Eliciting Beliefs through Rank Ordering

Greg Leo

October 28, 2020

Abstract

I present a novel procedure for eliciting beliefs using rank-ordering of payment propositions. This procedure can be used to elicit the probability of a binary outcome or any point on subjective cumulative distribution function [CDF] or inverse-CDF of a non-binary outcome. The procedure is simple, the incentives are transparent, require no calculation, and are robust to risk preferences.

1 Introduction

Beliefs play a key role in decision-making. Experimental economists have developed methods for incentivizing belief elicitation (Trautmann and van de Kuilen, 2015). Most of these methods condition a subject's payment on how their stated belief compares to a single draw from a random variable in question if the belief is about an inherently random prospect ("what is the probability that a randomly chosen person has brown hair?"), or to the true value if the topic of belief elicitation is non-random ("what is the distance from Los Angeles to Tucson?"). For payment rules of this type, designing the incentives for belief elicitation comes down to choosing a function that maps stated belief and a true outcome into payment.

Functions that provide incentive for subjects to provide their true belief about some function of their beliefs about a random variable or event are known as proper. For instance, a rule is proper for the probability of an event if a risk-neutral subject's expected utility is maximized by submitting their true belief about the probability. Proper scoring rules have been studied extensively in the statistical literature. A review is available in Gneiting and Raftery (2007). For instance, one rule that is proper for the probability of an event pays an amount of money decreasing linearly in the squared-error made in predicting the probability of that event. For instance, if a subject states that the probability of an even is 0.75 and the event is true, then the error made is $(1-0.75)^2 = 0.0625$. This "quadratic scoring rule" would pay $\$\alpha - \beta\,(0.0625)$ where α and β can be chosen arbitrarily.

The state-of-the-art in experimental economics uses binarization to make a scoring rule robust to risk preferences (Hossain and Okui, 2013; Harrison et al., 2014). In these payment rules, the probability of being paid some fixed amount, rather than the amount itself, is a function of stated beliefs and the true outcome. For the case of the quadratic scoring rule, instead of paying an amount of money that is linearly decreasing in the square of the probabilistic error, the subject is compensated by a binary lottery where the probability of winning some amount of money is linearly decreasing relative to this error. For instance, for the case of predicting a 0.75 chance of an event that does happen, the subject could be compensated with a lottery paying \$10 with a $1-(1-0.75)^2 = 0.9375$ chance.

There is nothing intuitive about these methods. They are difficult to understand. In response, researchers either carefully explain the mathematical function that determines subjects' payments or simply assure subjects that it is in their best interest to tell the truth. Both are insufficient responses to the complexity of these rules.

We can expect only the most motivated subjects to try and understand a complex scoring rule. Furthermore, assuring subjects that it is in their best interest to "tell the truth" about something like a probabilistic belief assumes subjects understand what that means in the first place. In our rain example, a subject might not think that "telling the truth" about their belief regarding rain would involve a probability at all. It either will rain or it will not.

Recent research further complicates this issue. Danz et al. (2020) have shown that subjects tend to bias their beliefs towards 0.5 whether the experimental instructions explicitly explain the binarized quadratic scoring rule or assure the subjects that it is in their best interest to tell the truth.

The remedy for these issues is to develop a way of eliciting beliefs that is both theoretically valid and intuitive for subjects. That is the goal of this paper. I propose an indirect methodology that implements discretized versions of binarized proper scoring rules. The methodology requires subjects to rank payment propositions. The following example provides intuition for the design of the methodology.

Suppose a subject is asked whether they would prefer to be paid \$10 with a 0.5 chance or \$10 if it rains tomorrow. If the subject has a probabilistic belief about the chance of rain tomorrow, then the fact that they prefer \$10 with a 0.5 probability is because they believe that there is less than a 0.5 chance of rain. Asking a series of these questions pinpoints the belief. If, for instance, they prefer being paid if it rains to being paid with a 0.4 chance, then their true belief about rain must be between 0.4 and 0.5.

The full methodology for eliciting a probability asks subjects to rank the proposition of being paid some fixed amount if the event is true among propositions involving being paid that same amount with a series of known probabilities. To incentivize correctly ranking these propositions, and thus truthfully revealing their belief, one probability is chosen at random. If the subject ranks the proposition involving the event higher than the proposition with known probability, they are paid according to the proposition contingent on the event and otherwise are paid according to the proposition with the known probability.

The incentive compatibility of this procedure is easy to explain. If the subject believes there is a 0.45 chance of rain tomorrow but ranks \$10 if it rains below being paid \$10 with a 0.4 chance, then if the 0.4 proposition is randomly chosen, they will be paid with a 0.4 chance. However, since they believe that there is a 0.45 chance of rain, they would prefer to be paid if it rains. It is thus in their best interest to correctly order all of these propositions.

Furthermore, the incentives underlying this procedure are a discrete version of the incentives offered by the binarized quadratic scoring rule. Thus, not only is this methodology theoretically valid, but it is much simpler for subjects to understand. Furthermore, it is possible to elicit several beliefs in one task by having subjects rank several uncertain propositions among payment with known probabilities.

The methodology can also extended to points on the subjective CDF and inverse CDF for a subject's beliefs about a non-binary event. For the distance is from Los Angeles to Tucson, eliciting points on the inverse CDF requires eliciting beliefs such as: "what is the probability you believe this distance is less than 500 miles?" Since this is simply a probabilistic belief, the procedure above applies. However, if one wants to elicit points on the CDF such as: "what is the number x that you think there is a 0.75 chance the distance between Los Angeles and Tucson is less than x?" A slight modification is needed.

To elicit this belief, consider asking the subject the following question. "Would you rather be paid \$10 with a 0.75 chance or \$10 if the distance is less than 500 miles?" If they prefer to be paid if the distance is less than 500 miles, then they must believe there is more than a 0.75 chance it is less than 500 miles. If, however, they prefer to be paid with a 0.75 chance to being paid if the distance is less than 400 miles, the 0.75 quantile of their subjective belief about this distribution must be between 400 and 500 miles. Asking the subject to rank the proposition about being paid with the known probability q among propositions about being paid according to a series of nested events (less than 1000 miles, less than 900 miles, ...) pinpoints the value of the q quantile in their subjective CDF.

2 Methodology

The decision-maker is assumed to have *complete* and *transitive* preferences $\succsim$ over objective lotteries $\mathcal{G}$. $\succsim$ is further assumed to meet the *monotonicity*, *Archimedean*, and *reduction of compound lottery* axioms. The decisions maker's beliefs about random variable R are given by subjective distribution F and $\succsim$ is extended onto lotteries involving contingencies about R through F. The set of subjective lotteries is denoted $\mathcal{G}_R$.

2.1 Eliciting Beliefs

The experimenter's goal is to learn about the subjective distribution F. This can be achieved by comparing revealed information about $\succsim$ to find differences of the form $g \sim \tilde{g}$, with $g \in \mathcal{G}$ and $\tilde{g} \in \mathcal{G}_R$.

Since it is assumed preferences over objective lotteries are extended into subjective lotteries through the subjective distribution F, these indifferences imply direct information about F. For instance, an indifference between the contingent lottery "\$10 if the distance from Los Angeles to Tucson is less than 500 miles" and the objective lottery "\$10 with probability 0.5" implies that the subjective belief F is consistent with the event "Los Angeles and Tucson are less than 500 miles apart" having probability 0.5.

2.2 Eliciting a Probability

To elicit the probability of event x in the range of random variable, select monetary outcomes $\$a$ and $\$b$ with $a \geq b$, and construct simple lotteries g_i over these events of the form $p_i \circ \$a, (1 - p_i) \circ \$b$ with $p_i > p_{i+1}$. Note that since the subject's preferences over simple lotteries are monotonic, then the subject's preferences over lotteries in the set of lotteries of the form $p_i \circ \$a, (1 - p_i) \circ \$b$ are a linear order with $g_i \succ g_{i+1}$. This linear order extends onto the contingent lottery g_x: "$\$a$ if x is true and $\$b$ if x is false". To elicit the position that this lottery falls in the linear order, display the simple objective lotteries in order and ask the subject to place g_x in its correct position among the objective lotteries. An example is given below.

To incentivize correct ranking, randomly choose one objective lottery g_i. If the subject ranks the g_x above g_i, pay according to g_i and otherwise pay according to g_i. Since $\succsim$ is assumed to obey *monotonicity* and *the reduction of compound lotteries,* this procedure is incentive compatible.

Fixed Contingent Lottery	Known-Probability Lotteries
\$10 if distance $\leq$ 500 miles	0% chance of \$10
	20% chance of \$10
	40% chance of \$10
	60% chance of \$10
	80% chance of \$10
	100% chance of \$10

Table 1: Example: Eliciting a Probability.

Let g_t and g_{t+1} be the lotteries such that $g_t \succsim g_x \succsim g_{t+1}$ according to the ranking of g_x, and let p_t and p_{t+1} be the two probabilities associated with those lotteries. By monotonicity and the *Archimedean* axiom, there is some $\hat{p} \in [p_t, p_{t+1}]$ such that the simple lottery $(\hat{p} \circ \$a, (1 - \hat{p}) \circ \$b) \sim g_x$. Since preferences over contingent lotteries are induced by preferences over simple lotteries and the distribution F, then $\hat{p}$ is the subject's belief about the probability of event x.

2.3 Eliciting a Quantile

To elicit the p quantile of subjective belief F about random variable X, select monetary outcomes $\$a$ and $\$b$ with $a \geq b$, and construct nested contingent lotteries g_i of the form "$\$a$ if $X \leq q_i$ and $\$b$ otherwise" with $q_i > q_i'$ for $s > s'$. Since preferences over these lotteries are induced by F, and the subject's preferences over simple lotteries are monotonic, then the subject's preferences over these contingent lotteries are a linear order with $g_i \succ g_{i+1}$. This linear order extends onto the simple lottery g_s: $(p \circ \$a, (1-p) \circ \$b)$.

To elicit the position that this lottery falls in the linear order, display the nested contingent lotteries in order and ask the subject to place g_s in its correct position among the objective lotteries. An example is given below.

To incentivize correct ranking, randomly choose one nested contingent lottery g_i. If the subject ranks the g_s above g_i, pay according to g_s and otherwise pay according to g_i. Moving g_s up in the ranking weakly increases the probability of g_s being chosen from the compound lottery. Since $\succsim$ is assumed to obey *monotonicity* and *the reduction of compound lotteries,* this procedure is incentive compatible.

Nested Contingent Lotteries	Fixed Known-Probability Lottery
\$10 if it snows less than 10 inches tomorrow.	50% chance of \$10
\$10 if it snows less than 8 inches tomorrow.	
\$10 if it snows less than 6 inches tomorrow.	
\$10 if it snows less than 4 inches tomorrow.	
\$10 if it snows less than2 inches tomorrow.	
\$10 if it snows less than 0 inches tomorrow.	

Table 2: Example: Eliciting a Quantile.

Let g_i and g_{i+1} be the lotteries such that $g_i \succsim g_s \succsim g_{i+1}$ according to the ranking of g_x, and let q_t and q_{t+1} be the two values in the domain of F associated with those lotteries. Since F is increasing and continuous, and preferences are *monotonic* and obey the *Archimedean* axiom, there is some $\hat{q} \in [q_t, q_{t+1}]$ such that the simple lottery $(p \circ \$a, (1-p) \circ \$b)$ is indifferent to "\$a if $X \leq \hat{q}$ and \$b otherwise." Since preferences over contingent lotteries are induced by preferences over simple lotteries and the distribution F, then $\hat{q}$ is the p-quantile of the subjective distribution F.

References

Danz, D., Vesterlund, L., Wilson, A.J., 2020. Belief elicitation: Limiting truth telling with information on incentives. Technical Report. National Bureau of Economic Research.

Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. Journal of the American statistical Association 102, 359–378.

Harrison, G.W., Martínez-Correa, J., Swarthout, J.T., 2014. Eliciting subjective probabilities with binary lotteries. Journal of Economic Behavior & Organization 101, 128–140.

Hossain, T., Okui, R., 2013. The binarized scoring rule. Review of Economic Studies 80, 984–1001.

Trautmann, S.T., van de Kuilen, G., 2015. Belief elicitation: A horse race among truth serums. The Economic Journal 125, 2116–2135.