

A Mechanism for the Elicitation of Second-Order Belief and Subjective Information Structure

Edi Karni*

Johns Hopkins University and University of Warwick

July 11, 2018

Abstract

This paper describes a direct revelation mechanism for eliciting (a) Decision makers' range of subjective priors under Knightian uncertainty and their second-order introspective belief and (b) Bayesian decision makers' range of posteriors and their subjective information structure.

Keywords: Second-order belief; subjective information structure; Knightian uncertainty; probability elicitation; revelation mechanism; robust Bayesian statistics; random choice

JEL classification numbers: D80, D82, D83

*I am grateful to Zvi Safra, Andrei Savochkin, Jay Lu, Christopher Chambers, Nicholas Lambert and an anonymous referee for their useful comments and suggestions.

1 Introduction

An expert’s assessment of the likelihood of an event in which he has no stake may be of interest to others. For example, a patient may want to obtain a second opinion about the likelihood of success of a treatment recommended by his physician. A venture capitalist may be interested in an engineer’s assessment of the chance of success of a new technology for generating electricity from sea waves. In some instances the expert’s beliefs may be represented by a *set of priors*, which makes it impossible for him to deliver a precise assessment of the probability of the event of interest. Such an expert may be able to provide a range of the probabilities instead.¹ In other instances, a Bayesian expert (that is, a subjective expected utility maximizer whose beliefs are represented by a unique prior) may anticipate receiving new, private, information that would affect his assessment. Such an expert could deliver a precise assessment of his prior of the event of interest. However, he entertains a *set of posteriors* corresponding potential information signals he anticipates receiving. In this case, the expert’s assessment takes the form of a range of posterior probabilities. Whether priors or posteriors, I refer to the expert’s set of subjective probabilities of the event of interest as *first-order beliefs*.

In both instances, the expert may also entertain beliefs about the likelihoods of the probabilities in the corresponding sets. In the case in which the first-order beliefs is represented by a set of priors, I refer to these likelihoods as *second-order belief*, and in the case in which the first-order beliefs is represented by a set of posteriors I refer to them as the *subjective information structure*.² In either case, the second-order belief or information structure

¹In robust Bayesian statistics, the elicitation of the set of priors is analogous to the elicitation of a single prior in Bayesian statistics (see Seidenfeld, Schervish and Kadane [1995]). Multi-prior models are part of the more general theory of imprecise probability, which allows for partial probability specifications. It is applicable when information is insufficient to identify a unique (prior) probability distribution (“it is useful for dealing with expert elicitation, because decision makers have a limited ability to determine their own subjective probabilities and might find that they can only provide an interval.” [Wikipedia, Imprecise probability]).

²The term “second-order belief” is defined in the literature dealing with interactive decision making to describe one player’s belief of another player belief. It is also employed in the theory of decision making under uncertainty in the sense of its use here, namely, to describe a decision maker’s belief over a set of priors (e.g., Seo [2009], Nascimento and

are represented by a unique probability measure.

The situation in which an expert or a decision maker entertains a set of priors arises when the expert's preference relation over the set of non-constant, state-contingent consequences (or *acts*), is incomplete. Bewley's (2002) Knightian uncertainty model characterizes this situation.³ A decision maker's second-order belief in the context of incomplete preferences is modeled in Karni and Safra (2016), according to whom, decision makers' display random choice among acts which is governed by their moods, beliefs, or states of mind. In the present context, states of mind correspond to the decision maker's beliefs and are represented by belief-contingent preference relations on acts. The fact that a decision maker's may entertain distinct beliefs à la Bewley, reflects the underlying incompleteness of her preference relation on the set of acts. According to Karni and Safra, decision makers entertain introspective belief about their likely beliefs.

The situation in which a Bayesian expert entertains a set of posterior beliefs corresponding to his subjective information structure was studied in Dillenberger, Lleras, Sadowski and Takeoka (2014) and Lu (2016). Both models describe Bayesian decision makers who anticipate receiving private signals before choosing an act from a set of acts (menu).

Dillenberger et al. characterize what they refer to as subjective learning representations of preference relations on menus. A subjective learning representation involves a unique information structure which takes the form of a probability measure on canonical signal space (that is, the set of distributions on the state space) representing the decision maker's subjective beliefs on the set of posteriors.

Using similar framework, Lu (2016) models an analyst who observes the decision maker's random choice but is not privy to the signal he receives and acts upon. Lu shows that, if the distribution on the canonical signal space corresponds to the observed random choice rule (that is, the empirical distribution of choices of elements from menus of acts), the analyst can identify the decision maker's private information structure by observing binary choices.

Despite the differences in their analytical frameworks, their axiomatic foundations, and their implied choice behavior, the works of Karni and Safra

Riella [2013] and Giraud [2014]).

³See also Galaabaatar and Karni (2013).

(2016), Dillenberger et al. (2014) and Lu (2016), share in common representations of the decision maker’s second-order belief or private information structure. Moreover, in all of these models decision makers exhibit random choice behavior. According to Karni and Safra, when facing a choice between two alternatives that are non-comparable (that is one alternative is preferred over another according to some priors while it is less preferred according to other priors in the set of priors), choice behavior is random. More specifically, some impulse triggers a choice of a prior from the set of priors according to the second-order belief, and this prior determines the alternative to be chosen according to the subjective expected utility criterion. According to Dillenberger et al. and Lu a decision maker receives, at an interim stage, a private signal on the basis of which he updates his beliefs. The posterior belief governs his eventual choice.

Suppose that an uninformed party (henceforth, the elicitor) would like to elicit the expert’s assessment of range of the prior or posterior probabilities of the event of interest as well as his subjective belief regarding the likelihoods of the different priors or the posteriors in the corresponding sets. In this paper I propose a direct revelation mechanism requiring the expert to submit a report that allows the simultaneous elicitation of the range of priors (or posteriors) and his subjective assessment of the probabilities that the priors (or posteriors) in the corresponding sets are true. The mechanism is a modified quadratic scoring rule. If the expert displays incomplete beliefs the mechanisms elicit the range of the expert’s priors of an event of interest, as well as his introspective second-order belief (à la Karni and Safra [2016]). If the expert is Bayesian, the mechanism can be used to elicits the range of the expert’s posterior probabilities of the event of interest and his subjective information structure (à la Dillenberger et al. [2014] and Lu [2016]).

With two notable exceptions, Chambers and Lambert (2014, 2015, 2017), and Prelec (2004), the mechanisms described in the literature on the elicitation of subjective probabilities require the conditioning of the expert’s payoff on the event of interest.⁴ This paper deals with the elicitation of the probabilities of events in a subjective state space (that is, probabilities on subjective beliefs). Consequently, the true state (or event) is never observed and hence

⁴Unlike the mechanism in this paper and the protocol of Chambers and Lambert, Prelec’s (2004) elicitation scheme is not designed for the elicitation of second-order beliefs and is, therefore, less pertinent for the problem under study.

cannot be used to condition the payoff of the expert. Chambers and Lambert deals with the elicitation of the private information structure that governs the evolution of an expert beliefs over time. Because the information acquired by the expert is private, the elicitor cannot condition the expert's payoff on this information.

Despite the similarity in their objectives, the mechanisms presented here and the protocol of Chambers and Lambert (henceforth the CL protocol) are distinct. The differences include the underlying decision models, which affects the nature of the information they are designed to elicit, the elicitation schemes themselves, and the preference relations of the experts they admit. A more detailed discussion of the relation of this work to that of Chambers and Lambert is put off until after the reader review the proposed elicitation mechanisms of this paper (see section 4.3).

2 The Elicitation Mechanism

2.1 The analytical framework

I adopt the analytical framework of Anscombe and Aumann (1963). Let S be a set of *states*, one of which is the true state. Subsets of S are *events*. An event is said to obtain if the true state belongs to it. Let $\Delta(X)$ be the set of simple probability distributions (that is, distributions with finite supports) on an interval X in \mathbb{R} , and denote by $H := \{h : S \rightarrow \Delta(X)\}$ the set of Anscombe-Aumann *acts*. I identify the set of constant acts with $\Delta(X)$.

A *bet* on an event E , denoted $x_E y$, is a mapping from S to X that pays x dollars if E obtains and y dollars otherwise, where $x > y$. Denote by $\ell(p; x, y) = [x, p; y, (1 - p)]$, $p \in [0, 1]$, the act that pays off x dollars with probability p and y dollars with probability $(1 - p)$ in every state. Henceforth, I refer to such constant act as *lottery*. Let $B := \{x_E y \mid E \subset S, x, y \in \mathbb{R}, x > y\}$ be the set of bets and $L := \{\ell(p; x, y) \mid p \in [0, 1], x, y \in \mathbb{R}, x > y\}$ be the set of lotteries. Clearly, $B \subset H$ and $L \subset \Delta(X)$.

2.2 The subject

Consider a subject whose assessment of the probability of the event E is of interest. The subject's assessment may not be unique for two reasons.

First, his belief might not be complete in which case the assessment involves a non-singleton set of subjective priors. Second, even if his prior assessment is represented by a unique subjective probability, the subject might anticipate receiving a private signal that would make him revise his initial assessment. In this case the subject's posterior beliefs are represented by a signal-contingent assessments of the posterior probability of the event of interest. These are distinct possible situations depicted by different models of the subject's beliefs.

The notion of incomplete prior beliefs was first axiomatized by Bewley (2002). According to Bewley, the subject's prior beliefs are represented by a subset, $\Pi_0 \subseteq \Delta(S)$, where $\Delta(S)$ the set of probability distributions on S , and he strictly prefers and act f over another act g if and only if $\mathbb{E}_\pi^u f > \mathbb{E}_\pi^u g$, for all $\pi \in \Pi_0$, where $\mathbb{E}_\pi^u f = \sum_{s \in S} [\sum_{x \in X} u(x) f(s)(x)] \pi(s)$ and $u : X \rightarrow \mathbb{R}$ is a strictly monotonic increasing function. I assume throughout that u is twice continuously differentiable.

In the model of Karni and Safra (2016) Knightian uncertainty corresponds to the special case in which possible states of mind are represented by first-order beliefs (that is, elements of Π_0).⁵ Let $M \subseteq H$ be a menu of acts and suppose that given a menu to start with, there is an interim period in which the decision maker acquire knowledge of his first-order beliefs and chooses an act from the menu based on these beliefs. According to Karni and Safra (2016), for any pair of menus in H , $M \succ M'$ if and only if $\sum_{\pi \in \Pi_0} \mu(\pi) [\mathbb{E}_\pi^u f_M - \mathbb{E}_\pi^u f_{M'}] > 0$, where $f_M : \Pi_0 \rightarrow H$ is defined as follows: For each $\pi \in \Pi_0$, $f_M(\pi)$ is a selection of a unique element from the set $\{f \in M \mid \mathbb{E}_\pi^u f \geq \mathbb{E}_\pi^u f', \forall f' \in M\}$ of undominated acts in M , and μ is a second-order belief on $\Pi_0(E)$. In other words, given a menu M , f_M assigns to each first-order belief, $\pi \in \Pi_0$, an act in M such that, give this belief, the assigned act, $f_M(\pi)$, is at least as preferred as any other act in the menu. Consequently, $\mathbb{E}_\pi^u f_M$ is the value of the menu M . Thus, given $x, y \in X$, $x > y$, and $M = \{x_E y, \ell(p; x, y)\}$, $f_M(\pi) = x_E y$ if $\pi(E) \geq p$ and $f_M(\pi) = \ell(p; x, y)$ otherwise.

The notion of multi-posterior beliefs figures in the representations of Dillenberger et al. (2014) or Lu (2016). Let \succsim be the subject's preference

⁵This is the case of a subjective expected utility maximizing decision-maker whose preference relation on the set of acts is incomplete, but restricted to the subset of constant acts the preference relation is complete.

relation on the set of compact subsets of H (called *menus*). Suppose that the subject is a Bayesian decision maker whose preference relation satisfies the axioms of Dillenberger et al. (2014). In this case, $\Pi_1 \subseteq \Delta(S)$ can be interpreted as the set of posterior beliefs.⁶

For each event E , the subject's prior beliefs under Knightian uncertainty are represented by the probabilities $\Pi_0(E) = \{\pi(E) \mid \pi \in \Pi_0\}$, and if the subject is Bayesian then posterior beliefs are represented by $\Pi_1(E) = \{\pi(E) \mid \pi \in \Pi_1\}$. In the former case, the subject believes that the prior probability of an event E is a random variable, $\tilde{\pi}_0$, taking values in the interval $\Pi_0(E) = [\underline{\pi}_0(E), \bar{\pi}_0(E)]$ and that the likelihood that the true prior probability is π is described by a cumulative distribution function μ_0 on $[\underline{\pi}_0(E), \bar{\pi}_0(E)]$, interpreted as the subject's second-order belief.⁷ In the latter case, the subject believes that the posterior probability of an event E is a random variable, $\tilde{\pi}_1$, taking values in the interval $\Pi_1(E) = [\underline{\pi}_1(E), \bar{\pi}_1(E)]$ and that the likelihood that the true posterior probability is π is described by a cumulative distribution function, μ_1 on $[\underline{\pi}_1(E), \bar{\pi}_1(E)]$, interpreted as the subject's subjective information structure.⁸

2.3 The elicitation mechanism

The mechanism described below is designed to elicit the range of subjective prior or posterior probabilities of an event E as well as the subject's corresponding second-order belief or the subjective information structure. The proposed elicitation scheme requires the subject to report, at time $t = 0$, a function $\alpha : [0, 1] \rightarrow [0, 1]$. Following the report α , a random number, r , is drawn from a uniform distribution on $[0, 1]$. The subject is awarded the choice between the bet $(y + \theta - \theta^2 \alpha(r))^2_E (y - \theta^2 \alpha(r))^2$ and the lottery $\ell(r; y + \theta - \theta^2(1 - \alpha(r))^2, y - \theta^2(1 - \alpha(r))^2)$, where $\theta > 0$. In the last period, $t = 2$, after it becomes clear whether or not the event E obtained and the outcome of the lottery are revealed, all payments are made.

⁶Kanri and Safra (2016) provide a detailed discussion of connection between their work and those of Dillenberger et al. (2014) and Lu (2016). They argue that, despite sharing some features, their work is different conceptually, methodologically, and structurally.

⁷The existence of μ_0 , representing the decision maker's second-order belief on the set of priors, is implied by the model of Karni and Safra (2016).

⁸The existence of μ_1 , representing the subject's subjective information structure, is implied by the models of Dillenberger et al. (2014) and Lu (2016).

A crucial aspect of the mechanism is the flexibility it affords in delaying the choice. If the subject's preferences display Knightian uncertainty the value of this delay is in allowing the subject more time to consider his beliefs in the hope that in the interim some, possibly subconscious, stimuli would “trigger” a selection of an element from the set of priors and that the selected prior is used to determine the choice.⁹

If the subject is Bayesian, the delayed choice allows him to receive new information, or a signal, before making up his mind. In either case, the subject reports $\alpha(r) = 1$ if he is confident that, for that r , the lottery dominates the bet and intends to choose the lottery. The subject reports $\alpha(r) = 0$ if he is confident that the bet dominated the lottery and intends to choose the bet. The subject reports $\alpha(r) \in (0, 1)$ if he is unsure about his preference between the bet and the lottery and prefers postponing her choice to the interim period. Put differently, the subject's preference for flexibility is manifested in his willingness to pay a price $\theta^2 \alpha(r)^2$ in case he decides to chose the bet or $\theta^2(1 - \alpha(r))^2$ in case he decides to choose the lottery, in order to preserve his right to choose from the menu

$\{(y + \theta - \theta^2 \alpha(r))^2\}_E (y - \theta^2 \alpha(r)^2), \ell(r; y + \theta - \theta^2(1 - \alpha(r))^2, y - \theta^2(1 - \alpha(r))^2)\}$
in the interim period, $t = 1$.¹⁰

2.4 The elicitation mechanism analyzed

To fix the ideas, suppose that the subject's preference relation on H displays incomplete beliefs.¹¹ Also, to simplify the notations, without loss of generality I assume that $y = 0$, and, fixing the event of interest, E , I denote the

⁹Kreps (1979) articulates this presumption as follows: “In many problems of individual choice, the choice is made in more than one stage. At early stages, the individual makes decisions which will constrain the choices that are feasible later. In effect, these early choices amount to choice of a subset of items from which subsequent choice will be made. This paper concerns choice among such opportunity sets, where the individual has a “desire for flexibility” which is “irrational” if the individual knows what his subsequent preferences will be” (Kreps [1979], p. 565). The focus of the discussion here is the subject's subsequent beliefs rather than subsequent tastes.

¹⁰In Dillenberger et al. (2014) and Lu (2016) the delay is built in as the interim period in which the decision maker receives the information signal. The willingness to delay the choice is a manifestation of the value of the anticipated information.

¹¹The same analysis pertains to Bayesian subjects with private information structures.

corresponding set of priors by $\Pi = [\underline{\pi}, \bar{\pi}]$ instead of $\Pi_0(E) = [\underline{\pi}_0(E), \bar{\pi}_0(E)]$ and the subject's second-order belief by μ instead of μ_0 .

Under the mechanism, the subject's optimal choice of $\alpha(r)$, permits the recovery of his second-order belief, μ . Moreover, it is in the subjects' best interest to truthfully reveal the range of his beliefs by announcing $\alpha(r) = 1$ of all $r \in [\bar{\pi}, 1]$ and $\alpha(r) = 0$ for all $r \in [0, \underline{\pi}]$. Formally, given θ denote by $\alpha^*(r; \theta)$ the optimal choice under the mechanism, then the next theorem asserts that $\lim_{\theta \rightarrow 0} \alpha^*(\cdot; \theta) = \mu(\cdot)$.

Theorem 1: *In elicitation scheme the optimal report of $\alpha^*(r; \theta)$ satisfies $\lim_{\theta \rightarrow 0} \alpha^*(r; \theta) = \mu(r)$, for all $r \in [0, 1]$. Moreover, $\alpha^*(r; \theta) = 1$, for all $r \in [\bar{\pi}, 1]$, and $\alpha^*(r; \theta) = 0$, for all $r \in [0, \underline{\pi}]$.*

Proof. In the interim period $t = 1$ the subject learns his first-order belief, π . Given his announcement $\alpha(r)$, the subject chooses the bet if

$$ru(\theta - \theta^2(1 - \alpha(r))^2) + (1 - r)u(-\theta^2(1 - \alpha(r))^2) \leq \\ \pi u(\theta - \theta^2\alpha(r))^2 + (1 - \pi)u(-\theta^2\alpha(r)^2),$$

otherwise he chooses the lottery. Thus, the subject will choose the bet if

$$\pi \geq rA(\theta, \alpha(r)) + B(\theta, \alpha(r))$$

where

$$A(\theta, \alpha(r)) = \frac{u(\theta - \theta^2(1 - \alpha(r))^2) - u(-\theta^2(1 - \alpha(r))^2)}{u(\theta - \theta^2\alpha(r))^2 - u(-\theta^2\alpha(r)^2)}$$

and

$$B(\theta, \alpha(r)) = \frac{u(-\theta^2(1 - \alpha(r))^2) - u(-\theta^2\alpha(r)^2)}{u(\theta - \theta^2\alpha(r)^2) - u(-\theta^2\alpha(r)^2)}.$$

The subject is a subjective expected utility maximizer, hence, anticipating his choice in the interim period, he reports a function $\alpha : [0, 1] \rightarrow [0, 1]$ so as to maximize

$$\int_0^1 \{\mu(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))) [ru(\theta - \theta^2(1 - \alpha(r))^2) + (1 - r)u(-\theta^2(1 - \alpha(r))^2)]\} \\ (1)$$

$$+ \int_{rA(\theta, \alpha(r)) + B(\theta, \alpha(r))}^{\bar{r}} [\pi u(\theta - \theta^2 \alpha(r)^2) + (1 - \pi) u(-\theta^2 \alpha(r)^2)] d\mu(\pi) dr.$$

The maximal value is attained by maximizing the integrand pointwise. For every $r \in [0, 1]$, the first-order condition is¹²

$$\begin{aligned} & \mu(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))) (1 - \alpha(r)) \left[ru'(\theta - \theta^2(1 - \alpha(r))^2) + (1 - r) u'(-\theta^2(1 - \alpha(r))^2) \right] + \\ & \mu'(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))) (rA_\alpha(\theta, \alpha(r)) + B_\alpha(\theta, \alpha(r))) \times \\ & \left[ru(\theta - \theta^2(1 - \alpha(r))^2) + (1 - r) u(-\theta^2(1 - \alpha(r))^2) \right] \\ & = (1 - \mu(rA(\theta, \alpha(r)) + B(\theta, \alpha(r)))) \alpha(r) \times \\ & \left[\hat{\pi}(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))) [u'(\theta - \theta^2 \alpha^*(r)^2) - u'(-\theta^2 \alpha^*(r)^2)] + u'(-\theta^2 \alpha^*(r)^2) \right] \\ & + \mu'(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))) (rA_\alpha(\theta, \alpha(r)) + B_\alpha(\theta, \alpha(r))) \times \\ & \left[\pi(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))) [u(\theta - \theta^2 \alpha(r)^2) - u(-\theta^2 \alpha(r)^2)] + u(-\theta^2 \alpha(r)^2) \right] \end{aligned}$$

where

$$\hat{\pi}(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))) = \int_{rA(\theta, \alpha(r)) + B(\theta, \alpha(r))}^{\bar{r}} \frac{\pi}{(1 - \mu(rA(\theta, \alpha(r)) + B(\theta, \alpha(r))))} d\mu(\pi),$$

$$A_\alpha(\theta, \alpha(r)) = \partial A(\theta, \alpha(r)) / \partial \alpha \text{ and } B_\alpha(\theta, \alpha(r)) = \partial B(\theta, \alpha(r)) / \partial \alpha.$$

Since u' and u are continuous, for all $\alpha(r)$, $\lim_{\theta \rightarrow 0} A(\theta, \alpha(r)) = 1$, $\lim_{\theta \rightarrow 0} B(\theta, \alpha(r)) = 0$, $\lim_{\theta \rightarrow 0} A_\alpha(\theta, \alpha(r)) = \lim_{\theta \rightarrow 0} B_\alpha(\theta, \alpha(r)) = 0$. Moreover,

$$\begin{aligned} & \lim_{\theta \rightarrow 0} \left[ru'(\theta - \theta^2(1 - \alpha(r))^2) + (1 - r) u'(-\theta^2(1 - \alpha(r))^2) \right] = u'(0) \\ & = \lim_{\theta \rightarrow 0} \left[\hat{\pi}(r) u'(\theta - \theta^2(1 - \alpha(r))^2) + (1 - \hat{\pi}(r)) u'(-\theta^2 \alpha(r)^2) \right] \end{aligned}$$

¹²It is easy to verify that the second-order condition is satisfied, so the first-order condition is necessary and sufficient for a maximum.

and

$$\begin{aligned} \lim_{\theta \rightarrow 0} \left[ru \left(\theta - \theta^2(1 - \alpha(r))^2 \right) + (1 - r) u \left(-\theta^2(1 - \alpha(r))^2 \right) \right] &= u(0) = \\ &= \lim_{\theta \rightarrow 0} \left[\pi(r) u \left(\theta - \theta^2(1 - \alpha(r))^2 \right) + (1 - \pi(r)) u \left(-\theta^2 \alpha(r)^2 \right) \right]. \end{aligned}$$

Let $\alpha^*(r; \theta)$ denote the optimal choice of the subject, then these conditions imply that,

$$\mu(r) \left(1 - \lim_{\theta \rightarrow 0} \alpha^*(r; \theta) \right) = \lim_{\theta \rightarrow 0} \alpha^*(r; \theta) (1 - \mu(r)), \quad (2)$$

for all $r \in [0, 1]$. Thus, $\lim_{\theta \rightarrow 0} \alpha^*(r; \theta) = \mu(r)$, for all $r \in [0, 1]$.

Suppose that $r \in [\bar{\pi}, 1]$ then

$$ru \left(\theta - \theta^2(1 - \alpha(r))^2 \right) + (1 - r) u \left(-\theta^2(1 - \alpha(r))^2 \right) \geq \pi u \left(\theta - \theta^2 \alpha(r)^2 \right) + (1 - \pi) u \left(-\theta^2 \alpha(r)^2 \right) \quad (3)$$

for all $\pi \in \Pi = [\underline{\pi}, \bar{\pi}]$. By first order stochastic dominance

$$ru(\theta) + (1 - r) u(0) > ru \left(\theta - \theta^2(1 - \alpha(r))^2 \right) + (1 - r) u \left(-\theta^2(1 - \alpha(r))^2 \right), \quad (4)$$

for all $\alpha(r) < 1$. Hence, $\alpha^*(r; \theta) = 1$ for all $r \in [\bar{\pi}, 1]$. By a similar argument $\alpha^*(r; \theta) = 0$ for all $r \in [0, \underline{\pi}]$. \blacksquare

Let $\lim_{\theta \rightarrow 0} \alpha^*(r; \theta) := \alpha^*(r)$, then $\alpha^*(r) = \mu(r)$, for all $r \in (\underline{\pi}, \bar{\pi})$ implies that $\alpha(\pi) = \mu(\pi)$, for all $\pi \in \Pi(E)$. Moreover, as θ tends to zero, the subject chooses the lottery if $\pi < r$ and the bet if $r \leq \pi$. Thus, the subject chooses the lottery with probability $\mu(r)$ and with probability $1 - \mu(r)$ she chooses the bet. Hence, the subject exhibits random choice behavior.

Remark: Consider a more general mechanism that allows the subject a choice, in the interim period, between the bet $(x - \theta \alpha(r))^2_E (y - \theta \alpha(r))^2$ and the lottery $\ell(r; x - \theta(1 - \alpha(r))^2, (y - \theta(1 - \alpha(r))^2)$, where $\theta > 0$ and $x > y$. If the expert is risk neutral then the marginal utilities are constant and, consequently, Theorem 1 holds in the limit as θ tends to zero, for all values of x and y . By the argument in the proof of Theorem 1, the necessary and sufficient condition for optimality of the report $\alpha^*(r; \theta) = \lim_{\theta \rightarrow 0} \alpha^*(r; \theta)$ is $\mu(r) (1 - \alpha^*(r)) = \alpha^*(r) (1 - \mu(r))$. Hence, for every given $x > y$, $\mu(r) = \alpha^*(r)$, for all $r \in [0, 1]$. If the expert is risk averse, then the taking the limit as θ tends to zero would induce a biased report that depends on x and y . To avoid this bias, it is necessary to take the limit letting the $x - y \rightarrow 0$.

Alternatively, it is possible to use the lottery payoff scheme described in Roth and Malouf (1979) to obtain an assessment of μ , for all values of $x > y$.

Because the mechanism requires the specification of a function, $\alpha(\cdot)$, it is difficult, if not impossible, to implement in practice. However, the analysis of the mechanism-induced choice behavior suggests a practical method of approximating the solution to any desired degree.

3 Implementation

3.1 The discrete elicitation schemes described

Consider a discrete version of the elicitation mechanism depicted in the preceding section. Assume that the subject is a subjective expected utility maximizer whose assessment of the probability of an event, E , is a random variable, $\tilde{\pi}$, taking values in the interval $\Pi(E) = [\underline{\pi}, \bar{\pi}]$. Let μ denote the subject's cumulative distribution function on $[0, 1]$ representing his beliefs about the distribution of $\tilde{\pi}$.

Fix n and let $r_i^n = i/n$, $i = 0, 1, \dots, n$. At $t = 0$ the subject is asked to report a number $\alpha(z) \in [0, 1]$ for each $z \in \{r_i^n \mid i = 0, 1, \dots, n\}$. A random number r is then selected from a uniform distribution on $\{r_0^n, \dots, r_n^n\}$. In the interim period $t = 1$, the subject is allowed to choose between the bet,

$$(\theta - \theta^2 \alpha(r)^2)_E (-\theta^2 \alpha(r)^2)$$

and the lottery

$$\ell(r; \theta - \theta^2(1 - \alpha(r))^2), (-\theta^2(1 - \alpha(r))^2),$$

$\theta > 0$. All the payoffs are affected at $t = 2$.

3.2 The discrete elicitation scheme analyzed

Consider the partition $\mathcal{P}_n = \{[r_j^n, r_{j+1}^n) \mid j = 0, \dots, n-2\} \cup [(n-1)/n, 1]$ of the unit interval. Suppose that $\underline{\pi} \in [r_i^n, r_{i+1}^n)$ and $\bar{\pi} \in [r_k^n, r_{k+1}^n)$, for some $0 \leq i \leq k \leq n-1$. For sufficiently small θ the subject's optimal choice of $\alpha(z)$, $z \in \{r_i^n \mid i = 0, 1, \dots, n\}$, denoted $\alpha^*(z; \theta)$, permits the recovery of $\mu(r_i^n)$, for all $r_i^n \in [0, 1]$. Formally,

Theorem 2: *In elicitation scheme the report of $\alpha^*(r; \theta)$ satisfies $\lim_{\theta \rightarrow 0} \alpha^*(z; \theta) = \mu(z)$, for all $z \in \{r_i^n \mid i = 0, 1, \dots, n\}$. Moreover, $\alpha^*(r_j^n; \theta) = 1$ for all $r_j^n \in [\bar{\pi}, 1]$, and $\alpha^*(r_j^n; \theta) = 0$ for all $r_j^n \in [0, \underline{\pi}]$.*

Proof. To simplify the notations, define $\bar{U}(\tau; \theta) = \tau u(\theta) + (1 - \tau)u(0)$, and $\bar{U}^\Delta(\tau; \theta) = \tau u'(\theta) + (1 - \tau)u'(0)$, $\tau \in [0, 1]$. By the argument in the proof of Theorem 1, in the interim period $t = 1$ the subject whose first-order belief is π chooses the bet if

$$\pi \geq r_j^n A(\theta, \alpha(r_j^n)) + B(\theta, \alpha(r_j^n)).$$

The subject chooses $\alpha(r_j^n)$ so as to maximize

$$\begin{aligned} & \mu(r_j^n A(\theta, \alpha(r_j^n)) + B(\theta, \alpha(r_j^n))) \times \left[r_j^n u(\theta - \theta^2(1 - \alpha(r_j^n))^2) + (1 - r_j^n) u(-\theta^2(1 - \alpha(r_j^n))^2) \right] \\ & + \int_{r_j^n A(\theta, \alpha(r_j^n)) + B(\theta, \alpha(r_j^n))}^{\bar{\pi}} \left[\pi u(\theta - \theta^2 \alpha(r_j^n)^2) + (1 - \pi) u(-\theta^2 \alpha(r_j^n)^2) \right] d\mu(\pi). \end{aligned} \quad (5)$$

By the same argument as in the proof of Theorem 1, the first-order condition evaluated at the limit as $\theta \rightarrow 0$ is

$$\mu(r_j^n) (1 - \alpha(r_j^n; \theta)) \bar{U}^\Delta(r_j^n; \theta) - \alpha(r_j^n; \theta) (1 - \mu(r_j^n)) \int_{r_j^n}^{\bar{\pi}} \bar{U}^\Delta(\pi; \theta) d\mu(\pi \mid \pi \geq r_j^n) = 0. \quad (6)$$

Let $Ex[\bar{U}^\Delta(\pi; \theta) \mid \pi \geq r_j^n] = \int_{\underline{\pi}}^{r_j^n} \bar{U}^\Delta(\pi; \theta) d\mu(\pi \mid \pi \geq r_j^n)$. Then, the first order condition implies that

$$\frac{\alpha(r_j^n; \theta)}{(1 - \alpha(r_j^n; \theta))} = \frac{\mu(r_j^n) \bar{U}^\Delta(r_j^n; \theta)}{(1 - \mu(r_j^n)) Ex[\bar{U}^\Delta(\pi; \theta) \mid \pi \geq r_j^n]}. \quad (7)$$

In the limit $\theta \rightarrow 0$, $\bar{U}^\Delta(\tau; \theta) = u'(0)$ is independent of τ . Hence, denoting the optimal solution by $\alpha^*(r_j^n; \theta)$, (7) implies $\lim_{\theta \rightarrow 0} \alpha^*(r_j^n; \theta) = \mu(r_j^n)$.

Moreover, by the same argument as in the proof of Theorem 1, for $r_j^n \geq \bar{\pi}$, $\mu(r_j^n) = 1$, hence $\alpha^*(r_j^n; \theta) = 1$ and for $r_j^n < \underline{\pi}$, $\mu(r_j^n) = 0$, hence, $\alpha^*(r_j^n; \theta) = 0$, for all $\theta > 0$. ■

In the limit, as n tends to infinity, these estimates coincide with the true values and α^* converge to μ .

Theorem 3: Let $\bar{\mu}(\pi) := \lim_{\theta \rightarrow 0} \alpha^*(r_j^n; \theta)$, for all $j = 0, \dots, 1$ and $\pi \in (r_{j-1}, r_j]$, then $\lim_{n \rightarrow \infty} \bar{\mu}(\pi) = \mu(\pi)$, for all $\pi \in [0, 1]$.

Proof. For each $\pi \in [0, 1]$ and $n = 1, 2, \dots$, let $[r_n^l, r_n^h)$ be the cell of the partition $\mathcal{P}_n = \{[r_j^n, r_{j+1}^n)\}_{j=0}^n$ such that $\pi \in [r_n^l, r_n^h)$. Consider the sequence $\{r_n^h\}$. Since $\pi < r_n^h$, for all $n = 1, \dots$, $\inf\{r_n^h \mid \pi < r_n^h, n = 1, 2, \dots\}$ exists and is equal to π . But μ is right continuous, hence $\lim_{n \rightarrow \infty} \bar{\mu}(r_n^h) = \mu(\pi)$. ■

Under the limit processes described above, the mechanism yields the range of the subject's beliefs about the likelihood of any event in the state space and of his introspective assessment of the likelihoods of his beliefs. In practice, an appropriate choice of the parameters n, θ , and the payoff y yields estimated values that approximate the true values of the subject's beliefs to any degree desired.

4 Extension, Variation and Related Literature

4.1 Elicitation of distributions of real-valued random variables

The procedure described above is designed to elicit a subject's beliefs about the likely realization of events. By extension, this method can also be used to elicit an entire distribution of a random variable. Consider the case in which the variable of interest is a subject's beliefs regarding the distribution of a real-valued random variable. Suppose that the subject entertains multiple such beliefs. To elicit the subject beliefs, the procedure described below combines the elicitation mechanism described in Section 2 with an elicitation procedure due to Qu (2012).

Consider the set, \mathcal{H} , whose elements are cumulative distribution functions (CDF) on \mathbb{R} .¹³ Suppose that a subject's beliefs regarding the distribution

¹³Assume that \mathcal{H} is endowed with the topology of weak convergence, and denote by Σ

of a real-valued random variable of interest are represented by $\mathcal{F} \subset \mathcal{H}$ and that his introspective beliefs about the likely realizations of elements of \mathcal{F} are depicted by a probability measure, μ on \mathcal{F} .¹⁴

The subject is asked to report a function $\alpha : \mathbb{R} \times [0, 1] \rightarrow [0, 1]$. The mechanism then draws a number k from a distribution with full support on the real line and a random number r from a uniform distribution on the unit interval. For each $k \in \mathbb{R}$, define $E_k = (-\infty, k]$.

For each possible realization of k , the subject is allowed to choose between the bet

$$(\theta - \theta\alpha(k, r)^2)_{E_k} (-\theta^2\alpha(k, r)^2)$$

and the lottery

$$\ell(r; \theta - \theta(1 - \alpha(k, r))^2, -\theta^2(1 - \alpha(k, r))^2).$$

The next theorem asserts that, for every value in the support of the random variable whose CDF is of interest, truthful reporting of the beliefs about the range and the likelihoods of values of the CDF is the unique best response.

Theorem 4: *For each $(k, r) \in \mathbb{R} \times [0, 1]$, in the limit as θ tend to zero, the optimal report satisfies $\alpha(k, r) = \mu\{F \in \mathcal{F} \mid F(k) \leq r\}$. Moreover, for each $k \in \mathbb{R}$, let $\bar{G}(k) = \sup_{F \in \mathcal{F}}\{F(k)\}$ and $\underline{G}(k) = \inf_{F \in \mathcal{F}}\{F(k)\}$, then $\alpha^*(k, r; \theta) = 1$ for all $r \in [\bar{G}(k), 1]$ and $\alpha^*(k, r; \theta) = 0$ for all $r \in [0, \underline{G}(k)]$.*

Proof. Applying the proof of Theorem 1 for each k the subject's optimal report, $\alpha^*(k, r; \theta)$ satisfies $\lim_{\theta \rightarrow 0} \alpha^*(k, r; \theta) = \mu\{F \in \mathcal{F} \mid F(k) \leq r\}$. Moreover, by the same argument as in the proof of Theorem 1, $\alpha^*(k, r; \theta) = 1$ for all $r \in [\bar{G}(k), 1]$ and $\alpha^*(k, r; \theta) = 0$ for all $r \in [0, \underline{G}(k)]$. ■

The same procedure can be employed to elicit a subject's beliefs about the distribution of a vector valued random variable.

4.2 Direct elicitation of the range of probabilities of an event

A direct elicitation of the range of priors is possible using an incentive scheme that combines a mechanism described in Grether (1981) and Karni (2009)

the Borel sigma algebra on \mathcal{H} .

¹⁴I assume that \mathcal{F} together with the trace of Σ on \mathcal{F} is the relevant measurable space.

for the elicitation of unique subjective prior and a modified proper scoring rule applied over a restricted set of measures.

Fix an event E and let $\Pi = [\underline{\pi}, \bar{\pi}]$ denote the rang of subjective probabilities representing the subjects beliefs about the likelihood of E . The modified scheme requires the subject to report, at time $t = 0$, two numbers, $\underline{r}, \bar{r} \in [0, 1]$ (intended to demarcate the range of his subjective prior or posterior probability assessments of the event E) and, for each $r \in (\underline{r}, \bar{r})$, to report a number $\alpha(r) \in (0, 1)$. A random number, r , is drawn from a uniform distribution on $[0, 1]$. In the interim period, $t = 1$, the subject is awarded the bet $x_E y$ if $r \leq \underline{r}$ and the lottery $\ell(r; x, y)$ if $r \geq \bar{r}$, where $x > y$. If $r \in (\underline{r}, \bar{r})$, then the subject is allowed to choose, at $t = 1$, between the bet $(x - \theta\alpha(r)^2)_E (y - \theta\alpha(r)^2)$ and the lottery $\ell(r; x - \theta(1 - \alpha(r))^2, y - \theta(1 - \alpha(r))^2)$, where $\theta > 0$. In the last period, $t = 2$, whether or not the event E obtained and the outcome of the lottery are revealed, and all payments are made.

Theorem 5: *In the modified scheme, in the limits as $\theta \rightarrow 0$, the subject's unique dominant strategy is to report $\underline{r} = \underline{\pi}$ and $\bar{r} = \bar{\pi}$.*

Proof. Fix $x > y$ and $\theta > 0$. Suppose that the subject reports $\bar{r} > \bar{\pi}$. If $r \leq \bar{\pi}$ or $r \geq \bar{r}$ the subject's payoff is the same regardless of whether he reports \bar{r} or $\bar{\pi}$. If $r \in (\bar{\pi}, \bar{r})$, the subject's payoff is a choice between the bet $(x - \theta\alpha(r)^2)_E (y - \theta\alpha(r)^2)$ and the lottery $\ell(r; x - \theta(1 - \alpha(r))^2, y - \theta(1 - \alpha(r))^2)$; had he reported $\bar{\pi}$ instead of \bar{r} his payoff would have been $\ell(r; x, y)$. By first-order stochastic dominance,

$$\ell(r; x, y) \succcurlyeq \ell(r; x - \theta(1 - \alpha(r))^2, y - \theta(1 - \alpha(r))^2),$$

for all $\alpha(r) \in [0, 1]$, with strict preference except when $\alpha = 1$. Since $r > \bar{\pi}$,

$$\ell(r; x, y) \succ x_E y \succ (x - \theta\alpha(r)^2)_E (y - \theta\alpha(r)^2),$$

for all $\alpha(r) \in (0, 1]$. Thus the subject is worse off reporting \bar{r} instead of $\bar{\pi}$.

Suppose that the subject reports $\underline{r} < \underline{\pi}$. If $r \leq \underline{r}$ or $r \geq \underline{\pi}$ the subject's payoff is the same regardless of whether he reports \underline{r} or $\underline{\pi}$. If $r \in (\underline{r}, \underline{\pi})$, the subject's payoff is a choice between $(x - \theta\alpha(r)^2)_E (y - \theta\alpha(r)^2)$ and the lottery $\ell(r; x - \theta(1 - \alpha(r))^2, y - \theta(1 - \alpha(r))^2)$; had he reported $\underline{\pi}$ instead of \underline{r} his payoff would have been $x_E y$. By stochastic dominance, $x_E y \succcurlyeq (x - \theta\alpha(r)^2)_E (y - \theta\alpha(r)^2)$, for all $\alpha(r) \in [0, 1]$, with strict preference except when $\alpha = 0$. Since $r < \underline{\pi}$,

$$x_E y \succ \ell(r; x, y) \succ \ell(r; x - \theta(1 - \alpha(r))^2, y - \theta(1 - \alpha(r))^2),$$

for all $\alpha(r) \in [0, 1]$. Thus the subject is worse off reporting \underline{r} instead of $\underline{\pi}$.

Suppose that the subject reports $\bar{r} \in (\underline{\pi}, \bar{\pi})$. If $r \in [\bar{r}, \bar{\pi}]$, the subject's payoff is $\ell(r; x, y)$, whereas had he reported $\bar{\pi}$ he would have the opportunity to choose between the bet $(x - \theta\alpha(r)^2)_E (y - \theta\alpha(r)^2)$ and the lottery $\ell(r; x - \theta(1 - \alpha(r))^2, y - \theta(1 - \alpha(r))^2)$. Thus, in the limit $\theta \rightarrow 0$, the subject's subjective expected utility is:

$$\mu(r) [ru(x) + (1 - r)u(y)] + \int_r^{\bar{\pi}} [\pi u(x) + (1 - \pi)u(y)] d\mu(\pi).$$

It is easy to verify that this expression exceeds the expected utility of the lottery $\ell(r; x, y)$, $ru(x) + (1 - r)u(y)$. Thus, reporting $\bar{r} < \bar{\pi}$ is dominated by reporting $\bar{\pi}$. By similar argument, $\underline{r} \not\geq \underline{\pi}$. ■

4.3 Related literature

Belief elicitation procedures have been the subject of inquiry for more than half a century, beginning with the work of Brier (1950) and Good (1952) followed by Savage (1971) Kadane and Winkler (1988) and others.¹⁵ Except for Prelec (2004), Chambers and Lambert (2014, 2015, 2017), and the mechanisms described in this paper, the elicitation schemes in the literature condition the subject's (expert's) reward on the event of interest. This requires that the event of interest be publicly observable. In this sense, the game mechanism of Prelec, the protocol of Chambers and Lambert, and the elicitation scheme described in this paper are unconventional.

Both the CL protocol and the elicitation scheme of this paper are designed to elicit the expert's subjective "belief over beliefs," interpreted, respectively, as subjective information structure and second-order introspective beliefs. Both schemes presume that the decision maker refines his belief over beliefs over time and, in both instances, the event of interest is a set "first-order beliefs," or subjective probabilities, that, by definition, is not observable and, consequently, cannot be used to condition the expert's reward. Despite these similarities the CL protocol and the elicitation scheme of this paper employ distinct procedures. According to a special CL protocol (i.e., the simple example) described in Chambers and Lambert (2017) the elicitor selects two

¹⁵For a recent review, see Chambers and Lambert (2017).

numbers, α and β , from uniform distributions on $[0, 1]$ and $[-1, 1]$, respectively. These numbers represent, respectively, the price of an option, dubbed α -option, to short-sell a security in the interim period ($t = 1$) at the price α , and the price of the α -option as of the initial period ($t = 0$). The mechanism requires the expert to announce a prior \hat{F} and, later, a posterior \hat{p} .¹⁶ The elicitor, acting on the expert's behalf, makes the optimal decision to buy or not to buy the option, then to exercise or not to exercise the right to sell the security. The elicitor must never inform the expert of which decision she has made until all uncertainty about the random variable is resolved. The expert is awarded the payoff resulting from the elicitor's choices.

Finally, it is noteworthy that, because the results of this work focus on approximate truthfulness (with strict truthfulness only in the limit) the conclusion of Chambers and Lambert (2017) that no protocol in the general class they define as "stage separate protocols" can be strategy proof does not contradict the result of this paper.

¹⁶Note that \hat{F} is a second-order belief and is the same as μ in this paper. Similarly, \hat{p} is a first-order belief and is the same as π in this paper. It is noteworthy, however, that the mechanism described in this paper does not require that the subject report his posterior belief.

References

- [1] Anscombe, Francis J. and Aumann, Robert J. (1963) “A Definition of Subjective Probability,” *Annals of Mathematical Statistics* 43, 199–205.
- [2] Bewley, Truman F. (2002) “Knightian Decision Theory: Part I,” *Decision in Economics and Finance*, 25, 79–110.
- [3] Brier, Glenn W. (1950) “Verification of Forecasts Expressed in Terms of Probabilities,” *Monthly Weather Review*, 78, 1–3.
- [4] Chambers, Christopher P. and Lambert, Nicholas S. (2014) “Dynamically Eliciting Unobservable Information,” extended abstract, *Proceedings of the 15th ACM Conference on Electronic Commerce*, 987–988.
- [5] Chambers, Christopher P. and Lambert, Nicholas S. (2015) “Introduction to Dynamic Belief Elicitation,” *ACM SIGecom Exchanges* 14, 80–83.
- [6] Chambers, Christopher P. and Lambert, Nicholas S. (2017) “Dynamic Belief Elicitation,” unpublished manuscript.
- [7] Dekel, Eddie, Lipman, Barton L. and Rustichini, Aldo (2001) “Representing Preferences with a Unique Subjective State Space,” *Econometrica* 69, 891–934.
- [8] Dillenberger, David, Lleras, Juan S., Sadowski, Philipp and Takeoka, Norio (2014) “A Theory of Subjective Learning,” *Journal of Economic Theory* 153, 287–312.
- [9] Galaabaatar, Tsogbadral and Karni, Edi (2013) “Subjective Expected Utility with Incomplete Preferences,” *Econometrica* 81, 255–284.
- [10] Giraud, Raphaël (2014) “Second Order Beliefs Models of Choice under Imprecise Risk,” *Theoretical Economics* 9, 779–816.
- [11] Good, Irving J. (1952) “Rational Decisions,” *Journal of the Royal Statistical Society, Series B (Methodological)*, 14, 107–114.
- [12] Grether, David M. (1981) “Financial Incentive Effects and Individual Decision Making,” California Institute of Technology, Social Science Working Paper 401.

- [13] Gul, Faruk and Pesendorfer, Wolfgang (2006) "Random Expected Utility," *Econometrica* 74, 121-146.
- [14] Kadane, Joseph B., and Winkler Robert L. (1988) "Separating Probability Elicitation from Utility," *Journal of the American Statistical Association*, 83, 357-363.
- [15] Karni, Edi (2009) "A Mechanism Design for Probability Elicitation," *Econometrica* 77, 603-606.
- [16] Karni, Edi and Safra, Zvi (2016) "A Theory of Stochastic Choice under Uncertainty," *Journal of Mathematical Economics*, (forthcoming)
- [17] Kreps, David M. (1979) "A Representation Theorem for 'Preference for Flexibility'," *Econometrica* 47, 565-576.
- [18] Lu, Jay (2016) "Random Choice and Private Information," *Econometrica* 84, 1983-2027.
- [19] Minardi, Stefania and Savochkin, Andrei (2015) "Preferences with Grades of Indecisiveness," *Journal of Economic Theory* 155, 300-331.
- [20] Nascimento, Leandro and Gil Riella (2013) "Second Order Ambiguous Beliefs," *Economic Theory* 52, 1005 -1037.
- [21] Prelec, Dražen (2004) "A Bayesian Truth Serum for Subjective Data," *Science* 306, 462-466.
- [22] Qu, Xiangyu (2012) "A Mechanism for Eliciting a Probability Distribution," *Economic Letters* 115, 399-400.
- [23] Roth, Alvin E., and Malouf, Michael W. K. (1979) "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review* 86, 574-594.
- [24] Savage, Leonard J. (1971) "Elicitation of Personal Probabilities and Expectations," *Journal of American Statistical Association*, 66, 783-801.
- [25] Seidenfeld, Teddy, Schervish, Mark J. and Kadane, Joseph B. (1995) "A Representation of Partially Ordered Preferences," *Annals of Statistics* 23, 2168-2217.

- [26] Seo, Kyoungwon (2009) “Ambiguity and Second-Order Belief,” *Econometrica* 77, 1575-1605.