A Mechanism for the Elicitation of Second-Order Beliefs and Subjective Information Structures

Edi Karni*
Johns Hopkins University and University of Warwick

November 2, 2016

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 beliefs and (b) Bayesian decision makers' range of posteriors and their subjective information structures.

Keywords: Second-order beliefs; 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, and Nicholas Lambert 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 (that is, unique) 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 improve 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 the potential information signals he anticipate receiving. In this case, the expert's assessment takes the form of a range of the posterior probabilities.

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

The situation in which an expert – or, more generally, a decision maker – entertains a set of priors arises when the expert's preference relation over non-constant, state-contingent consequences (or *acts*), is incomplete. Bewley's (2002) Knightian uncertainty model characterizes this situation.[2] 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' choice behavior is determined by preference relations representing their

[1] In robust Bayesian statistics, the elicitation of the set of priors is analogous to the elicitation of a single prior in Bayesian statistics (see Seidenfld, 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]).

[2] See also Galaabaatar and Karni (2013).

moods, beliefs, or states of mind. Decision makers entertain introspective belief about their likely states of mind, which is manifested in random choice behavior. In the special case in which states of mind correspond to a decision maker's beliefs, Karni and Safra characterize the set of priors and the decision maker's second-order 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 received and acts upon. Lu shows that, assuming that 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 work of Karni and Safra (2016), Dillenberger et al. (2014) and Lu (2016), share in common representations of the decision maker's second-order beliefs or private information structure. Moreover, in all of these models decision makers exhibit random choice behavior. According to Karni and Safra, a decision maker acts as if responding to a random selection of a prior from his subjective set of priors and the selected prior governs his choice. According to Dillenberger et al. and Lu a decision maker receives, at some interim stage, a private signal on the basis of which he updates his beliefs. The posterior belief governs his eventual choice. If the model Karni and Safra is applied to a Bayesian decision maker, it can be interpreted as giving rise to a set of posteriors, and the second-order belief can be viewed as a subjective information structure.

Suppose that the uniformed 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 beliefs regarding the likelihoods of the different priors or the posteriors in the corresponding sets. If the expert displays incomplete preferences, 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, a second 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]).

In this paper I propose two direct revelation mechanisms that require 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 first mechanism combines two incentive schemes that were discussed in the literature on the elicitation of a unique prior (a) a procedure described in Grether (1981) and Karni (2009) and (b) a modified proper scoring rule applied over a restricted set of measures.[3] The second mechanism is a modified proper scoring rule.

With the exception of Chambers and Lambert (2016), 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. 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 cannot be used to condition the payoff of the expert. The work of Chambers and Lambert (2016) 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 payoff on this information.

Despite the similarity in their objectives, the mechanism presented here and the protocol described and analyzed in Chambers and Lambert are quite different. 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.

[3] Savage (1971) and Kadane and Winkler (1988) proved a discussion of proper scoring rules.

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 4.2).

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 $H := \{h : S \to \Delta(X)\}$ the set of Anscombe-Auamnn *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 $\mathbb{R}$ that pays x dollars if E obtains and y dollars otherwise, where $x > y$. Denote by $\ell(p; x, y) = [x, p; y, (1-p)]$, $p \in [0,1]$, a lottery that pays off x dollars with probability p and y dollars with probability $(1-p)$ in every state. 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. Note that $B \subset H$ and $L \subset \Delta(X)$.

2.2 The subject

Consider a subject whose assessment of the probability of the event E is of interest. Let $\succcurlyeq$ denote the subject's preference relation on H. Assume that $\succcurlyeq$ has expected utility representation on $\Delta(X)$, the set of constant acts, and Knightian uncertainty representation on H, the set of acts.[4] Denote by $\Delta(S)$ the set of probability distributions on S. Then the subject's prior beliefs are represented by a subset, $\Pi_0 \subseteq \Delta(S)$, of prior probability distributions on S.

Let $\succsim$ be the subject's preference 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).

[4]This assumption is consistent with Bewley's (2002) model, which assumes that the preference relation is complete on the subset of constant acts.

In this case, $\Pi_1 \subseteq \Delta(S)$ can be interpreted as the set of posterior beliefs that figure in the representations of Dillenberger et al. (2014) or Lu (2016).

For each event E, the subject's prior beliefs under Knightian uncertainty are represented by the probabilities $\Pi_0(E) = \{\pi(E) \mid \pi \in \Pi_0\}$. If the subject is Bayesian and his information structure is private, 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, $\widetilde{\pi}_0$, taking values in the interval $\Pi_0(E) = [\underline{\pi}_0(E), \bar{\pi}_0(E)]$ and that the likelihood that the true 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 beliefs.[5] In the latter case, the subject believes that the posterior probability of an event E is a random variable, $\widetilde{\pi}_1$, taking values in the interval $\Pi_1(E) = [\underline{\pi}_1(E), \bar{\pi}_1(E)]$ and that the likelihood that the true 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.[6]

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 beliefs or the subjective information structure. The 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 $\alpha(r) \in (0, 1)$ (intended to capture his second-order beliefs or his subjective information structure). A random number, r, is drawn from a uniform distribution on $[0, 1]$.[7] 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 $r \geq \bar{r}$. If $r \in (\underline{r}, \bar{r})$, then the subject is allowed to choose, at $t = 1$, between the bet $\left(x - \theta\alpha(r)^2\right)_E \left(y - \theta\alpha(r)^2\right)$ 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

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

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

[7]The uniformity of the distrbution of r is not necessary. It is used here to simplify the exposition.

the lottery are revealed, and all payments are made.

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 determine his beliefs before making his choice.[8] If the subject is Bayesian, it allows him to receive an informational signal before making up his mind. In either case, the subject's preference for flexibility is manifested in his willingness to pay a price in order to preserve his right to choose from the menu

$$\left\{\left(x-\theta\alpha\left(r\right)^{2}\right)_{E}\left(y-\theta\alpha\left(r\right)^{2}\right),\ell\left(r;x-\theta(1-\alpha\left(r\right))^{2},y-\theta(1-\alpha\left(r\right))^{2}\right)\right\}$$

in the interim period, $t=1$, rather than between the bet $x_{E}y$ and the lottery $\ell\left(r;x,y\right)$ at time $t=0$.[9]

2.4 The elicitation mechanism analyzed

To fix the ideas, suppose that the subject's preferences are incomplete and the incompleteness is due to his beliefs.[10] Also, to simplify the notation, fixing the event of interest, E, I denote the corresponding set of priors by $\Pi=[\underline{\pi},\bar{\pi}]$ instead of $\Pi_{0}\left(E\right)=\left[\underline{\pi}_{0}\left(E\right),\bar{\pi}_{0}\left(E\right)\right]$ and denote the subject's second-order belief by μ instead of μ_{0}.

Under the mechanism, the subject's unique dominant strategy is to truthfully reveal the range of his beliefs, and that if the parameters θ, x and y are chosen appropriately, the subject's choice of $\alpha\left(r\right)$ permits the recovery of his second-order belief, μ. Formally,

[8] Kreps (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.

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

[10] The same analysis pertains to Bayesian subjects with private information structures.

Theorem 1: *In the mechanism, for all* $\alpha : (\underline{r}, \bar{r}) \to [0,1]$; *(a) In the limits as* $\theta \to 0$, *the subject's unique dominant strategy is to report* $\underline{r} = \underline{\pi}$ *and* $\bar{r} = \bar{\pi}$. *(b) In the limit as* $\theta \to 0$, *and* $(x-y) \to 0$, *the choice of* $\alpha(\cdot)$ *satisfies* $\alpha(r) = \mu(r)$, *for all* $r \in (\underline{r}, \bar{r})$.

Proof. (a) 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 $\left(x - \theta\alpha(r)^2\right)_E \left(y - \theta\alpha(r)^2\right)$ 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, $\ell(r; x, y) \succ \ell(r; x - \theta(1-\alpha(r))^2, y - \theta(1-\alpha(r))^2)$, for all $\alpha(r) \in [0,1)$, and, since $r > \bar{\pi}$, $\ell(r; x, y) \succ x_E y \succ \left(x - \theta\alpha(r)^2\right)_E \left(y - \theta\alpha(r)^2\right)$. 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 $\left(x - \theta\alpha(r)^2\right)_E \left(y - \theta\alpha(r)^2\right)$ 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, $x_E y \succ \left(x - \theta\alpha(r)^2\right)_E \left(y - \theta\alpha(r)^2\right)$ and, since $r < \underline{\pi}$, $x_E y \succ \ell(r; x, y) \succ \left(x - \theta\alpha(r)^2\right)_E \left(y - \theta\alpha(r)^2\right)$. 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 $\left(x - \theta\alpha(r)^2\right)_E \left(y - \theta\alpha(r)^2\right)$ and the lottery $\ell(r; x - \theta(1-\alpha(r))^2, y - \theta(1-\alpha(r))^2)$. Clearly, in the limit $\theta \to 0$, the subject would choose the bet if $\pi \in [r, \bar{\pi}]$ and the lottery otherwise. Thus, his subjective expected utility is:

$$\mu(r)\left[ru\left(x - \theta(1-\alpha(r))^2\right) + (1-r)\,u\left(y - \theta(1-\alpha(r))^2\right)\right] +$$

$$\int_r^{\bar{\pi}} \left[\pi u\left(x - \theta\alpha(r)^2\right) + (1-\pi)\,u\left(y - \theta\alpha(r)^2\right)\right] d\mu(\pi).$$

It is easy to verify that, in the limit as $\theta \to 0$, 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, $\underline{r} \not> \underline{\pi}$.

(b) The subject is a subjective expected utility maximizer. Hence he chooses a function $\alpha : (\underline{r}, \bar{r}) \to (0,1)$ so as to maximize

$$\int_{\underline{r}}^{r} \{\mu(r)\left[ru\left(x-\theta(1-\alpha(r))^2\right)+(1-r)\,u\left(y-\theta(1-\alpha(r))^2\right)\right]+ \tag{1}$$

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

The maximal value is attained by maximizing the integrand pointwise. For every $r \in (\underline{r}, \bar{r})$, the first-order condition is[11]

$$\begin{aligned} &\mu(r)(1-\alpha(r))\left[ru'\left(x-\theta(1-\alpha(r))^2\right)+(1-r)\,u'\left(y-\theta(1-\alpha(r))^2\right)\right] = \\ &\alpha(r)(1-\mu(r))\left[u'(x-\theta\alpha(r)^2)\bar{\pi}(r)+u'(y-\theta\alpha(r)^2(1-\bar{\pi}(r))\right], \end{aligned} \tag{2}$$

where $\bar{\pi}(r) = \int_{r}^{\bar{r}} \frac{\pi}{(1-\mu(r))} d\mu(\pi)$. Evaluated at the limits as $\theta \to 0$, and $(x-y) \to 0$ these conditions become

$$\mu(r)(1-\alpha(r)) = \alpha(r)(1-\mu(r)), \tag{3}$$

for all $r \in (\underline{r}, \bar{r})$.[12] Thus in the limit, $\alpha(r) = \mu(r)$, for all $r \in (\underline{r}, \bar{r})$. ■

Since $\underline{r} = \underline{\pi}$ and $\bar{r} = \bar{\pi}$, $\alpha(r) = \mu(r)$, for all $r \in (\underline{r}, \bar{r})$ 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 $1-\mu(r)$ he chooses the bet. Hence, his choice is random.

Since 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.

---

[11] I assume throughtout that the utility function is twice continuously differentiable. It is easy to verify that the second-order condition is satisfied, so the first-order condition is necessary and sufficient for a maximum.

[12] It is easy to verify that the order of in which the limits are taken does not matter.

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, $\widetilde{\pi}$, taking values in the interval $\Pi(E) = [\underline{\pi}, \bar{\pi}]$. Let μ denote the subject's cumulative distribution function on $[\underline{\pi}, \bar{\pi}]$ representing his beliefs about the distribution of $\widetilde{\pi}$.

Fix n and let $r_i^n = i/n$, $i = 0, 1, ..., n$. The subject is asked to report two numbers, $\underline{r}^n, \bar{r}^n \in \{r_0^n, ..., r_n^n\}$, not necessarily distinct, and, for each $z \in \{r_i^n \mid \underline{r}^n < r_i^n < \bar{r}^n\}$, to indicate a number $\alpha(z) \in (0,1)$. A random number r is selected from a uniform distribution on $\{r_0^n, ..., r_n^n\}$. If $r \leq \underline{r}^n$, then the subject is awarded the bet $x_E y$; if $r \geq \bar{r}^n$, then the subject is awarded the lottery $\ell(r; x, y)$; and if $\underline{r}^n < r < \bar{r}^n$, then the subject is allowed to choose between the bet, $(x - \theta\alpha(r)^2)_E(y - \theta\alpha(r)^2)$, and the lottery $\ell(r; x - \theta(1 - \alpha(r))^2, (y - \theta(1 - \alpha(r))^2)$, $\theta > 0$. All the payoffs are affected at the same time.

3.2 The discrete elicitation scheme analyzed

Consider the partition of the unit interval $\{[r_j^n, r_{j+1}^n) \mid j = 0, ..., n-2\} \cup [(n-1)/n, 1]$. Without loss of generality, let $\underline{\pi} \in [r_i^n, r_{i+1}^n)$ 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 unique dominant strategy is to report $\underline{r}^n = r_i^n$ and $\bar{r}^n = r_{k+1}^n$; and if the parameters θ, x and y, are chosen appropriately, the subject's choice of $\alpha(r_i^n)$ permits the recovery of $\mu(r_i^n)$, for all $r_i^n \in (\underline{r}^n, \bar{r}^n)$. Formally;

Theorem 2: *In mechanism I, (a) If $\underline{\pi} \in [r_i^n, r_{i+1}^n)$ and $\bar{\pi} \in [r_k^n, r_{k+1}^n)$, for some $i \leq k$, then for some $\varepsilon > 0$ and all $\theta \in (0, \varepsilon)$, the subject's unique dominant strategy is to report $\underline{r}^n = r_i^n$ and $\bar{r}^n = r_{k+1}^n$. (b) In the limits $\theta \to 0$ and $(x - y) \to 0$, $\alpha(r_j^n) = \mu\left(r_j^n\right)$, for all $r_j^n \in \{\underline{r}^n < r_{i+1}^n, ..., r_k^n < \bar{r}^n\}$.*

Proof. (a) Let $\bar{\pi} \in [r_k^n, r_{k+1}^n)$ and compare reporting $\bar{r}^n = r_{k+1}^n$ and reporting $\bar{r}^n > r_{k+1}^n$. If $r \geq \bar{r}^n$, then in either case the subject is awarded the lottery $\ell(r; x, y)$. However, if the subject reported $\bar{r}^n > r_{k+1}^n$ and the random draw is $r = r_{k+1}^n$, he is allowed to choose between the bet, $(x - \theta\alpha\left(r_{k+1}^n\right)^2)_E(y - \theta\alpha\left(r_{k+1}^n\right)^2)$ and the lottery $\ell(r_{k+1}^n; x - \theta(1 - \alpha\left(r_{k+1}^n\right))^2, \left(y - \theta(1 - \alpha\left(r_{k+1}^n\right)\right)^2)$, whereas had he reported $\bar{r}^n = r^n_{k+1}$ he would have been awarded the lottery $\ell\left(r^n_{k+1}, x, y\right)$. By first-order stochastic dominance

$$\ell\left(r^n_{k+1}, x, y\right) \succ \ell(r^n_{k+1}; x - \theta(1-\alpha\left(r^n_{k+1}\right))^2, \left(y - \theta(1-\alpha\left(r^n_{k+1}\right)\right)^2),$$

for all $\theta > 0$, and, since $r^n_{k+1} > \bar{\pi}$, the lottery $\ell\left(r^n_{k+1}, x, y\right)$ is strictly preferred over the bet $(x - \theta\alpha\left(r^n_{k+1}\right)^2)_E(y - \theta\alpha\left(r^n_{k+1}\right)^2)$, for all $\pi \in [\underline{\pi}, \bar{\pi}]$. Hence $\bar{r}^n = r^n_{k+1}$ dominates $\bar{r}^n > r^n_{k+1}$.

Next compare reporting $\bar{r}^n = r^n_k$ and reporting $\bar{r}^n = r^n_{k+1}$. If $r < \bar{r}^n = r^n_k$ there is no difference between reporting $\bar{r}^n = r^n_{k+1}$ and reporting $\bar{r}^n = r^n_k$ (in both cases the subject is allowed to choose between the lottery $\ell(r; x - \theta(1-\alpha\left(r\right))^2, (y - \theta(1-\alpha\left(r\right))^2)$ and the bet $(x - \theta\alpha\left(r\right)^2)_E(y - \theta\alpha\left(r\right)^2)$). If $r = r^n_k$ then reporting $\bar{r}^n = r^n_{k+1}$ would allow the subject to choose between the lottery $\ell(r^n_k; x - \theta(1-\alpha\left(r^n_k\right))^2, (y - \theta(1-\alpha\left(r^n_k\right))^2)$ and the bet $(x - \theta\alpha\left(r^n_k\right)^2)_E(y - \theta\alpha\left(r^n_k\right)^2)$. In the limit, as θ tends to zero, the subject will choose the lottery if $\pi < r^n_k$ and the bet, otherwise. Had he reported $\bar{r}^n = r^n_k$, he would have received the lottery $\ell\left(r^n_k; x, y\right)$ and forgone the right to choose. Hence with probability $\mu\left[r^n_k, \bar{\pi}\right]$ (according to his own beliefs), the subject gets the lottery $\ell\left(r^n_k; x, y\right)$ although he would have strictly preferred the bet. Formally, the subject's expected utility when given the right to choose is:

$$\left[r^n_k u\left(x - \theta(1-\alpha\left(r^n_k\right))^2\right) + \left(1 - r^n_k\right) u(y - \theta\left((1-\alpha\left(r^n_k\right))^2\right)\right] \mu[\underline{\pi}, r^n_k) + \tag{4}$$

$$\int_{r^n_k}^{\bar{\pi}} \left[\pi u(x - \theta\alpha\left(r^n_{i+1}\right)^2) + (1-\pi)\, u(y - \theta\alpha\left(r^n_{i+1}\right)^2)\right] d\mu\left(\pi\right),$$

and the expected utility of the lottery $\ell\left(r^n_k; x, y\right)$ is

$$r^n_k u\left(x\right) + \left(1 - r^n_k\right) u(y). \tag{5}$$

Thus, in the limit as θ tends to zero, expression (4) is strictly larger than expression (5). Consequently, given the subject's beliefs, for $\varepsilon > 0$ sufficiently close to zero and all $\theta \in (0, \varepsilon)$, reporting $\bar{r}^n = r^n_{k+1}$ is a dominant strategy.

Let $\underline{\pi} \in [r^n_i, r^n_{i+1})$ and compare reporting $\underline{r}^n = r^n_i$ and $\underline{r}^n = r^n_{i+1}$. If $r \leq r^n_i$ or $r > r^n_{i+1}$, then, in both cases the subject is awarded the same payoff.[13]

[13]That is, the bet $x_E y$ if $r \leq \underline{r}^n$, the right to choose between the lottery $\ell(r; x - t(1-\alpha\left(r\right))^2, y - t\left((1-\alpha\left(r\right)\right)^2)$ and the bet $(x - t\alpha\left(r\right)^2)_E(y - t\alpha\left(r\right)^2)$ if $r \in (\underline{r}^n, \bar{r}^n)$, and the lottery $\ell\left(r; x, y\right)$ if $r \geq \bar{r}^n$.

If the subject reports $\underline{r}^n = r_i^n$ and $r = r_{i+1}^n$, he is allowed to choose between the lottery $\ell(r_{i+1}^n; x - \theta(1-\alpha(r_{i+1}))^2, y - \theta\left((1-\alpha(r_{i+1}))^2\right)$ and the bet $(x - \theta\alpha\left(r_{i+1}^n\right)^2)_E(y - \theta\alpha\left(r_{i+1}^n\right)^2)$; had he reported $\underline{r}^n = r_{i+1}^n$ instead, he would have been awarded the lottery $\ell(r_{i+1}^n; x, y)$. Given the right to choose, the subject would choose the bet if $\pi \in [r_{i+1}^n, \bar{\pi}]$ and the lottery if $\pi \in [\underline{\pi}, r_{i+1}^n)$. Hence his expected utility corresponding to this strategy is

$$\left[r_{i+1}^n u\left(x - \theta(1-\alpha\left(r_{i+1}^n\right))^2\right) + \left(1 - r_{i+1}^n\right) u(y - \theta\left((1-\alpha\left(r_{i+1}^n\right))^2\right)\right] \mu\left([\underline{\pi}, r_{i+1}^n)\right) + \tag{6}$$
$$\int_{r_{i+1}^n}^{\bar{\pi}} \left[\pi u(x - \theta\alpha\left(r_{i+1}^n\right)^2) + (1-\pi)\, u(y - \theta\alpha\left(r_{i+1}^n\right)^2)\right] d\mu(\pi).$$

The subject's expected utility of the lottery $\ell(r_{i+1}; x, y)$ is

$$r_{i+1}^n u(x) + \left(1 - r_{i+1}^n\right) u(y). \tag{7}$$

In the limit, as θ tends to zero, expected utility of the choice option is:

$$\left[r_{i+1}^n u(x) + \left(1 - r_{i+1}^n\right) u(y)\right] \mu\left([\underline{\pi}, r_{i+1}^n)\right) + \int_{r_{i+1}^n}^{\bar{\pi}} \left[\pi u(x) + (1-\pi)\, u(y)\right] d\mu(\pi). \tag{8}$$

Hence the option to choose is strictly preferred to the lottery (expression (8) is strictly larger than expression (7)). Consequently, there exist $\varepsilon > 0$ such that for all $\theta \in (0, \varepsilon)$ reporting $\underline{r}^n = r_i^n$ dominates reporting $\underline{r}^n = r_{i+1}^n$.

(b) Next I show that the data supplied by the subject in the context of the mechanism is sufficient to obtain an estimate of $\mu(\pi)$, for $\pi = r_j^n$, $j = 1, ..., n-1$. To simplify the notations, I define $\bar{U}(\tau) = \tau u(x) + (1-\tau) u(y)$, and $\bar{U}^{\Delta}(\tau) = \tau u'(x) + (1-\tau) u'(y)$, $\tau \in [0,1]$.

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

$$\left[r_j^n u\left(x - \theta(1-\alpha\left(r_j^n\right))^2\right) + \left(1 - r_j^n\right) u\left(y - \theta(1-\alpha\left(r_j^n\right))^2\right)\right] \mu\left(r_j^n\right) + \tag{9}$$
$$\int_{r_j^n}^{\bar{\pi}} \left[\pi u(x - \theta\alpha\left(r_j^n\right)^2) + (1-\pi)\, u(y - \theta\alpha\left(r_j^n\right)^2)\right] d\mu(\pi).$$

The first-order condition evaluated at the limit as $\theta \to 0$ is

$$\mu\left(r_j^n\right)\left(1-\alpha(r_j^n)\right)\bar{U}^{\Delta}\left(r_j^n\right)-\alpha(r_j^n)\left(1-\mu\left(r_j^n\right)\right)\int_{r_j^n}^{\bar{\pi}}\bar{U}^{\Delta}(\pi)\,d\mu\left(\pi \mid \pi \geq r_j^n\right)]=0. \tag{10}$$

For $r_j^n \geq \bar{r}^n = r_{k+1}^n$, $\mu\left(r_j^n\right) = 1$, hence $\alpha(r_j^n) = 1$. For $r_j^n < \underline{r}^n = r_{i+1}^n$ $\mu\left(r_j^n\right) = 0$, hence, $\alpha(r_j^n) = 0$.

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

$$\frac{\alpha(r_j^n)}{\left(1-\alpha(r_j^n)\right)} = \frac{\mu\left(r_j^n\right)\bar{U}^{\Delta}\left(\bar{r}^n\right)}{\left(1-\mu\left(r_j^n\right)\right)Ex\left[\bar{U}^{\Delta}(\pi) \mid \pi \geq r_j^n\right]}. \tag{11}$$

In the limit $x - y \to 0$, $\bar{U}^{\Delta}(\tau) = \tau u'(x) + (1-\tau)u'(y)$ is independent of τ. Hence, by (11) $\alpha(r_j^n) = \mu\left(r_j^n\right)$. ■

As n increases, the intervals around $\bar{\pi}$ and $\underline{\pi}$ shrink and the corresponding estimates $\bar{r}^n$ and $\underline{r}^n$ become more accurate. In the limit, as n tends to infinity, these estimates coincide with the true values and α converge to μ.

Theorem 3: *In mechanism I, (a)* $\lim_{n\to\infty}\bar{r}^n = \bar{\pi}$ and $\lim_{n\to\infty}\underline{r}^n = \underline{\pi}$. *(b) Let* $\bar{\mu}(\pi) := \alpha(r_j^n)$, *for all* $j = k+1, k, ..., i+1$ *and* $\pi \in (r_{j-1}, r_j]$, *then* $\lim_{n\to\infty}\bar{\mu}(\pi) = \mu(\pi)$, *for all* $\pi \in [\underline{\pi}, \bar{\pi}]$.

Proof. (a) For each n there are $i(n)$,$k(n) \in \{0, 1/n, ..., (n-1), 1\}$ such that $\underline{\pi} \in [r_{i(n)}^n, r_{i(n)+1}^n)$ and $\bar{\pi} \in [r_{k(n)}^n, r_{k(n)+1}^n)$. By Theorem 2, $\bar{r}^n = r_{k(n)+1}^n$, and $\underline{r}^n = r_{i(n)}^n$. Hence, $\bar{\pi} \in \cap_{n=1}^{\infty}[r_{k(n)}^n, r_{k(n)+1}^n)$ and $\underline{\pi} \in \cap_{n=1}^{\infty}[r_{i(n)}^n, r_{i(n)+1}^n)$. But $\lim_{n\to\infty} r_{k(n)}^n = \lim_{n\to\infty} r_{k(n)+1}^n$ and $\lim_{n\to\infty} r_{i(n)}^n = \lim_{n\to\infty} r_{i(n)+1}^n$. Thus $\lim_{n\to\infty} r_{i(n)}^n = \underline{\pi}$ and $\lim_{n\to\infty} r_{k(n)+1}^n = \bar{\pi}$.

(b) For each $\pi \in [\underline{\pi}, \bar{\pi}]$ and $n = 1, 2...$, 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 $n = 1, ...$, $\inf\{r_n^h \mid \pi < r_n^h, n = 1, 2, ...\}$ exists and is equal to π. But μ is right continuous, hence $\lim_{n\to\infty}\mu\left(r_n^h\right) = \mu(\pi)$. ■

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 payoffs x and

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 Mechanism I, 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 $\mathbb{R}$.[14] Suppose that a subject's beliefs regarding the distribution 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 $\mathcal{F}$.[15]

The subject is asked to report two functions, $\bar{G}$ and $\underline{G}$ mapping $\mathbb{R}$ to $[0,1]$ and a function $\alpha : \mathbb{R}\times [0,1] \to [0,1]$. 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's payoff is the bet $x_{E_k}y$ if $r \leq \underline{G}(k)$, the lottery $\ell(r; x, y)$ if $r \geq \bar{G}(k)$ and $r \in (\underline{G}(k), \bar{G}(k))$ then the subject is allowed to choose between the bet $\left(x - \theta\alpha(k,r)^2\right)_{E_k}\left(y - \theta\alpha(k,r)^2\right)$ and the lottery $\ell\left(r; x - \theta(1-\alpha(k,r))^2, y - \theta(1-\alpha(k,r))^2\right)$.

[14] Assume that $\mathcal{H}$ is endowed with the topology of weak convergence, and denote by Σ the Borel sigma algebra on $\mathcal{H}$.

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

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 dominant strategy.

Theorem 4: *For each* $(k,r) \in \mathbb{R}\times[0,1]$*: (a) reporting* $\bar{G}(k) = \sup_{F\in\mathcal{F}}\{F(k)\}$ *and* $\underline{G}(k) = \inf_{F\in\mathcal{F}}\{F(k)\}$ *is a unique dominant strategy. (b) In the limit as* θ *and* $x-y$ *tend to zero,* $\alpha(k,r) = \mu\{F\in\mathcal{F} \mid F(k)\leq r\}$.

Proof. By Theorem 1, for each k the subject's unique dominant strategy is to report $\underline{G}(k) = \inf_{F\in\mathcal{F}}\{F(k)\}$, $\bar{G}(k) = \sup_{F\in\mathcal{F}}\{F(k)\}$ and $\alpha(k,r) = \mu\{F\in\mathcal{F} \mid F(k)\leq r\}$. ■

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

4.2 A variation on the mechanism

A variation of the proposed elicitation mechanism (henceforth elicitation scheme II) dispenses with the direct elicitation of $\underline{r}$ and $\bar{r}$. It requires the subject to report a function $\alpha : [0,1] \to [0,1]$. Following the report of α, a random number, r, is drawn from a uniform distribution on $[0,1]$. The subject is awarded the choice between the bet $\left(x-\theta\alpha(r)^2\right)_E\left(y-\theta\alpha(r)^2\right)$ and the lottery $\ell\left(r; x-\theta(1-\alpha(r))^2, y-\theta(1-\alpha(r))^2\right)$, where $\theta>0$.

The truthful revelation of μ is an immediate implication of Theorem 1.

Corollary 1: *In elicitation scheme II, in the limit as* $\theta \to 0$, *and* $(x-y)\to 0$, *the choice of* $\alpha(\cdot)$ *satisfies* $\alpha(r) = \mu(r)$, *for all* $r\in[0,1]$.

The proof is an immediate implication of the proof of part (b) of Theorem 1 where $\underline{r}$ and $\bar{r}$ are replaced with 0 and 1, respectively. An implication of Corollary 1 below is that the support of μ can be recovered from that α. Formally, let $\underline{r} = \sup\{r\in[0,1] \mid \alpha(r)=0\}$ and $\bar{r} = \inf\{r\in[0,1] \mid \alpha(r) = 1\}$, then $Supp\mu = [\underline{r},\bar{r}]$.

Consider next a discrete variation of mechanism*. Fix an integer n and partition of the unit interval $\{[(i-1)/n, i/n) \mid i=1,...,n-1\} \cup [(n-1)/n, 1]$. The subject is asked to report a number $\alpha(i/n)\in[0,1]$, for each $i=0,1,...,n$. A random number r is selected from a uniform distribution on $\{i/n \mid n=0,1,...,n\}$. The subject is allowed to choose between the bet, $(x-\theta\alpha(r)^2)_E(y-\theta\alpha(r)^2)$, and the lottery $\ell(r; x-\theta(1-\alpha(r))^2, (y-\theta(1-\alpha(r))^2)$, $\theta>0$. By part (b) of Theorem 2 we have,

Corollary 2: *In mechanism*, in the limits* $\theta \to 0$ *and* $(x - y) \to 0$, $\alpha(i/n) = \mu(i/n)$, *for all* $i = 0, ..., n$.

By part (b) of Theorem 3, we have that in the limit, as n tends to infinity α converge to μ. Formally,

Corollary 3: *In mechanism*, let* $\bar{\mu}(\pi) := \alpha(i/n)$, *for all* $i = 1, ..., n$, *and* $\pi \in ((i-1)/n), i/n]$, *then* $\lim_{n\to\infty} \bar{\mu}(\pi) = \mu(\pi)$, *for all* $\pi \in [\underline{\pi}, \bar{\pi}]$.

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).[16] Except for Chambers and Lambert (2016) and the mechanisms described here, 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 protocol proposed by Chambers and Lambert (2016) (henceforth the CL protocol) and the elicitation scheme described in this paper are unconventional. These elicitations procedures are designed to elicit the expert's subjective "probability on probabilities," interpreted as second-order introspective beliefs or subjective information structure. Consequently, the event of interest is a set of subjective probabilities that, by definition, are not publicly observable.

Despite this similarity, the CL protocol and the mechanism of this paper are quite different. The mechanism proposed here is designed to elicit the range of expert's set of priors (or posteriors). The CL protocol is designed to elicit the expert's unique posterior. This difference stems from the choice theories underlying the mechanisms, as explained below.

To highlight the similarities and differences, I compare the two procedures in the context of the simple example introduced by Chambers and Lambert to explain the working of their protocol.[17] In this example, a random variable X taking values in $\{0, 1\}$ is an indicator of a publicly observed event, and there are three periods, $t = 0, 1, 2$.[18] In the last period, $t = 2$, the value of X is publicly observed. In the interim period, $t = 1$, the expert receives a private

[16] For a recent review, see Chambers and Lambert (2016).

[17] See Chambers and Lambert (2016) section 2.

[18] In terms of the procedure described here, the random variable X can be interpreted as the indicator of the event of interest, E.

signal of the values of a random variable, P, representing a distribution of X. Because X is an indicator function, P takes values in the unit interval and each realization of P, $p \in [0,1]$, is the probability of $X = 1$. In the initial period, $t = 0$, the expert entertains beliefs about the signal to be received in the interim period. These beliefs are captured by a cumulative distribution function (CDF) F on $[0,1]$. Consequently, the expert prior is given by $\bar{p}$, the mean of P under F. Chambers and Lambert refer to F as a prior and P as the posterior and their protocol is design to elicit both.

In the model of this paper, the sets of the expert's beliefs, Π_0 or Π_1 according to the two interpretations, are analogous to the posterior P, and the subject's random choice of π can be interpreted as response to unspecified impulse or information signal. Hence, the selection of π is analogous to the signal p. The prior F is analogous to μ_0 or μ_1, the cumulative distribution function on the sets Π_0 or Π_1, respectively, that represents the expert's second-order beliefs (or subjective information structure) about the distribution of $\widetilde{\pi}$. Unlike the CL protocol, *the mechanism of this paper admits the possibility, but does not require, that there be a unique prior.*

The elicitation mechanism proposed here and the CL protocol are designed to elicit the expert's second-order belief, μ, and his belief regarding information structure summarized by the CDF F, respectively. Both elicitation mechanisms are incentives schemes intended to induce the expert to reveal his true information. However, as mentioned above, the objectives of the proposed elicitation schemes are not identical. The objective of the CL protocol is: "to motivate the expert to reveal his prior at the outset, and subsequently the posterior he observes at the interim period" (Chambers and Lambert [2016], p. 8). By contrast, depending on the interpretation, the objective of the mechanisms described in this paper are the elicitation of the expert second-order beliefs and the ranges of the set of priors (or his subjective information structure and the range of the posteriors). In what follows, I use the example to further highlight the differences between the CL protocol and the elicitation scheme proposed in this paper.

The CL protocol is based on two randomly selected parameters, α and β, from the uniform distribution on $[0,1]$ and $[-1,1]$, respectively. These parameters represent, respectively, the price of an option, dubbed $\alpha-$option, to short-sell a security in the interim period ($t = 1$) at the price α, and the price of the $\alpha-$option as of the initial period ($t = 0$). The mechanism requires that "the expert announces a prior $\hat{F}$ and, later, a posterior $\hat{p}$. The elicitor randomly selects the option and its price, and, 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 eventually gets the final value that results from the elicitor's choices." (Chambers and Lambert [2016], p. 11)

The CL protocol presumes that the expert is an expected-value maximizer. In particular, the expert's *preferences are complete.* The completeness of the preference relation renders the expert's choice behavior predictable, which makes it possible for the elicitor to act on the expert's behalf. Thus, the CL protocol, designed to elicit the expert's information structure and posterior beliefs, applies to experts whose behavior is depicted in the models of Dillenberger et al. (2014) and Lu (2016). If the expert's preference are incomplete, his choices become impossible to predict and the CL protocol does not apply.

By contrast, the elicitation schemes proposed in this paper is designed to elicit the expert's range of priors and second-order beliefs. It is therefore applicable when the expert is a subjective expected-utility maximizer whose preference relation is incomplete, as in the model of Karni and Safra (2016). Unlike the CL protocol, the scheme proposed in this paper works even if the expert exhibits unpredictable (that is, random) choice behavior. To grasp the significance of this observation, suppose that an expert whose preference relation is incomplete participates in the CL protocol. Upon receiving a private signal, this expert updates his *set of priors*. Therefore, he does not have a single posterior representing his updated beliefs that he can report. If the expert reports a set of posteriors, the elicitor might find that the α-option is worth purchasing according to some posteriors but not others. Hence, the elicitor cannot predict what the expert would choose to do and, consequently, cannot act on his behalf. This lack of predictability makes the price of the $\alpha-$option in the first period impossible to determine and causes the CL protocol to break down.

If the elicitation scheme proposed here is applied to a Bayesian decision maker it would work as follows. At time $t = 0$, the expert reports two numbers, $\underline{r}$, $\bar{r} \in [0, 1]$, and, for each $r \in (\underline{r}, \bar{r})$, 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 expert is awarded the bet $x_E y$ if $r \leq \underline{r}$ and the lottery $\ell(r; x, y)$ if $r \geq \bar{r}$. If $r \in (\underline{r}, \bar{r})$ then the expert is allowed to choose between the bet $\left(x-\theta\alpha\left(r\right)^{2}\right)_{E}\left(y-\theta\alpha\left(r\right)^{2}\right)$ and the lottery $\ell\left(r;x-\theta(1-\alpha\left(r\right))^{2},y-\theta(1-\alpha\left(r\right))^{2}\right)$, where $\theta>0$. By Theorem 1, as θ and $x-y$ tend to zero, the expert's best response is to report his subjective assessment of the true range of the posteriors probabilities of the event E and his true second-order beliefs, which in this case correspond to his information structure. Note, however, that, unlike the CL protocol, this scheme is not designed to elicit the expert's unique posterior. Rather, it elicits the a priori range of the expert's posteriors.

In terms of implementation, the CL protocol "relies on the elicitor's ability to *commit* to making the best choices on the expert's behalf *without* the expert having the ability to verify the elicitor's actions" (Chambers and Lambert [2016], p. 11) The elicitor's commitment to act in the best interest of the expert might be problematic when there is potential conflicts of interests. The mechanisms of this paper do not require the elicitor's to act as an agent of the expert. Moreover, unlike the scheme proposed here, which does not restrict the expert's risk attitudes, the CL protocol stipulates that the elicitor maximize the expected value corresponding to the expert's announcements. Thus, the scheme applies naturally to risk-neutral experts. If the expert's is not risk neutral, the application of the CL protocol is still possible, provided that the expert is paid with lottery tickets (or "probability currency"), which adds another, unnatural, complication to the experimental design.

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. (2016) "Dynamic Belief Elicitation," unpublished manuscript.

[5] Dekel, Eddie, Lipman, Barton L. and Rustichini, Aldo (2001) "Representing Preferences with a Unique Subjective State Space," *Econometrica* 69, 891-934.

[6] Dillenberger, David, Lleras, Juan S., Sadowski, Philipp and Takeoka, Norio (2014) "A Theory of Subjective Learning," *Journal of Economic Theory* 153, 287-312.

[7] Galaabaatar, Tsogbadral and Karni, Edi (2013) "Subjective Expected Utility with Incomplete Preferences," *Econometrica* 81, 255-284.

[8] Good, Irving J. (1952) "Rational Decisions," *Journal of the Royal Statistical Society, Series B (Methodological)*, 14, 107-114.

[9] Grether, David M. (1981) "Financial Incentive Effects and Individual Decision Making," California Institute of Technology, Social Science Working Paper 401.

[10] Gul, Faruk and Pesendorfer, Wolfgang (2006) "Random Expected Utility," *Econometrica* 74, 121-146.

[11] Kadane, Joseph B., and Winkler Robert L. (1988) "Separating Probability Elicitation from Utility," *Journal of the American Statistical Association*, 83, 357-363.

[12] Karni, Edi (2009) "A Mechanism Design for Probability Elicitation," *Econometrica* 77, 603-606.

[13] Karni, Edi and Safra, Zvi (2016) "A Theory of Stochastic Choice under Uncertainty," *Journal of Mathematical Economics,* (forthcoming)

[14] Kreps, David M. (1979) "A Representation Theorem for 'Preference for Flexibility'," *Econometrica* 47, 565-576.

[15] Lu, Jay (2016) "Random Choice and Private Information," *Econometrica* (forthcoming)

[16] Minardi, Stefania and Savochkin, Andrei (2015) "Preferences with Grades of Indecisiveness," *Journal of Economic Theory* 155, 300-331.

[17] Qu, Xiangyu (2012) "A Mechanism for Eliciting a Probability Distribution," *Economic Letters* 115, 399-400.

[18] Savage, Leonard J. (1971) "Elicitation of Personal Probabilities and Expectations," *Journal of American Statistical Association*, 66, 783-801.

[19] Seidenfeld, Teddy, Schervish, Mark J. and Kadane, Joseph B. (1995) "A Representation of Partially Ordered Preferences," *Annals of Statistics* 23, 2168-2217.