ELICITING SUBJECTIVE REAL-VALUED BELIEFS†

GREG LEO* & SAM STELNICKI**

ABSTRACT. We present a simple and robustly incentive-compatible price-list methodology to elicit quantiles of a subjective real-valued belief. These elicited quantiles can be employed to approximate a subject's complete subjective distribution, and we establish that the distribution maximizing entropy while adhering to the elicited quantiles is piecewise linear. Using this approach, our methodology extends to estimating arbitrary unobserved attributes of the subjective distribution, such as mean and variance, which are otherwise challenging to elicit. We provide a proof-of-concept for our framework through an experiment involving the elicitation of participants' beliefs regarding the mathematical abilities of their peers.

Keywords: Belief Elicitation; Experimental Methods, Incentive Compatibility; Maximum Entropy

JEL Classification: C90, C91, D83.

February 4, 2024

†The authors thank Paul J. Healy, Kristine Koutout, Andrew Dustan, and seminar participants at UC Santa Barbara, Loyola Marymount, and the 2023 ESA North American Meetings for their valuable comments and feedback.

*Dept. of Economics, Vanderbilt University; g.leo@vanderbilt.edu.

**Dept. of Economics, The Ohio State University; stelnicki.4@buckeyemail.osu.edu.

1. Introduction

Subjective beliefs play a crucial role in economic decision-making. Relevant beliefs are often about real-valued random variables.[1] The decision to buy an asset depends on beliefs about the future price. The decision to install farm irrigation depends on beliefs about future rainfall. In these cases, beliefs take the form of a density function over the range of possible values.

Decisions can depend in complex ways on the subjective belief distribution. For example, the value of an asset to a risk-neutral buyer depends on the mean of the price distribution. The value of irrigation for a farmer depends on the likelihood and severity of dry conditions, which may be related to the lower tail of the rainfall distribution. Because of this, it is important that a belief elicitation methodology in this environment is flexible in allowing a researcher to capture the relevant information about the subjective distribution.

In this paper, we introduce a methodology for eliciting quantiles of a subjective real-valued belief using a price list methodology: the quantile price list. Our methodology is simple, incentive compatible under very general conditions, and can be extended to elicit any quantile. This allows researchers to pinpoint the information they require or to get a more comprehensive view of a participant's underlying beliefs by eliciting several quantiles. We also demonstrate how the elicited quantiles can be used to approximate an entire belief distribution and estimate unobserved properties, such as the mean.

We demonstrate our methodology using an experiment on beliefs about student math test performance. Twenty students from The Ohio State University took a math test and were marked as "successful" if they answered at least thirteen questions correctly. We elicited the beliefs of the participants about the number of students who were successful, specifically the 0.25, 0.50 (median), and 0.75 quantiles of their subjective belief distributions, and use these beliefs to approximate the entire cumulative distribution function for each participant.

Additionally, we elicited the probability each participant believes that a randomly chosen student passed the math task using a price list methodology for probabilities proposed in Holt and Smith (2016). We compare the approximated CDF's from our methodology to the implied binomial distributions induced from this elicited probability and find that the actual subjective CDFs are much "flatter" than the implied binomial distributions, demonstrating how our methodology can provide a more nuanced understanding of beliefs.

Section 2 outlines the theory of quantile price lists. Section 3 compares our methodology with other methods. Section 4 describes our experimental design. Section 5 reports the results of the experiment, and Section 6 concludes with a discussion focusing on potential applications.

[1]Either truly random or random from the perspective of the decision maker

2. THEORY

We begin this section by describing our procedure informally and discussing why and under what conditions it elicits quantiles of a subjective belief. In the next section, we formalize this discussion and prove incentive compatibility.

To make these examples concrete and to show that a subjective belief need not be "about" an inherently random quantity, suppose that we want to learn about a participant's belief about the distance between Los Angeles and San Diego. Let us start with a participant who is somewhat familiar with the geography of California. They know that these cities are both in southern California and are not *that* far apart.

Suppose that we ask this participant if they would rather have \$10 with a 75% chance or \$10 if the distance between the two cities is actually less than 1000 miles. This participant believes that the distance is almost certain to be less than 1000 miles and chooses the latter option, since they believe that it will yield the \$10 with a nearly 100% chance. We can conclude that whatever the distribution of their belief F about the distance between these two cities, it is the case that $F(X \leq 1000) > 0.75$.

Now suppose that we continue asking questions like this, but where the second option pays conditional on the distance being below $900, 800, 700, 600, 500, 400$ and 300 respectively. In each case, the participant chooses the second option each time. However, they conclude that the chance that the distance is below 200 miles is less than 75% and, when asked to compare being paid \$10 with a 75% chance and being paid conditional on the distance being below 200 miles, they now choose the first option. From these choices, we can conclude that for their belief F, $F(X \leq 300) > 0.75$ but $F(X \leq 200) < 0.75$. Thus, whatever number x solves $F(X \leq x) = 0.75$ (the 75-th quantile) must be between 200 and 300.

To find a different quantile, we repeat this exercise with a different objective lottery. For example, if the same participant were asked to compare \$10 with a 25% chance and \$10 if the distance is below 200, they may well choose the latter option but switch to choosing the objective lottery when asked to compare it to \$10 if the distance is below 100. These pairs of choices imply that $F(X \leq 200) > 0.25$ but $F(X \leq 100) < 0.25$. We can conclude the 0.25 quantile of their subjective belief is between 100 and 200.

Notice that in these cases, we use the choices to bound the number x that would create an indifference between some objective lottery that pays with probability p and an act that pays if the random variable is below x. We interpret this x as being the $p-$th quantile of the participant's belief. To interpret this indifference as a quantile of the belief, we need to assume a few things about preferences:

If a participant is indifferent between a 75% chance of \$10 and \$10 if the random variable is below 250, but strictly prefers a 75% chance of \$20 to \$20 if the random variable is below 250 then we cannot interpret 250 as the 0.75 quantile because the number we would infer as the 0.75 quantile appears to depend on whether \$10 or \$20 is the outcome. In this case, we cannot reliably infer a belief using some arbitrarily chosen outcome for incentivizing the revelation of that belief. Thus, to interpret the outcome of our procedure as identifying the quantile of a belief, we need to assume that participants' willingness to substitute between acts and objective lotteries does not depend on the arbitrarily chosen outcome. This is the *replacement* axiom.

Another way our procedure can fail to deliver a result that can be interpreted as the quantile of a belief is if preferences over acts or objective lotteries are not monotone. For example, if a participant is indifferent between a 75% chance of \$10 and \$10 if the random variable is below 250, but also indifferent between a 50% chance of \$10 and \$10 if the random variable is below 300, then it is hard to interpret either of these indifferences as generated by a coherent belief, since it implies that F would have to be downward sloping. One way to get preferences like this is if preferences over acts do not respect the monotonicity of events, so that an act can be preferred to another even if it pays conditional on an event that is a strict subset of another (such as $X \leq 250$ vs $X \leq 300$). Another way to obtain preferences like this is if preferences over objective lotteries are not monotonic in probabilities (for instance, \$10 with 50% chance is preferred to \$10 with 75%). On the other hand, if both of these monotonicity conditions are met, then the indifferences between acts that pay conditional on larger (in terms of inclusion) events will occur at higher objective probabilities, and thus these indifferences can be interpreted as belonging to some consistent distribution F. We refer to this pair of assumptions as the *act monotonicity* and *objective lottery monotonicity* axioms, respectively.

The replacement axiom and the two monotonicity axioms are enough to ensure that we can use indifferences to identify structure in preferences that can reasonably be interpreted as the quantiles of an underlying belief distribution. However, to actually find that indifference, we will need one more assumption, which has been latent in the discussion up until now. When we ask these participants to choose their favorite option from the relevant pair and incentivize it by randomly implementing one of these choices, how do we know that their choice in each menu is their favorite? This requires weak structure on the participant's preferences over compound lotteries over the types of acts and objective lotteries used in this procedure. This axiom is known as *statewise monotonicity*, which is an assumption required in any experimental economic procedure that asks participants to make multiple choices and then randomly chooses one to implement (Azrieli et al., 2018).

2.1. *Framework*

This section makes use of the following notation. The set of **outcomes** is $\mathscr{X} = \{a, b\}$. Throughout, it is assumed $a \succ b$. X is a real-valued random variable with state space Ω. A **tail event** E_x is an event of the form $X \leq x$. **Simple lotteries** are objective lotteries of the form $S = (a \circ p, b \circ (1-p))$, where $p \in [0,1]$. **Binary acts** are acts of the form $(a \circ E, b \circ E^c)$ where E is an event in X and E^c is the complement of that event. **Simple mixtures** are compound lotteries that mix (potentially) objective and subjective risk of the form $M = (L_1 \circ p_1, L_2 \circ p_2, ..., L_n \circ p_n)$ where $p_i \in [0,1]$ with $\sum_{i=1}^{n} p_i = 1$, and each L_i is a simple lottery or a binary act. **The set of all simple mixtures** is $\mathscr{M}$.

A **Quantile price list** is a probability p (the quantile to be elicited), along with a sequence of n values $(x_1, ..., x_n)$ with $x_i > x_{i+1}$. The **range** of the quantile price list is $[x_n, x_1]$.

The price list is constructed by pairing a constant objective lottery $o_p = (a \circ p, b \circ 1-p)$ with a sequence of increasing binary acts $A_{x_i} = a \circ E_{x_i}, b \circ E^c_{x_i}$ (recall that tail events E_{x_i} are events $X < x_i$) to create n **menus** of the form $\{o_q, A_{x_i}\}$. To implement the price list, participants are asked to choose either the lottery or the act from each menu. A menu is randomly chosen (according to any fixed distribution), and the participant is rewarded with the lottery chosen from that menu.

Assume participants have preference relation $\succsim$ over $\mathscr{M}$ that is complete, transitive, and also meets the additional axioms below, then participants behave "as-if" their preferences are generated by a well-formed subjective distribution F, and that quantiles of this subjective belief are revealed by switching-points in their choices from a quantile price list.

Axiom (1): *Objective Lottery Monotonicity:*
For $a \succ b,\ p \geq p' \Leftrightarrow (a \circ p, b \circ 1-p) \succsim (a \circ p', b \circ 1-p')$.

Axiom (2): *Act Monotonicity:*
For $a \succ b,\ E' \subseteq E \Rightarrow (a \circ E, b \circ E^c) \succsim (a \circ E', b \circ E'^c)$.

Axiom (3): *Continuity:*
$\forall E \in \Omega,\ \exists p \in [0,1] : (a \circ E, b \circ E^c) \sim (a \circ p, b \circ 1-p)$.

Axiom (4): *Replacement*
For all a, a', b, b': $(a \circ E, b \circ E^c) \sim (a \circ p, b \circ 1-p) \Leftrightarrow$
$(a' \circ E, b' \circ E^c) \sim (a' \circ p, b' \circ 1-p)$

Axiom (5): *Statewise Monotonicity*
$L^*_i \succ L_i \Leftrightarrow (L_1 \circ p_1, ..., L^*_i \circ p_i, ..., L_n \circ p_n) \succ (L_1 \circ p_1, ..., L_i \circ p_i, ..., L_n \circ p_n)$

2.2. *Belief Consistency*

In this section, we demonstrate that under *objective lottery monotonicity*, *act monotonicity*, *continuity*, and *replacement*, participants behave "as-if" their preferences are generated by a well-formed subjective distribution F.

Proposition 1. Under *axioms 1-4*, there exists a unique CDF function F (invariant to the choice of a and b) that solves $\left(a \circ F(x), b \circ 1 - F(x)\right) \sim \left(a \circ E_x, b \circ E^c_x\right)$.

Proof. By *continuity*, for any x there exists a probability $F(x) \in [0,1]$ that solves:

$$\left(a \circ F(x), b \circ 1 - F(x)\right) \sim \left(a \circ E_x, b \circ E^c_x\right) \tag{1}$$

By *replacement* $F(x)$ is invariant to the choices of a and b. We now prove that F is a CDF. We start by showing that the range of F is $[0,1]$:

$$\lim_{x\to\infty} (a \circ E_x, b \circ E^c_x) = (a \circ \Omega, b \circ \emptyset) = a = (a \circ 1, b \circ 0) \tag{2}$$

Thus, $lim_{x\to\infty} F(x) = 1$.

$$\lim_{x\to-\infty} (a \circ E_x, b \circ E^c_x) = (a \circ \emptyset, b \circ \Omega) = b = (a \circ 0, b \circ 1) \tag{3}$$

Thus, $lim_{x \to -\infty} F(x) = 0$.

We now show that F is non-decreasing in x. For $x \geq x'$, $E'_x \subset E_x$. By *act monotonicity*, $\left(a \circ E_x, b \circ E_x^c\right) \succsim \left(a \circ E_{x'}, b \circ E_{x'}^c\right)$. By transitivity and indifference equation 1:

$$\left(a \circ F(x), b \circ 1 - F(x)\right) \succsim \left(a \circ F(x'), b \circ 1 - F(x')\right) \tag{4}$$

By *objective lottery monotonicity* the preference in Equation 4 is true if and only if $F(x) \geq F(x')$. Thus, $F(x)$ is non-decreasing in x. We now show that $F(x)$ is unique. Suppose otherwise, then for some E_x there exists $p, p' \in [0,1]$ with $p \neq p'$ such that:

$$\left(a \circ E_x, b \circ E_x^c\right) \sim (a \circ p, b \circ 1 - p) \tag{5}$$

$$\left(a \circ E_x, b \circ E_x^c\right) \sim \left(a \circ p', b \circ 1 - p'\right) \tag{6}$$

By *transitivity*, $\left(a \circ p', b \circ 1 - p'\right) \sim (a \circ p, b \circ 1 - p)$. Without loss of generality, assume $p > p'$. By *objective lottery monotonicity*, $\left(a \circ p', b \circ 1 - p'\right) \succ (a \circ p, b \circ 1 - p)$, which contradicts the previous indifference. Thus, $F(x)$ is unique. □

2.3. *Incentive Compatibility*

In this section we show that under the addition of *statewise monotonicity*, if a participant's subjective distribution has a quantile within the range of the quantile price list, then the participant's choices in the quantile price list will have a switching-point and that switching-point reveals an interval that must contain a quantile p of the subjective distribution F. Call x_i a *switching-point for quantile* p if the participant chooses A_{x_i} from $\{o_p, A_{x_i}\}$ but o_p from $\{o_p, A_{x_{i+1}}\}$. Below, let $F(x)$ be the CDF that solves $(a \circ F(x), b \circ 1 - F(x)) \sim \left(a \circ E_x, b \circ E_x^c\right)$ under the conditions of *proposition 1* and let $[q_l, q_h]$ be the set of p quantiles of F. That is, the set of values q that solve $F(q) \geq p$ and $F(q) \leq p$. The set of values that solves this must be an interval since F is increasing.

Proposition 2. Under *axiom 5* (*statewise monotonicity*) if $q_h < x_1$ and $q_l > x_n$ then there is at least one switching-point x_i in the sequence $(x_1, ..., x_n)$ and for any switching-point, there is some x in $[x_i, x_{i+1}]$ that is a p-th quantile of F.

Proof. Since $q_h < x_1$, $F(x_1) > p$. Similarly, since $x_n < q_l$, $F(x_n) < p$. By *objective lottery monotonicity:*

$$\left(a \circ F(x_1), b \circ 1 - F(x_1)\right) \succ (a \circ p, b \circ 1 - p) \tag{7}$$

$$(a \circ p, b \circ 1 - p) \succ \left(a \circ F(x_n), b \circ 1 - F(x_n)\right) \tag{8}$$

Since $F(x)$ is the CDF that solves $\left(a \circ F(x), b \circ 1 - F(x)\right) \sim (a \circ E_x, b \circ E_x^c)$ under the conditions of *proposition 1*, we have the following two conditions:

$$\left(a \circ E_{x_1}, b \circ E_{x_1}^c\right) \succ (a \circ p, b \circ 1 - p) \tag{9}$$

$$(a \circ p, b \circ 1 - p) \succ \left(a \circ E_{x_n}, b \circ E_{x_n}^c\right) \tag{10}$$

Thus, in the the quantile price list, A_{x_1} is chosen from $\{o_q, A_{x_1}\}$ and o_q is chosen from $\{o_q, A_{x_n}\}$. Thus, there must be some first x_i in $(x_2, ..., x_n)$ such that o_q is chosen from

$\{o_q, A_{x_i}\}$. x_{i-1} is a switching-point. We now show that any switching-point provides a range of values that must contain some p quantile of F.

By *statewise monotonicity*, the choice of A_{x_i} from $\{o_q, A_{x_i}\}$ and o_q from $\{o_q, A_{x_{i+1}}\}$ implies the following pair of preferences:

$$\left(a \circ E_{x_i}, b \circ E^c_{x_i}\right) \succsim (a \circ p, b \circ 1-p) \tag{11}$$

$$\left(a \circ E_{x_{i+1}}, b \circ E^c_{x_{i+1}}\right) \precsim (a \circ p, b \circ 1-p) \tag{12}$$

By *subject/objective replacement* and *transitivity*, these are true if and only if:

$$\Big(a \circ F(x_i), b \circ 1-F(x_i)\Big) \succsim (a \circ p, b \circ 1-p) \tag{13}$$

$$\Big(a \circ F(x_{i-1}), b \circ 1-F(x_{i-1})\Big) \precsim (a \circ p, b \circ 1-p) \tag{14}$$

By *objective lottery monotonicity* this pair of preference inequalities is true if and only if:

$$F(x_i) \geq p \geq F(x_{i-1}) \tag{15}$$

If the left inequality is weak, then $F(x_i) = p$ and so x_i is a p-th quantile of F. Similarly, if the right inequality is weak, then x_{i-1} is a p-th quantile of F. If both are strict, then $F(x_i) > p > F(x_{i-1})$ since F is increasing, there is some smallest $x \in [x_{i-1}, x_i]$ such that $F(x) > p$. This point is a p-th quantile of F.

□

2.4. *Approximating Beliefs via Maximum Entropy*

Our procedure allows researchers to collect any set of quantiles. These quantiles can be interpreted as points on the CDF of the participant's underlying subjective belief distribution, or more precisely, as intervals where these points must lie. However, we envision that researchers may want to extrapolate from the information collected by our methodology to calculate unobserved properties of these distributions (such as moments or unelicited quantiles).

Researchers often have parametric assumptions in mind that would allow them to construct a distribution from the elicited information or use a methodology like maximum likelihood to estimate the best-fitting parameters/distribution. Since these assumptions and procedures would depend on the particular research questions being studied, we will not elaborate on them here. Instead, in this section, we offer a procedure to approximate the entire distribution of subjective beliefs in a non-parametric way *using only the information provided by our methodology*.

To do this, we apply the principle of maximum entropy.[2] The principle of maximum entropy is closely related to the principle of insufficient reason. A distribution that maximizes entropy subject to constraints is the distribution that is least informative beyond the information encoded by those constraints. More informally, the process determines which

[2]The principle of maximum entropy was first formalized in Jaynes (1957a,b). We are unaware of the use of maximum entropy in a similar context (estimating a subjective belief distribution), but the principle has been used for inference elsewhere in economics (see Scharfenaker and Yang (2020) for a review) and in models of decision-making (see Kirman et al. (2023) for a recent example).

distribution uses known information and *only* the known information.[3]

2.5. *Characterization of the Maximum Entropy Approximation*

Suppose that a random variable X is known to be continuously distributed on $\left[\underline{q}, \overline{q}\right]$. Furthermore, there are n quantile values $\{q_{p_1}, q_{p_2}, ..., q_{p_n}\}$ associated with the probabilities $\{p_1, p_2, ..., p_n\}$ with each quantile value q_{p_i} constrained to be within an interval: $\left[l_{p_i}, h_{p_i}\right]$. Without loss of generality, assume $p_i > p_{i-1}$. For convenience, let $p_0 = 0$ and $p_{n+1} = 1$. The fact that the distribution has support $\left[\underline{q}, \overline{q}\right]$ can be represented by $l_{p_0} = h_{p_0} = \underline{q}$ and $l_{p_n} = h_{p_n} = \overline{q}$. The other values of l_{p_i}, h_{p_i} are determined by our elicitation methodology.

We wish to find the distribution $\tilde{F}$ and the associated density $\tilde{f}$ that maximizes the entropy function $h(f) = \int_{\underline{q}}^{\overline{q}} -f(x)ln(f(x))dx$ subject to the n restrictions $l_{p_1} \leq q_{p_i} \leq h_{p_i}$. These constraints can be interpreted geometrically to restrict F to be in the set of all increasing functions from $\left[\underline{q}, \overline{q}\right]$ to $[0, 1]$ that pass through the horizontal line segments represented by each pair $[l_{p_i}, h_{p_i}]$. This is shown in figure I.

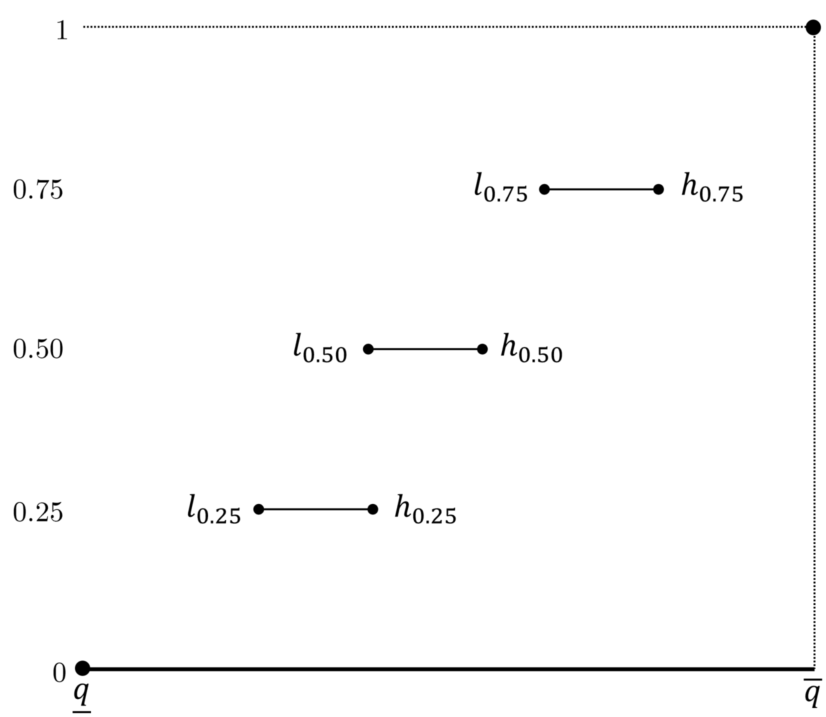

FIGURE I. Constraints represented elicited quantile intervals.

[3]We note that the principle of insufficient reason might be applied at a higher level, directly to the researchers knowledge of the participant's potential subjective belief, and instead propose that all potential subjective distributions be treated equally. However, since application yields a distribution over distributions rather than a distribution, which is our goal, we sidestep the issue and apply maximum entropy to the subjective distribution itself.

Of course, the maximum entropy distribution must have some value for each q_{p_i}. Proposition 3 shows that between these points, whatever they are, the maximum entropy distribution $\tilde{F}$ is a simple piecewise linear interpolation of the quantiles q_{p_i}. If the quantiles were known exactly, rather than constrained to be within some interval, this would be a full characterization of the maximum entropy distribution. An example of this is shown in Figure II for known quantiles $\{q_{0.25}, q_{0.50}, q_{0.75}\}$.

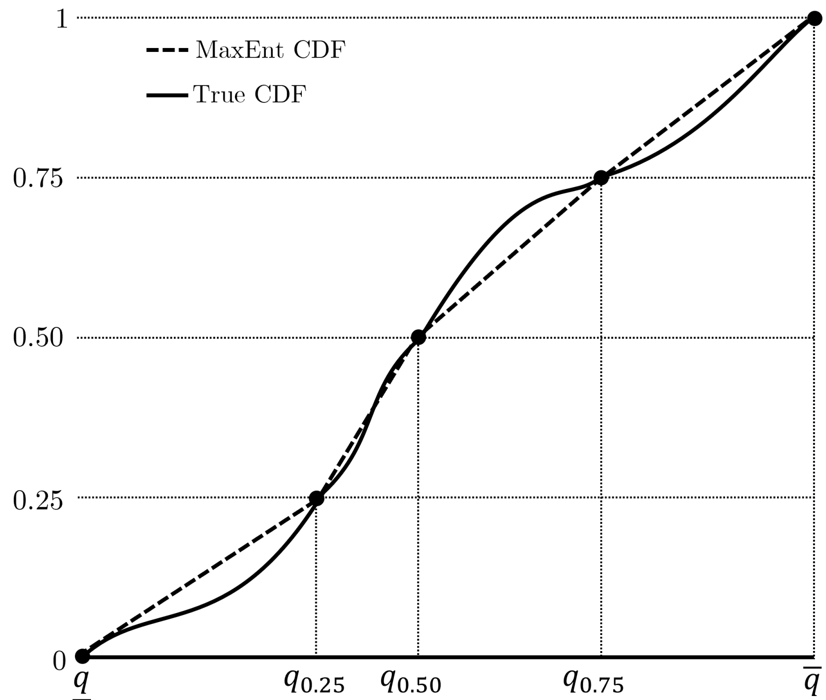

FIGURE II. Maximum Entropy CDF with known quantiles at 25%,50%,75%.

Proposition 3. The distribution $\tilde{F}$ over the range $[\underline{q}, \overline{q}]$ that maximizes entropy subject to quantile restrictions $\{q_{p_1}, q_{p_2}, ..., q_{p_n}\}$ is a piecewise linear CDF connecting the points $\left(\left(\underline{q}, 0\right), \left(q_{p_1}, p_1\right), \left(q_{p_1}, p_1\right), \left(q_{p_2}, p_2\right), ..., \left(q_{p_n}, p_n\right), \left(\overline{q}, 1\right)\right)$.

Proof. The CDF must contain the points $\left\{\left(\underline{q}, 0\right), \left(q_{p_1}, p_1\right), \left(q_{p_1}, p_1\right), \left(q_{p_2}, p_2\right), ..., \left(q_{p_n}, p_n\right), \left(\overline{q}, 1\right)\right\}$ by the constraints. To show that F is piecewise linear on the intervals between these points, it is sufficient to show that f is constant on these intervals. The range and quantile restrictions can be written as follows:

$$\int_{\underline{q}}^{\overline{q}} f(x)dx = 1, \int_{\underline{q}}^{q_{p_1}} f(x)dx = p_1, \int_{q_{p_1}}^{q_{p_2}} f(x)dx = p_2 - p_1, ..., \int_{q_{p_n}}^{\overline{q}} f(x)dx = 1 - p_n \tag{16}$$

By Theorem 12.1.1 of Cover and Thomas (2006), the f that maximizes the entropy function $h(f) = \int_{\underline{q}}^{\overline{q}} -f(x)ln(f(x))dx$ has the following form:

$$f(x) = e^{\lambda_0 - 1 + \sum_{i=1}^{n+1} \lambda_i I_{x \in (q_{p_{i-1}}, q_{p_i}]}} \tag{17}$$

We can solve for $\lambda_i i \in [1, n+1]$ using this form and the restrictions above. For each of those restrictions, we have:

$$\int_{q_{p_{i-1}}}^{q_{p_i}} e^{\lambda_0 - 1} e^{\lambda_i} dx = p_i - p_{i-1} \tag{18}$$

This simplifies to:

$$\lambda_i = ln\left(\frac{p_i - p_{i-1}}{q_{p_i} - q_{p_{i-1}}}\right) - (\lambda_0 - 1) \tag{19}$$

Plugging these back into the distribution, eliminating the term $\lambda_0 - 1$, and simplifying the resulting expression yields:

$$f(x) = \sum_{i=1}^{n+1} \left(\frac{p_i - p_{i-1}}{q_{p_i} - q_{p_{i-1}}}\right) I_{x \in (q_{p_{i-1}}, q_{p_i}]} \tag{20}$$

Thus, $f(x)$ is constant on every interval $[q_{p_i}, q_{p_{i+1}}]$, completing the proof. □

Notice that *proposition 3* does not fully characterize the maximum entropy distribution subject to the quantile intervals elicited by our methodology. However, it greatly simplifies the problem of finding this distribution. Although maximizing entropy in general is an infinite-dimensional optimization problem, *proposition 3* simplifies the problem to one of finding the values of $q_{p_2}, ..., q_{p_{n-1}}$ for which F constructed from the piecewise linear interpolation of the points $\left(\left(\underline{q}, 0\right), \left(q_{p_1}, p_1\right), \left(q_{p_1}, p_1\right), \left(q_{p_2}, p_2\right), ..., \left(q_{p_n}, p_n\right), \left(\overline{q}, 1\right)\right)$ maximizes the entropy function $h(f) = \int_{\underline{q}}^{\overline{q}} -f(x) ln(f(x)) dx$.

The problem can be solved efficiently with generic nonlinear optimization packages. Figure III shows the maximum entropy distribution subject to the quantile intervals shown in Figure I. In this case, $q_{0.25}$ and $q_{0.50}$ are maximized at corner solutions, while $q_{0.75}$ is maximized at an interior. The first-order conditions on each quantile ensure that whenever the maximum entropy function involves an interior solution for some quantile, the density is constant and equal on both sides of that quantile and therefore $\tilde{F}$ does not have a kink at that quantile.

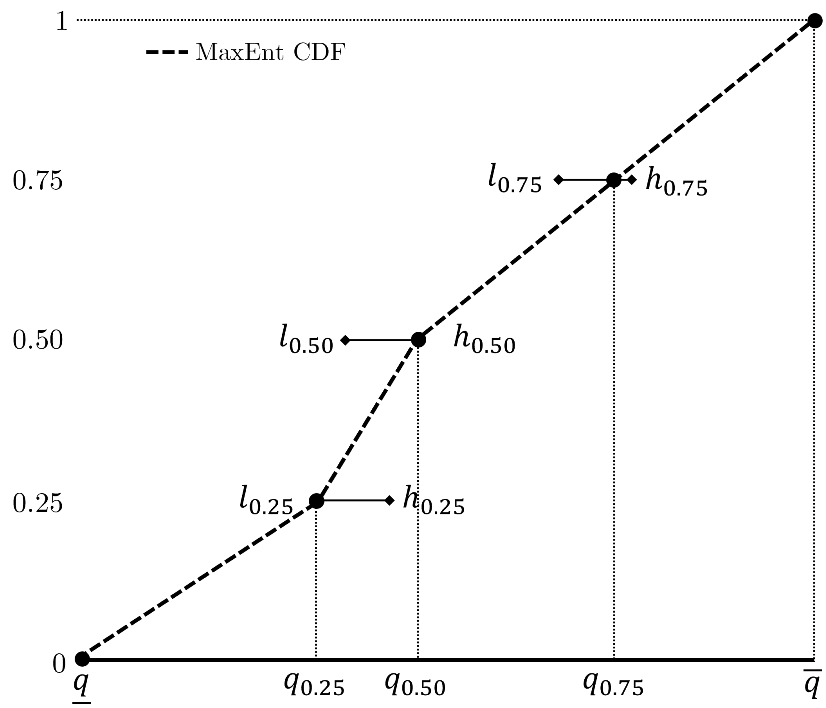

FIGURE III. Maximum Entropy CDF from elicited quantile intervals at 0.25, 0.50, 0.75.

2.6. *Approximating Moments*

Once we have approximated beliefs using the maximum entropy distribution $\tilde{F}$, it is possible to estimate properties of a belief that have not been observed. As an example, suppose that we want to calculate the belief about the mean of X. This can be done using the CDF and the fact that $E(X) = \int_{-\infty}^{\infty} 1 - F(x)dx$ or using the approximate density $\tilde{f}$ which has the density $\frac{p_i - p_{i-1}}{q_{p_i} - q_{p_{i-1}}}$ in each interval. Thus, the mean of the approximated distribution can be calculated by the following sum.

$$
\begin{aligned}
(21) \qquad & \int_{\underline{q}}^{q_{p_1}} x\left(\frac{p_1}{q_{p_1} - \underline{q}}\right)dx + \int_{q_{p_1}}^{q_{p_2}} x\left(\frac{p_2 - p_1}{q_{p_2} - q_{p_1}}\right)dx + \ldots \\
& + \int_{q_{p_{n-1}}}^{q_{p_n}} x\left(\frac{p_n - p_{n-1}}{q_{p_n} - q_{p_{n-1}}}\right)dx + \int_{q_{p_n}}^{\overline{q}} x\left(\frac{1 - p_n}{\overline{q} - q_{p_n}}\right)dx
\end{aligned}
$$

As an example, suppose we have quantiles $q_{0.25}$, $q_{0.50}$, and $q_{0.75}$. This sum simplifies to: $\sum_{i=1}^{n+1} \frac{1}{8}(q_i + q_{i-1})$. This provides the following expression for $\tilde{\mu}$ in terms of a, b and the quantiles:

$$
(22) \qquad \tilde{\mu} = \frac{1}{8}\left(\underline{q} + \overline{q}\right) + \frac{1}{4}(q_{0.25} + q_{0.50} + q_{0.75})
$$

3. Comparison to Other Methodologies

The choice of methodology to elicit subjective beliefs involves two main decisions: what to elicit and how to elicit it. It is possible to elicit any aspect of beliefs (Lambert et al., 2008).[4] The decision of what to elicit ultimately depends on the information required by the research question. We propose eliciting quantiles as a flexible way to gather information about a real-valued distribution when the required information is about the distribution itself rather than about the probabilities of particular events. For example, if a researcher wants to know the probability that a participant believes the value of a random variable is greater than some fixed quantity, eliciting a probability is appropriate. If a researcher wants to know the values that a participant believes are in the upper and lower 5% tails of a distribution, then quantiles are more appropriate.[5]

It is also possible to elicit moments of a real-valued subjective distribution. However, in contrast to quantiles, moments are not a fundamental part of a distribution but are derived from it. Furthermore, moments may not exist for certain extreme distributions. More practically, because moments are not event-based but rather a summary of the entire distribution, the incentives involved in eliciting moments are more complex than those for eliciting quantiles. This also makes them harder to communicate to participants when researchers want to tell participants exactly what they are eliciting.

Moments, unlike quantiles, cannot always be elicited in isolation. For example, it is impossible to elicit the variance of a subjective distribution without learning about the mean (Lambert et al., 2008). Finally, information about moments can be approximated without parametric assumptions given information about quantiles, but the reverse is not true. For these reasons, we believe that it is almost always more appropriate to elicit quantiles rather than moments when research questions depend on information about real-valued subjective distributions.

For research questions that *specifically require the elicitation of moments* and where the researcher is not satisfied with estimating these from the elicited quantiles, scoring rules can be used. Scoring rules are a flexible way to elicit quantities of a distribution. Participants provide their estimate of the target value, such as the mean, which is then compared to the truth or a sample from the true distribution. It is possible to tailor the scoring rules to elicit a wide range of information about a distribution (Gneiting and Raftery, 2007; Lambert, 2019). In experimental economics, it is common to "binarize" scoring rules to eliminate the effect of risk preferences on biasing elicited beliefs. In binarized scoring rules, the probability of a reward, rather than the value of the reward, changes with the accuracy of the reported beliefs (Savage, 1971; Hossain and Okui, 2013; Harrison et al., 2014).

[4] Not all aspects can be elicited *directly* and must either be inferred through directly elicited quantities or elicited along with other quantities used in their calculation.

[5] When a research question involves eliciting probabilities of specific events, scoring rules for probabilities or price list methods can be used. We note that these methods can also be used to provide a comprehensive view of a belief distribution by eliciting probabilities throughout the range of the random variable. This is akin to eliciting the PDF rather than the CDF of a distribution. This may provide a reasonable and simple alternative to eliciting several quantiles when general information about a distribution is required.

Researchers convinced by our arguments on eliciting quantiles may still be skeptical that eliciting quantiles using price lists is the best solution. Like moments and probabilities, quantiles can also be elicited using scoring rules. However, unlike our price list methodology, scoring rules require participants to provide their beliefs about the target quantities directly. This may be more cumbersome for participants since the scoring rules do not provide any guidance on how to discover these values. For example, Eyting and Schmidt (2021) propose eliciting quantiles with the following language (in this case, to elicit the 0.75 quantile): "What do you say is X if underestimation is four times more costly than overestimation?" Although this is incentive compatible, participants are left to determine how to arrive at X that maximizes their utility under this rule. Even if researchers tell participants what they are eliciting, such as in Dustan et al. (2022), who tell participants their incentives are designed to elicit the median of their belief, participants must still arrive at what their belief about the median is without any direction. In contrast, our price list methodology leverages the event-based nature of a quantile and allows participants to “discover” their belief through a series of simpler questions about event probabilities.

In conclusion, we believe that eliciting quantiles with price lists is a simple, intuitive, and practical approach, especially when the research focus is on understanding particular aspects of participants’ real-valued subjective belief distributions. The flexibility of quantile elicitation, in contrast to the complexities and limitations associated with moments, makes it a superior choice in most research scenarios. Furthermore, the use of price list methodologies, as opposed to more complex scoring rules, further simplifies the elicitation process, making it more accessible to participants in complex scenarios. These benefits highlight the potential “behavioral incentive compatibility” of our methodology in addition to the theoretical incentive compatibility (Danz et al., 2022).

4. Experimental Design

The purpose of this experiment is to demonstrate one possible application of our methodology and simultaneously investigate how subjective belief distributions are affected by eliciting both quantiles and mean of the distribution.

The event that participants considered in our experiment is whether individuals pass or fail a math task. Before the experiment, 20 Ohio State students completed a math task. Each student was shown a sequence of five two-digit numbers. Students had ten minutes to correctly calculate the sum of as many of these sets of two-digit numbers as they could. For each correctly answered math problem, they received \$1. The students also received a \$5 show-up fee. The performance of these students on the math task is the basis of the main experiment.

We categorized each of the 20 students according to how many math problems they correctly answered. If a student correctly answered at least 13 math problems, we categorized them as successful in the math task. If they answered any number less than thirteen of the math problems correctly, we categorized them as failing the math task. All students who participated in the math task knew that their answers would be seen later by other participants but did not know about the categorization of success. We chose not to disclose the categorization to the students in order to not change their performance goal on the task.

For the main experiment, we used the performance of the 20 students on the math task as our subjective event of interest. Participants in the main experiment were also recruited through and participated in the experiment at the Ohio State Experimental Economics Laboratory. Anyone who participated in the previous math task was excluded from participation in the experiment. We chose to do both the math task and the experiment in person at the Ohio State laboratory so that the participants in the experiment would have some baseline facts they could use to form beliefs over the students' performance in the math task. Performance in the math task is a subjective event because before knowing the actual results, the participants in the experiment cannot find the correct answer, and there is no way to look up the results. Furthermore, we wanted participants to be able to form their beliefs about math task performance using facts they knew so that their beliefs were not completely uninformed. Information like how well students did in their math classes or how smart they believe the general population of the university to be are facts that can be used by our participants to form their beliefs about other students' performances on the prior math task.

The participants were first told about the math task and given an example of the math problems. The participants then told us their beliefs about the students' performance on the math task. Their beliefs were elicited through their switching-points on four separate Multiple Price Lists (MPLs). Three of the four MPLs were *quantile price lists* targeted at the 0.25, 0.50, and 0.75 quantiles of their beliefs. The fourth MPL elicited the mean of their subjective belief distribution via the probability a randomly chosen student passed. This MPL is not part of our methodology. Instead, it is based on a popular methodology for eliciting *subjective probabilities* first introduced by Holt and Smith (2016). Each Multiple Price List consisted of 21 rows. In each row, there were two options, Option A and Option B. Participants decided which of the two options they preferred in each row.

In each of the quantile elicitations, participants were asked how many students they believed had succeeded in the math task. The right-hand side of each quantile price list, or Option B, was constant. For the 0.25 quantile, Option B was "Win \$10 with a 25% chance" in each row. For the 0.5 quantile, Option B was "Win \$10 with a 50% chance" in each row. And for the 0.75 quantile, Option B was "Win \$10 with a 75% chance" in each row. The left-hand side, or Option A, of all three quantile price lists were identical. Option A was changed in each row. Option A in the first row of all three lists was "Win \$10 if ≤ 20 people succeeded in the math task." In each row, the number of people who succeeded in the math task decreased by 1 down to "Win \$10 if 0 people succeeded in the math task" in the 21st row. An example of the 0.25 quantile price list is shown in Figure IV.

The switching-point of each of these lists elicits a value of different points along a participant's belief distribution. From the three quantile price lists, we get bounds on the 0.25, 0.5, and 0.75 quantiles of each participant's belief distribution. These three quantiles allow us to build a approximation of each participant's entire subjective belief distribution using our maximum entropy methodology. In all MPLs, a single switching-point was enforced.

Would you rather have:

	Option A	or	Option B
Q1.	\$10 if ≤ 20 people passed	or	\$10 with a 25% chance
Q2.	\$10 if ≤ 19 people passed	or	\$10 with a 25% chance
Q3.	\$10 if ≤ 18 people passed	or	\$10 with a 25% chance
Q4.	\$10 if ≤ 17 people passed	or	\$10 with a 25% chance
Q5.	\$10 if ≤ 16 people passed	or	\$10 with a 25% chance
Q6.	\$10 if ≤ 15 people passed	or	\$10 with a 25% chance
Q7.	\$10 if ≤ 14 people passed	or	\$10 with a 25% chance
Q8.	\$10 if ≤ 13 people passed	or	\$10 with a 25% chance
Q9.	\$10 if ≤ 12 people passed	or	\$10 with a 25% chance
Q10.	\$10 if ≤ 11 people passed	or	\$10 with a 25% chance
Q11.	\$10 if ≤ 10 people passed	or	\$10 with a 25% chance
Q12.	\$10 if ≤ 9 people passed	or	\$10 with a 25% chance
Q13.	\$10 if ≤ 8 people passed	or	\$10 with a 25% chance
Q14.	\$10 if ≤ 7 people passed	or	\$10 with a 25% chance
Q15.	\$10 if ≤ 6 people passed	or	\$10 with a 25% chance
Q16.	\$10 if ≤ 5 people passed	or	\$10 with a 25% chance
Q17.	\$10 if ≤ 4 people passed	or	\$10 with a 25% chance
Q18.	\$10 if ≤ 3 people passed	or	\$10 with a 25% chance
Q19.	\$10 if ≤ 2 people passed	or	\$10 with a 25% chance
Q20.	\$10 if ≤ 1 people passed	or	\$10 with a 25% chance
Q21.	\$10 if = 0 people passed	or	\$10 with a 25% chance

FIGURE IV. MPL that elicits the 0.25 quantile.

The fourth MPL elicited participants' beliefs about the probability that a randomly chosen student passed the math task. The right-hand side of this list, or Option B, was constant across all rows: "Win \$10 if one randomly selected person succeeded on the math task." The left-hand side, or Option A, changed in each row. Option A in the first row was "Win \$10 with 100% chance." In each row, the probability that participants received \$10 in Option A decreased by 5% down to "Win \$10 with 0% chance" in the 21st row. The exact MPL used is shown in Figure V.

Participants were randomly assigned to one of two treatments that only differed in the order of the four MPLs. In the *mean-first* treatment, the mean eliciting MPL was shown first. In the *mean-last* treatment, the three quantile price lists were shown first. In both treatments, the order of the three quantile price lists was randomized at the participant level.

One of the four MPLs was randomly chosen to determine payment. For the randomly chosen MPL, a row was then randomly chosen. The participant's choice in that row was used to determine final payment. If the chosen option was winning \$10 with a certain probability (Option B for quantile price list and Option A for mean MPL), a 100-sided die was rolled. If the number on the die was lower than or equal to the probability stated in that row, the

Would you rather have:

	Option A	or	Option B
Q1.	$10 with a 100% chance	or	$10 if one randomly chosen subject passed
Q2.	$10 with a 95% chance	or	$10 if one randomly chosen subject passed
Q3.	$10 with a 90% chance	or	$10 if one randomly chosen subject passed
Q4.	$10 with a 85% chance	or	$10 if one randomly chosen subject passed
Q5.	$10 with a 80% chance	or	$10 if one randomly chosen subject passed
Q6.	$10 with a 75% chance	or	$10 if one randomly chosen subject passed
Q7.	$10 with a 70% chance	or	$10 if one randomly chosen subject passed
Q8.	$10 with a 65% chance	or	$10 if one randomly chosen subject passed
Q9.	$10 with a 60% chance	or	$10 if one randomly chosen subject passed
Q10.	$10 with a 55% chance	or	$10 if one randomly chosen subject passed
Q11.	$10 with a 50% chance	or	$10 if one randomly chosen subject passed
Q12.	$10 with a 45% chance	or	$10 if one randomly chosen subject passed
Q13.	$10 with a 40% chance	or	$10 if one randomly chosen subject passed
Q14.	$10 with a 35% chance	or	$10 if one randomly chosen subject passed
Q15.	$10 with a 30% chance	or	$10 if one randomly chosen subject passed
Q16.	$10 with a 25% chance	or	$10 if one randomly chosen subject passed
Q17.	$10 with a 20% chance	or	$10 if one randomly chosen subject passed
Q18.	$10 with a 15% chance	or	$10 if one randomly chosen subject passed
Q19.	$10 with a 10% chance	or	$10 if one randomly chosen subject passed
Q20.	$10 with a 5% chance	or	$10 if one randomly chosen subject passed
Q21.	$10 with a 0% chance	or	$10 if one randomly chosen subject passed

FIGURE V. MPL that elicits probability a randomly chosen participant passed the math task.

participant won $10. If the chosen option was winning $10 dependent on how students performed in the math task (Option A for quantile price lists and Option B for mean MPL), the actual performance of the students was used to determine payment. For the three quantile price lists, if the number of participants who actually succeeded was less than or equal to the number stated in the row, participants won $10. For the mean MPL, a student was randomly chosen from the math task. If this student succeeded, participants won $10.

The experimental design not only tests the methodology proposed in this paper, but also allows us to test whether answering the quantile price lists helped participants "discover" their subjective distribution. Using the two treatments, we can determine whether reporting mean belief before or after quantile beliefs allowed for more consistency across the four elicitations. Additionally, we test whether the order of the MPLs made participants' belief distributions "more binomial". These are just some possible applications of the methodology that we find particularly interesting, but much more could be done using it.

All sessions were recruited through and run at the Experimental Economics Laboratory at The Ohio State University. Twenty students participated in the math task. Each student received a $5.00 show-up fee and $1.00 for every question answered. The math sessions took approximately 15 minutes and the average payment was $7.40. 158 participants were recruited for the main experiment. 71 participants were randomly assigned to mean-first and 87 participants were randomly assigned to mean-last. The experiment took 10 minutes and the participants received a \$5.00 show-up fee plus average payment of \$3.29.

5. Results

As our main motivation of the experiment is to demonstrate our methodology, we first report results for the three quantile price lists. As twenty students participated in the math task participant's beliefs are bounded between 0 and 20. In order for our methodology to work appropriately, the quantile beliefs of the participants must be monotonic. This means that the number of students participants believe passed at the 0.25 quantile must be less than or equal to the 0.50 quantile, which must be less than or equal to the 0.75 quantile. These restrictions ensure an increasing approximated cumulative distribution function. We exclude participants who do not have monotonic beliefs from all analyses.

73% of the participants have monotonic beliefs. Participants in the mean-first or mean-last treatment do not exhibit differences in whether their beliefs are monotonic. 69.01% of the mean-first and 75.86% of the mean-last participants have monotonic beliefs. A chi-squared test for the difference of these proportions is not significant with a p-value of 0.3360. It does not seem that allowing participants to report their mean first or last changes their ability to report beliefs that are consistent with an increasing subjective CDF.

Next, we show the actual beliefs the participants reported for each of the quantile price lists. A switching-point provides a range of possible beliefs. In our quantile price lists, the difference between each row is 1 student. Thus, a switching-point in the 10th row represents a range of 10-11 for the relevant quantile. Instead of plotting the full range, we have chosen a single point within the range to represent the beliefs of the participants. The belief chosen is the point in the range that maximizes the entropy of the CDF approximation. [6] Figure VI shows a summary of the beliefs of all participants for each elicited quantile. The shaded region provides a 50% probability interval for approximated CDFs. Put another way, for any quantile, 50% of participants' approximated CDFs fall within the shaded region.

[6] As our ranges are quite small, the choice of what point to use does not have a big impact on our results. Using the midpoint or upper/lower bounds of this range does not result in notable differences in this analysis.

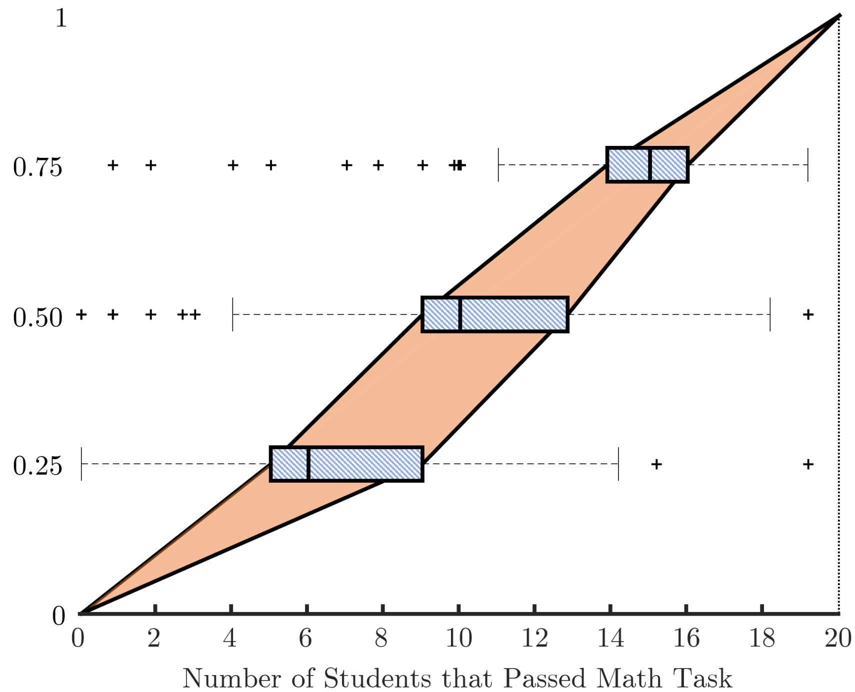

FIGURE VI. Box plots of each of the elicited quantile price lists. The shaded region represents CDFs approximated for the aggregate data.

The 50% probability region for each quantile is quite distinct, with very little overlap between the three elicited quantiles. There is a substantial amount of consistency, especially at the aggregate level. Furthermore, only 4% of the participants have the same switching-point in all three quantile price lists. This further supports the existence of an underlying distribution over this domain rather than a point-precise belief. Interestingly, the 0.75 quantile has a tighter range than the other quantiles. Most of the participants have a belief in the range of 14 to 16 for the 0.75 quantile. The interquartile ranges of beliefs elicited for the other two quantiles are double this size, with much wider whiskers.

Using the participants' reported beliefs, their entire subjective belief distributions about the performance of students on the math task are approximated using *Proposition 3*. An example (in this case for participant number 3) of an approximated CDF using our methodology is shown in Figure VII.

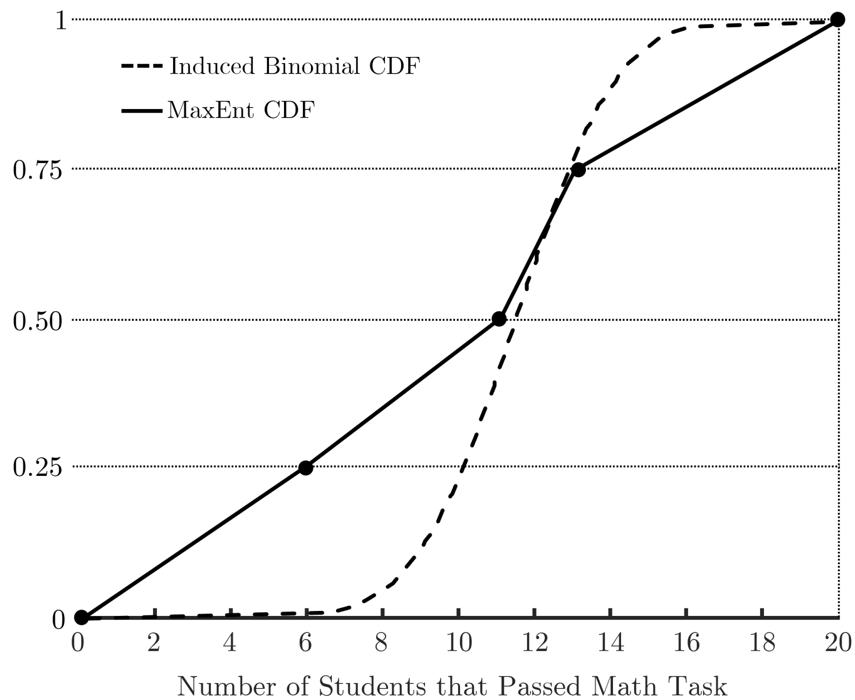

FIGURE VII. Participant 3's approximated subjective belief distribution using the quantile price lists.

Using this approximation, we can recover different moments of the distribution. Here, we focus on the mean. After finding the mean from the approximated CDFs, we determine how the directly elicited means differ from these approximated means.

The directly elicited means are determined from the mean MPL by assuming that participants believe that student performances are independent. Under this assumption, the mean number of students who pass is twenty times the probability a randomly chosen student passes. We call this the *elicited mean*. The *approximated mean* is instead calculated from the approximated subjective belief distributions using equation (18). Figure VIII shows a histogram of the difference between the elicited mean and approximated mean for each participant in the two treatments.

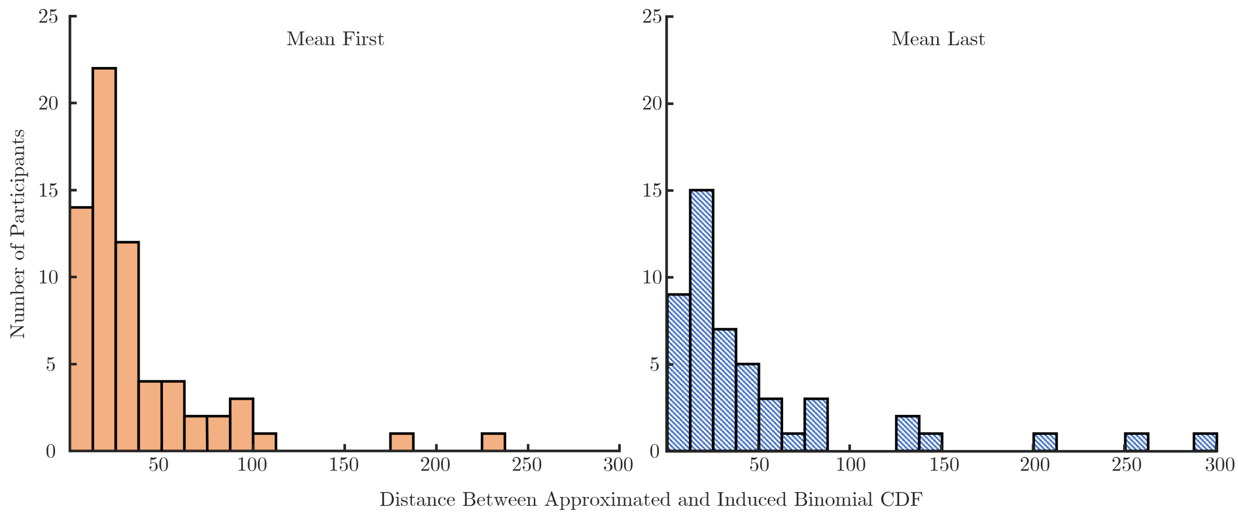

FIGURE VIII. Distance of elicited mean from approximated mean for each participant by treatment.

We find that the elicited means and the approximated means are not too different at the aggregate level. For the mean-first treatment, the average difference is -1.71 and for the mean-last treatment, the average difference is -0.92. The participant's elicited means are on average slightly lower than the approximated means, although a Wilcoxon rank-sum test for the differences in these distributions finds no significant difference (p-value 0.5526). The mean-first treatment, where the mean MPL is answered prior to the quantiles, is skewed slightly left compared to the mean-last treatment. This provides weak evidence that eliciting means after quantiles leads to greater consistency with approximated CDFs.

In addition to the estimation of moments, we can also use our methodology to test hypotheses about how participants form their beliefs. For instance, assuming that participants view the performance of each student on the math task as independent and have a point belief about the probability that a randomly chosen student passed, the distribution for the number who passed is binomial. However, if participants have higher-level uncertainty about the probability that a randomly chosen student will pass, the approximated CDFs would be flatter than the binomial distribution.[7]

Using our data, we can compare the approximated CDFs of each participant with the induced binomial distributions using the reported beliefs about the probability a randomly chosen student passed. An example comparing the approximated CDF with the induced binomial CDF is shown in Figure VII.

To compare the relative flatness of the approximated CDF with the induced binomial distributions, we compare the 0.25 and 0.75 quantiles of the two distributions for each participant. Figure IX shows the histograms of these differences.

[7]For instance, if participant has a beta-distributed belief about the probability a randomly chosen student passed, then their belief about the number of students who passed would have a beta-binomial distribution.

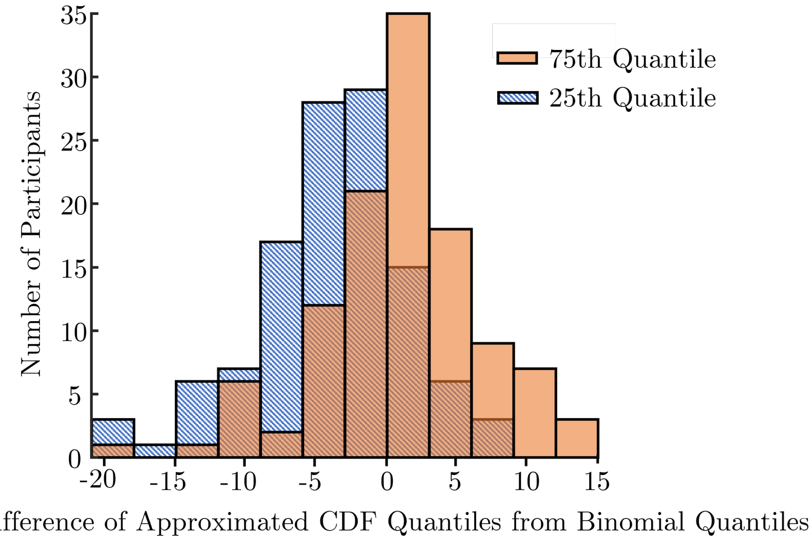

FIGURE IX. Differences between the approximated CDF 0.25 and 0.75 quantiles and the respective induced binomial 0.25 and 0.75 quantiles for each participant.

We find that the 0.25 quantile is below and the 0.75 quantile is above the binomial baseline for 44% of the participants, indicating a strong tendency for participants to have "flatter" distributions than can be explained by a binomial belief about the number of students who pass. This pattern also holds between treatments. For both mean-first and mean-last, we find that differences in the elicited 0.25 and 0.75 quantiles from the binomial distribution are indicative of participants having underlying distributions that are not binomial in nature. There is no significant difference in the differences of the 0.25 quantile or the 0.75 quantile between treatments.[8]

To additionally test the difference between the induced binomial distributions and the approximated distributions, we can look at the distance between the binomial CDF and the approximated CDF. To measure this distance, we numerically integrate the squared difference between the distributions over the domain $[0,20]$. If the approximated CDFs are close to binomial, the distance between the binomial CDF and the approximated CDF will be close to 0. For reference, a perfectly linear approximated CDF with the three elicited quantiles at 5, 10, 15 respectively, is distance 13.2592 from a binomial distribution with mean 10.

We do not find evidence that the timing of the mean MPL affects our distance measure between the approximated and induced binomial distributions.[9] Figure X shows a histogram of this distance measure for each participant (pooled treatments).

[8]Wilcoxon rank-sum test result in p-values of 0.5394 and 0.5621 respectively.

[9]Wilcoxon rank-sum test for the difference in these distributions has a p-value of 0.3113.

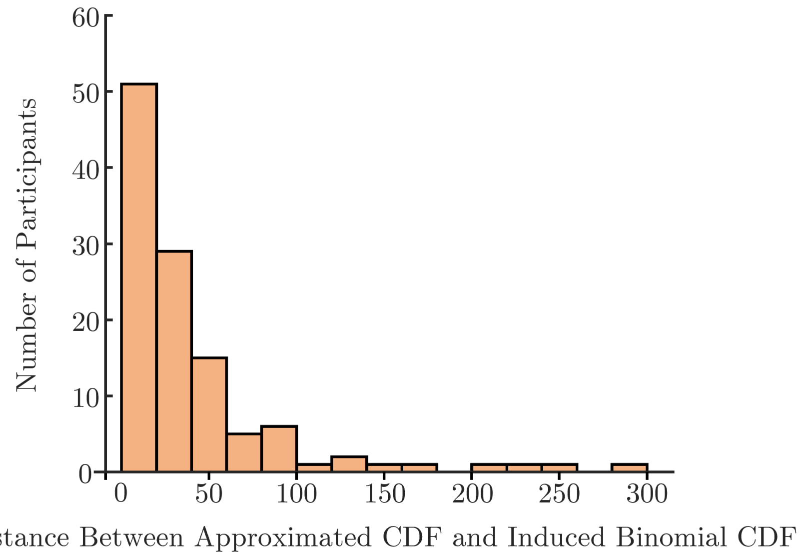

FIGURE X. Distance between the approximated CDF and the respective binomial CDF for each subject (pooled treatments).

We find that the the majority of participants have a distance between their approximated and induced binomial distributions that is less than 35. Additionally, we find a large spike between 10 and 14, suggesting that the approximated CDFs are close to linear.[10] This would suggest that instead of binomial distributions underlying beliefs, participants have flatter, more linear distributions. This is consistent with participants having higher-level uncertainty about the difficulty of the math task.

Together, these results show that eliciting just a mean belief may not be sufficient to understand the entire subjective CDF. Ultimately, eliciting different quantiles of the belief distribution gives a fuller understanding of the participant's beliefs, and we hope our analysis and discussion in this section demonstrates how our methodology can help researchers peer deeper into the minds of participants.

6. Discussion

In this paper, we propose a robust and straightforward methodology for eliciting subjective real-valued beliefs. Our methodology allows practitioners to pin-point aspects of beliefs that are relevant to their research question and to build a picture of participants entire belief distributions by eliciting several quantiles.

Given the flexibility of this methodology, there are potential applications in many areas of economics. One potential application is to obtain the tails of a belief distribution. In many situations, models predict similar central-tendencies but large differences in tails. Using our methodology, the tails of the distributions can be easily elicited to test hypotheses generated from models of this type. For example, beliefs about men and women are often compared on

[10]Although only 6% of participants have perfectly linear beliefs.

average, where important differences may be in terms of the shape or variances of distributions. The greater male variability hypothesis suggests that men are often more variable across traits than women (Thöni and Volk, 2021; Thöni et al., 2021). Since variance (as a moment) is not elicitable on its own (see Lambert et al. (2008)), our methodology likely represents the simplest way to test hypotheses in terms of beliefs about this type of difference.

Another potential application is to obtain richer beliefs that can be used for forecasting. For example, two experts might have beliefs about future inflation that are similar in terms of their mean but significantly different in terms of shape. Our methodology allows researchers to capture the full knowledge and uncertainty of experts on important indicators. Similarly, the potential of extreme events often drives decision-making in agriculture and insurance, and our methodology allows researchers to easily learn about beliefs about extremes by focusing on quantiles in the tail.

References

Azrieli, Y., Chambers, C.P., Healy, P.J., 2018. Incentives in experiments: A theoretical analysis. Journal of Political Economy 126, 1472–1503.

Cover, T.M., Thomas, J.A., 2006. Elements of Information Theory 2nd Edition (Wiley Series in Telecommunications and Signal Processing). Wiley-Interscience.

Danz, D., Vesterlund, L., Wilson, A.J., 2022. Belief elicitation and behavioral incentive compatibility. American Economic Review 112, 2851–2883.

Dustan, A., Koutout, K., Leo, G., 2022. Second-order beliefs and gender. Journal of Economic Behavior Organization 200, 752–781. URL: https://www.sciencedirect.com/science/article/pii/S0167268122002190, doi:https://doi.org/10.1016/j.jebo.2022.06.023.

Eyting, M., Schmidt, P., 2021. Belief elicitation with multiple point predictions. European Economic Review 135, 103700. URL: https://www.sciencedirect.com/science/article/pii/S0014292121000532, doi:https://doi.org/10.1016/j.euroecorev.2021.103700.

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.

Holt, C.A., Smith, A.M., 2016. Belief elicitation with a synchronized lottery choice menu that is invariant to risk attitudes. American Economic Journal: Microeconomics 8, 110–139.

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

Jaynes, E.T., 1957a. Information theory and statistical mechanics. Physical Review 106, 620–630. doi:10.1103/PhysRev.106.620.

Jaynes, E.T., 1957b. Information theory and statistical mechanics ii. Physical Review 108, 171–190. doi:10.1103/PhysRev.108.171.

Kirman, A., Laisney, F., Pezanis-Christou, P., 2023. Relaxing the symmetry assumption in participation games: a specification test for cluster-heterogeneity. Experimental Economics , 1–29.

Lambert, N.S., 2019. Elicitation and evaluation of statistical forecasts. Preprint.

Lambert, N.S., Pennock, D.M., Shoham, Y., 2008. Eliciting properties of probability distributions, in: Proceedings of the 9th ACM Conference on Electronic Commerce, pp. 129–138.

Savage, L.J., 1971. Elicitation of personal probabilities and expectations. Journal of the American Statistical Association 66, 783–801.

Scharfenaker, E., Yang, J., 2020. Maximum Entropy Economics: Foundations and Applications. Springer.

Thöni, C., Volk, S., 2021. Converging evidence for greater male variability in time, risk, and social preferences. Proceedings of the National Academy of Sciences 118, e2026112118.

Thöni, C., Volk, S., Cortina, J.M., 2021. Greater male variability in cooperation: Meta-analytic evidence for an evolutionary perspective. Psychological Science 32, 50–63.