for representations of ambiguity-sensitive preference axiomatised by Klibanoff, Marinacci, and Mukerji (2005) (KMM), Nau (2006), Ergin and Gul (2009), Seo (2009). Consider an act $a_E b$ which pays a monetary amount a if event E happens, and an amount b otherwise. Suppose the agent has a finite set of beliefs $\{\mu_1, \dots, \mu_m\}$ about E , and suppose f_i is the second order belief on μ_i , $i \in \{1, \dots, m\}$. Here $\mu_i \in [0, 1]$, $f_i \in (0, 1]$ and $\sum_i f_i = 1$. In the KMM representation, an agent evaluates this act according to

$$\sum_{i=1}^m \phi(\mu_i u(a) + (1 - \mu_i)u(b)) f_i$$

where ϕ is an increasing function and u is the vN-M utility function. The function ϕ captures attitude towards ambiguity: the agent’s preference exhibits ambiguity aversion, neutrality or fondness according as ϕ is strictly concave, linear or strictly convex. Note that the first-order beliefs μ_i are the belief-states¹ with the second-order belief as the (unique) prior over these states. Representations based on different axiomatisations by Nau (2006), Ergin and Gul (2009), Seo (2009) have a similar structure, with attitude towards dispersion of second-order probabilities capturing ambiguity attitude.

In this paper we use this representation class and elicit² beliefs μ_i and f_i , given that the agent is sensitive to ambiguity. The precise attitude towards ambiguity is immaterial:

¹In the literature, the belief-states are sometimes called second-order states, with the payoff relevant states (E and E^c , the complement of E) being the first-order states. We do not use this terminology in the paper - we refer to the first-order beliefs μ_i as belief states and to f_i as the second order beliefs.

²By “elicit” we mean “to make truthful reporting of beliefs the *unique* best response of the agent.”

we elicit the beliefs for any increasing function ϕ that satisfies a certain “smoothness” property (explained later).³ The idea behind our methods is quite general, and could be adapted to other settings featuring beliefs on unobservable states.

The problem of eliciting (first and second order) beliefs of an ambiguity-sensitive agent throws up challenges that are qualitatively different from those encountered when dealing with beliefs on observable events. Since the belief-states are fundamentally unobservable,⁴ standard elicitation techniques fail to work. Indeed, as discussed in the literature review below, *all available mechanisms in the literature* depend on a reward based on an observable event and/or require the agent to be an SEU maximizer. Therefore solving the current problem requires new techniques.

It is worth noting that even though the belief-states are unobservable, if the set of belief-states of an agent (i.e. the support of the second-order beliefs) were known, we could elicit the second-order beliefs in a relatively straightforward manner using a variation of known mechanisms (described in section 3.1). However, since these are not observable and, except for some special cases (for example, for synthetic ambiguous events created in a laboratory) unknown to the mechanism designer, elicitation presents a significant challenge. In response, we introduce new techniques and conceptual innovations in designing the belief-elicitation mechanism. We describe a novel direct revelation mechanism in which the agent is asked to report ambiguous beliefs (the first order beliefs and the second-order belief distribution on these belief-states) and show that the unique best response of the agent is to report beliefs truthfully. The literature on belief elicitation on observable states compares *levels* of utility from acts and lotteries. Here, with unobservable states and ambiguity-sensitive agents, we show that we need to consider how utility *changes* as we vary the prizes. Our mechanism constructs a test of beliefs based on report-dependent variations in lottery prizes with the property that it is only under the truthful report that the agent receives the most improving variation always.

³Of course, our mechanism for belief elicitation is not meant for agents who are ambiguity-neutral; the smoothness property on ϕ implies that the agent cannot have ambiguity-neutral preferences. It is intuitively obvious that when the agent is ambiguity-neutral the agent chooses as if the agent has a unique belief regarding the likelihood of the event E (the “predictive prior”) which is what an elicitation mechanism should be able to elicit. The standard elicitation techniques of the literature elicit this (unique) predictive prior.

⁴Suppose an agent believes E will happen with probability either $1/3$ or $2/3$ and that these two belief-states are equally likely. While it is possible to observe whether E happens, it is not possible to ascertain whether E happened with probability $1/3$ or $2/3$ if it does happen.

As the literature shows, beliefs on observable events can be elicited without requiring knowledge of the utility function of the agent. However, as noted above, our mechanism relies on report-dependent variations in lottery prizes, and we do need to know how such variations change the agent's utility. Therefore, it is no longer possible to be agnostic about the evaluation functions (in this case ϕ and u), and we elicit these as part of our mechanism.

The analysis proceeds in four steps. First, we elicit the vN-M utility function u using a scheme similar to the well-known mechanism of [Becker, Degroot, and Marschak \(1964\)](#). This allows us to write all monetary payoffs as vN-M utilities, so that we can write expected utility from acts with money prizes as simply expected payoffs over corresponding vN-M utilities.

Second, we present a mechanism with prize variations and show some of its properties. We then assume ϕ is known and elicit beliefs (belief-states and second-order distribution) using the mechanism. This is the main part of the paper and comprises our principal contribution. Since our methods are new, before formally presenting our methods and results, we explain the idea behind our belief elicitation technique in some detail in [section 3](#).

Third, we elicit the ϕ function by creating a synthetic ambiguous event (so that the set of belief-states for that event are controlled by the mechanism designer, and the realized belief-state is known as it is the designer's choice) and then employing a variation of the mechanism in [Karni \(2009\)](#). Finally, we set up a grand mechanism to elicit ϕ and beliefs simultaneously.

1.1 Literature review

Let us now describe the relevant literature briefly. A large literature on scoring rules addresses the problem of eliciting beliefs on observable events. See [Gneiting and Raftery \(2007\)](#) for a discussion of this literature. Recent surveys by [Schotter and Trevino \(2014\)](#) and [Schlag, Tremewan, and van der Weele \(2015\)](#) provide superb accounts of theoretical and practical issues with procedures used in experiments to elicit subjective beliefs. The literature is unrelated to our belief-elicitation procedure since we elicit unobservable belief-states and beliefs on such states.

The literature on eliciting beliefs of agents who have non-standard preferences is relatively small. [Chambers \(2008\)](#) studies belief-elicitation of agents with maxmin expected utility preference using proper scoring rules. The main result is that the agent announces a single probability belonging to her set of priors. [Bose and Daripa \(2016\)](#) address the problem of belief-elicitation under the more general α -maxmin preferences, and, in contrast to the above, elicit the entire set of beliefs as well as the relevant preference parameter for α -maxmin (elicitation of the entire set of beliefs under maxmin is a special case).

The papers mentioned above do not face the question of eliciting beliefs over unobservable states. This question does arise in [Karni \(2016\)](#), who considers elicitation of beliefs of decision makers who have non-standard preferences but different from the one we focus on.⁵ In Karni's work, the agent faces a multi-stage environment. The agent faces Knightian uncertainty in period 0, but is an SEU maximizer in period 1 when this uncertainty is resolved.⁶ The mechanism elicits the beliefs in period 0 by allowing the choice between acts and lotteries to be delayed to period 1, allowing SEU payoff comparisons across choices. The fact that agents are SEU when making a decision is crucial for elicitation of beliefs in this work. In contrast, we have the more standard problem where the agent is sensitive to ambiguity, and must make choices while facing ambiguity. For this problem, the procedure in [Karni \(2016\)](#) cannot be applied, and our methods are quite different.

To summarise, all existing belief-elicitation procedures depend on observability of the belief-relevant event and/or agents being SEU maximizers for making choices. The problem we study is a departure from this, and we present novel belief-elicitation methods to address it.

2 Preliminaries

We are interested in eliciting the agent's beliefs regarding an event E .

Let a finite *non-singleton* set of probabilities represent the agent's (first-order) beliefs about

⁵See also [Karni \(2017\)](#).

⁶Either the agent has incomplete preferences (à la [Bewley \(2002\)](#)) in period 0 and beliefs are determined in period 1, or the agent is a Bayesian decision maker who in period 0 entertains the possibility of a range of possible posteriors for period 1 with a prior in period 0 over (information signals corresponding to) the set of the posteriors. In either case, the agent is an SEU maximizer (with a unique prior) in period 1.

E. The agent’s beliefs involve a second-order belief which is a probability distribution on the first-order beliefs. Any probability in the support of the second-order belief is called a belief-state.

Let the agent’s belief be denoted by $B \equiv ((\mu_1, f_1), \dots, (\mu_m, f_m))$ where (μ_1, \dots, μ_m) are the belief-states, $\mu_i \in [0, 1]$ and $\mu_i \neq \mu_j$ for all $i, j \in \{1, \dots, m\}$. Further, (f_1, \dots, f_m) are the second-order beliefs, where $f_i > 0$ is the belief associated with the state μ_i , $i \in \{1, \dots, m\}$ and $\sum_i f_i = 1$.

The task is to elicit B , i.e. elicit both the belief-states and the second-order belief attached to each state. We use direct revelation mechanisms for this task and “elicit X ” means the agent’s *unique* best response is to truthfully report X when facing the mechanism.

We consider eliciting beliefs of an agent with preference represented by the **KMM** model of smooth ambiguity. Consider an act $\hat{x}_E \hat{y}$ which pays a monetary amount \hat{x} if E happens, and an amount \hat{y} otherwise. An agent with KMM preference evaluates this according to

$$\sum_{i=1}^m \phi(\mu_i u(\hat{x}) + (1 - \mu_i)u(\hat{y})) f_i$$

where ϕ is an increasing function and u is the vN-M utility function.

Expressing the payoffs using vN-M numbers simplifies our presentation. Of course the agent’s vN-M utility function $u(\cdot)$ may be unknown to the mechanism designer but there are known mechanisms for eliciting agent’s risk preference and for the sake of completeness, we discuss in the next section a mechanism to elicit the vN-M utility function u . The agent’s risk preference (as represented by the function $u(\cdot)$) plays no role in our analysis and our mechanisms for belief elicitation would work irrespective of the risk attitude of the agent.

2.1 Eliciting the vN-M utility function

The agent’s risk preference is reflected in the vN-M utility function $u(\cdot)$. This function can be elicited using a mechanism similar to the well-known scheme of **Becker et al. (1964)**. Appendix **A.1** describes the mechanism to elicit $u(\cdot)$ ⁷. Since $u(\cdot)$ can be elicited and does not play any other role in the analysis, we reduce notation by adopting the convention

⁷More precisely, the agent has preference over risky alternatives with $u(\cdot)$ being a numerical representation; any affine transformation of $u(\cdot)$ is also a representation. Our use of the phrase “elicit $u(\cdot)$ ” is really

that all prizes for objective lotteries or ambiguous acts are vN-M utility numbers rather than money prizes. So, for example, if the money prize is \hat{x} , we write the prize as $x = u(\hat{x})$.

In common with most of the literature, we assume that *all* departures from expected utility is due to ambiguity aversion. In particular, this means that for all types of ambiguity-averse preferences, the agent's payoffs from *objective compound lotteries* is exactly the same as an expected utility maximizer's with same risk preferences.

2.2 Notation

In what follows, x, y, z etc denote vN-M numbers generated using $u(\cdot)$ from monetary payoffs. We assume the agent strictly prefers more money; in other words, if the money prize \hat{x} generating vN-M utility x is strictly higher than the money prize \hat{y} generating vN-M utility y , it follows that $x > y$.

$x_E y$ denotes the (subjective) act that pays x if event E happens and y if event E does not happen (i.e. event E -complement happens). An agent with KMM preference evaluates this according to

$$\phi \circ B = \sum_{i=1}^m \phi(\mu_i x + (1 - \mu_i) y) f_i$$

As noted before, using vN-M payoff numbers simplifies the presentation. We sometimes use the phrase KMM EU to refer to $\phi \circ B$.

3 Belief elicitation: an outline

This section provides an informal outline of our belief elicitation procedure. For subsections 3.1 and 3.2 we assume that the function ϕ is known to the mechanism designer.

First, consider a simpler case when the set of belief states are known. Of course this would not be true for most situations. Nevertheless, we consider this case since (i) in

a short form for saying elicit any member of this equivalence class of numerical representations. Put differently, by the phrase "a mechanism elicits u " we mean the agent's unique best response when facing this mechanism is to report a function such that the certainty equivalent of any lottery calculated by using this reported function is the same as the agent's true certainty equivalent of the lottery.

some situations the beliefs states may be known and it would be useful to describe the (simpler) mechanism that can be used in such cases, and (ii) more importantly, discussing this case allows us to explain better the complexities that arise when the mechanism has to elicit both the support (i.e. belief states) and the distribution of second order beliefs.

3.1 A simpler case: set of belief-states known

Here we elicit second-order beliefs $F = (f_1, \dots, f_m)'$ under the assumption that the set of belief-states (μ_1, \dots, μ_m) is known to the mechanism designer. Suppose the agent reports $\hat{F} = (\hat{f}_1, \dots, \hat{f}_m)'$. Using a simple Vickrey scheme, it is easy to elicit (as unique best response) the certainty equivalent (CE) of KMM EU from the act $x_E y$.⁸ Using this, we can back out F by creating enough acts as follows.

The agent is asked to report F and reports \hat{F} . The mechanism designer announces m prize vectors $(x_i, y_i), i \in \{1, \dots, m\}$ and chooses one prize vector randomly, calculates the associated report of CE, and runs the CE-elicitation scheme (see footnote 8).

Let H denote the following $m \times m$ matrix:

$$H = \begin{pmatrix} \phi(\mu_1 x_1 + (1 - \mu_1) y_1) & \dots & \phi(\mu_m x_1 + (1 - \mu_m) y_1) \\ \vdots & & \vdots \\ \phi(\mu_1 x_m + (1 - \mu_1) y_m) & \dots & \phi(\mu_m x_m + (1 - \mu_m) y_m) \end{pmatrix}$$

Note that H is known to the designer. The agent faces the problem of ensuring that reported and actual CE coincide for each prize vector, implying $H\hat{F} = HF$.

If H has full rank, it is clear that $\hat{F} = F$ is the only possible solution. So the only question is whether it is possible to ensure H has full rank. This is in fact a non-trivial question. However, we answer this question for a very similar matrix (where elements are derivatives of ϕ and the μ -values are reports) later, and prove that so long as the function ϕ has non-zero derivatives of a high enough order, this question can be translated into a question about the rank of a Vandermonde matrix, which is known to have full rank. A very similar proof can be used to show that it is indeed possible to construct a full-rank

⁸The CE is given by $\phi^{-1}(\phi \circ B)$. The agent is asked to report the CE, and reports r . The designer draws a value q randomly (say, using the uniform distribution) from $[y, x]$ (recall that $x > y$) and awards the agent the act if $r > q$ and the prize q otherwise. It is straightforward to see that the agent's best response is to report CE truthfully.

H . This then completes belief-elicitation if the set of belief-states were known.

3.2 An outline of the elicitation procedure

Unlike the (possibly unrealistic) case above, in the actual procedure we need to elicit the belief-states as well as the second-order probability weights attached to them. Let us see why the above procedure no longer works in this case.

Suppose the agent announces n values (n could be different from m , the actual number of belief-states) $(\hat{\mu}_1, \dots, \hat{\mu}_n)$ and corresponding second-order beliefs $(\hat{f}_1, \dots, \hat{f}_n)$. The mechanism designer would then announce n prize vectors, so that H is an $n \times n$ matrix. Since true H is no longer known, the designer can only construct \hat{H} based on reported μ -values. If we now used the same procedure as in section 3.1, the agent would need to ensure $\hat{H}\hat{F} = HF$. The possibility that this equality holds for reported μ and f values that are not truthful cannot be ruled out.

How can we force the agent to truthfully reveal both belief-states and second-order weights? Using prize-variations based on the reports received, we develop a test of the belief report. We explain this test in section 3.3 below. Passing the test for n different prize vectors requires the agent to ensure *two* conditions involving announced beliefs and true beliefs. We then show that so long as \hat{A} , the $n \times n$ matrix below, spans n dimensions, satisfying these conditions is impossible if the agent lies in any way: misreports any belief-state(s) or any second-order probabilities. It follows that truth-telling is uniquely optimal.

$$\hat{A} \equiv \begin{pmatrix} \phi'(\hat{\mu}_1 x_1 + (1 - \hat{\mu}_1) y_1) & \dots & \phi'(\hat{\mu}_n x_1 + (1 - \hat{\mu}_n) y_1) \\ \vdots & & \vdots \\ \phi'(\hat{\mu}_1 x_n + (1 - \hat{\mu}_1) y_n) & \dots & \phi'(\hat{\mu}_n x_n + (1 - \hat{\mu}_n) y_n) \end{pmatrix} \quad (3.1)$$

The remaining question then is whether it is possible to construct the matrix \hat{A} to have full rank. We show this depends on whether the function ϕ has enough non-vanishing derivatives. In particular, so long as the first to n -th derivative is non-zero (almost everywhere on the set of possible lotteries), the question can be transformed into one asking whether the $n \times n$ Vandermonde matrix is full rank. From standard results in linear algebra, we know that Vandermonde matrices are indeed full rank, which then implies \hat{A} is full rank.

Since n , the cardinality of the set of belief-states announced by the agent, can be chosen

by the agent, a sufficient condition for belief elicitation to work is that ϕ has non-zero derivatives (almost everywhere) of all orders. In the paper, we assume this property. This class includes exponential and log functions, power functions such as x^t where $t < 0$ or t is a positive fraction, functions such as x^x , trigonometric functions. Essentially the only class excluded is polynomials of positive integer degree. For this class, our proof works only so long as the degree exceeds the cardinality of the announced set of belief-states.

3.3 Prize variation: an introduction

We now explain the idea of creating a test of belief-reports through prize variations. Consider a prize vector (x, y) where $x > y$ and the act $x_E y$. Pick a value of μ from the unit interval. Call this μ_c . Consider the following prize variation: increase x and reduce y so as to keep the agent indifferent at the belief-state μ_c , i.e. keep $\mu_c x + (1 - \mu_c)y$ the same. In this case,

$$dy = -\frac{\mu_c}{1 - \mu_c} dx.$$

Would the variation make the agent better or worse off? Note that at all belief-states $\mu > \mu_c$ the agent is better off, while the opposite is true for $\mu < \mu_c$. Suppose $\mu_c \leq \mu_1$. Then for all μ values in the support of the second-order beliefs distribution, the agent is better off - so he is definitely better off overall. Similarly, the agent would definitely be worse off if $\mu_c \geq \mu_m$. As μ_c increases, the overall payoff continuously (and monotonically) decreases. It follows that there is a unique $\mu_c^* \in [\mu_1, \mu_m]$ such that if we were to change x and y to keep the agent indifferent at belief μ_c^* , the agent would be overall indifferent between receiving the variation and not receiving it. Figure 1 plots the change in KMM EU (of the act $x_E y$) from the variation.

Formally, starting from the KMM EU of the act $x_E y$, the change in KMM EU from the variation is

$$\sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) \left(\mu_i + (1 - \mu_i) \frac{dy}{dx} \right) f_i.$$

Using the value of dy/dx from above, the cutoff μ_c^* is then given implicitly by equating the change to 0:

$$\sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) \left(\frac{\mu_i - \mu_c^*}{1 - \mu_c^*} \right) f_i = 0. \quad (3.2)$$

Before we continue, we remind the reader that the exposition up to this point was done




Figure 1: The figure is drawn assuming $\phi(t) = \frac{1 - e^{-\alpha t}}{\alpha}$ and $B = ((\mu_1 = 1/3, f_1 = 1/3), (\mu_2 = 1/2, f_2 = 1/3), (\mu_3 = 2/3, f_3 = 1/3))$. Suppose \hat{B} is such that $\hat{\mu}_c^* < \mu_c^*$. At any μ_c between these two points, the agent's utility would improve from the variation but he does not get it. Similarly, if $\hat{\mu}_c^* > \mu_c^*$, at any μ_c in between, the agent would rather not get the variation, but does receive it.

as if the mechanism designer knows the true function ϕ . Of course that is typically not the case and our mechanism asks the agent to report both ϕ and B .

Suppose the agent reports $(\hat{\phi}, \hat{B})$. Above we calculated the “true” cutoff value μ_c^* using (ϕ, B) . Let $\hat{\mu}_c^*$ be the value of the cutoff calculated from $(\hat{\phi}, \hat{B})$.

Let us show that we can elicit μ_c^* , i.e. we can construct a mechanism such if the report generates a cutoff $\hat{\mu}_c^* \neq \mu_c^*$, the report would be suboptimal. This idea – that we can provide incentives to the agent to maintain the cutoffs generated through variations at their true values – forms a crucial part of the mechanism to elicit beliefs. Consider the following scheme.

- Ask the agent to announce ϕ, B . The agent announces $\hat{\phi}, \hat{B}$. Calculate $\hat{\mu}_c^*$.
- Pick a μ_c randomly (e.g. using a uniform distribution) from $(0, 1)$. If $\mu_c \leq \hat{\mu}_c^*$, offer the variation. Otherwise offer no variation (i.e. simply offer the original act).

Suppose that the agent's reports are such that $\hat{\mu}_c^* < \mu_c^*$. If $\mu_c \leq \hat{\mu}_c^*$, the agent gets the variation whether his report leads to cutoff value $\hat{\mu}_c^*$ or μ_c^* . If $\mu_c > \mu_c^*$, the agent gets no variation in either case. If $\hat{\mu}_c^* < \mu_c \leq \mu_c^*$, the agent's utility improves from the variation, but he does not get it under the report, while he would get it if his report had led to the true cutoff value. This can be seen from figure 1. A similar argument shows that any report leading to $\hat{\mu}_c^* > \mu_c^*$ is suboptimal. Therefore reporting in such a way that maintains

μ_c^* at the true value is the weakly dominant strategy.

The above shows that by introducing report-dependent variations, we can generate a cutoff value and create incentives for the agent to ensure that his report leads to the true cutoff value. If we can now show that only the truthful report on beliefs would lead to the true cutoff value, we would be done. However, the mechanism as described so far gives the agent too high a “degree of freedom” - it is perfectly possible for the agent to misreport ϕ and/or B in such a way that the calculated cutoff nevertheless coincides with the true cutoff.

An obvious way to try to fix this problem is to “curtail” the degree of freedom - for example the mechanism can try to create incentive such that the agent will always want to report ϕ truthfully. This is in fact part of what we do: the actual (grand) mechanism elicits ϕ (using a mechanism that does not ask for any report on beliefs). So for eliciting beliefs, we can assume, without loss of generality, that ϕ is known. It follows that the mechanism can use the true function ϕ in its calculation of the cutoff. This curtails degree of freedom to an extent, but this is not enough. Even if ϕ cannot be misreported, the mechanism with just one single value of cutoff still allows for the possibility that the agent might misreport beliefs B , yet the calculated cutoff coincides with the true one. To solve this problem, we need to generate multiple (at least two) cutoffs with the same property that the agent’s best response is to ensure all report-generated cutoffs are equal to their true values. With (properly chosen) multiple cutoffs, we can show that the only way to ensure these equalities is for the agent to report beliefs truthfully.

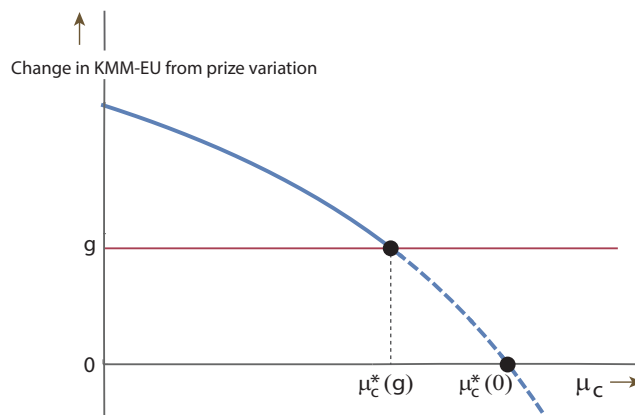


Figure 2: For $g = 0$, the agent is overall indifferent under $\text{Var}(x, y)$ at $\mu_c = \mu_c^*(0)$. If now we set $g > 0$, we get a lower cutoff $\mu_c^*(g)$.

How can we generate other cutoffs? If we could introduce a further variation that would improve the agent's KMM EU at a rate g that is independent of the agent's beliefs (true or reported), we could find a cutoff – call it $\mu_c^*(g)$ – such that if we were to change x and y to keep the agent indifferent at belief $\mu_c^*(g)$, the agent would be overall indifferent between receiving this (x, y) variation and receiving the belief-independent variation. Since $\mu_c^*(g)$ depends on g , by varying g we could find different values for the cutoff. This is shown in figure 2.

The mechanism we set up in section 5 finds a way to introduce such a belief-independent variation to compare with the (x, y) variation we introduced before, allowing us to generate multiple cutoffs. We ensure it is optimal for the agent to maintain all such cutoff values at their true values. We then show this is only possible if the report is truthful.

The rest of the paper is organized as follows. Section 4 formally introduces the mechanism with prize variations. Section 5 shows that B can be elicited assuming that ϕ is known using an augmented form of the mechanism introduced in the previous section. Next, section 6 sets up a mechanism to extract ϕ (this does not require any knowledge of or reports about B). Finally, section 7 constructs a grand mechanism combining the previously-introduced ones and shows that ϕ and B can be elicited simultaneously.

4 A mechanism with prize variations

This section introduces a mechanism with prize variations and derives certain properties. We make use of these in later sections to elicit beliefs.

We use the following two acts and prize variations on these acts.

$x_E y$ and $\text{Var}(x, y)$: Consider a prize vector (x, y) where $x > y > 0$ and the act $x_E y$. Let $\text{Var}(x, y)$ denote the following prize variation: Pick a value $\mu_c \in (0, 1)$, then raise x and reduce y so as to keep $\mu_c x + (1 - \mu_c)y$ the same. In this case, $dy = -\frac{\mu_c}{1-\mu_c} dx$.

$z_E z$ and $\text{Var}(z)$: Next, consider a constant act $z_E z, z > 0$. Let $\text{Var}(z)$ denote the following prize variation: raise z at the same rate as x in the other variation, i.e. $dz = dx$.

Recall that the agent's second-order belief is given by $B \equiv ((\mu_1, f_1), \dots, (\mu_m, f_m))$ where (μ_1, \dots, μ_m) are the belief-states or second-order states, $\mu_i \in [0, 1]$ and $\mu_i \neq \mu_j$ for all $i, j \in \{1, \dots, m\}$. Further, (f_1, \dots, f_m) are the second-order beliefs, where $f_i > 0$ is the belief associated with the state $\mu_i, i \in \{1, \dots, m\}$ and $\sum_i f_i = 1$.

Let $\hat{\phi}$ denote the agent's report on ϕ and let the belief report be $\hat{B} \equiv \{(\hat{\mu}_1, \hat{f}_1), \dots, (\hat{\mu}_n, \hat{f}_n)\}$

Let $\hat{\mu}_c^*(p)$ be given by the solution to

$$p \sum_{i=1}^n \hat{\phi}'(\hat{\mu}_i x + (1 - \hat{\mu}_i)y) \left(\frac{\hat{\mu}_i - \hat{\mu}_c^*(p)}{1 - \hat{\mu}_c^*(p)} \right) \hat{f}_i = (1 - p) \hat{\phi}'(z) \quad (4.1)$$

Let

$$\underline{p} \equiv \frac{\hat{\phi}'(z)}{\hat{T}_1 + \hat{\phi}'(z)} \quad (4.2)$$

where $\hat{T}_1 \equiv \sum_{i=1}^n \hat{\phi}'(\hat{\mu}_i x + (1 - \hat{\mu}_i)y) \hat{\mu}_i \hat{f}_i$. Note that $\underline{p} \in (0, 1)$, so that the interval $(\underline{p}, 1]$ is non-empty. It is straightforward to verify that so long as $\underline{p} < p \leq 1$, there is a unique solution $\hat{\mu}_c^*(p) \in (0, 1)$.⁹

It is worth noting at the outset that in all mechanisms specified, any randomization is done before the event E is realized.

⁹Let $\hat{T}_2 \equiv \sum_{i=1}^n \hat{\phi}'(\hat{\mu}_i x + (1 - \hat{\mu}_i)y) \hat{f}_i$. Solving from equation (4.1), $\hat{\mu}_c^*(p) = (\hat{T}_1 - \frac{1-p}{p} \hat{\phi}'(z)) / (\hat{T}_2 - \frac{1-p}{p} \hat{\phi}'(z))$. Since $\hat{T}_1 < \hat{T}_2$, $\hat{\mu}_c^*(p) < 1$. The value of \underline{p} follows from $\hat{\mu}_c^*(p) = 0$.

A mechanism with prize variations

The agent is asked to report B and ϕ . The agent reports $\widehat{B} \equiv \{(\widehat{\mu}_1, \widehat{f}_1), \dots, (\widehat{\mu}_n, \widehat{f}_n)\}$ and $\widehat{\phi}$. Note that the cardinality of the set of reported belief-states is n . After receiving the report the mechanism designer follows the procedure below, also depicted in Figure 3.

1. Announce n prize vectors $x_i, y_i, i \in \{1, \dots, n\}$, and randomly choose one: (x_k, y_k) .¹⁰ To simplify notation, we will denote the chosen (x_k, y_k) as (x, y) in the rest of this section. (Since only one prize vector is chosen there is no scope for confusion if we drop the subscript.)
2. Choose a $p \in (\underline{p}, 1]$ randomly (e.g. using a uniform distribution). (\underline{p} is defined by (4.2).)
3. Choose the act $x_E y$ with probability p and the constant act z_{EZ} with probability $(1 - p)$.
4. Calculate $\widehat{\mu}_c^*(p)$ (using equation (4.1) above).
5. Next, choose a μ_c randomly from $[0, 1]$ (e.g. using a uniform distribution). Adjust the act as follows.
 - If and only if (a) $\mu_c \leq \widehat{\mu}_c^*(p)$, and (b) $x_E y$ is picked in step 3, award the variation $\text{Var}(x, y)$ at μ_c .
 - If and only if (a) $\mu_c > \widehat{\mu}_c^*(p)$, and (b) z_{EZ} is picked in step 3, award the variation $\text{Var}(z)$.
 - In any other case, award no variation - i.e. simply award the original act picked in step 3.

§

¹⁰Each of the prize vectors are chosen from some subset of \mathbb{R}^2 according to some known distribution; the prize (x_k, y_k) is similarly chosen according to some known rule. The exact nature of the subset and the distribution is unimportant. So, for notational simplicity, we do not explicitly mention the distribution nor write the agent's payoff below as expectation with respect to the distribution. What is crucial is that the set of the n prize vectors as well as the chosen prize vector (x_k, y_k) are determined according to known rules that are part of the description of the mechanism - in particular these should not add any additional source of ambiguity to the choice problem that the agent faces. For belief elicitation, there is one crucial property that the chosen prizes must satisfy. This is discussed later in section 5.




Figure 3: The mechanism with prize variations. The expressions at the terminal nodes show not the level of KMM EU but the change in KMM EU (of the relevant act) arising from the variations. Note that only the variation received by the agent depends on the agent's announcement (through determination of $\mu_c^*(p)$).

Let us first calculate the (true) cutoff $\mu_c^*(p)$ given the mechanism above. For any choice of p , the agent calculates KMM EU as $\sum_{i=1}^m \phi(\mu_i w_1 + (1 - \mu_i) w_2) f_i$, where $w_1 = px + (1 - p)z$ and $w_2 = py + (1 - p)z$. Now suppose the agent is awarded variation $\text{Var}(x, y)$. The agent's utility should then change as:

$$\sum_{i=1}^m \phi'(\cdot) p \left(\mu_i + (1 - \mu_i) \frac{dy}{dx} \right) f_i dx$$

which can be written as

$$p \sum_{i=1}^m \phi'(\cdot) \left(\frac{\mu_i - \mu_c}{1 - \mu_c} \right) f_i dx$$

using the fact that $dy = -\frac{\mu_c}{1 - \mu_c} dx$.

To evaluate $\phi'(\cdot)$, the agent conditions on the information that variation $\text{Var}(x, y)$ can only be awarded if $x_E y$ is chosen in the first place. Therefore $\phi'(\cdot)$ is evaluated at $w_1 = x$ and

$w_2 = y$. It follows that the rate of change of KMM EU from $\text{Var}(x, y)$ is

$$p \sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) \left(\frac{\mu_i - \mu_c}{1 - \mu_c} \right) f_i dx. \quad (4.3)$$

Similarly, if the agent is awarded variation $\text{Var}(z)$, utility changes as: $\sum_{i=1}^m \phi'(\cdot)(1 - p)f_i dz$ where $\phi'(\cdot)$ is evaluated at $w_1 = z$ and $w_2 = z$ since the agent conditions on the information that this variation can be awarded only if initially z_{EZ} was chosen. Thus the rate of change of KMM EU from $\text{Var}(z)$ is given by $\sum_{i=1}^m \phi'(z)(1 - p)f_i dz$ which simplifies to

$$(1 - p)\phi'(z)dz. \quad (4.4)$$

From expressions (4.3) and (4.4), and using $dz = dx$, it follows that $\mu_c^*(p)$ is given implicitly by

$$p \sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) \left(\frac{\mu_i - \mu_c^*(p)}{1 - \mu_c^*(p)} \right) f_i = (1 - p)\phi'(z). \quad (4.5)$$

Note that when calculating $\hat{\mu}_c^*$ from the agent's report, we have used equation (4.1), which is the same equation but uses $\hat{\phi}$ and \hat{B} - the reported values for ϕ and beliefs.

We are now ready to present the first important result. This shows that the agent does not want the cutoff $\hat{\mu}_c^*(p)$ calculated from his report to be different from the true cutoff.

Theorem 1. *Under the mechanism with prize variations stated above, if the cutoff $\hat{\mu}_c^*(p)$ (derived using the report $(\hat{\phi}, \hat{B})$) differs from the "true" value $\mu_c^*(p)$ (derived using the true (ϕ, B)) for any $p \in [\underline{p}, 1]$, the report is suboptimal.*

While the formal proof is in the appendix, the discussion in section 3.3 explains the idea of the proof, which relies on a Vickrey-type dominance argument. Further, note that the act received by the agent does not depend on the agent's report. The report only influences the calculation of the cutoff $\hat{\mu}_c^*(p)$ (and therefore only influences the variation received by the agent). This, combined with the theorem above, shows that any report that preserves the cutoffs at their true value is an optimal report. It follows that truth-telling is an optimal strategy. In other words, the agent cannot do better than telling the truth by lying, and is at best indifferent between reporting truthfully or lying. We note this in the following corollary, which is useful for later arguments.

Corollary 1. *Under the mechanism with prize variations, truth-telling is a best response. In other words, any agent who submits a report $(\hat{\phi}, \hat{B})$ where either $\hat{\phi} \neq \phi$ and/or $\hat{B} \neq B$, must be indifferent between submitting this and submitting (ϕ, B) .*

5 Eliciting Beliefs

Corollary 1 above says that truth-telling is *an* optimal strategy. We now proceed to show that if we assume ϕ is known, then under the mechanism with prize variations presented in the previous section augmented by an additional property, reporting B truthfully is the *uniquely* optimal strategy. Later we elicit ϕ , and then construct a grand mechanism that elicits B and ϕ simultaneously.

Before we state the main result of this section on eliciting B , we need to introduce an assumption about the function ϕ . In constructing the proof, it is crucial that the $n \times n$ matrix \hat{A} (given by (3.1)), can be constructed to be of full rank by choosing the n prize vectors appropriately. We show next that a sufficient condition for \hat{A} to have full rank is that the function ϕ has non-zero derivatives (almost everywhere) of order at least n . Since the agent can report any number of belief-states, a sufficient condition is that ϕ has non-zero derivatives of all orders almost everywhere. We assume this:

Assumption 1. *The function $\phi(\cdot)$ has non-zero derivatives of all orders almost everywhere on its domain (set of all possible expected utilities from the set of lottery prize vectors).*

We assume that any report $\hat{\phi}(\cdot)$ has the same property.¹¹

Note that this class includes exponential and log functions, power functions such as x^t where $t < 0$ or t is a positive fraction, functions such as x^x , trigonometric functions. Essentially the only class excluded is polynomials of positive integer degree. For this latter class, our proof works only so long as the degree exceeds the cardinality of the announced set of belief-states.

Theorem 2. *Under assumption 1, it is possible to construct the matrix \hat{A} to have full rank for any finite n .*

The proof shows that the question can be translated into one about the rank of a Vandermonde matrix and then using its properties. We relegate this to the appendix. For the matrix \hat{A} to have full rank, the set of n prize vectors (x_i, y_i) have to be chosen in a certain way (see also footnote 10). One procedure is as follows. For any $\varepsilon > 0$, small, let X denote

¹¹Alternatively, if a report $\hat{\phi}$ does not have the property, the seller would withdraw from participation, leaving the agent with a zero payoff. By construction, the worst (realized) payoff that the agent can receive by participating in the mechanism is strictly positive.

the set $[\varepsilon, 1] \times [\varepsilon, 1]$. The first prize vector (x_1, y_1) is chosen from X randomly (for example according to the uniform distribution). By definition, the smaller of the two number is called y_1 (if the vector happens to be from the 45 degree line, it is discarded and a new vector is chosen). Let $M(k)$ denote the $k \times k$ square sub-matrix formed by the first k rows and first k columns of \hat{A} . The vector (x_1, y_1) determines the first row of \hat{A} . Given (x_1, y_1) (and the agent's belief report) calculate $X_2 \subset X$ such that if the second row of \hat{A} is constructed by using any vector from X_2 then $M(2)$ has rank 2, implying \hat{A} has rank at least 2. The proof shows this set is non-empty. The prize vector (x_2, y_2) is chosen randomly (according to conditional uniform) from X_2 . Similarly, (x_3, y_3) is chosen to ensure $M(3)$ has rank 3 and this process is continued till all the n prize vectors are chosen such that the constructed matrix \hat{A} has full rank.

Consider the following mechanism.

Γ_B : A mechanism to elicit belief states and second-order beliefs B

This is the same as the mechanism with prize variations outlined in the previous section with the following augmentation: in step 1, the mechanism designer chooses the n prize vectors to ensure that the matrix \hat{A} has full rank (we know this is possible from Theorem 2 above).

§

We now state the main belief-elicitation result.

Theorem 3. *Suppose ϕ is known, and assumption 1 holds. Then under the mechanism Γ_B , the agent's unique best response is to report B truthfully, i.e. to report truthfully the belief states and the second-order belief attached to each state.*

The idea of the proof is as follows: suppose we calculate $\hat{\mu}_c^*(p)$ using ϕ (the true function) and \hat{B} (the reported beliefs). Using Theorem 1 above, we know optimality requires $\hat{\mu}_c^*(p) = \mu_c^*(p)$. Using this condition, we derive two matrix equality conditions involving B and \hat{B} . We then show these can only be satisfied if $\hat{B} = B$.

5.1 Proof of Theorem 3

5.1.1 Step 1

We know from Theorem 1 that any optimal report must ensure $\widehat{\mu}_c^*(p) = \mu_c^*(p)$ for any $p \in (\underline{p}, 1]$. Using this, and $\widehat{\phi} = \phi$ (since we assume ϕ is known) in equation (4.1) and rewriting (4.5), we get the following:

$$\begin{aligned} \sum_{i=1}^n \phi'(\widehat{\mu}_i x + (1 - \widehat{\mu}_i)y) (\widehat{\mu}_i - \mu_c^*(p)) \widehat{f}_i &= \left(\frac{1-p}{p} \right) (1 - \mu_c^*(p)) \phi'(z) \\ \sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) (\mu_i - \mu_c^*(p)) f_i &= \left(\frac{1-p}{p} \right) (1 - \mu_c^*(p)) \phi'(z) \end{aligned}$$

Since the right hand sides are equal, we can equate the left hand sides:

$$\begin{aligned} \sum_{i=1}^n \phi'(\widehat{\mu}_i x + (1 - \widehat{\mu}_i)y) \widehat{\mu}_i \widehat{f}_i - \mu_c^*(p) \sum_{i=1}^n \phi'(\widehat{\mu}_i x + (1 - \widehat{\mu}_i)y) \widehat{f}_i \\ = \sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) \mu_i f_i - \mu_c^*(p) \sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) f_i \end{aligned}$$

Now, recall that the mechanism designer announces n prize vectors $x_i, y_i, i \in \{1, \dots, n\}$, and randomly chooses one: (x_k, y_k) . The above equation must be satisfied for any choice of (x_k, y_k) , so we get n conditions, which we can write in matrix notation as follows.

Let A be the following $n \times m$ matrix.

$$A \equiv \begin{pmatrix} \phi'(\mu_1 x_1 + (1 - \mu_1)y_1) & \dots & \phi'(\mu_m x_1 + (1 - \mu_m)y_1) \\ \vdots & & \vdots \\ \phi'(\mu_1 x_n + (1 - \mu_1)y_n) & \dots & \phi'(\mu_m x_n + (1 - \mu_m)y_n) \end{pmatrix}$$

\widehat{A} is as defined by (3.1). Next, let M be an $m \times m$ diagonal matrix of belief-states:

$$M \equiv \begin{pmatrix} \mu_1 & 0 & \dots & 0 \\ 0 & \mu_2 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & \mu_m \end{pmatrix}$$

Let \widehat{M} be the similar $n \times n$ diagonal matrix of reported belief-states. Next, let F denote the $m \times 1$ vector of second-order beliefs

$$F \equiv \begin{pmatrix} f_1 \\ \vdots \\ f_m \end{pmatrix}$$

and similarly define \widehat{F} as the $n \times 1$ vector of reported second-order beliefs. Finally, let $\mu_c^*(p)_k$ denote the cutoff value for the k -th prize vector (x_k, y_k) , $k = 1, \dots, n$. Let C_p be the following $n \times n$ diagonal matrix of cutoff values:

$$C_p \equiv \begin{pmatrix} \mu_c^*(p)_1 & 0 & \dots & 0 \\ 0 & \mu_c^*(p)_2 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & \mu_c^*(p)_n \end{pmatrix}$$

The n conditions above can now be written as

$$\widehat{A} \widehat{M} \widehat{F} - C_p \widehat{A} \widehat{F} = A M F - C_p A F$$

Rewriting,

$$\widehat{A} \widehat{M} \widehat{F} - A M F = C_p \left(\widehat{A} \widehat{F} - A F \right).$$

Note that the only term that involves p is C_p . Further, for any prize vector (x, y) , the corresponding cutoff $\mu_c^*(p)$ changes with p .¹² Thus every diagonal element in C_p changes with p . It follows that the right hand side varies with p . Therefore the above equality can hold only if

$$\widehat{A} \widehat{F} = A F. \tag{5.1}$$

This also implies

$$\widehat{A} \widehat{M} \widehat{F} = A M F.$$

Solving for \widehat{F} from each equation and setting them equal we have

$$\widehat{A}^{-1} A F = \widehat{M}^{-1} \widehat{A}^{-1} A M F.$$

¹²This is intuitively obvious: as p rises, we are raising the value of variation $\text{Var}(x, y)$, so $\mu_c^*(p)$ should rise with p . This can be seen formally by solving for $\mu_c^*(p)$ and taking the derivative. Let $T_1 \equiv \sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) \mu_i f_i$ and $T_2 \equiv \sum_{i=1}^m \phi'(\mu_i x + (1 - \mu_i)y) f_i$. Note that $T_1 < T_2$. Then, solving from equation (4.5) and differentiating, $\frac{\partial \mu_c^*}{\partial p} = \frac{\phi'(z)(T_2 - T_1)}{(T_2 p - (1-p)\phi'(z))^2} > 0$.

which implies that

$$\widehat{A}^{-1}A = \widehat{M}^{-1}\widehat{A}^{-1}AM. \quad (5.2)$$

This proves the following result.

Lemma 1. *Any report $\widehat{B} \equiv \{(\widehat{\mu}_1, \widehat{f}_1), \dots, (\widehat{\mu}_n, \widehat{f}_n)\}$ that does not satisfy equation (5.2) fails to satisfy $\mu_c^*(p) = \widehat{\mu}_c^*(p)$.*

5.1.2 Step 2

Lemma 1 above, combined with theorem 1, implies that any report not satisfying equation (5.2) is suboptimal. The proof now proceeds to show that it is not possible to satisfy equation (5.2) unless $\widehat{F} = F$ and $\widehat{M} = M$, implying that truth-telling is uniquely optimal.

Step 2.1 First, let us show that we cannot have $n \neq m$. We consider the two cases, $n > m$ and $n < m$, separately.

To clarify the idea behind the proof, here we use two simple examples for the two cases below. The general proof of step 2.1 – which uses the same logic but is perhaps less transparent – is in the appendix (section A.4).

Case A: $n > m$ Suppose the agent has a unique prior μ_1 , but reports 2 priors $\widehat{\mu}_1$ with weight \widehat{f}_1 and $\widehat{\mu}_2$ with weight \widehat{f}_2 , where $\widehat{f}_i > 0$ and $\sum_i \widehat{f}_i = 1$.

Let z_{ij} denote a typical element of the matrix $\widehat{A}^{-1}A$. Since the matrix is $n \times m$, we have $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$.

Applied to this case, equation (5.2) implies

$$\begin{pmatrix} z_{11} \\ z_{21} \end{pmatrix} = \begin{pmatrix} \frac{\mu_1}{\widehat{\mu}_1} z_{11} \\ \frac{\mu_1}{\widehat{\mu}_2} z_{21} \end{pmatrix} \quad (5.3)$$

Now, recall that $\widehat{F} = \widehat{A}^{-1}AF$. Here F is simply 1, so

$$\begin{pmatrix} \widehat{f}_1 \\ \widehat{f}_2 \end{pmatrix} = \begin{pmatrix} z_{11} \\ z_{21} \end{pmatrix}$$

Since $\hat{f}_i > 0$, both z_{11} and z_{21} must be positive. But then equation (5.3) cannot possibly hold, since μ_1 cannot be equal to both $\hat{\mu}_1$ and $\hat{\mu}_2$.

In the general case of m values and $n > m$ reports, exactly the same problem emerges. All n rows of $\hat{A}^{-1}A$ must contain positive elements. This then implies that for equation (5.2) to hold, some μ values would need to be equal to more than one $\hat{\mu}$ value, which is impossible.

Case B: $n < m$

Consider an example where the agent reports two belief states $\hat{\mu}_1$ and $\hat{\mu}_2$ (with positive weights \hat{f}_1 and \hat{f}_2 , respectively), while the agent has 4 belief states $\mu_i, i \in \{1, \dots, 4\}$ with second-order beliefs f_i . In this case, equation (5.2) implies

$$\begin{pmatrix} z_{11} & z_{12} & z_{13} & z_{14} \\ z_{21} & z_{22} & z_{23} & z_{24} \end{pmatrix} = \begin{pmatrix} \frac{\mu_1}{\hat{\mu}_1}z_{11} & \frac{\mu_2}{\hat{\mu}_1}z_{12} & \frac{\mu_3}{\hat{\mu}_1}z_{13} & \frac{\mu_4}{\hat{\mu}_1}z_{14} \\ \frac{\mu_1}{\hat{\mu}_2}z_{21} & \frac{\mu_2}{\hat{\mu}_2}z_{22} & \frac{\mu_3}{\hat{\mu}_2}z_{23} & \frac{\mu_4}{\hat{\mu}_2}z_{24} \end{pmatrix}$$

First, note that $\hat{\mu}_i$ must be equal to some μ_j . Otherwise, it is impossible to have $z_{ij} = \frac{\mu_j}{\hat{\mu}_i}z_{ij}$ for any choice of i, j where $i \in \{1, 2\}$ and $j \in \{1, \dots, 4\}$. Suppose, in this case,

$$\hat{\mu}_1 = \mu_2 \quad \text{and} \quad \hat{\mu}_2 = \mu_4.$$

Then \hat{A} is the matrix formed by columns 2 and 4 of A . It follows that the second column of $\hat{A}^{-1}A$ is the vector $(1, 0)'$ and the fourth column is the vector $(0, 1)'$ (these can be verified by direct calculation, but the idea should be obvious). In other words, the matrix formed by columns 2 and 4 of $\hat{A}^{-1}A$ is the 2×2 identity matrix.

Thus, equation (5.2) implies

$$\begin{pmatrix} z_{11} & 1 & z_{13} & 0 \\ z_{21} & 0 & z_{23} & 1 \end{pmatrix} = \begin{pmatrix} \frac{\mu_1}{\hat{\mu}_1}z_{11} & 1 & \frac{\mu_3}{\hat{\mu}_1}z_{13} & 0 \\ \frac{\mu_1}{\hat{\mu}_2}z_{21} & 0 & \frac{\mu_3}{\hat{\mu}_2}z_{23} & 1 \end{pmatrix}$$

Clearly, this is possible only if all z_{ij} terms in columns 1 and 3 of $\hat{A}^{-1}A$ are 0, so we would have

$$\hat{A}^{-1}A = \hat{M}^{-1}\hat{A}^{-1}AM = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

But $\hat{F} = \hat{A}^{-1}AF$. Therefore,

$$\begin{pmatrix} \hat{f}_1 \\ \hat{f}_2 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} f_1 \\ \vdots \\ f_4 \end{pmatrix}$$

But that implies $\hat{f}_1 = f_2$ and $\hat{f}_2 = f_4$. This is impossible since then $\sum \hat{f}_i = f_2 + f_4 < 1$.

The proof extends immediately to the general case of $n < m$.

Step 2.2 From step 2.1 above, it follows that we must have $n = m$. In this case, $z_{ij} = \frac{\mu_j}{\hat{\mu}_i} z_{ij}$ requires each $\hat{\mu}_i$ to be equal to a distinct μ_j . Since $n = m$, the reported vector must coincide with the true vector. In other words, we have $\hat{M} = M$. But then $\hat{A} = A$. Therefore, $\hat{A}^{-1}A$ is the $m \times m$ identity matrix, and $\hat{F} = \hat{A}^{-1}AF = F$. This completes the proof. ||

6 Eliciting ϕ

So far we have described the belief elicitation mechanism under the assumption that the function ϕ is known. However, this would be a fairly unrealistic assumption for most situations and we now discuss the mechanism that elicits ϕ . This should be of some independent interest as well (for example for policy making purposes where the policy maker may need to know ϕ .)

We first describe the construction of an ambiguous act and a (random) second-order act, which play crucial roles in the mechanism that elicits ϕ .

Ambiguous Act: For prizes $x > y$, consider a lottery that pays x with probability p (and hence y with probability $1 - p$) and it is announced that p is chosen from a non-singleton set \mathcal{P} . We can assume, without loss of generality, that \mathcal{P} has two elements, p_1 and p_2 . No further information as to how p is chosen from the set \mathcal{P} is provided. The lottery therefore constitutes an ambiguous act¹³ and letting H denote the agent's second-order belief over \mathcal{P} , the agent's KMM EU is $\phi \circ H$. With a slight abuse of notation, we use H to represent the agent's (second-order) belief that $p = p_1$; hence the number $1 - H$ represents the belief that $p = p_2$.

Second-Order Act (SOA): For some prizes $w > z$, the second-order act pays w if p is chosen to be p_1 and z if the chosen p is equal to p_2 . A **Random SOA (RSOA)**, chooses w and z independently according to some *announced* distribution. For example, two numbers are drawn independently from the uniform distribution over $[0, 1]$ with the larger of the two called w . Once w and z have been chosen, the rest of RSOA is then exactly as in SOA.

Objective Lottery: For $q \in [0, 1]$, let $\ell(q; w, z)$ denote the objective lottery that pays w with probability q and z with probability $1 - q$.

¹³The actual, *physical* construction of the ambiguous act can be done following techniques well-known in the literature and used in laboratory experiments. For example, a *two coloured urn* is constructed, consisting of red and white balls of unknown proportion and the act pays x if the colour of the ball drawn from the urn is red and pays y if the colour is white. It is announced that the proportion of red and white balls is such that the probability of drawing a red ball is either p_1 or p_2 but no further information is given regarding the probability of choosing p_1 or p_2 . There are well-known practices followed in laboratory experiments to make the process clear and credible to the experiment subjects and we assume some such method is employed.

Γ_ϕ : A mechanism for eliciting ϕ ¹⁴

The agent is asked to report ϕ and H . Let $\hat{\phi}$ and \hat{H} denote the corresponding reports.

1. The mechanism chooses prizes w, z as in RSOA.
2. With (objective) probability 1/2 each, the mechanism selects either **scheme A** or **scheme B**.
 - Under **Scheme A**, the mechanism chooses a number q randomly according to the uniform distribution over $[0, 1]$.
 - If $q \geq \hat{H}$ the agent is given the objective lottery $\ell(q; w, z)$.
 - If $q < \hat{H}$, the agent is awarded the SOA with prizes w and z .
 - Under **Scheme B**, the mechanism chooses a number randomly from the uniform distribution over $[z, w]$. Call the randomly chosen number t . The mechanism calculates, according to the reports of the agent, the certainty equivalent (recall this is in terms of vN-M utils, not money) of the ambiguous act with prizes w and z . That is, the mechanism calculates the certainty equivalent of $\hat{H} \hat{\phi}(p_1 w + (1 - p_1)z) + (1 - \hat{H}) \hat{\phi}(p_2 w + (1 - p_2)z)$. Let \widehat{CE} denote this calculated certainty equivalent.¹⁵ The mechanism allocates as follows.
 - If $t \geq \widehat{CE}$, the agent is awarded the prize t .
 - If $t < \widehat{CE}$, the agent is awarded the ambiguous act.

§

We now show why this mechanism elicits ϕ . Consider first scheme B. Let CE be the agent's *true* certainty equivalent. It is easy to see (using the standard weak-dominance argument) that the agent's best response is to submit \hat{H} and $\hat{\phi}$ such that the calculated certainty equivalent \widehat{CE} is equal to CE. Clearly, truthful reporting is a best response; however, it may not be the unique best response. On the other hand if scheme A is chosen

¹⁴The usual caveats apply. The mechanism elicits a function that is a representation of the agent's preference over second-order acts. If the agent reports an affine transformation of this function then that obviously is also a truthful elicitation (See **KMM**).

¹⁵This is given by $\hat{\phi}^{-1}\left(\hat{H} \hat{\phi}(p_1 w + (1 - p_1)z) + (1 - \hat{H}) \hat{\phi}(p_2 w + (1 - p_2)z)\right)$.

(which happens with strictly positive probability), the reported $\hat{\phi}$ is irrelevant and reporting \hat{H} to be the same as the true H is the unique best response. (Scheme A is essentially the mechanism in [Karni \(2009\)](#) and its chief advantage is that it elicits beliefs while allowing for the mechanism to remain agnostic about ϕ .) Given that the agent reports H truthfully, reporting $\hat{\phi} = \phi$ is then the unique best response under scheme B. To see this, note that if, say, $\hat{\phi}(\cdot) > \phi(\cdot)$ at any point in the domain of prizes, by continuity there must be an open interval in the domain on which $\hat{\phi}(\cdot) > \phi(\cdot)$. But then with positive probability $p_i w + (1 - p_i)z$ belongs to this interval for $i \in \{1, 2\}$, implying that $\widehat{CE} \neq CE$. Similarly, we can rule out having $\hat{\phi}(\cdot) < \phi(\cdot)$ at any point in the domain of prizes. This proves the following result.

Theorem 4. *Under the mechanism Γ_ϕ , the agent's unique best response is to report ϕ truthfully.*

It might be useful to compare Γ_ϕ - the mechanism to elicit ambiguity-preference - with procedures that elicit risk-preference $u(\cdot)$ (see [appendix A.1](#) for one such procedure). Elicitation of risk-preference involves use of constructed/engineered objective lotteries and of calculated (monetary) certainty equivalent of the lotteries using reports submitted by the agent. As one would expect, a constructed/engineered ambiguous act - and the calculated certainty equivalent of this act based on the agent's report - is used in Γ_ϕ . There is however an important and subtle difference. The certainty equivalent of a risky prospect does not depend on whether the prospect is subjective or objective. But $\phi(\cdot)$, which reflects the agent's ambiguity-preference, comes into play only if the act is subjective. Furthermore, the associated second-order belief H - without the knowledge of which the certainty equivalent of the ambiguous act cannot be calculated by the mechanism - is the agent's private information. This explains the need for scheme A. Its sole purpose is to elicit the second-order belief H which can then be used for calculating the certainty equivalent used in scheme B.

7 A grand mechanism to elicit B and ϕ concurrently

We have so far constructed the mechanism Γ_B in [section 5](#) to elicit B assuming ϕ is known and Γ_ϕ in [section 6](#) to elicit ϕ (the latter does not require any knowledge of B). In this section we use these to construct a grand mechanism to elicit B and ϕ together.

The grand mechanism is a straightforward combination of the two mechanisms discussed

so far:

Γ : A grand mechanism to elicit B and ϕ concurrently

Ask the agent to report B and ϕ . The agent reports \widehat{B} and $\widehat{\phi}$. Using these as the relevant reported values, run

- the mechanism Γ_B with probability $1/2$, and
- the mechanism Γ_ϕ with probability $1/2$.

§

Corollary 1 shows that under the mechanism with prize variations introduced in section 4 – and therefore also under mechanism Γ_B – reporting both ϕ and B truthfully is an optimal strategy. It follows that at worst the agent is indifferent between lying and telling the truth.

Next, under mechanism Γ_ϕ , reporting ϕ truthfully is uniquely optimal (Theorem 4). Also note that this mechanism does not seek reports on B . Since mechanism Γ_ϕ is offered with positive probability, truthful report of ϕ is uniquely optimal.

It follows that we need to only consider reports of the type (ϕ, \widehat{B}) , and then rule out $\widehat{B} \neq B$. This is accomplished by Theorem 3, which shows that if ϕ is known and assumption 1 holds, reporting B truthfully is uniquely optimal under mechanism Γ_B . Since Γ_B is offered with positive probability, reporting B truthfully is uniquely optimal.

It follows that under the grand mechanism Γ , it is uniquely optimal to report both ϕ and B truthfully. This proves the result below.

Theorem 5. *Suppose assumption 1 holds. Under the grand mechanism Γ above, the uniquely optimal response of the agent is to report B and ϕ truthfully.*

8 Conclusion

We study elicitation of subjective beliefs of an agent facing model uncertainty or ambiguity. The agent has non-singleton (first-order) priors on an event. Each prior (belief-state)

refers to a different underlying data generating process and the corresponding probability model. The agent then has a second-order prior on these first-order belief-states. The agent's preference exhibits non-neutrality to the second-order distribution. As discussed at the outset, such a two-stage decomposition of uncertainty and non-reduction of compound lotteries resulting from non-neutrality to the second-order distribution plays an important role in the literature in resolving the Ellsberg Paradox.

The large belief-elicitation literature focuses on eliciting beliefs on events that are observable. The entire scoring rule literature fits into this category. Here, on the other hand, we elicit (second-order) beliefs on unobservable belief-states. A few recent papers do also address problems pertaining to belief-elicitation on unobservable states, but these require the agent to be SEU maximizer at some decision making stage. In this paper, in contrast, we consider the more standard problem of eliciting first and second-order beliefs of agents who are ambiguity-sensitive and do in fact face ambiguity when making a decision. Here we adopt the smooth ambiguity representation of **KMM** for specificity. However, as discussed, several other axiomatisations result in a similar representation structure.

The literature on belief elicitation on observable states compares *levels* of utility from acts and lotteries. Here, with unobservable states and ambiguity-sensitive agents, we show that we need to consider how utility *changes* under report-dependent prize variations. Using such variations, we construct a novel direct revelation mechanism that induces truthful reporting of the first-order belief states as well as the second-order distribution on the belief-states as the unique best response. The mechanism requires knowledge of the sensitivity function to second-order distribution (capturing ambiguity attitude) and the vN-M utility function. We construct variations on established techniques to elicit these aspects of the agent's preferences.

While we consider a smooth ambiguity representation to elicit second-order beliefs, our idea of using report-dependant prize-variations is quite general, and should be applicable for eliciting beliefs on unobservable states more broadly.

A Appendix

A.1 Eliciting $u(\cdot)$

For $x > y$, where x and y are in *monetary* units, consider the objective lottery $\ell(p; x, y)$ that awards x with probability p and y with probability $(1 - p)$. A **mechanism** that elicits $u(\cdot)$ is as follows. (See footnote 7.)

- The agent is asked to report u . Let \hat{u} denote the report.
- The mechanism selects two numbers (call the larger of them x and the smaller one y) randomly from some interval according to independent draws. Without loss of generality let the interval be $[0, 1]$ with the random selections made according to the uniform distribution.
- The mechanism calculates the expected utility of $\ell(p; x, y)$ using reported function \hat{u} . Let \hat{V} denote the certainty equivalent.
- Next, the mechanism picks a number r randomly - for example using the uniform distribution - from the set $[y, x]$. If $r > \hat{V}$, the agent is awarded the certain amount r . If $r \leq \hat{V}$, the agent receives the lottery $\ell(p; x, y)$.

Let us show that this mechanism elicits u . Let V denote the true certainty equivalent of $\ell(p; x, y)$. Suppose $\hat{V} > V$. The misreporting is irrelevant if $r > \hat{V}$ or $r < V$. If $\hat{V} > r > V$, the report of \hat{u} earns $u(V)$ while the true report earns $u(r)$, which is higher. Similar arguments show that truthful reporting is weakly dominant if $V > \hat{V}$. Therefore reporting u such that $\hat{V} = V$ is the weakly dominant strategy under the mechanism. Now, if, say, $\hat{u}(\cdot) > u(\cdot)$ at any point in the domain of prizes, by continuity there must be an open interval in the domain on which $\hat{u}(\cdot) > u(\cdot)$. But then with positive probability $\hat{V} \neq V$. It follows that reporting $u(\cdot)$ truthfully is optimal.

A.2 Proof of Theorem 1

Recall that μ_c^* (we drop the argument p here as it plays no role in this proof - the arguments are valid for any $p \in (\underline{p}, 1]$) is given by equation (4.5). The expression on the left

hand side is the change in KMM EU from variation $\text{Var}(x, y)$ and that on the right hand side is the change in KMM EU from variation $\text{Var}(z)$. The cutoff μ_c^* is such that the two are equal. Note that the change from $\text{Var}(x, y)$ is preferred for $\mu_c < \mu_c^*$ and that from $\text{Var}(x, y)$ is preferred for $\mu_c > \mu_c^*$.

Case 1. Suppose $\hat{\mu}_c^* < \mu_c^*$.

If the μ_c chosen is such that $\mu_c < \hat{\mu}_c^*$, reporting (ϕ, B) or $(\hat{\phi}, \hat{B})$ does not make any difference - in either case, the agent receives the change in KMM EU from prize variation $\text{Var}(x, y)$ at μ_c . If $\mu_c > \hat{\mu}_c^*$, in either case the agent receives the change from $\text{Var}(z)$. Finally, suppose μ_c is such that $\hat{\mu}_c^* < \mu_c < \mu_c^*$. Then under a truthful report the agent receives the change in KMM EU from $\text{Var}(x, y)$ at μ_c and under the false report, the agent receives the change from $\text{Var}(z)$. But in this case, the change in KMM EU is higher from $\text{Var}(x, y)$.

It follows that reporting $(\hat{\phi}, \hat{B})$ under which $\hat{\mu}_c^* = \mu_c^*$ weakly dominates reporting some other $(\hat{\phi}, \hat{B})$ under which $\hat{\mu}_c^* < \mu_c^*$.

Case 2. Suppose $\mu_c^* < \hat{\mu}_c^*$.

If $\mu_c > \hat{\mu}_c^*$ or $\mu_c < \mu_c^*$, whether the report is true or false makes no difference to the agent's payoff. Finally, suppose $\mu_c^* < \mu_c < \hat{\mu}_c^*$. Then under the truthful report the agent receives the change in KMM EU from $\text{Var}(z)$, while under the false report the agent receives the change from $\text{Var}(x, y)$ at μ_c . However, in this case the agent prefers the change in KMM EU from $\text{Var}(z)$.

It follows that reporting $(\hat{\phi}, \hat{B})$ under which $\hat{\mu}_c^* = \mu_c^*$ weakly dominates reporting some other $(\hat{\phi}, \hat{B})$ under which $\mu_c^* < \hat{\mu}_c^*$.

The two cases above imply that reporting $(\hat{\phi}, \hat{B})$ under which $\hat{\mu}_c^* = \mu_c^*$ is the weakly dominant strategy. Since, under any report that leads to $\hat{\mu}_c^* \neq \mu_c^*$, there is a strictly positive probability that the choice of μ_c is such that the agent gets the lower change in KMM EU, reporting $(\hat{\phi}, \hat{B})$ under which $\hat{\mu}_c^* \neq \mu_c^*$ is suboptimal. ||

A.3 Proof of Theorem 2

We show here that under assumption 1, it is possible to find prize vectors $(x_1, y_1), \dots, (x_n, y_n)$ such that \hat{A} (given by equation (3.1)) has full rank.

In what follows, all μ -values are reports (i.e. they all have a “hat” on top). Since there is no possibility of confusion, for economy of notation, we remove all “hat” symbols in the proof that follows. Further, since we need to repeatedly use expressions of the form $\phi'(\hat{\mu}_i x_k + (1 - \hat{\mu}_i) y_k)$, we shorten it to $\phi'(\mu_i, x_k, y_k)$ (again, note that we are also removing the hat on μ_i).

With these notation changes, the matrix \hat{A} in equation (3.1) is now written as

$$\mathcal{A} = \begin{pmatrix} \phi'(\mu_1, x_1, y_1) & \dots & \phi'(\mu_n, x_1, y_1) \\ \vdots & \ddots & \vdots \\ \phi'(\mu_1, x_n, y_n) & \dots & \phi'(\mu_n, x_n, y_n) \end{pmatrix}$$

We need to construct \mathcal{A} to have full rank n .

Let $M(k)$ be the following $k \times k$ minor of \mathcal{A} , where $2 \leq k \leq n$.

$$M(k) = \begin{pmatrix} \phi'(\mu_1, x_1, y_1) & \phi'(\mu_2, x_1, y_1) & \dots & \phi'(\mu_k, x_1, y_1) \\ \phi'(\mu_1, x_2, y_2) & \phi'(\mu_2, x_2, y_2) & \dots & \phi'(\mu_k, x_2, y_2) \\ \vdots & \vdots & \dots & \vdots \\ \phi'(\mu_1, x_k, y_k) & \phi'(\mu_2, x_k, y_k) & \dots & \phi'(\mu_k, x_k, y_k) \end{pmatrix}$$

Note that a sufficient condition for the rank of \mathcal{A} to be at least k is that $M(k)$ is full rank.

To show that it is possible to construct \mathcal{A} so that it has full rank n , we use a proof by induction. Let $R(X)$ denote the rank of matrix X . We show that for $k = 2$, the minor $M(2)$ has full rank, establishing that $R(\mathcal{A}) \geq 2$. We then show that if $M(k-1)$ has full rank, then we can construct $M(k)$ so that the latter has full rank k , for any $k \in \{3, \dots, n\}$. In other words, $R(\mathcal{A}) \geq k-1$ implies $R(\mathcal{A}) \geq k$. Since $M(2)$ has rank 2, this shows that $M(3) \dots M(n)$ can be constructed to have full rank. But $M(n)$ is simply the matrix \mathcal{A} , which proves that $R(\mathcal{A}) = n$.

A.3.1 Step 1: $M(2)$ has full rank

We first show that $M(2)$ can be constructed to have rank 2.

Suppose not. Suppose for all possible choices of (x_2, y_2) , $M(2)$ has rank 1.

It follows that there are numbers β_1 and β_2 not both zero such that

$$\begin{pmatrix} \phi'(\mu_1, x_1, y_1) & \phi'(\mu_2, x_1, y_1) \\ \phi'(\mu_1, x_2, y_2) & \phi'(\mu_2, x_2, y_2) \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

Since

$$\beta_1 \phi'(\mu_1, x_2, y_2) + \beta_2 \phi'(\mu_2, x_2, y_2) = 0 \quad (\text{A.1})$$

for all possible choices of x_2, y_2 , it follows that a necessary condition for equation (A.1) to hold is that the derivative of the expression on the left hand side with respect to x_2 must also be zero:

$$\beta_1 \mu_1 \phi''(\mu_1, x_2, y_2) + \beta_2 \mu_2 \phi''(\mu_2, x_2, y_2) = 0. \quad (\text{A.2})$$

Further, suppose $\beta_1 + \beta_2 \neq 0$. We can write equation (A.1) as

$$\beta_1 \left(\phi'(\mu_1, x_2, y_2) - \phi'(\mu_2, x_2, y_2) \right) = -(\beta_1 + \beta_2) \phi'(\mu_2, x_2, y_2).$$

But then by choosing y_2 arbitrarily close to x_2 , we can make the left hand side arbitrarily close to zero, while the right hand side is bounded away from zero, thus violating equation (A.1), which is a contradiction.

It follows that for equation (A.1) to hold for all possible choices of x_2, y_2 , a necessary condition is

$$\beta_1 + \beta_2 = 0 \quad (\text{A.3})$$

Starting from equation (A.2) and using exactly the same argument, it follows that a further necessary condition for equation (A.1) to hold for all possible choices of x_2, y_2 is given by

$$\mu_1 \beta_1 + \mu_2 \beta_2 = 0. \quad (\text{A.4})$$

However, it is easy to see that the necessary conditions (A.3) and (A.4) can only be satisfied if both β_1 and β_2 are 0. To see this, write the necessary conditions (A.3) and (A.4) as

$$\begin{pmatrix} 1 & 1 \\ \mu_1 & \mu_2 \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

Since the first matrix on the left hand side has rank 2 (the determinant is $\mu_2 - \mu_1 > 0$), the only solution is $\beta_1 = \beta_2 = 0$.

It follows that equation (A.1) cannot be satisfied for all possible choices of x_2, y_2 unless all β -values are 0, implying that it must be possible to have some value x_2, y_2 for which $M(2)$ has rank 2.

A.3.2 Step 2: $M(k)$ has full rank if $M(k - 1)$ has full rank

Suppose $M(k - 1)$ is full rank but $M(k)$ has rank $k - 1$ for all possible values of the prize vector (x_k, y_k) . Then there exists $(\beta_1, \dots, \beta_k)$ not all zero such that for every value of (x_k, y_k) ,

$$\sum_{t=1}^k \beta_t \vec{c}_t = 0, \quad (\text{A.5})$$

where \vec{c}_t is the t -th column vector

$$\begin{pmatrix} \phi'(\mu_t, x_1, y_1) \\ \vdots \\ \phi'(\mu_t, x_k, y_k) \end{pmatrix}$$

Since $\sum_{t=1}^k \beta_t \phi'(\mu_t, x_k, y_k) = 0$ for all possible values of (x_k, y_k) where not all β_t values are 0, we get the following result.

Lemma 2. Consider any $k \in \{1, \dots, n\}$. The assumption that equation (A.5) holds for all values of the prize vector (x_k, y_k) implies that for any $\ell \in \{1, 2, \dots, k\}$,

$$\sum_{t=1}^k \beta_t \mu_t^{\ell-1} \phi^{(\ell)}(\mu_t, x_k, y_k) = 0, \quad (\text{A.6})$$

where $\phi^{(\ell)}(\cdot)$ denotes the ℓ -th partial derivative of $\phi(\cdot)$ with respect to x_k , and where not all values of β_t are zero.

Proof: Note that we have assumed that all derivatives of $\phi(\cdot)$ of orders up to k exist and are non-zero almost everywhere.

Also note that since we assume equation (A.5) holds, the equation is true for $\ell = 1$ by assumption.

Fix any y_k . Let

$$g(x, \ell) \equiv \sum_{t=1}^k \beta_t \mu_t^{\ell-1} \phi^{(\ell)}(\mu_t, x, y_k).$$

and suppose that $g(x, 1) = 0$ for all values of x .

Consider the case of $\ell = 2$. Suppose $g(x, 2) \neq 0$ for some value of x . Specifically, suppose $g(x, 2) > 0$ for some x . By continuity, there is some non-empty interval $[a, b]$ of values of x such that $g(x, 2) > 0$ on (a, b) . In this case if we choose $x \in (a, b)$, then by raising x slightly $g(x, 1)$ can be raised. Therefore it cannot be true that $g(x, 1) = 0$ for all values of x , which is a contradiction.

Now, for any $\ell \in 2, \dots, k$, suppose $g(x, \ell - 1) = 0$ for all values of x . This implies that $g(x, \ell) = 0$ for all values of x . To see this, suppose $g(x, \ell) > 0$ for some value of x . As above, this implies that $g(x, \ell) > 0$ for values of x in some interval (a, b) . Choosing x in the interval and then raising x would then raise $g(\cdot, \ell - 1)$ above 0, which is a contradiction.

This proves that if $g(x, \ell - 1) = 0$ for all values of x , then $g(x, \ell) = 0$ for all values of x .

Since we have shown that $g(x, 1) = g(x, 2) = 0$ for all values of x , it follows that $g(x, \ell) = 0$ for all values of $\ell = 1, \dots, k$. This completes the proof. ||

The next result derives some necessary conditions for equation (A.5) to hold for all prize vectors.

Lemma 3. *Consider any $k \in \{2, \dots, n\}$. The following conditions are necessary for equation (A.5) to hold for all prize vectors (x_k, y_k) . For any $\ell \in \{1, \dots, k\}$,*

$$\sum_{t=1}^k \beta_t \mu_t^{\ell-1} = 0, \tag{A.7}$$

where not all values of β_t are 0.

Proof: Fix any $\ell \in \{1, \dots, k\}$ and suppose

$$\sum_{t=1}^k \beta_t \mu_t^{\ell-1} \neq 0. \tag{A.8}$$

From Lemma 2, we know that a necessary condition for equation (A.5) to hold for all prize vectors (x_k, y_k) is equation (A.6):

$$\sum_{t=1}^k \beta_t \mu_t^{\ell-1} \phi^{(\ell)}(\mu_t, x_k, y_k) = 0.$$

The equation above can be written as

$$\sum_{t=1}^{k-1} \beta_t \mu_t^{\ell-1} \left(\phi^{(\ell)}(\mu_t, x_k, y_k) - \phi^{(\ell)}(\mu_k, x_k, y_k) \right) = - \sum_{t=1}^k \beta_t \mu_t^{\ell-1} \phi^{(\ell)}(\mu_k, x_k, y_k)$$

It follows from equation (A.8) that the right hand side is not 0.

By choosing y_k very close to x_k we can make the left hand side as small as we like, while the right hand side gets close to

$$- \sum_{t=1}^k \beta_t \mu_t^{\ell-1} \phi^{(\ell)}(\mu_k, x_k, x_k)$$

which is bounded away from zero. It follows that equation (A.6) is violated. This is a contradiction. This proves the result. ||

Writing the necessary conditions given by lemma 3 in matrix form, we get

$$\begin{pmatrix} 1 & 1 & \dots & 1 \\ \mu_1 & \mu_2 & \dots & \mu_k \\ \mu_1^2 & \mu_2^2 & \dots & \mu_k^2 \\ \vdots & \vdots & \dots & \vdots \\ \mu_1^{k-1} & \mu_2^{k-1} & \dots & \mu_k^{k-1} \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_k \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

Now, let $V(k)$ denote the first matrix on the left hand side. The transpose of $V(k)$ is a $k \times k$ Vandermonde matrix. It follows that¹⁶

$$\det(V(k)^T) = \prod_{1 \leq i < j \leq k} (\mu_j - \mu_i).$$

Since $\mu_j \neq \mu_i$ for any $j \neq i$, the determinant is non-zero. Therefore the rank of $V(k)^T$ is k . Since taking a transpose does not change rank (or, indeed, the determinant), $V(k)$ has full rank as well.

¹⁶This is a standard result in matrix algebra. See, for example, [Horn and Johnson \(2013\)](#), chapter 0.9.11.

Since $V(k)$ has full rank, there is a unique solution for the β -values. Since $\beta_t = 0$ for all $t = 1, \dots, k$ is a solution, this must be the only solution.

This implies that the null-space of the minor $M(k)$ contains only the zero-vector, implying that $M(k)$ has k linearly independent columns, indicating that $R(M(k)) = k$.

Thus, starting from the assumption that $M(k-1)$ has full rank, we have shown that we can find a prize vector x_k, y_k such that $M(k)$ has full rank. We have also shown that $M(2)$ has full rank. It follows by induction that $M(k)$ can have full rank for all values of $k \in \{2, \dots, n\}$. Since $M(n)$ is simply the matrix \mathcal{A} , this proves that we can find n prize vectors such that \mathcal{A} has full rank.

A.4 Step 2.1 in the proof of Theorem 3

In the paper, we showed that it is not possible to satisfy equation (5.2) if $n \neq m$ using examples. We present the general proof here, which uses the same ideas as in the examples.

Note that since $\hat{F} = \hat{A}^{-1}AF$, and \hat{F} is $n \times 1$, $\hat{A}^{-1}A$ must have at least one non-zero element in each row. It follows that for equation (A.9) to hold, the following must hold:

$$\text{each } \hat{\mu}_i \text{ must be equal to some } \mu_j. \quad (*)$$

Case A: $n > m$: If $n > m$, it is not possible to satisfy condition (*) for all $\hat{\mu}_i$ since there are more values $\hat{\mu}_i$ than values μ_j .

Case B: $n < m$: Let Z be a $n \times (m-n)$ matrix with typical element denoted by z_{ij} . Let I_n denote the $n \times n$ identity matrix. Let $k_i, i \in \{1, \dots, n\}$ be n numbers such that $1 \leq k_1 < k_2 \dots < k_n \leq m$. Let $Z \oplus I_n$ denote an augmented matrix that inserts columns from I_n in the Z matrix as follows: the first column of I_n is inserted after the $(k_1 - 1)$ -th column of Z , second column of I_n inserted after the $(k_2 - 2)$ -th column of Z and so on until the last column of I_n inserted after the $(k_n - n)$ -th column of Z .

Let Z_1 denote the matrix formed by the first $(k_1 - 1) \geq 0$ columns of Z , Z_2 denote the matrix formed by the next $(k_2 - k_1 - 1) \geq 0$ columns of Z and so on until Z_{n+1} formed by the last $(k_{m-n} - k_{m-n-1} - 1) \geq 0$ columns of Z . The augmented matrix has the following

form:

$$Z \oplus I_n \equiv \left(Z_1 \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} Z_2 \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix} \cdots Z_n \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix} Z_{n+1} \right)$$

k_1 -th col k_2 -th col k_n -th col

Recall that condition (5.2) requires $\hat{A}^{-1}A = \hat{M}^{-1}\hat{A}^{-1}AM$. Note that $\hat{A}^{-1}A$ is $n \times m$ with typical element z_{ij} , where $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. The matrix $\hat{M}^{-1}\hat{A}^{-1}AM$ is also $n \times m$ with typical element $\frac{\mu_j}{\hat{\mu}_i}z_{ij}$. Thus condition (5.2) requires

$$z_{ij} = \frac{\mu_j}{\hat{\mu}_i}z_{ij} \quad (\text{A.9})$$

for all $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. Note that this either requires $z_{ij} = 0$, or if $z_{ij} \neq 0$, then $\mu_j = \hat{\mu}_i$.

The proof proceeds through the following Lemma. Let Z_0 denote an $n \times (m - n)$ matrix where each element in the matrix is 0.

Lemma 4. *To satisfy equation (5.2), a necessary condition is that the matrix $\hat{A}^{-1}A$ is of the form $Z_0 \oplus I_n$.*

Proof: To satisfy condition (*), let $k_i, i \in \{1, \dots, n\}$ be such that $\hat{\mu}_i = \mu_{k_i}$, where $1 \leq k_1 < k_2 < \dots < k_n \leq m$. In this case, the matrix \hat{A} is simply the matrix formed by columns k_1, k_2, \dots, k_n of matrix A . It follows that $\hat{A}^{-1}A$ is of the form $Z \oplus I_n$.

Given this, each row has one element 1. Consider row i . The element in the k_i -th column is 1. Suppose some $z_{ij} \neq 0$ where $j \neq k_i$. Then equation (A.9) would require $\hat{\mu}_i = \mu_j$. This is impossible since we already have $\hat{\mu}_i = \mu_{k_i}$. It follows that row i must have $z_{ij} = 0$ for all $j \neq k_i$. This implies that all elements of the matrix Z must be 0. It follows that $\hat{A}^{-1}A$ is of the form $Z_0 \oplus I_n$. This completes the proof of Lemma 4. ||

Since $\hat{F} = \hat{A}^{-1}AF$, and $\hat{A}^{-1}A$ is of the form $Z_0 \oplus I_n$ (from Lemma 4 above), it follows that

$$\hat{f}_i = f_{k_i}.$$

But since $n < m$, $\{k_1, \dots, k_n\}$ is a strict subset of $\{1, \dots, m\}$, it follows that $\sum_{i=1}^n \hat{f}_i = \sum_{i=1}^n f_{k_i} < 1$, which is impossible. Therefore we can rule out the case $n < m$.

Above we ruled out $n > m$ and $n < m$. This proves the claim in step 2.1 that we cannot have $n \neq m$.||

References

- BECKER, G. M., M. H. DEGROOT, AND J. MARSCHAK (1964): "Measuring utility by a single-response sequential method," *Behavioral Science*, 9, 226–232. 3, 5
- BEWLEY, T. F. (2002): "Knightian decision theory: Part I," *Decisions in Economics and Finance*, 25, 79–110. 4
- BOSE, S. AND A. DARIPA (2016): "Eliciting Ambiguous Beliefs Under α -Maxmin Preference," Working Paper, <https://ssrn.com/abstract=2749424>. 4
- CHAMBERS, C. P. (2008): "Proper scoring rules for general decision models," *Games and Economic Behavior*, 63, 32–40. 4
- ELLSBERG, D. (1961): "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics*, 75, 643–669. 1
- ERGIN, H. AND F. GUL (2009): "A Subjective Theory of Compound Lotteries," *Journal of Economic Theory*, 144, 899–929. 1
- GNEITING, T. AND A. E. RAFTERY (2007): "Strictly Proper Scoring Rules, Prediction, and Estimation," *Journal of the American Statistical Association*, 102, 359–378. 3
- HORN, R. A. AND C. R. JOHNSON (2013): *Matrix Analysis*, Cambridge University Press, 2nd ed. 35
- KARNI, E. (2009): "A Mechanism for Eliciting Probabilities," *Econometrica*, 77, 603–606. 3, 26
- (2016): "A Mechanism for the Elicitation of Second-Order Beliefs and Subjective Information Structures," Working Paper. 4
- (2017): "A Mechanism for Eliciting Second-Order Beliefs and the Inclination to Choose," Working Paper. 4
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73(6), 1849–1892. 1, 5, 25, 28
- NAU, R. F. (2006): "Uncertainty Aversion with Second-Order Utilities and Probabilities," *Management Science*, 52, 136–145. 1

- SCHLAG, K. H., J. TREMEWAN, AND J. J. VAN DER WEELE (2015): "A penny for your thoughts: a survey of methods for eliciting beliefs," *Experimental Economics*, 18, 457–490. 3
- SCHOTTER, A. AND I. TREVINO (2014): "Belief elicitation in the laboratory," *Annual Review of Economics*, 6, 103–128. 3
- SEGAL, U. (1987): "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 28, 175–202. 1
- (1990): "Two-Stage Lotteries without the Reduction Axiom," *Econometrica*, 58, 349–377. 1
- SEO, K. (2009): "Ambiguity and Second-Order Belief," *Econometrica*, 77, 1575–1605. 1