Learning from Survey Density Forecasts*

Martín Almuzara† Richard K. Crump§
Dennis Kongoli‡ Gizem Koşar¶

March 2026

[Link to the latest version]

Abstract

We propose a new framework to recover nonparametric upper and lower bounds on functionals of survey forecast densities—e.g., the mean, median, variance and interquartile range. By imposing only mild smoothness constraints on the unknown density, our approach delivers informative bounds that automatically accommodate features such as skewness and multimodality. This circumvents the need to assume a specific parametric family of densities when working with density forecast data. We also show how to incorporate the bounds into a subsequent regression model in the role of a dependent variable, an independent variable or both. We provide highly efficient algorithms to compute the bounds and adjust regression estimates and inference procedures. We illustrate our methodology in an application to the consumption Euler equation using data from the NY Fed's Survey of Consumer Expectations.

Keywords: Survey density forecasts, expectations, density elicitation.

*We have greatly benefited from discussions with Keshav Dogra, Marco Del Negro, Mikkel Plagborg-Møller, Giorgio Topa and Wilbert Van Der Klaauw. All errors are our own. The views expressed in this paper do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

†FRBNY: martin.almuzara@ny.frb.org
‡FRBNY: dennis.kongoli@ny.frb.org
§FRBNY: richard.crump@ny.frb.org
¶FRBNY and CESIfo: gizem.kosar@gmail.com

1 Introduction

The use of survey data in economics has expanded considerably in recent years (Bachmann, Topa, and van der Klaauw, 2022). A key feature of many household, firm, and professional forecaster surveys is that they elicit not only respondents' point forecasts but also information about their subjective forecast distributions.[1] The appeal of density responses over point forecasts is that they not only provide measures of central tendency but also characterize aspects of the full distribution of beliefs which enables, for example, tests of model predictions about uncertainty and disagreement. Moreover, a probabilistic response can be used to quantify subjective uncertainty via the variance or interquartile range, and to detect features such as skewness and multimodality. Point forecasts, on the other hand, are inherently ambiguous to interpret since it is unclear what target the respondent is reporting (e.g., a mean, median, or mode?).

In this paper, we introduce a novel approach to obtain nonparametric upper and lower bounds of functionals of survey-elicited forecast densities. Rather than committing to a parametric family, our approach only requires that the unknown density be sufficiently smooth, which is generally a baseline assumption in nonparametric estimation. This single mild restriction defines a convex class of densities that is consistent with the reported survey response probabilities, and lower and upper bounds on any linear functional of the density are then obtained by solving the extremal problem over this class. The resulting bounds are tighter than the bounds proposed by Engelberg, Manski, and Williams (2009) (hereafter, EMW), which are obtained without imposing any structure, and can be recovered as a special case of our approach, while remaining strictly more flexible than a parametric point estimate. In particular, our bounds automatically accommodate features such as multimodality and skewness; they are always consistent with the reported probabilities exactly; and they can be modified in a straightforward way to accommodate rounding and measurement error. We cover a broad range of functionals: measures of central tendency (mean, median), measures of uncertainty (variance, interquartile range), the full CDF, and quantile functions. We prove that the infinite-dimensional extremal problem can be approximated arbitrarily well by a finite-dimensional linear program obtained by discretizing the support, and that the

[1] The Survey of Professional Forecasters (SPF), one of the earliest macroeconomic surveys, has long asked respondents to assign probabilities to intervals of future inflation and output growth. The NY Fed's Survey of Consumer Expectations (SCE) and the Survey of Primary Dealers (SPD) similarly include density questions about inflation, income, and spending growth. Similar formats appear in household surveys such as the Bank of Italy's Survey of Household Income and Wealth.

discrete solution converges to the continuous-problem solution as the grid refines, providing a rigorous foundation for the numerical approach. Efficient linear programming solvers make the method scalable to large survey samples.

Density forecasts are typically elicited through one of two formats: respondents either assign probabilities to a fixed partition of the support into a small number of intervals, the "bin-based" format, or they report the value of their subjective CDF at a set of specified quantiles. In both cases, the response fixes the subjective CDF at only a finite number of points, to keep the cognitive load to the respondent limited. This implies that the full density, or any smooth functional of it, is not identified without further assumptions. The dominant approach to operationalizing the responses to density forecast questions has been to fit a parametric density with a small number of parameters to the reported bin probabilities of each individual. Following EMW, relying on a generalized beta family of distributions has become the de facto standard in applied work. While fitting a parametric density achieves point identification, the fitted density may not closely approximate the true survey responses (see Section 2 for empirical examples). Moreover, even if the fitted density matches the reported bin probabilities exactly, there may be many other densities with very different properties that are *equally* consistent with the observed responses. This non-uniqueness is fundamental: a survey response with K bins generally fixes the CDF at only $K + 1$ points, leaving an infinite-dimensional family of densities compatible with the data. Any specific parametric choice thus imposes untestable restrictions on the shape of the density, and functionals such as the variance or higher-order moments can differ substantially across members of the identified set.

We also show how to incorporate the resulting interval-valued estimates into subsequent regression analyses. Many empirical applications involve regressing one density-based moment on another or on a directly observed outcome to recover a structural parameter. When both the dependent variable and the regressor are only known to lie within an interval, standard regression estimators are inconsistent. We study three cases in a linear panel data model: (i) the outcome is interval-valued, and the regressor is observed; (ii) the regressor is interval-valued, and the outcome is observed; and (iii) both are interval-valued. In each case, we characterize the identified set for the OLS coefficient and propose computationally feasible algorithms to recover the set estimate and construct valid confidence intervals, extending existing partial identification results to accommodate individual fixed effects and additional controls.

Finally, we illustrate our approach with an empirical application of estimating the con-

sumption Euler equation using the SCE Household Spending Survey. Survey-based subjective moments have attractive properties for estimating Euler equations because they circumvent the challenging identification problems, such as the need to instrument realized consumption growth, that arise with realized panel data (Christelis, Georgarakos, Jappelli, and van Rooij, 2020; Crump, Eusepi, Tambalotti, and Topa, 2022). Estimation of the Euler equation requires both the first and second moments of the subjective spending growth distribution, so both the dependent variable and the key regressor are density-based, placing us in the third case of our regression framework. Prior work has either relied on realized consumption data (Dynan, 1993; Merrigan and Normandin, 1996) or assumed a parametric density, such as the triangular or beta, for the subjective distribution (Bertola, Guiso, and Pistaferri, 2005; Christelis et al., 2020). Our nonparametric approach sidesteps the misspecification concerns associated with these assumptions. We obtain set estimates for the coefficient of relative prudence, ρ, ranging from approximately 5.4 to 8.5 across our main specifications, suggesting a strong precautionary savings motive for U.S. households. These estimates exceed those reported by Christelis et al. (2020) and Bertola et al. (2005), and are more in line with the calibration value of 4 used by Hubbard, Skinner, and Zeldes (1994), suggesting that parametric assumptions can meaningfully distort inference in this setting.

Our paper is most closely related to the seminal work of EMW, who studied the information content of bin-probability responses in the SPF. EMW proposes two approaches. The first, which we refer to as "hands-off" bounds, uses only the reported bin probabilities to bound the mean, median, and mode of the subjective distribution without any distributional assumptions. These bounds place the functional of interest between the extreme values attainable by moving all probability mass within each bin to the top or bottom of the interval. However, when the outermost bins are open-ended, the hands-off bounds on the mean are uninformative. Our contribution is to find a middle ground between these two approaches: we exploit smoothness conditions on the density to obtain bounds that are strictly tighter than the hands-off bounds while avoiding the untestable restrictions inherent in any parametric specification. In addition, our framework covers a broader set of functionals, including the variance, the interquartile range, and the full CDF, and extends to accommodate rounding and measurement error. Several subsequent papers build on the EMW parametric approach. Becker, Duersch, Eife, and Glas (2021) extend the generalized beta fitting procedure to surveys with varying interval widths, such as the ECB's SPF. Bellemare, Bissonnette, and Kröger (2012) propose a semiparametric alternative based on cubic spline interpolation of the subjective CDF, which is more flexible than fitting a single para-

metric family but still requires smoothing choices that can affect the estimated moments. Bissonnette, De Bresser, van Soest, and Vonkova (2017) study the estimation of subjective distributions from probability questions in the context of consumption uncertainty.

A separate strand of the literature focuses on the panel of density forecast responses, aggregating information across respondents and time periods. For example, Bassetti, Casarin, and Del Negro (2024) propose a Bayesian nonparametric approach that embeds a parametric model in a Dirichlet process mixture and conducts joint inference across the cross-section of SPF forecasters. In contrast, our framework focuses on single survey responses in isolation, without pooling across respondents or time. This makes our bounds directly applicable as inputs to any analysis without requiring assumptions on the cross-sectional distribution of beliefs.

On the regression side, the problem of inference with interval-valued data has a large literature following Manski and Tamer (2002). Beresteanu, Molchanov, and Molinari (2011) study identification and inference with set-valued regressors using random set theory, and Molinari (2020) provides a comprehensive survey of partial identification methods. We contribute to this literature by deriving sharp bounds on the regression coefficient when the outcome is interval-valued and computationally feasible bounds when the regressor, or both variables, are interval-valued, with inference adapted to the panel data setting with individual fixed effects.

Our paper is structured as follows. Section 2 motivates the problem under study and discusses previous approaches using some examples from actual survey responses. Section 3 presents our main results in a general setup and discusses the implementation of our approach. Section 4 shows how to utilize our new nonparametric bounds in subsequent regression analysis. Section 5 provides an empirical analysis of the consumption Euler equation using our approach, and Section 6 concludes. Proofs can be found in Appendix A.A `MATLAB` code repository is available online.[2]

2 Setup and motivation

The general structure of most survey density questions comprises of a target variable and a mutually exclusive partition of the support of this variable. Respondents are then asked to provide their assessed probabilities of the target variable realizing in each of the partitions.

[2] https://github.com/TinchoAlmuzara/SurveyDensities.

To make things concrete, consider the survey density question on year-ahead inflation[3] given in the Survey of Consumer Expectations (SCE):

Question Q9. Now we would like you to think about the different things that may happen to inflation **over the next 12 months**. We realize that this question may take a little more effort.

In your view, what would you say is the percent chance that, **over the next 12 months** ...

increased by 12% or more	_____	percent chance
increased by 8% to 12%	_____	percent chance
increased by 4% to 8%	_____	percent chance
increased by 2% to 4%	_____	percent chance
increased by 0% to 2%	_____	percent chance
decreased by 0% to 2%	_____	percent chance
decreased by 2% to 4%	_____	percent chance
decreased by 4% to 8%	_____	percent chance
decreased by 8% to 12%	_____	percent chance
decreased by 12% or more	_____	percent chance
Total	100	

A completed response to Question Q9 clearly provides *some* information about a respondent's subjective probability distribution. Unfortunately, it does not provide a complete picture. To see this, let $F(\cdot)$ be the true subjective forecast cumulative distribution function (CDF). In addition, define π_1 as the probability that "the rate of inflation will be 12% or higher", π_2 the probability that "the rate of inflation will be between 8% and 12%", and so on. We can summarize the information provided in the response as,

$$\begin{aligned} F(0.12) &= 1 - \pi_1 \\ F(0.08) &= 1 - \pi_1 - \pi_2 \\ F(0.04) &= 1 - \pi_1 - \pi_2 - \pi_3 \\ &\vdots \\ F(-0.08) &= 1 - \textstyle\sum_{i=1}^{9} \pi_i \\ F(-0.12) &= \pi_{10} \end{aligned}$$

[3]The questionnaire for the main SCE survey is available here.

Question Q9 thus provides the subjective CDF evaluated at nine points across the whole support of future inflation. As pointed out by Engelberg et al. (2009), hereafter EMW, without further assumptions, we can only draw limited conclusions about objects of interest such as the mean or median of this subjective distribution.

EMW suggested "hands off" upper and lower bounds for the mean and median. These are best introduced by example. The top chart in Figure 1 presents a visualization of a single SCE response where each bar represents a specific value of π_i. If the median, $F^{-1}(0.5)$ coincides with either endpoint of the partition, it is known directly. If not, upper and lower bounds for the subjective median can be found using the endpoints of the bin in which the subjective CDF takes on a value of 0.5. In this particular response we have that $F(0) = 0.37$ and $F(2) = 0.92$ so that we can form the lower bound, LB^{med}_{EMW}, and the upper bound, UB^{med}_{EMW}, as

$$\mathsf{LB}^{med}_{EMW} = 0\% \quad \text{and} \quad \mathsf{UB}^{med}_{EMW} = 2\%$$

To obtain analogous bounds for the subjective mean, EMW suggested to obtain the lower bound by assuming all probability mass in the interval is concentrated on the left endpoint of the bin, and an upper bound. The middle chart of Figure 1 visualizes this approach. The blue vertical lines represent the probability mass function associated with the lower bound so that

$$\mathsf{LB}^{mean}_{EMW} = (-0.08)\pi_8 + (-0.04)\pi_7 + (-0.02)\pi_6 + (0.02)\pi_4 + (0.04)\pi_3 = -0.78\%.$$

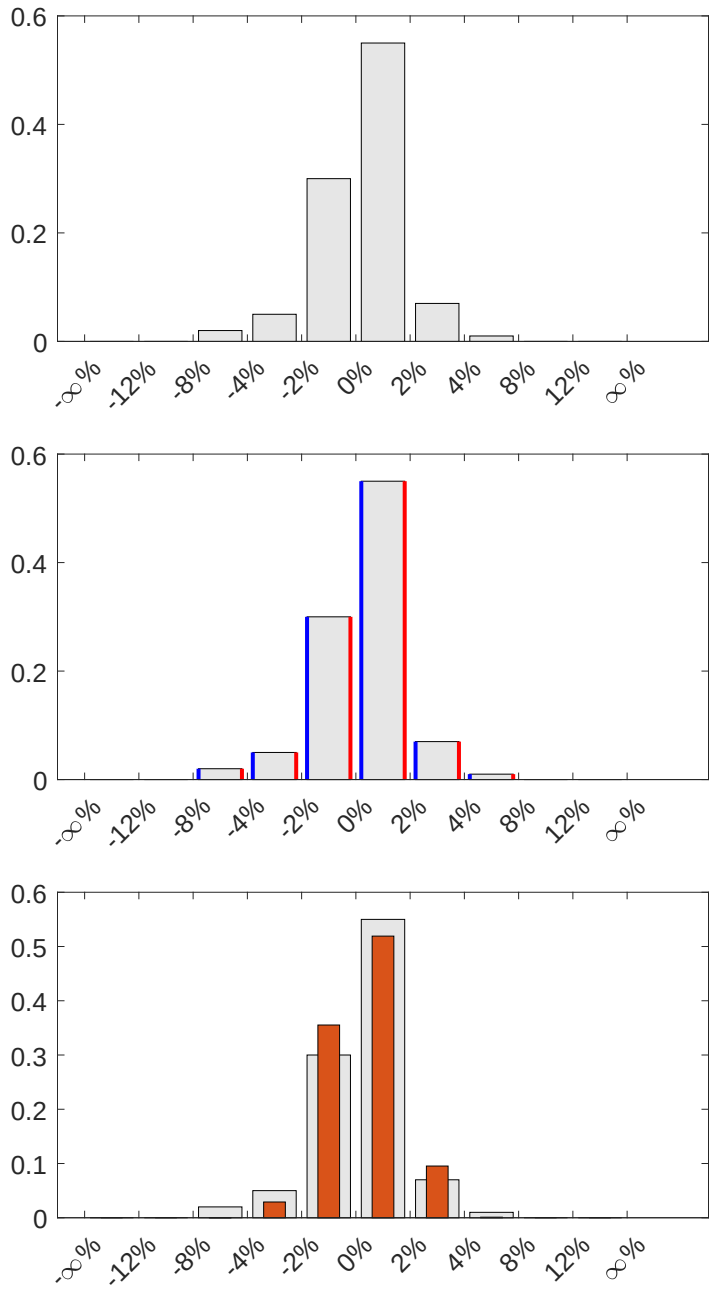
Similarly, the probability mass function represented by the red vertical lines gives the upper bound,

$$\mathsf{UB}^{mean}_{EMW} = (-0.04)\pi_8 + (-0.02)\pi_7 + (0.02)\pi_5 + (0.04) * \pi_4 + (0.02)\pi_3 = 1.28\%.$$

Although LB^{med}_{EMW}, UB^{med}_{EMW}, LB^{mean}_{EMW}, and UB^{mean}_{EMW} have the advantage that the bounds are obtained without imposing any structure on the distribution, a clear downside is that these bounds are generally quite wide—about 2 percentage points for both the mean and median in the example above. Indeed, the bounds on the mean straddle zero so that the sign is ambiguous. Moreover, following similar steps to obtain upper and lower bounds on measures of uncertainty such as the subjective variance or subjective interquartile (IQR) range, would lead to even wider bounds.[4] As an alternative, EMW also proposed fitting

[4]For example, using the example in Figure 1 the 25th percentile is in the interval $[-.02\%, 0.0\%]$ and the

FIGURE 1. **Example SCE Response with Zero Tail Probabilities.** This figure presents an example survey response from the SCE where $\pi_1 = \pi_{10} = 0$. The top chart presents the raw responses. The middle chart adds the probability mass functions associated with the lower bound (blue) and upper bound (red) proposed by Engelberg et al. (2009). The bottom chart adds the fitted probabilities (orange bars) using the generalized beta family of distributions as suggested by Engelberg et al. (2009).



75th percentile is in the interval $[0.0\%, 0.02\%]$. Thus, the hands-off lower bound for the IQR is 0% and the upper bound is 4%.

the survey responses using a class of parametric distributions—the generalized beta family of distributions.[5] An advantage of this approach is that it provides a completely specified density which allows the direct calculation of any functional including the mean, median, mode, IQR, etc. This is especially important when trying to quantify uncertainty as the hands-off bounds can become very wide. However, a downside to this approach is that the parametric class of densities may not contain a density which well-approximates the true density. This can be true even if the fitting error between the survey probabilities and those implied by the candidate density are small, since there are generally many possible densities (including those outside the parametric class) that are consistent with each survey response.[6]

The bottom chart of Figure 1 overlays the implied probabilities using the best fit from the generalized beta family on the survey responses. We can clearly observe that the generalized beta distribution misses key aspects of the response. The fitted density has only negligible mass in the $[-8\%, -4\%]$, $[0\%, 2\%]$, and $[4\%, 8\%]$ intervals but places excess mass in the $[-2\%, 0\%]$ and $[2\%, 4\%]$ intervals. The associated mean and median from the generalized beta distribution are both about 0.37, which lies between the hands-off lower and upper bounds from above. However, the substantial fitting errors suggest that caution should be taken when utilizing any function of the fitted density.

In the previous example, the respondent placed zero probability in both of the extreme intervals ($[-\infty, -12\%]$ and $[12\%, \infty]$), which simplified the analysis. Figure 2 presents two examples from the SCE where each respondent has assigned nonzero probability to at least one of the extreme intervals. In both cases, we are still able to obtain bounds for the subjective median. In Figure 2(a), $F(0.02) = 0.45$ and $F(0.04) = 0.70$ so that

$$\mathsf{LB}^{med}_{\mathrm{EMW}} = 2\% \quad \text{and} \quad \mathsf{UB}^{med}_{\mathrm{EMW}} = 4\%.$$

Similarly, in Figure 2(a) we have that $F(0) = 0.39$ and $F(0.02) = 0.69$ so that

$$\mathsf{LB}^{med}_{\mathrm{EMW}} = 0\% \quad \text{and} \quad \mathsf{UB}^{med}_{\mathrm{EMW}} = 2\%.$$

However, with nonzero probability in the extreme intervals, $\mathsf{LB}^{mean}_{\mathrm{EMW}}$ and $\mathsf{UB}^{mean}_{\mathrm{EMW}}$ cannot

[5] EMW only use the generalized beta family when the respondent reports nonzero probabilities in three or more consecutive intervals. When the response has nonzero probability in only one or two consecutive intervals, a triangular distribution is used. Here, we focus on the generalized beta family but the same drawbacks we discuss apply to any parametric class of densities.

[6] A related issue is that the choice of density, from a parametric family, that best fits a survey response will be affected by the choice of criterion function used.

be defined without further restrictions on the subjective density. In particular, with no additional assumptions, the lower and upper bounds of the density represented in Figure 2(b) become $-\infty$ and ∞, respectively.[7] Furthermore, the bottom charts of Figure 2 show that the generalized beta family of densities is not sufficiently flexible to match the survey responses—especially the tail probabilities in the extreme intervals.

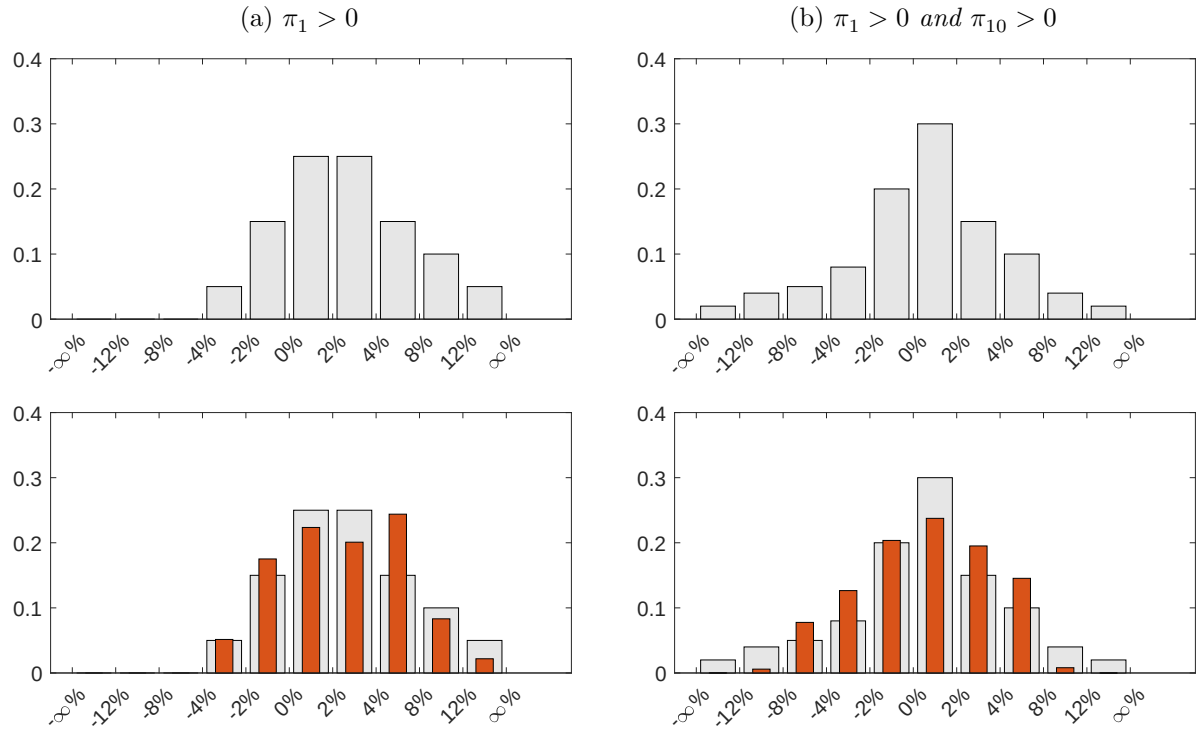
Motivated with these examples, in the next section we introduce a new approach to construct lower and upper bounds on objects such as the subjective mean and median and measures of uncertainty. Our approach strikes a balance between the above-described hands-off bounds and the point identification from assuming a parametric form for the density. We achieve this balance by imposing smoothness restrictions on the underlying density. This retains the flexibility of the hands-off bounds as we can avoid taking a stand on the parametric family of densities, but is still able to provide substantially more informative bounds compared to the hands-off bounds approach. The bounds, $\mathsf{LB}^{mean}_{\text{EMW}}$ and $\mathsf{UB}^{mean}_{\text{EMW}}$, based on the data in Figure 1 provides an illustrative example of how smoothness can in fact tighten nonparametric bounds. The two probability mass functions (red and blue lines) in the middle chart of Figure 1 underpin the hands-off upper and lower bounds. Importantly, they assign *zero* probability to all values in the interior of each interval implying that the probabilities jump from zero to a nonzero value at the end of each interval. If we were to instead restrict the class of distributions to have some degree of smoothness—ruling out jumps such as these—then this will raise the lower bound and lower the upper bound. Moreover, we can impose these restrictions while still limiting ourselves to densities that satisfy the reported probabilities for each interval given by the respondent. This eliminates the concerns that arise from the fitting error generated by using parametric families such as the generalized beta distribution. We will formalize these steps in the next section.

3 Bounds on functionals of survey forecast densities

We now state the problem formally. We consider an economic unit who is asked about the (subjective) probability distribution of a random variable X—the rate of inflation over the next 12 months in Section 2 or household spending in the empirical analysis of Section 5.

[7] We can construct the lower bound, but not the upper bound, for the response shown in Figure 2(a).

FIGURE 2. **Example SCE Response with Non-Zero Tail Probabilities.** This figure presents two example survey responses from the SCE where either $\pi_1 > 0$ or $\pi_{10} > 0$. The top chart presents the raw responses. The bottom chart adds the fitted probabilities (orange bars) using the generalized beta family of distributions as suggested by Engelberg et al. (2009).



The unit reports for a class of known measurable sets $\{A_k\}_{k=1}^K$ the probabilities

$$\pi_k = P_{f^\star}(\, X \in A_k) = \int_{A_k} f^\star(x)\, dx, \tag{1}$$

where $f^\star$ is the density of X. Throughout this section, we assume that X is continuously distributed and scalar, but most of what we do can be extended to a multivariate setup and to discrete distributions with known support.[8]

Next, suppose we have collected the probabilistic data $D = \{A_k, \pi_k\}_{k=1}^K$ and are interested in the quantity $\theta^\star = E_{f^\star}[\, G(X)\,] = \int G(x) f^\star(x)\, dx$ for a known real-valued function G. The

[8] A natural question concerns elicitation error: the survey response reflects the respondent's reported distribution, which may differ from their true latent subjective distribution. We view our framework as providing the observables in a measurement system for the latent density, as our bounds characterize the information content of the reported responses under only smoothness assumptions. Connecting these observables to the latent density requires additional modeling of the elicitation process, which we leave for future work.

true density $f^\star$ is unknown to us but we know that it satisfies (1). We also have additional information about $f^\star$ which may include constraints on the smoothness, tails, location of the mode, etc. We represent the available information about $f^\star$ by a set $\mathcal{F}_D$ of all candidate densities that satisfy (1) and the additional constraints.

The core insight from our paper is that, given this information, we can summarize our knowledge of $\theta^\star$ with a pair of bounds, $\underline{\theta} \leq \theta^\star \leq \overline{\theta}$, where

$$\underline{\theta} = \inf_{f \in \mathcal{F}_D} \left(\int G(x) f(x)\, dx \right) \quad \text{and} \quad \overline{\theta} = \sup_{f \in \mathcal{F}_D} \left(\int G(x) f(x)\, dx \right). \tag{2}$$

The bounds $\underline{\theta}$ and $\overline{\theta}$ are appealing because they impose the probabilistic data D exactly, in the spirit of the hands-off approach of EMW, thus eliminating the fitting error concerns that plague fully parametric approaches. At the same time, our bounds make efficient use of all the information in $\mathcal{F}_D$—notably the information on the properties of $f^\star$. This feature can sharpen estimates of $\theta^\star$ quite substantially, as we will illustrate.

The practical question is how to compute the bounds. We will show below that many restrictions of interest take the form of linear equality and inequality constraints on $f^\star$. Hence, $\mathcal{F}_D$ will often be a closed convex set and (2) a pair of functional linear programs. With this structure at hand, we can provide a highly efficient algorithm to approximate $\underline{\theta}$ and $\overline{\theta}$ to any desired level of accuracy.

This section has three parts. First, we illustrate the wide applicability of our framework by describing the classes of target functions G and candidate density spaces $\mathcal{F}_D$ that can be handled (Section 3.1). Second, we introduce our computational strategy to approximate $\underline{\theta}$ and $\overline{\theta}$, and discuss its mathematical properties (Section 3.2). Finally, we show that our approach also provides valuable information about objects (such as centered moments or quantiles) that cannot be directly written as functional linear programs (Section 3.3). We give a step-by-step implementation guide in Section 3.4.

In this section, our focus is on learning about features of $f^\star$ for a single unit. Given data $\{D_i\}_{i=1}^N$ for a sample of N units, each with true unknown density $f_i^\star \in \mathcal{F}_{D_i}$, one can use our proposed method one unit at a time to compute a sample of bounds $\{\underline{\theta}_i, \overline{\theta}_i\}_{i=1}^N$. How to exploit the sample of bounds for regression analysis is the topic of Section 4.

3.1 Target functions and candidate density spaces

3.1.1 Target function G

Many objects of interest that arise in applications can be expressed in terms of functionals $\theta^\star$ of a density $f^\star$. Two leading examples of objects that can be mapped to $\theta^\star$ *directly* are uncentered moments and probabilities of events.

The uncentered moment $E_{f^\star}\left[X^j\right]$ with j a positive integer corresponds to $\theta^\star$ with the target function $G(x) = x^j$. This covers the subjective mean in the illustration of Section 2. The probability $P_{f^\star}(A)$ with A a measurable set corresponds to $\theta^\star$ with the target function $G(x) = \mathbf{1}\{x \in A\}$. This covers the CDF at a point $F(\bar{x}) = P_{f^\star}((-\infty, \bar{x}])$.

Other objects of interest are centered moments, such as the variance or skewness, and functions of the quantiles of X, such as the median or interquartile range. Although these cannot be cast directly in the form of (2), our method delivers valuable information about them, as we show in Section 3.3.[9]

3.1.2 Candidate density space $\mathcal{F}_D$

As our baseline, we ask that any $f \in \mathcal{F}_D$ (i) is a proper density, (ii) fits the probabilistic data (1), and (iii) satisfies a smoothness constraint. To be precise, let $\mathcal{X}$ be the support of X, let $\mathcal{C}^{p-1,1}(\mathcal{X})$ be the space of $(p-1)$-times continuously differentiable functions on $\mathcal{X}$ that have a Lipschitz continuous p-th derivative, and let D^p be the p-th derivative operator. Also let $\|\cdot\|_\infty$ by the sup-norm. For a fixed (and small) $\phi > 0$, we define

$$\mathcal{F}_D(\phi) = \left\{ f \in \mathcal{C}^{p-1,1}(\mathcal{X}) : f \geq 0,\ \int_{\mathcal{X}} f(x)\,dx = 1,\ \int_{A_k} f(x)\,dx = \pi_k,\ \text{all } k,\ \|D^p f\|_\infty \leq \phi \right\} \quad (3)$$

We now explain each element in definition (3). First, each $f \in \mathcal{F}_D(\phi)$ is a proper density as it is nonnegative and integrates to one. Second, we impose the probabilistic restrictions (1); that is, the set $\mathcal{F}_D(\phi)$ only contains densities that share the probabilistic implications of the true unknown density $f^\star$ for $\{A_k\}_{k=1}^K$. In practice, the most common situation is one where $\{A_k\}_{k=1}^K$ is a partition of (a set that covers) $\mathcal{X}$, that is, $A_k = (a_{k-1}, a_k]$ for increasing $\{a_k\}_{k=0}^K$ with $\sum_{k=1}^K \pi_k = 1$—see our examples in Sections 2 and 5. In that case, the condition

[9] As mentioned, we assume a scalar X, but our methods can be extended to a multivariate X. Uncentered cross-moments and joint probabilities are covered directly while covariances, correlations and coefficients in least-square or quantile regressions are covered indirectly.

$\int_{\mathcal{X}} f(x)\,dx = 1$ may be omitted.[10]

Third, we require that f be sufficiently smooth. Given that $f \in \mathcal{C}^{p-1,1}(\mathcal{X})$, the candidate density f has p continuous derivatives. Moreover, its p-th derivative is uniformly bounded by a small number ϕ that controls (together with p) the degree of smoothness. As this is an important parameter for problem (2), we make the dependence of $\mathcal{F}_D$ on ϕ explicit. We note that the condition $\|D^p f\|_\infty \leq \phi$ is equivalent to

$$D^p f(x) \leq \phi \text{ and } -D^p f(x) \leq \phi, \text{ for all } x \in \mathcal{X},$$

which is an infinite set of linear inequality constraints on f. Since the requirement that f is a proper density and satisfies the probabilistic restrictions also involves linear constraints, $\mathcal{F}_D(\phi)$ is a closed convex subset of $\mathcal{C}^{p-1,1}(\mathcal{X})$ and problem (2) is a pair of functional linear programs. Our computational strategy exploits this observation.

Where does information about smoothness come from? Constraints on derivatives are often used in the nonparametric density estimation literature (Stone, 1980). Although we study a conceptually different problem, a lesson from that literature is that there is limited scope for data-driven choice of smoothness class, particularly with regards to inference (Low, 1997; Cai and Low, 2004; Giné and Nickl, 2010; Bull, 2012). In our case, we can fit the probabilistic data D perfectly for many different choices of p and ϕ. This clearly indicates that smoothness constraints must be formulated with reference to a priori knowledge or external data. We discuss practical recommendations below.

3.1.3 Additional structure

In some cases, the restrictions on $\mathcal{F}_D(\phi)$ may not be enough to obtain meaningful bounds. In other cases, the researcher may have extra information that can be used to refine the bounds. Our method can easily accommodate additional structure when it takes the form of linear equality or inequality constraints on f. We discuss a few examples next.

Tail constraints. As noted in Section 2, the hands-off bounds for the mean may be infinite if the unit puts positive probability on the tail bins. When the support $\mathcal{X}$ is unbounded,

[10] Some surveys ask respondents about their CDF at different points—$A_k = (-\infty, a_k]$ for increasing $\{a_k\}_{k=1}^K$ with $\pi_{k-1} \leq \pi_k$—but one can always map this to the more standard partition format. If $\tilde{A}_k = (-\infty, \tilde{a}_k]$ has probability $\tilde{\pi}_k$ for $k = 1, \ldots, K-1$ (in increasing order), we set $A_1 = \tilde{A}_1$, $A_k = \tilde{A}_k \backslash \tilde{A}_{k-1}$ for $1 < k < K$, and $A_K = \mathcal{X} \backslash \tilde{A}_{K-1}$ together with $\pi_1 = \tilde{\pi}_1$, $\pi_k = \tilde{\pi}_k - \tilde{\pi}_{k-1}$ for $1 < k < K$ and $\pi_K = 1 - \tilde{\pi}_{K-1}$.

the same issue affects any unbounded target G, and in particular, all uncentered moments. This can be prevented by imposing restrictions on the tail behavior of $\mathcal{F}_D$. For example, for a dominating function γ with $|G\gamma|$ integrable and a set $\mathcal{T}$ with $\mathcal{X}\backslash\mathcal{T}$ compact,

$$f(x) \leq \gamma(x), \text{ for all } x \in \mathcal{T}.$$

Clearly, this is an infinite set of linear inequality constraints on f.

Monotonicity. Although more restrictive, it is straightforward to include information about the monotonicity of f or its derivatives. For example, if the q-th derivative of f is known to be non-increasing over a set $\mathcal{M}$, this can be written as

$$D^q f(x) - D^q f(x') \geq 0, \text{ for all } x, x' \in \mathcal{M} \text{ such that } x < x'.$$

Setting $q = 1$ and $\mathcal{M} = \mathcal{X}$, the restriction imposes that every $f \in \mathcal{F}_D$ is unimodal (allowing for flat areas). Accordingly, monotonicity can also be expressed as an infinite set of linear inequality constraints.

Location of the mode. Some surveys ask units to report the mode of the forecast density. That is the value $\widetilde{x}$ such that

$$f(\widetilde{x}) - f(x) \geq 0, \text{ for all } x \in \mathcal{X}.$$

This is another example of linear inequality constraint on f.

Rounding bias. Our setup can handle cases where some or all of the bin probabilities reported by the unit are subject to rounding bias by an amount r known to the researcher. This involves replacing the equality constraint (1) by

$$\pi_k - r \leq \int_{A_k} f(x)\, dx \leq \pi_k + r,$$

for those bins k for which the researcher knows that rounding has occurred. Again, these are linear inequality constraints on f.

Our implementation (described below and in Section 3.4) automatically adopts a tail condition when $\mathcal{X}$ is unbounded to ensure the problem is well defined. The monotonicity,

mode and rounding bias restrictions are instead less standard. We do not include them in our default implementation, but it would be straightforward to do so.

3.2 Computational strategy and main result

We now develop our computational approach to approximate the bounds $\underline{\theta}$ and $\overline{\theta}$ in (2). To explain the idea, we assume initially a bounded support of the form $\mathcal{X} = [\underline{x}, \overline{x}]$ and a candidate density space $\mathcal{F}_D(\phi)$ that imposes smoothness as in (3). We discuss how to relax these assumptions and incorporate more structure into $\mathcal{F}_D$ later in the section.

Our method relies on discretization of the support of X and the set $\mathcal{F}_D$. Specifically, we discretize $\mathcal{X} = [\underline{x}, \overline{x}]$ by a grid $\mathcal{X}_M = \{x_0, x_1, \ldots, x_M\}$ with step size $h = (\overline{x} - \underline{x})/M$, so that $x_m = \underline{x} + mh$ for each $m = 0, \ldots, M$. Next, we define the p-th discrete derivative operator by $\Delta_h^p = \Delta_h \circ \Delta_h^{p-1}$ where, for each vector $u = (u_0, u_1, \ldots, u_M)$,

$$(\Delta_h u)_m = \frac{u_m - u_{m-1}}{h}.$$

This is the discrete counterpart to the derivative operator D^p. We then discretize the closed convex set $\mathcal{F}_D(\phi)$ by replacing each function f with a vector $\bar{f} = (\bar{f}_0, \bar{f}_1, \ldots, \bar{f}_M)$, where $\bar{f}_m$ can be viewed as the probability of a discrete random variable taking the value x_m, subject to the probabilistic and smoothness restrictions in Equation (3):

$$\mathcal{F}_{D,M}(\phi) = \left\{ \bar{f} \in \mathbb{R}^{M+1} : \bar{f} \geq 0,\ \sum_{m=0}^{M} \bar{f}_m = 1,\ \sum_{x_m \in A_k} \bar{f}_M = \pi_k,\ \text{all } k,\ \left\| \Delta_h^p \bar{f} \right\|_\infty \leq h\phi \right\}. \quad (4)$$

As in our discussion of (3), the condition $\sum_{m=0}^{M} \bar{f}_m = 1$ may be omitted if $\{A_k\}_{k=1}^K$ is a partition of $\mathcal{X}_M$. More importantly, the discrete smoothness constraint can be written as

$$(\Delta_h^p \bar{f})_m \leq h\phi \text{ and } -(\Delta_h^p \bar{f})_m \leq h\phi, \text{ for } m = 0, \ldots, M,$$

which is a finite set of linear inequality constraints on $\bar{f}$. Finally, we note that discrete derivatives are bounded, not by ϕ, but by $h\phi$ (which is smaller than ϕ for large M). This is because the natural counterpart to $\bar{f}_m$ in the continuous case is $\int_{[x_m \pm h/2]} f(x)\, dx \approx f(x_m)h$.

We can now introduce our approximate bounds:

$$\underline{\theta}_M(\phi) = \min_{\bar{f} \in \mathcal{F}_{D,M}(\phi)} \left(\sum_{m=0}^{M} G(x_m) \bar{f}_m \right) \quad \text{and} \quad \overline{\theta}_M(\phi) = \max_{\bar{f} \in \mathcal{F}_{D,M}(\phi)} \left(\sum_{m=0}^{M} G(x_m) \bar{f}_m \right). \quad (5)$$

For a fixed M, the two problems in (5) are finite-dimensional linear programs and, provided that $\mathcal{F}_{D,M}(\phi)$ is nonempty (which requires a not too small ϕ), they can be solved by efficient numerical routines even for large M. We propose to approximate the original bounds $(\underline{\theta}, \overline{\theta})$ by $(\underline{\theta}_M(\phi), \overline{\theta}_M(\phi))$ for a reasonably large M.

The justification for our approach is the following proposition. In (2), when $\mathcal{F}_D = \mathcal{F}_D(\phi)$ as defined in Equation (3), we will write $\underline{\theta}(\phi)$ and $\overline{\theta}(\phi)$ to emphasize the dependence on the smoothness coefficient ϕ. Let $\mathrm{BV}(\mathcal{X})$ be the class of functions of bounded variation.

Proposition 1. *Suppose $\mathcal{X} = [\underline{x}, \overline{x}]$ is compact, $\{A_k\}_{k=1}^K$ is a partition of $\mathcal{X}$, $G \in \mathrm{BV}(\mathcal{X})$, and the spaces $\mathcal{F}_D(\phi)$ and $\mathcal{F}_{D,M}(\phi)$ in (3) and (4) are nonempty. Then, as $M \to \infty$*

$$\underline{\theta}_M(\phi) \longrightarrow \underline{\theta}(\phi) \quad \text{and} \quad \overline{\theta}_M(\phi) \longrightarrow \overline{\theta}(\phi).$$

Proof. See Appendix A. □

As mentioned above, that $\{A_k\}_{k=1}^K$ partitions $\mathcal{X}$ is without loss of generality in practice, and the spaces $\mathcal{F}_D(\phi)$ and $\mathcal{F}_{D,M}(\phi)$ are generally nonempty if ϕ (and M) are not too small. The assumption that $G \in \mathrm{BV}(\mathcal{X})$ is also broad as $\mathrm{BV}(\mathcal{X})$ includes the (restrictions to $\mathcal{X}$ of the) functions $G(x) = x^j$ for j positive and $G(x) = \mathbf{1}\{x \in A\}$ for A given by a finite union of intervals. Thus, Proposition 1 covers subjective means, probabilities and several other objects of applied interest under very general conditions. We examine the performance of our approach in numerical experiments in Section 3.4. We now make a few remarks.

Remark 1. Given M, $\mathcal{F}_{D,M}(\phi)$ is a polytope—it is the intersection of an affine space and a finite (but possibly large) number of half-spaces—contained in the $(M+1)$-dimensional simplex. Hence, it is convex and compact and, by Weierstrass theorem, the minimum and maximum in (5) are attained. Moreover, seen as a function of ϕ, $\mathcal{F}_{D,M}(\phi)$ is a continuous compact-valued correspondence and, by Berge's maximum theorem, $\underline{\theta}_M(\phi)$ and $\overline{\theta}_M(\phi)$ are continuous in ϕ. In fact, they are piecewise linear with $\underline{\theta}_M(\phi)$ decreasing and convex, and $\overline{\theta}_M(\phi)$ increasing and concave in ϕ.

In the functional problem (2), $\mathcal{F}_D(\phi)$ (with the sup-norm) is also a continuous compact-valued correspondence and the objective function $\int_{\mathcal{X}} G(x)f(x)\,dx$ is continuous in f. Thus, $\underline{\theta}(\phi)$ and $\overline{\theta}(\phi)$ are continuous (and monotonic) in ϕ.

Remark 2. If M is such that $A_k \cap \mathcal{X}_M \neq \emptyset$ for all k, there is always a finite value of ϕ for which $\mathcal{F}_{D,M}(\phi)$ is nonempty. The minimum such ϕ, which will be denoted ϕ^*, delivers the tightest bounds $\underline{\theta}_M(\phi)$ and $\overline{\theta}_M(\phi)$ for $\theta^\star$. In the experiments reported in Section 3.4, the

tightest bounds $\underline{\theta}_M(\phi^*)$ and $\overline{\theta}_M(\phi^*)$ work remarkably well for a large class of true densities $f^\star$ and we therefore advocate their use.

An important property is that ϕ^* is independent of G and continuous in $\{\pi_k\}_{k=1}^K$. This is because ϕ^* is given by $\phi^* = \min_{\tilde{f}\in\mathcal{F}^*_{D,M}} h^{-1} \left\|\Delta_h^p \tilde{f}\right\|_\infty$ where

$$\mathcal{F}^*_{D,M} = \left\{ \bar{f} \in \mathbb{R}^{M+1} : \bar{f} \geq 0, \sum_{x_m \in A_k} \bar{f}_m = \pi_k, \text{ all } k \right\}$$

is a continuous compact-valued correspondence of $\{\pi_k\}_{k=1}^K$. This simplifies implementation as ϕ^* can be computed efficiently prior to solving (5) and can be reused for different G.

Remark 3. The set of minimizer and maximizer densities that attain the bounds in (5) can also be computed by linear programming techniques. However, one has to be cautious in interpreting them as they are generally not unique, they depend on the target function G, and they convey little direct information about the true density $f^\star$. Our recommendation is to focus on the bounds $\underline{\theta}_M(\phi)$ and $\overline{\theta}_M(\phi)$, possibly for various choices of target function G, as the means to summarize our knowledge about $f^\star$.[11]

Remark 4. The essence of the proof of Proposition 1 is that (a) any function in $\mathcal{F}_D(\phi)$ can be discretized to a vector that approximately satisfies the constraints in $\mathcal{F}_{D,M}(\phi)$ and (b) the smoothness constraints in $\mathcal{F}_{D,M}(\phi)$ are preserved under weak convergence. Extending our result to the monotonicity, mode and rounding restrictions introduced in Section 3.1 can be done with few changes to the argument.

Proposition 1 assumes a bounded support $\mathcal{X}$ but under suitable tail constraints we can extend the convergence result to the unbounded case. For economy of space, we only give a sketch of the argument. Suppose $A_1 = (-\infty, \underline{a})$ and $A_K = (\overline{a}, \infty)$ with $-\infty < \underline{a} < \overline{a} < \infty$; $\{A_k\}_{k=2}^{K-1}$ is a partition of $[\underline{a}, \overline{a}]$. We define the truncated support $\mathcal{X}_R = [\underline{a} - R, \overline{a} + R]$ and its discretized version $\mathcal{X}_{R,M}$. We also introduce the dominating function γ that satisfies $\int_{A_1 \cup A_K} |G(x)\gamma(x)|\, dx < \infty$. The idea is to re-define $\mathcal{F}_D(\phi)$ so that $\mathcal{X} = \mathbb{R}$ and $f \in \mathcal{F}_D(\phi)$ implies $f(x) \leq \gamma(x)$ for all $x \in A_1 \cup A_K$. We then let $\mathcal{F}_{D,R}(\phi)$ be the truncated version of $\mathcal{F}_D(\phi)$ with $\mathcal{X}$ replaced by $\mathcal{X}_R$ and A_k by $A_k \cap \mathcal{X}_R$, and $\mathcal{F}_{D,R,M}(\phi)$ the discrete counterpart to $\mathcal{F}_{D,R}(\phi)$. Denote by $\underline{\theta}(\phi)$, $\underline{\theta}_R(\phi)$ and $\underline{\theta}_{R,M}(\phi)$ the lower bounds of the problems associated to $\mathcal{F}_D(\phi)$, $\mathcal{F}_{D,R}(\phi)$ and $\mathcal{F}_{D,R,M}(\phi)$ (the argument for the upper bounds is identical).

[11] When interest is in the probability that X lies in a neighborhood around a point $\widetilde{x}$, computing bounds on $\theta^\star$ with $G(x) = \mathbf{1}\{|x - \widetilde{x}| \leq \nu\}$ for a small but nonzero ν is a more sensible approach.

For fixed truncation R, a slight variation on the proof of Proposition 1 (to accommodate the dominating function) delivers $\underline{\theta}_{R,M}(\phi) \longrightarrow \underline{\theta}_R(\phi)$ as $M \to \infty$ and, using the tail bound, one can show that $\underline{\theta}_R(\phi) \longrightarrow \underline{\theta}(\phi)$ as $R \to \infty$. It follows from double-array techniques that there is a sequence $\{R_M\}$ with $R_M \to \infty$ as $M \to \infty$ such that $\underline{\theta}_{R_M,M}(\phi) \longrightarrow \underline{\theta}(\phi)$. That sequence needs $h = (\overline{a} - \underline{a} + 2R_M)/M \to 0$ which, in particular, requires $R_M/M \to 0$; that is, the range has to grow but slowly compared to the gridsize. The argument highlights that to approximate $\underline{\theta}(\phi), \overline{\theta}(\phi)$ one needs a truncated support large enough to include most of the "action" in Gf and, given the choice of R, a fine enough grid.

In our numerical experiments (Section 3.4) we find an exponentially decaying γ and a simple rule to calibrate R are a good choice for many densities and target functions.

3.3 Bounding other features of the distribution

As mentioned in Section 3.1, sharp bounds for centered moments (such as the variance) and quantiles cannot be characterized directly using functional linear programs. However, our approach offers valuable information about them, as we explain next.

Centered moments. For any two densities $f_1, f_2 \in \mathcal{F}_D$, define the auxiliary function

$$m_j(f_1, f_2) = E_{f_2}\Big[\,(X - E_{f_1}[\,X\,])^j\,\Big].$$

The sharp bounds for the j-th centered moment $\theta^\star = m_j(f^\star, f^\star)$ are

$$\underline{\theta}^s = \inf_{f \in \mathcal{F}_D} m_j(f, f) \quad \text{and} \quad \overline{\theta}^s = \sup_{f \in \mathcal{F}_D} m_j(f, f).$$

Unfortunately, these are functional polynomial programs that, even after discretization, can sometimes be computationally intractable.[12] Instead of a single polynomial program, we propose an approach that consists of solving multiple linear programs. Specifically,

$$\underline{\theta} = \inf_{f_1 \in \mathcal{F}_D} \left\{ \inf_{f_2 \in \mathcal{F}_D} m_j(f_1, f_2) \right\} \quad \text{and} \quad \overline{\theta} = \sup_{f_1 \in \mathcal{F}_D} \left\{ \sup_{f_2 \in \mathcal{F}_D} m_j(f_1, f_2) \right\}.$$

Given f_1, the inner problems fall under the linear programming framework of our paper. The outer problems can then be implemented by bounding $E_{f_1}[\,X\,]$ (linear programs) and

[12] For example, in the case of the variance, the upper bound can be obtained by solving a convex quadratic program but this is not the case for the lower bound.

solving the inner problem for a grid of values between those bounds. This approach is valid but conservative, that is, $\underline{\theta} \leq \underline{\theta}^s \leq \theta^\star \leq \overline{\theta}^s \leq \overline{\theta}$.[13] Moreover, the implied bounds for even moments are guaranteed to be non-negative.

A similar sequential approach can be developed for ratios $m_j(f_1, f_3)/\{m_2(f_1, f_2)\}^{j/2}$, such as the skewness ($j = 3$) and kurtosis ($j = 4$) coefficients.

Quantiles. In order to bound the quantile function $Q(\tau) = \inf\{x : F(x) \geq \tau\}$, we begin by constructing lower and upper bounds, $\underline{F}(x)$ and $\overline{F}(x)$, for $F(x) = P_{f^\star}(X \leq x)$ for each $x \in \mathcal{X}$ using our approach. Valid (but possibly conservative) bounds for $Q(\tau)$ are given by

$$\underline{Q}(\tau) = \inf\{x : \overline{F}(x) \geq \tau\} \quad \text{and} \quad \overline{Q}(\tau) = \inf\{x : \underline{F}(x) \geq \tau\}.$$

In words, a lower bound for the quantile function is the left inverse of an upper bound for the CDF, and viceversa. These calculations are straightforward to implement (e.g., by using a grid of values of x).

For functionals of Q, such as the interquartile range or the Kelley's measure of skewness, a sequential approach like the one outlined for centered moments is applicable.

3.4 Numerical illustration and practical implementation

In this section we demonstrate the performance of our bounds in an empirical exercise using one-year ahead inflation expectations data from the SCE and compare the resultant bounds to the outcomes from existing approaches.[14] As part of this exercise, we also detail our recommendations for how to implement the bounds in practice.

3.4.1 SCE inflation expectations

Just as in Section 2 we use Question Q9 from the SCE which asks about the anticipated inflation rate over the next twelve months. Like most density questions, respondents are allowed to place mass in bins with unbounded support (inflation or deflation which is "12%

[13] In numerical experiments where both problems can be solved, we found the degree of conservativeness of the alternative approach to be small but the computational gains to be large.

[14] A widely used simplification assigns uniform probability within each bin, which corresponds to assuming the density is piecewise constant over the partition. We do not include comparisons to this approach because the assumption that a density jumps to zero at bin boundaries is typically adopted for computational convenience, not as a genuine belief about the shape of the distribution. Our framework nests this as a special case and, by imposing smoothness restrictions, always delivers weakly tighter bounds.

or higher"). To accommodate these responses we need to place some restrictions on the tail behavior of the true density.

We take the following approach for the tail bound parameters $(\mathcal{T}, \gamma)$. We impose that f is dominated by a function strictly increasing on $\mathcal{T}_{\text{lo}} = (-\infty, x_{\text{lo}}]$ and strictly decreasing on $\mathcal{T}_{\text{up}} = [x_{\text{up}}, \infty)$ with exponential rate of decay λ. Let $\tilde{f}_{\text{lo}} = f(x_{\text{lo}})$ and $\tilde{f}_{\text{up}} = f(x_{\text{up}})$. We require that for all m such that $x_m \leq x_{\text{lo}}$,

$$\tilde{f}_m \leq \tilde{f}_{\text{lo}} \cdot e^{-\lambda(x_{\text{lo}} - x_m)}$$

and for all m such that $x_m \geq x_{\text{up}}$,

$$\tilde{f}_m \leq \tilde{f}_{\text{up}} \cdot e^{-\lambda(x_m - x_{\text{up}})}$$

This transforms the problem of choosing $\mathcal{T} = \mathcal{T}_{\text{lo}} \cup \mathcal{T}_{\text{up}}$ and γ into selecting $(x_{\text{lo}}, x_{\text{up}}, \lambda)$. Given values for these three parameters, the tail bound restrictions can be written as linear inequality constraints in the format

$$\Lambda \tilde{f} \leq 0_{M_{\mathcal{T}} \times 1}$$

where $M_{\mathcal{T}}$ is the cardinality of the set $\mathcal{X}_M \cap \mathcal{T}$ and the entries of the matrix Λ are either 0, 1, $-e^{-\lambda(x_{\text{lo}} - x_m)}$ or $-e^{-\lambda(x_m - x_{\text{up}})}$.

To simplify the implementation of the tail bound, rather than choose λ directly we instead choose two points on the support $(x_{\mathsf{LT}}, x_{\mathsf{UT}})$ along with associated upper bounds for the density evaluated at these points, $f(x_{\mathsf{LT}})$ and $f(x_{\mathsf{UT}})$. Based on existing work using the SCE (Armantier, Giorgio, van der Klauw, and Zafar, 2016), we choose $(x_{\mathsf{LT}}, x_{\mathsf{UT}}) = (-38, 38)$ along with the bounds $f(-38) \leq 10^{-4}$ and $f(38) \leq 10^{-4}$. Then, we can utilize the following choice of λ for the lower and upper tail,

$$\lambda_{\mathsf{LT}} = -\frac{\log(10^{-4})}{x_{\text{lo}} - (-38)}, \qquad \lambda_{\mathsf{UT}} = -\frac{\log(10^{-4})}{38 - x_{\text{up}}}.$$

Of course, the choice of tail bound parameters will play a larger role when the target function G is unbounded. When the researcher asks for bin probabilities (case D-1) and the bins are all bounded or if some intervals are unbounded but the respondent places zero probability on them, we recommend to omit the tail bound restrictions.

We choose an equally-spaced grid, with the distance between gridpoints equal to 0.076,

so that $M + 1 \leq 1,000$. This choice strikes a balance between reducing discretization error through a larger M but also limiting the size of M to eliminate any numerical problems when implementing the linear program. When the respondent places mass in both the first and last interval, this choice of gap produces $M + 1 = 1,000$ equally spaced points on $[-38, 38]$.

Here we recommend taking (p, ϕ) to allow our approach to accommodate a large class of standard distributions (e.g., the normal, beta, and gamma) with a reasonable grid size M. To ensure adequate smoothness in the class of densities we choose $p = 4$. For ϕ we follow a common prescription in the nonparametric literature. We first fit a pilot density to the respondent's bin probabilities. We choose two candidate pilot densities: a generalized beta density, as in EMW, and a skew normal density, as in García and Manzanares (2007). We then take the pilot density which has the smallest root mean-square error and calculate the corresponding discrete 4th derivative and take its maximum absolute value as ϕ. We will denote this choice as ϕ_{pilot}. We will also report some results based on a choosing ϕ to be the minimum value for which the programs in equation (5) have a solution. We will denote this choice as ϕ_{min}.

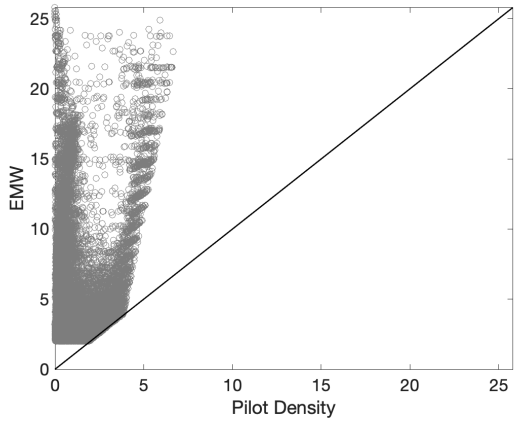
To construct our dataset we use the publicly available data from the SCE over the period from June 2013 through June 2024. We omit observations where the respondent reports 100% probability in a single interval. We also omit observations where the respondent reports 0% probability in an interval which is adjacent to intervals with non-zero probability assigned. This results in a dataset comprising 140,942 observations of which 65,419 responses place positive probability in at least one of the two extreme bins.

We first study the mean on the respondent's forecast density. To provide a comparison, we also construct the "hands off" bounds for the mean using the procedure from EMW. When the respondent places non-zero probability in the first or last bin we construct the EMW bound by placing all the mass at -38 or 38, as appropriate. The top left chart in Figure 3 presents a scatter plot comparing our new bounds using ϕ_{pilot} and the EMW bounds. By construction, all of the data points are above the 45 degree line; however, the graph shows that the improvement in precision can be substantial. This is clearly illustrated in the top right chart of Figure 3 which summarizes the relative interval length between the two bounds. The median improvement is more than one percentage point and for 25% of the respondents the improvement is at least 3.7 percentage points.

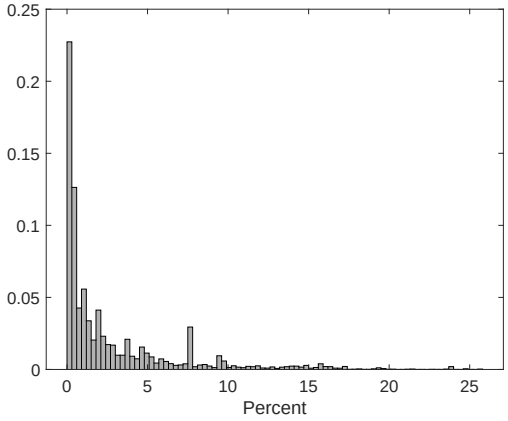
We can investigate the role of the pilot density choice by comparing to the bounds derived using ϕ_{min}. The bottom left chart shows that this choice produces much tighter bounds than either those using ϕ_{pilot} or those using the approach of EMW. That said, the

FIGURE 3. **Bounds on the Mean: SCE One-Year Ahead Inflation Expectations.** This figure presents the properties of different bounds for the mean of the forecast density. The bounds introduced in Section 3 are designated $\phi_{\min}$ or ϕ_{pilot} denoting the choice of ϕ. The nonparametric bounds from Engelberg et al. (2009) are denoted by EMW. Based on data from the Survey of Consumer Expectations over the sample period 2013m6–2024m6.

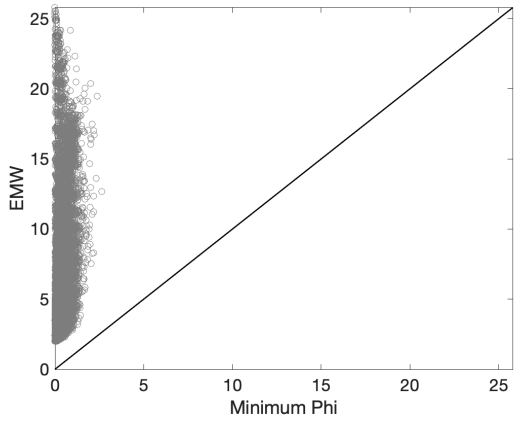
(a) *Bounds: EMW vs.* ϕ_{pilot}



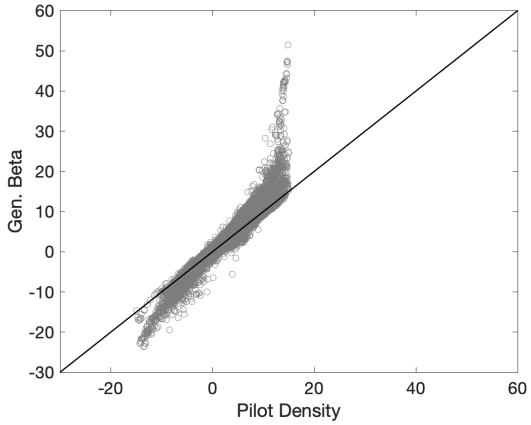
(b) *Interval Length: EMW less* ϕ_{pilot}



(c) *Bounds: EMW vs.* $\phi_{\min}$



(d) *Parametric vs. Midpoint of Bound*

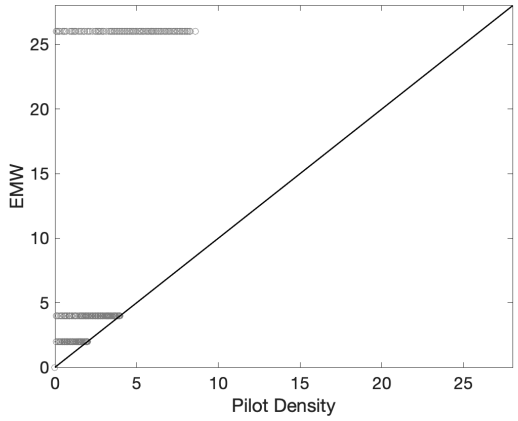


$\phi_{\min}$ bounds do not compress to a point and can be as large as 2.6 percentage points. This result underscores the importance of utilizing bounds on functionals of the density. Even the tightest possible nonparametric bounds, subject to the smoothness assumptions, result in meaningful uncertainty about the mean of the forecast density. The bottom right chart of Figure 3 provides a similar note of caution. This chart is a scatter plot between the mean based on a fitted generalized beta distribution and the midpoint of the nonparametric bounds for the mean using ϕ_{pilot}. The cloud of points deviates meaningfully from the 45 degree line with especially steeper slopes for more extreme values. Thus, the parametric approach can produce point estimates of the mean that may be far away from the center of

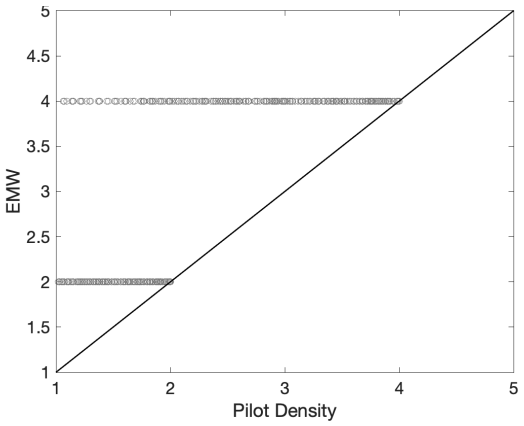
the nonparametric bound.

FIGURE 4. **Bounds on the Median: SCE One-Year Ahead Inflation Expectations.** This figure presents the properties of different bounds for the median of the forecast density. The bounds introduced in Section 3 are designated ϕ_{min} or ϕ_{pilot} denoting the choice of ϕ. The nonparametric bounds from Engelberg et al. (2009) are denoted by EMW. Based on data from the Survey of Consumer Expectations over the sample period 2013m6–2024m6.

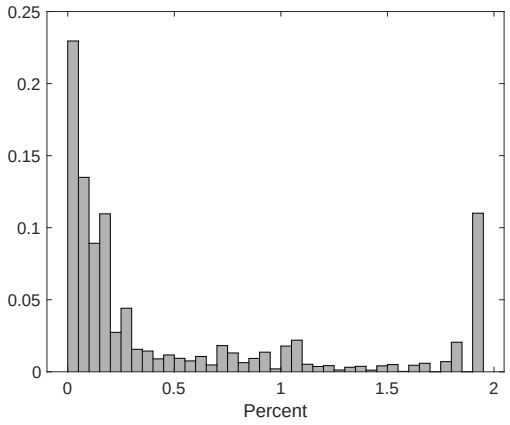
(a) *Bounds: EMW vs.* ϕ_{pilot}



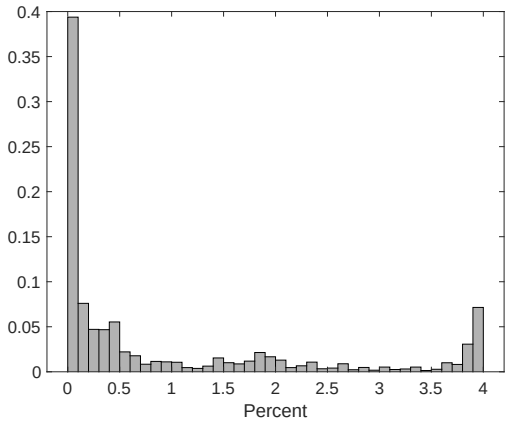
(b) *Bounds: EMW vs.* ϕ_{pilot}



(c) *Interval Length: EMW less* ϕ_{pilot} *(2%)*



(d) *Interval Length: EMW less* ϕ_{pilot} *(4%)*



We next provide the corresponding results for the median of the respondent's forecast distribution. Given ϕ, we can construct lower and upper bounds for the median via the following steps:

Step 1: Set $G(x) = \mathbf{1}\{x \leq x_0\}$ and obtain lower and upper bounds for $F(x_0)$ using the procedure from Section 3. Denote the lower and upper bound as $\underline{F}_0$ and $\overline{F}_0$, respectively.

Step 2: Repeat Step 1 for each $G(x) = \mathbf{1}\Big\{x \leq x_j\Big\}$ with $j = 1, \ldots, M$ and denote the lower and upper bound as $\underline{F}_j$ and $\overline{F}_j$, respectively.

Step 3: Construct the index set $\left\{j : \underline{F}_j \leq 0.5 \leq \overline{F}_j\right\} = \{j^* + 1, j^* + 2, \ldots, j^* + m\}$ and set the lower bound for the median as x_{j^*} and the upper bound as x_{j^*+m+1}.

Figure 4 presents the the bounds for the median using ϕ_{pilot} as compared to the bounds from EMW. The top left chart of Figure 4 displays a scatter plot of the two bounds. Since the EMW bounds can only take on the value of the length of the bin interval, the data are constant at $\{2, 4, 26\}$. In contrast, the bounds using ϕ_{pilot} almost always improve upon the EMW bounds. The top right chart focuses on respondents whose median is outside of the two extreme bins. This chart shows that the precision improvement can be substantial. This is further illustrated in the two bottom charts of Figure 4 which shows the difference in interval lengths for bins of length 2% (left chart) and length 4% (right chart). The improvement can be as large as 1.9% and 3.9%, respectively. Thus, just as for the case of the mean, we have shown that substantial gains in precision may be obtained by using the bounds introduced in Section 3.

4 Regression analysis of bounded data

A popular use of expectations measures is to use them as inputs to subsequent regression analyses. Consider the following linear regression,

$$y_i^* = \beta x_i^* + \gamma' w_i + \varepsilon_i, \qquad i = 1, \ldots, N, \tag{6}$$

where y_i^* is an outcome variable, x_i^* is the variable of interest, and w_i are a set of control variables including a constant. We will always assume we observe w_i; however, we will assume that y_i^* or x_i^* (or both) are based on survey responses and so are interval valued, i.e.,

$$\underline{y}_i \leq y_i^* \leq \overline{y}_i, \qquad \underline{x}_i \leq y_i^* \leq \overline{x}_i.$$

These lower and upper bounds could be, for example, obtained from the approach introduced in Section 3.

Typically, we would estimate β via ordinary least squares

$$\hat{\beta} = \frac{\sum_{i=1}^N x_i^*(y_i^* - \hat{\pi}_y' w_i)}{\sum_{i=1}^N x_i^*(x_i^* - \hat{\pi}_x' w_i)} \tag{7}$$

where $\hat{\pi}_y = \left(\sum_{i=1}^N w_i w_i'\right)^{-1} \sum_{i=1}^N w_i y_i^*$ and $\hat{\pi}_x = \left(\sum_{i=1}^N w_i w_i'\right)^{-1} \sum_{i=1}^N w_i x_i^*$. Because we cannot

observe both y_i^* and x_i^*, this is an infeasible estimator but serves as a natural target for our estimation approach. In particular, we will aim to find informative bounds in the guise of estimators $\underline{\hat{\beta}}$ and $\overline{\hat{\beta}}$ which satisfy,

$$\underline{\hat{\beta}} \leq \hat{\beta} \leq \overline{\hat{\beta}}.$$

Before proceeding it will be convenient to stack the elements of equation (6). Let $y = (y_1, \ldots, y_N)'$, $x = (x_1, \ldots, x_N)'$ and $P = I_N - w(w'w)^{-1}w'$ with $w = (w_1, \ldots, w_N)'$. Then, we can rewrite the OLS estimator of equation (7) as,

$$\hat{\beta} = \frac{x'Py}{x'Px},$$

which suggests choosing $\underline{\hat{\beta}} = B_{\mathsf{min}}$ as

$$B_{\mathsf{min}} = \min_{x,y \in \mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad y_i \in [\underline{y}_i, \overline{y}_i], \; x_i \in [\underline{x}_i, \overline{x}_i] \quad \forall i \tag{8}$$

and $\overline{\hat{\beta}} = B_{\mathsf{max}}$ as

$$B_{\mathsf{max}} = \max_{x,y \in \mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad y_i \in [\underline{y}_i, \overline{y}_i], \; x_i \in [\underline{x}_i, \overline{x}_i] \quad \forall i. \tag{9}$$

In general, both equations (8) and (9) represent non-convex optimization problems which are computationally infeasible—especially for large values of N. In some special cases we can obtain sharp bounds. In all other cases we suggest simple, computationally-convenient bounds. In Section 5 we show that these bounds are informative in practice providing new information about an important structural parameter for understanding households' consumption decisions.

We separate the general problem into three different cases which we treat accordingly:

Case 1: x_i^* is Observed In this case we aim to solve

$$B_{1\mathsf{min}} = \max_{y \in \mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad y_i \in [\underline{y}_i, \overline{y}_i], \qquad i = 1, \ldots, N$$

and

$$B_{1\mathsf{max}} = \min_{y \in \mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad y_i \in [\underline{y}_i, \overline{y}_i], \qquad i = 1, \ldots, N.$$

This is the simplest case as both optimization problems are linear programs and can be solved by standard methods. Further, these are the best available bounds.

Case 2: y_i^* is Observed In this case we aim to solve

$$B_{2\mathsf{min}} = \min_{x \in \mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i], \qquad i = 1, \ldots, N,$$

and

$$B_{2\mathsf{max}} = \max_{x \in \mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i], \qquad i = 1, \ldots, N.$$

Under some circumstances we can solve these fractional programs directly by applying the method of Dinkelbach (1967) to obtain $B_{2\mathsf{min}}$ and $B_{2\mathsf{max}}$. In particular, if $b_{2\mathsf{min}}^{xy} < 0$ where

$$b_{2\mathsf{min}}^{xy} = \min_{x \in \mathbb{R}^N} x'Py, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i] \quad \forall i,$$

then $B_{2\mathsf{min}}$ can be found using Dinkelbach's method, and similarly if $b_{2\mathsf{max}}^{xy} > 0$ where

$$b_{2\mathsf{max}}^{xy} = \max_{x \in \mathbb{R}^N} x'Py, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i] \quad \forall i,$$

then $B_{2\mathsf{max}}$ can also be found using Dinkelbach's method. When these conditions fail to hold we can still find simple-to-implement bounds using the following result.

Proposition 2. *Let D_P be a diagonal matrix with diagonal entries the same as P. Further, let $P - D_P$ have the spectral decomposition, $P - D_P = HLH'$, with the (i,j) element of H denoted by H_{ij}. Define,*

$$\underline{z}_j = \sum_{i=1}^N \left(\underline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} > 0\} + \overline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} < 0\} \right),$$

and

$$\overline{z}_j = \sum_{i=1}^N \left(\overline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} > 0\} + \underline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} < 0\} \right).$$

Then,

(*i*) $\underline{B}_{\mathsf{2min}} \leq B_{\mathsf{2min}}$ *where* $\underline{B}_{\mathsf{2min}}$ *is the unique solution to* $\underline{\mathsf{C}}_2^\star(\underline{B}_{\mathsf{2min}}) = 0$ *where*

$$\begin{aligned}\underline{\mathsf{C}}_2^\star(\ell) = &\max_{x\in\mathbb{R}^N}\left\{-x'D_Py - \ell\cdot x'D_Px\right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i] \quad \forall i \\ &+ \max_{z\in\mathbb{R}^N}\left\{-z'L(H'y) - \ell\cdot z'Lz\right\}, \quad s.t. \quad z_i \in [\underline{z}_i, \overline{z}_i] \quad \forall i,\end{aligned}$$

(*ii*) $\overline{B}_{\mathsf{2max}} \geq B_{\mathsf{2max}}$ *where* $\overline{B}_{\mathsf{2max}}$ *is the unique solution to* $\overline{\mathsf{C}}_2^\star(\overline{B}_{\mathsf{2max}}) = 0$ *where*

$$\begin{aligned}\overline{\mathsf{C}}_2^\star(\ell) = &\max_{x\in\mathbb{R}^N}\left\{x'D_Py - \ell\cdot x'D_Px\right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i] \quad \forall i \\ &+ \max_{z\in\mathbb{R}^N}\left\{z'L(H'y) - \ell\cdot z'Lz\right\}, \quad s.t. \quad z_i \in [\underline{z}_i, \overline{z}_i] \quad \forall i\end{aligned}$$

Proposition 2 provides lower and upper bounds, $\underline{B}_{\mathsf{2min}}$ and $\overline{B}_{\mathsf{2max}}$, which can be calculated quickly and efficiently. This is because the maximization problems which underpin $\underline{\mathsf{C}}_2^\star(\ell)$ and $\overline{\mathsf{C}}_2^\star(\ell)$ are easy to solve (see Lemmas 1 and 2 in the Appendix) and $\underline{\mathsf{C}}_2^\star(\ell)$ and $\overline{\mathsf{C}}_2^\star(\ell)$ are monotonic functions of ℓ which enables straightforward calculation of the optimum. The result exploits two key properties of our problem. First, the optimization problem depends on a projection matrix, P, and so is heavily influenced by its diagonal elements as the sample size grows. To see this, recall that by idempotency we have that $P_{ii} = \sum_{j=1}^N P_{ij}^2$. Second, we can easily solve non-convex optimization of linear and quadratic programs when the associated matrix is diagonal. We combine these two properties to construct our bounds. We decompose P as D_P and $(P - D_p)$. Since $(P - D_p)$ is not diagonal we use a spectral decomposition and map the optimization problem to $z = H'x$. It can then be shown that taking these two steps results in valid bounds for B_{2min} and B_{2max}.

Case 3: Neither y_i^* nor x_i^* is Observed In this case we aim to solve

$$\min_{x,y\in\mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad y_i \in [\underline{y}_i, \overline{y}_i],\ x_i \in [\underline{x}_i, \overline{x}_i], \qquad i = 1, \ldots, N,$$

and

$$\max_{x,y\in\mathbb{R}^N} \frac{x'Py}{x'Px}, \quad s.t. \quad y_i \in [\underline{y}_i, \overline{y}_i],\ x_i \in [\underline{x}_i, \overline{x}_i], \qquad i = 1, \ldots, N.$$

If we fix $\ell \in \mathbb{R}$ then we have,

$$x'Py - \ell \cdot x'Px = \begin{pmatrix} x & y \end{pmatrix}' \begin{bmatrix} -\ell P & P/2 \\ P/2 & 0 \end{bmatrix} \begin{pmatrix} x & y \end{pmatrix},$$

which is non-definite in general and so this is a non-convex quadratic program. Although there are some existing algorithms that may apply in this setting they are infeasible for anything but very small values of N. To circumvent these numerical issues we follow a similar approach as in Proposition 2 exploiting the special structure of our problem.

Proposition 3. *Let D_P be a diagonal matrix with diagonal entries the same as P. Further, let $P - D_P$ have the spectral decomposition, $P - D_P = HLH'$, with the (i,j) element of H denoted by H_{ij}. Define,*

$$\begin{aligned}
\underline{z}_{1j} &= \sum_{i=1}^{N} \Big(\underline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} > 0\} + \overline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} < 0\} \Big), \\
\overline{z}_{1j} &= \sum_{i=1}^{N} \Big(\overline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} > 0\} + \underline{x}_i \cdot H_{ij} \mathbf{1}\{H_{ij} < 0\} \Big), \\
\underline{z}_{2j} &= \sum_{i=1}^{N} \Big(\underline{y}_i \cdot H_{ij} \mathbf{1}\{H_{ij} > 0\} + \overline{y}_i \cdot H_{ij} \mathbf{1}\{H_{ij} < 0\} \Big),
\end{aligned}$$

and

$$\overline{z}_{2j} = \sum_{i=1}^{N} \Big(\overline{y}_i \cdot H_{ij} \mathbf{1}\{H_{ij} > 0\} + \underline{y}_i \cdot H_{ij} \mathbf{1}\{H_{ij} < 0\} \Big).$$

Then,

(i) $\underline{B}_{3\mathsf{min}} \leq B_{3\mathsf{min}}$ *where* $\underline{B}_{3\mathsf{min}}$ *is the unique solution to* $\underline{\mathsf{C}}_3^{\star}(\underline{B}_{3\mathsf{min}}) = 0$ *where*

$$\begin{aligned}
\underline{\mathsf{C}}_3^{\star}(\ell) = &\max_{x,y \in \mathbb{R}^N} \Big\{ -x'D_P y - \ell \cdot x'D_P x \Big\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i],\ y_i \in [\underline{y}_i, \overline{y}_i] \quad \forall i \\
&+ \max_{z_1, z_2 \in \mathbb{R}^N} \Big\{ -z_1'Lz_2 - \ell \cdot z_1'Lz_1 \Big\}, \quad s.t. \quad z_{1i} \in [\underline{z}_{1i}, \overline{z}_{1i}],\ z_{2i} \in [\underline{z}_{2i}, \overline{z}_{2i}] \quad \forall i
\end{aligned}$$

(ii) $\overline{B}_{3\mathsf{max}} \geq B_{3\mathsf{max}}$ *where* $\overline{B}_{3\mathsf{max}}$ *is the unique solution to* $\overline{\mathsf{C}}_3^{\star}(\overline{B}_{3\mathsf{max}}) = 0$ *where*

$$\begin{aligned}
\overline{\mathsf{C}}_3^{\star}(\ell) = &\max_{x \in \mathbb{R}^N} \Big\{ x'D_M y - \ell \cdot x'D_M x \Big\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i],\ y_i \in [\underline{y}_i, \overline{y}_i] \quad \forall i \\
&+ \max_{z \in \mathbb{R}^N} \Big\{ z_1'Lz_2 - \ell \cdot z_1'Lz_1 \Big\}, \quad s.t. \quad z_{1i} \in [\underline{z}_{1i}, \overline{z}_{1i}],\ z_{2i} \in [\underline{z}_{2i}, \overline{z}_{2i}] \quad \forall i
\end{aligned}$$

Proposition 3 follows the same reasoning as for Proposition 2. Again, we have that the

objective functions of interest, $\underline{\mathsf{C}}_3^\star(\ell)$ and $\overline{\mathsf{C}}_3^\star(\ell)$, are easy to compute and monotonic in ℓ. We use the bounds in Proposition 3 in our empirical application in Section 5.

5 Empirical application

In this section we demonstrate the appeal of our new procedures in an application using survey responses to estimate the consumption Euler equation. As argued in Christelis et al. (2020) and Crump et al. (2022), survey data have appealing properties for the estimation of Euler equations because they circumvent challenging statistical issues which arise using realized data.

The theoretical reference point for our empirical exercise is the standard intertemporal Euler equation, which encapsulates the optimal consumption and saving choice of a household that can freely borrow and lend in an asset with a given rate of return. In its most basic, and widely used, form, this equilibrium condition can be written as,

$$\left(\frac{1+r_i}{1+\delta_{i,t}}\right) \cdot E_i\left[U'(C_{i,t+1})\right] = U'(C_{i,t}) \tag{10}$$

where $U(\cdot)$ is the utility function, r_i is the relevant interest rate, and $\delta_{i,t}$ is the time preference rate. As shown in Dynan (1993) and Christelis et al. (2020) we can utilize a second-order Taylor series expansion of $U'(C_{i,t+1})$ around $U'(C_{i,t})$ which yields

$$\frac{U'(C_{i,t+1})}{C_{i,t}U''(C_{i,t})} \approx -\frac{1}{\xi} + GC_{i,t+1} - \frac{\rho}{2}GC_{i,t+1}^2, \qquad GC_{i,t+1} = \frac{C_{i,t+1}-C_{i,t}}{C_{i,t}}, \tag{11}$$

where ξ, the coefficient of relative risk aversion, and ρ, the coefficient of relative prudence are defined as

$$\xi = -\frac{C_{i,t}\cdot U''(C_{i,t})}{U'(C_{i,t})}, \qquad \rho = -\frac{C_{i,t}\cdot U'''(C_{i,t})}{U''(C_{i,t})}.$$

Plugging equation (11) into equation (10) yields,

$$E_i\left[GC_{i,t+1}\right] \approx \frac{1}{\xi}\left(\frac{r_i-\delta_{i,t}}{1+r_i}\right) + \frac{\rho}{2}E_i\left[GC_{i,t+1}^2\right]. \tag{12}$$

We follow Christelis et al. (2020) and use Equation (12) as the basis of our empirical exercise with the aim of estimating ρ. The value of ρ is important as it governs the strength of the precautionary savings motive. This is an ideal setting to highlight the strengths of our methods because it relies on regressions of first moments on second moments from the same

distribution. More specifically, we have the following regression specification,

$$ExpGC_{i,t+1} = \alpha + \beta \cdot ExpGC^2_{i,t+1} + \gamma \cdot w_{i,t} + \epsilon_{i,t+1}, \tag{13}$$

where $ExpGC_{i,t+1}$ is expected spending growth for $t+1$ formed at time t, $ExpGC^2_{i,t+1}$ is expected squared spending growth, and $w_{i,t}$ are additional control variables included to accommodate the presence of differential interest rates and discount rates across respondents.

To be able use survey data to estimate equation (13), we require both the first *and* second moments of the predictive distribution of spending growth. This will, in general, require using the same survey question(s) to generate both the dependent variable and the regressor of interest. This induces common variation in both the dependent and independent variables. If a parametric density is used, and it is correctly specified for each respondent, then this commonality does not pose any problems. However, under misspecification of the parametric density the estimates of ρ will be biased. Christelis et al. (2020) estimate equation (13) using data from the CentER Internet Panel and assuming the underlying density for each respondent is either a simple triangular distribution or a split triangle distribution. In either case, the assumption of a triangular density is restrictive and at risk for misspecification relative to the true density. In contrast, we can circumvent the misspecification issue entirely by using our nonparametric bounds and associated regression approach, introduced in Sections 3 and 4, as they do not rely on any specific parametric assumptions.

Our primary data source is the SCE Household Spending Survey[15] which is fielded every four months, since August 2015, as a rotating module as part of the of the monthly SCE. We merge the Household Spending Survey data with the core monthly SCE to obtain demographic and other characteristics for each respondent. We restrict the sample to the pre-COVID period of August 2015 through December 2019. To obtain information on the subjective distribution of $GC_{i,t+1}$ we rely on the following question (Question QSP7dens) in the SCE Household Spending Survey which queries households on their anticipated spending growth over the following 12 months:

[15] See https://www.newyorkfed.org/microeconomics/sce/household-spending#/ for more information.

Question QSP7dens Now we would like you to think about the different things that may happen to the total spending of all members of your household (including you) **over the next 12 months**.
What do you expect will happen to the total spending of all members of your household (including you) **over the next 12 months**?

increase by 12% or more	____ percent chance
increase by 8% to 12%	____ percent chance
increase by 4% to 8%	____ percent chance
increase by 2% to 4%	____ percent chance
increase by 0% to 2%	____ percent chance
decrease by 0% to 2%	____ percent chance
decrease by 2% to 4%	____ percent chance
decrease by 4% to 8%	____ percent chance
decrease by 8% to 12%	____ percent chance
decrease by 12% or more	____ percent chance

Based on this density question we are able to construct lower and upper bounds for each respondent's first and second moments of spending growth using the results in Section 3. For this empirical application we rely on $\phi_{\min}$ for implementation. Since we cannot directly observe either $ExpGC_{i,t+1}$ or $ExpGC^2_{i,t+1}$ then we must rely on the regression bounds approach summarized in Case 3 of Section 4. We report results based on four different specifications for $w_{i,t}$:

spec1 $w_{i,t} = 1$ (in other words, no controls);

spec2 SPEC1 plus month effects for each survey month in the sample;

spec3 SPEC2 plus a range of respondent characteristics[16];

spec4 Survey respondent fixed effects only.

[16]In particular, we include indicator variables for two gender categories, three age categories ("under 40", "40 to 60", and "over 60"), two numeracy categories ("high" and "low" numeracy), four geographic categories ("Midwest", "Northeast", "South", and "West"), three education categories ("high school", "some college", and "college and beyond"), and three income categories ("Under 50K", "50K to 100K", and "Over 100K").

TABLE 1. **Set Estimates of the Coefficient of Relative Prudence** This table presents the set estimate and corresponding confidence interval for the coefficient of relative prudence based on equation (13). Each column reports results for one of the four specifications described in the text. Bootstrap-based confidence intervals are based on $B = 399$ bootstrap draws.

	spec1	**spec2**	**spec3**	**spec4**
Set Estimate for ρ	$[5.40, 8.51]$	$[5.38, 8.52]$	$[4.65, 11.02]$	$[1.45, 12.80]$
99% Set Confidence Interval	$[4.86, 9.06]$	$[4.77, 9.17]$	$[3.94, 12.22]$	$[0.40, 14.98]$

Table 1 presents the set estimates for the coefficient of relative prudence across these four different specifications. The set estimates for ρ in the first three specification all indicate a value of the prudence parameter of at least 4.65. This implies that households display a strong precautionary motive, contrary to the findings by Dynan (1993), who estimates a consumption Euler equation using realized consumption growth and the realized consumption growth variability from the Consumer Expenditure Survey. Merrigan and Normandin (1996) applies a similar methodology to the UK Family Expenditure Survey—a repeated cross-section—but instrument realized consumption growth and the variability of consumption growth by projecting consumption $t + 1$ to t based on observable characteristics. They find relative prudence parameters between 1.8 and 2.3 for UK households, in contrast to the findings by Dynan (1993). Refining the estimation approach by Dynan (1993), Bertola et al. (2005) use subjective variance of earnings growth, assuming a triangular distribution for the probability of earnings, from the Survey of Household Income and Wealth as an instrument for the realized variability of consumption growth and find a relative prudence parameter of 2, indicating a strong precautionary motive. Similarly, Christelis et al. (2020) find a relative prudence parameter around 2 for Dutch households, providing empirical support for households' precautionary motives.

The estimates presented in Table 1 suggest an even stronger precautionary savings motive for the US households compared to the estimates reported by Merrigan and Normandin (1996), Bertola et al. (2005), and Christelis et al. (2020). Instead, our estimates are more in line with those reported by Hubbard et al. (1994), who review the empirical estimates and pick a value of 4 to calibrate the relative prudence parameter in their dynamic model.

Note that, the last column of Table 1 displays a much larger set (and thus, a larger confidence interval) for the estimated relative prudence parameter. This is expected, as we include specification 4 (SPEC4) only as a comparison to better understand our new approach. In practice, we have only 2–3 time-series observations per respondent due to the rotating

panel nature of the survey[17] and since the SCE Household Survey is only implemented three times per year. This feature of the data would be expected to result in imprecise estimates and less informative bounds on the relative prudence parameter. This is, in fact, the case as our set estimate and corresponding confidence interval becomes substantially wider.

To characterize the uncertainty around our set estimates we implement a panel nonparametric bootstrap. Specifically, we sample with replacement from the collection of respondents which produces unbalanced bootstrapped panel datasets of varying size. We then construct set estimates (i.e., lower and upper bounds) across 399 bootstrap samples. We then form the 99% set confidence interval by choosing the $(.01/2)$th quantile of the bootstrapped lower bounds and the $(1-.01/2)$th quantile of the bootstrapped upper bounds. In all four columns, the 99% confidence intervals exclude zero indicating that these data are compatible with a (strong) precautionary saving motive.

6 Conclusion

The elicitation of forecast distributions from survey respondents is now commonplace in surveys of consumers, firms, and professional forecasters. Almost exclusively, respondents are prompted to provide their forecast distribution in the form of a small number of probabilities assigned to intervals which partition the support of the target variable. To utilize these forecast density responses, it is now ubiquitous to rely on the best fit from a chosen family of parametric density functions. With this fitted density, functionals such as the mean or standard deviation can be obtained easily and used in any subsequent analyses.

However, there are clear deficiencies of this popular approach. The fitted density may not provide a close approximation to the actual survey responses. Furthermore, even if the fitting error is zero, there may be multiple densities which are each consistent with the reported probabilities. We circumvent these shortcomings by introducing a new framework which recovers nonparametric upper and lower bounds on functionals of the underlying density relying only on the imposition of a mild smoothness assumption. We provide specific guidance on implementation including efficient algorithms to compute the bounds. Using data from the Survey of Consumer Expectations, we show that our new approach delivers informative bounds for the mean and median of the distribution with pronounced improvements relative to existing approaches (Engelberg et al., 2009). We also show how to use the

[17] SCE respondents rotate off the survey after a maximum of a twelve month tenure.

resultant bounds in second-stage regression models providing valid approaches for estimation and inference. We illustrate this methodology in an application to the consumption Euler equation and provide set estimates of the coefficient of relative prudence.

A Proofs

Section 3

In addition to the continuous and discrete derivative operators, D^p and Δ_h^p, we introduce notation for the scaled backward finite difference operator acting on functions. For any function $U : \mathcal{X} \to \mathbb{R}$, we let $\nabla_h^p = \nabla_h \circ \nabla_h^{p-1}$ where

$$(\nabla_h U)(x_m) = \frac{U(x_m) - U(x_{m-1})}{h},$$

If follows that if $u_m = U(x_m)$ for each $x_m \in \mathcal{X}_M$, then $(\nabla_h^p U)(x_m) = (\Delta_h^p u)_m$ for all $p \geq m$.

We prove Proposition 1 for the lower bound and omit the proof for the upper bound as it is entirely analogous. The argument relies on three technical lemmas that we establish in Supplemental Appendix B.

Proof of Proposition 1. We show that $\lim_{M\to\infty} \underline{\theta}_M(\phi) = \underline{\theta}(\phi)$ in two steps by showing that $\limsup_{M\to\infty} \underline{\theta}_M(\phi) \leq \underline{\theta}(\phi)$ and that $\liminf_{M\to\infty} \underline{\theta}_M(\phi) \geq \underline{\theta}(\phi)$.

Limit superior. Take any $f \in \mathcal{F}_D$ and define its scaled discretized version $\tilde{f}_m^{(M)} = hf(x_m)$. By Lemma 1, $\tilde{f}^{(M)}$ satisfies the smoothness constraint

$$\left\|\Delta_h^p \tilde{f}^{(M)}\right\|_\infty \leq h\phi.$$

The vector $\tilde{f}^{(M)}$ need not satisfy the probabilistic constraints $\sum_{x_m \in A_k} \tilde{f}_m^{(M)} = \pi_k$, all k. Thus, create a new vector $\hat{f}^{(M)}$ by $\hat{f}_m^{(M)} = \tilde{f}_m^{(M)} \times \pi_{k(m)} / (\sum_{x_j \in A_{k(m)}} \tilde{f}_j^{(M)})$ where $k(m)$ is the index of the set A_k to which x_m belongs (there is a single index because $\{A_k\}_{k=1}^K$ is a partition of $\mathcal{X}$). Now, $\hat{f}^{(M)}$ satisfies the probabilistic constraints and, by linearity of Δ_h^p, Lemmas 2 and 3 and a tedious calculation, for a constant C_1 independent of f,

$$\left\|\Delta_h^p \hat{f}^{(M)}\right\|_\infty \leq (1 + hC_1)\, h\phi.$$

That is, $\hat{f}^{(M)} \in \mathcal{F}_{D,M}(\phi_h)$ where $\phi_h = (1 + hC_1)\phi = \phi + O(h)$.

Since G is bounded over $\mathcal{X}$, a useful result is the following:

$$0 \leq \underline{\theta}_M(\phi) - \underline{\theta}_M(\phi_h) \leq \frac{2hC_1}{1 + hC_1} \|G\|_\infty,$$

where the upper bound is obtained by taking some $\hat{f}^* \in \mathcal{F}_{D,M}(\phi_h)$ that attains the minimum in the ϕ_h-problem and defining $(\phi/\phi_h)\hat{f}^* + (1-\phi/\phi_h)\cdot 1_{(M+1)\times 1}/(M+1) \in \mathcal{F}_{D,M}(\phi)$.

By assumption, we have $Gf \in \mathrm{BV}(\mathcal{X})$, and by the product rule for total variation and Lemma 3, $V_{\mathcal{X}}(Gf) \leq \|G\|_\infty V_{\mathcal{X}}(f) + \|f\|_\infty V_{\mathcal{X}}(G) \leq C_2$ for some constant C_2. Then,

$$\begin{aligned}\underline{\theta}_M(\phi) \leq \underline{\theta}_M(\phi_h) + O(h) \leq \sum_{m=0}^{M} G(x_m)\hat{f}_m^{(M)} + O(h) &= h\sum_{m=0}^{M} G(x_m)f(x_m) + O(h) \\ &= \int_{\underline{x}}^{\overline{x}} G(x)f(x)\,dx + O(h).\end{aligned}$$

Taking supremum over $f \in \mathcal{F}_D(\phi)$ and letting $M \to \infty$, we get $\limsup_{M\to\infty} \underline{\theta}_M(\phi) \leq \underline{\theta}(\phi)$.

Limit inferior. For each M, we let $\hat{f}^{(M)} \in \mathcal{F}_{D,M}(\theta)$ be a minimizer and define the discrete measure $\mu_M = \sum_{m=0}^{M} \hat{f}_m^{(M)}\delta_{x_m}$ where δ_{x_m} denotes Dirac mass on x_m. We will show below that, over subsequences, μ_M converges weakly to a measure μ with density $f \in \mathcal{F}_D(\phi)$. Since $G \in \mathrm{BV}(\mathcal{X})$, the set of points where G is discontinuous has zero Lebesgue measure. It will then follow that, again over subsequences, as $M \to \infty$,

$$\underline{\theta}_M(\phi) = \sum_{m=0}^{M} \hat{f}_m^{(M)} G(x_m) \longrightarrow \int_{\underline{x}}^{\overline{x}} f(x)G(x)\,dx \geq \underline{\theta}(\phi),$$

by the Portmanteau theorem, and that $\liminf_{M\to\infty} \underline{\theta}_M(\phi) \geq \underline{\theta}(\phi)$ as desired.

Since the set of probability measures on $\mathcal{X}$ is compact in the topology of weak convergence, for any subsequence of $\{\mu_M\}$ we can extract a further subsequence with limit μ. Our calculations below are along this (arbitrary) convergent subsequence, although we keep the index M for notational simplicity. Moreover, it is immediate that $\mu(A_k) = \pi_k$, all k, so all we need to show is that μ has a density that satisfies the smoothness restriction of the space $\mathcal{F}_D(\phi)$.

Let $g^{(M)}$ be the piecewise-constant interpolant of $h^{-1}(\Delta_h^p \hat{f}^{(M)})$ on $\mathcal{X}$. For any $x \in \mathcal{X}$,

$$\left|g^{(M)}(x)\right| \leq \max_{0\leq m\leq M} \left|g^{(M)}(x_m)\right| = \frac{\left\|(\Delta_h^p \hat{f}^{(M)})_m\right\|_\infty}{h} \leq \phi,$$

from where it follows that $\{g^{(M)}\}$ is a bounded sequence on $\mathcal{L}^\infty(\mathcal{X})$. By the Banach-Alaoglu theorem, $g^{(M)}$ converges weakly to g (along a further subsequence) with $\|g\|_\infty \leq \phi$.

Take an arbitrary test function $\varphi \in \mathcal{C}_c^\infty(\mathcal{X})$ and do summation by parts p times,

$$\begin{aligned}\int \varphi \, d\mu_M = h \sum_{m=0}^{M} \hat{f}_m^{(M)} \varphi(x_m) &= (-1)^p h \sum_{m=p}^{M} g^{(M)}(x_m) \cdot (\nabla_h^{-p}\varphi)(x_{m-p}) + O(h) \\ &= (-1)^p \int g^{(M)}(x)(S_h^p \varphi)(x)\, dx + O(h).\end{aligned}$$

where ∇_h^{-p} is the p-times discrete antiderivative operator (i.e., $\nabla_h^{-p}(\nabla_h^p u) = \nabla_h^p(\nabla_h^{-p} u) = u$), $S_h^p\varphi$ is the piecewise constant interpolant of $\nabla_h^{-p}\varphi$ on $\mathcal{X}$ and the $O(h)$ term is uniformly bounded by $C_3 h$ with a constant C_3 that only depends on φ and p.

Define the p-times continuous antiderivative of φ recursively: $\varphi^{(-p)}(x) = \int_{\underline{x}}^{x} \varphi^{(-p+1)}(\xi)\, d\xi$ starting at $\varphi^{(0)}(x) = \varphi(x)$. By convergence of Riemann sums, $\left\| S_h^p\varphi - \varphi^{(-p)} \right\|_\infty \to 0$ as $M \to \infty$ and, therefore,

$$\lim_{M\to\infty} \int \varphi \, d\mu_M = (-1)^p \int_{\underline{x}}^{\overline{x}} g(x)\varphi^{(-p)}(x)\, dx = \int_{\underline{x}}^{\overline{x}} f(x)\varphi(x)\, dx = \int \varphi \, d\mu,$$

where the first equality uses weak convergence of $g^{(M)}$ to g and uniform convergence of $S_h^p\varphi$ to $\varphi^{(-p)}$, the second equality uses integration by parts p times (integrating g while differentiating $\varphi^{(-p)}$), and the third uses weak convergence of μ_M to μ. In particular, the second equality implies that $D^p f = g$ and, in consequence, $\|D^p f\|_\infty \leq \phi$. Therefore, μ has density $f \in \mathcal{F}_D(\phi)$ which concludes the proof. □

Section 4

We start with two lemmas we utilize for the proofs of the propositions. The first lemma is used in the case where y_i is observed but x_i is interval-valued.

Lemma 1. *Let D be an $n \times n$ diagonal matrix with ith diagonal element $d_i \neq 0$. Also, let $a = (a_1, \ldots, a_n)' \in \mathbb{R}^n$, $b \in \mathbb{R}$, and $x = (x_1, \ldots, x_n)'$ where $\underline{x}_i \leq x_i \leq \overline{x}_i$ for $i = 1, \ldots, n$. Then, the solution to*

$$\max_x \Big\{ x'Da - b \cdot x'Dx \Big\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i], \qquad i = 1, \ldots, n,$$

is given by $x^* = (x_1^*, \ldots, x_n^*)'$ *with*

$$x_i^* = \begin{cases} \frac{a_i}{2b} & \text{if } A_i \text{ holds} \\ \underline{x}_i & \text{if } A_i \text{ fails to hold and } d_i\left((\underline{x}_i - \overline{x}_i)\, a_i - b\left(\underline{x}_i^2 - \overline{x}_i^2\right)\right) \geq 0 \\ \overline{x}_i & \text{if } A_i \text{ fails to hold and } d_i\left((\underline{x}_i - \overline{x}_i)\, a_i - b\left(\underline{x}_i^2 - \overline{x}_i^2\right)\right) < 0 \end{cases}$$

where A_i *is the condition that* $bd_i > 0$ *and* $\underline{x}_i \leq \frac{a_i}{2b} \leq \overline{x}_i$.

Proof. Since the objective function is quadratic, either the maximum occurs at the global maximum or it occurs at either of the endpoints: $\underline{x}_i$ or $\overline{x}_i$. Because D is diagonal we can work with each i at a time. Taking the first derivative yields, $d_i(a_i - 2bx_i)$, so that the candidate optimum is $\frac{a_i}{2b}$ which is a global maximum if $-2bd_i < 0$. Thus when the condition A_i holds then $x_i^* = \frac{a_i}{2b}$. When the condition A_i fails to hold then we need only evaluate the endpoints and choose which endpoint provides the highest value, i.e., comparing $d_i(a_i\underline{x}_i - b\underline{x}_i^2)$ and $d_i(a_i\overline{x}_i - b\overline{x}_i^2)$ □

The second lemma is used in the case where y_i and x_i are interval-valued.

Lemma 2. *Let* D *be an* $n \times n$ *diagonal matrix with* i*th diagonal element* $d_i \neq 0$. *Also, let* $b \in \mathbb{R}$, $x = (x_1, \ldots, x_n)'$, *and* $y = (y_1, \ldots, y_n)'$ *where* $\underline{x}_i \leq x_i \leq \overline{x}_i$ *and* $\underline{y}_i \leq y_i \leq \overline{y}_i$ *for* $i = 1, \ldots, n$. *Then, the solution to*

$$\max_{x,y} \left\{ x'Dy - b \cdot x'Dx \right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i],\ y_i \in [\underline{y}_i, \overline{y}_i], \qquad i = 1, \ldots, n,$$

is given by $x^* = (x_1^*, \ldots, x_n^*)'$ *and* $y^* = (y_1^*, \ldots, y_n^*)'$ *described by the following.*

$$(x_i^*, y_i^*) = \text{argmax}_{(x_i,y_i) \in \{(\overline{x}_i, \overline{y}_i), (\overline{x}_i, \underline{y}_i), (\underline{x}_i, \overline{y}_i), (\underline{x}_i, \underline{y}_i)\}} \left\{ d_i(x_i y_i - b x_i^2) \right\},$$

if $\underline{A}_i$ *and* $\overline{A}_i$ *fail to hold where* $\underline{A}_i$ *is the condition that* $bd_i > 0$ *and* $\underline{x}_i \leq \frac{\underline{y}}{2b} \leq \overline{x}_i$ *and* $\overline{A}_i$ *is the condition that* $bd_i > 0$ *and* $\underline{x}_i \leq \frac{\overline{y}}{2b} \leq \overline{x}_i$.

$$(x_i^*, y_i^*) = \text{argmax}_{(x_i,y_i) \in \{(\frac{\underline{y}}{2b}, \underline{y}_i), (\frac{\overline{y}}{2b}, \overline{y}_i)\}} \left\{ d_i(x_i y_i - b x_i^2) \right\},$$

if $\underline{A}_i$ *and* $\overline{A}_i$ *hold.*

$$(x_i^*, y_i^*) = \text{argmax}_{(x_i,y_i) \in \{(\frac{\underline{y}}{2b}, \underline{y}_i), (\overline{x}_i, \overline{y}_i), (\underline{x}_i, \overline{y}_i)\}} \left\{ d_i(x_i y_i - b x_i^2) \right\},$$

if $\underline{A}_i$ *holds and* $\overline{A}_i$ *fails to hold.*

$$(x_i^*, y_i^*) = \text{argmax}_{(x_i,y_i)\in\{(\frac{\overline{y}}{2b},\overline{y}_i),(\overline{x}_i,\underline{y}_i),(\underline{x}_i,\underline{y}_i)\}}\left\{d_i(x_i y_i - b x_i^2)\right\},$$

if $\underline{A}_i$ *fails to hold and* $\overline{A}_i$ *holds.*

Proof. Because D is diagonal we can work with each i at a time. Taking the partial derivatives with respect to x and y we obtain $d_i(y_i - 2bx_i)$ and dx_i. The second partial derivatives are then $-2bd_i$ and 0. Thus, no global maximum exists. We can then proceed by contradiction. Suppose that (x_i^*, y_i^*) are such that $x_i \in [\underline{x}_i, \overline{x}_i]$ and $\underline{y}_i < y_i < \overline{y}_i$; however, for fixed x_i^* the function is linear in y_i so that the criterion function can be increased by moving to either $\underline{y}_i$ or $\overline{y}_i$. Thus, the upper and lower bounds are the only candidates for y_i^*. Now, fixing $y_i^* = \underline{y}_i$ or $y_i^* = \overline{y}_i$, the function is quadratic and so, following similar steps as in the proof or Lemma 2, then, if $bd_i > 0$, $(\frac{\overline{y}}{2b}, \overline{y}_i)$ and $(\frac{y}{2b}, \underline{y}_i)$ are candidate solutions so long as $\underline{x}_i \leq \frac{y}{2b} \leq \overline{x}_i$ and $\underline{x}_i \leq \frac{\overline{y}}{2b} \leq \overline{x}_i$. □

Proof of Proposition 2. This proof follows by similar steps as the proof of Proposition 3 (using Lemma 1 instead of Lemma 2) and so is omitted.

□

Proof of Proposition 3. First define $\overline{\mathsf{C}}_3(\ell)$ and $\underline{\mathsf{C}}_3(\ell)$ as,

$$\overline{\mathsf{C}}_3(\ell) = \max_{x,y}\left\{x' D_P y - \ell \cdot x' D_P x\right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i],\ y_i \in [\underline{y}_i, \overline{y}_i], \qquad i = 1, \ldots, N,$$

and

$$\underline{\mathsf{C}}_3(\ell) = \max_{x,y}\left\{-x' D_P y - \ell \cdot x' D_P x\right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i],\ y_i \in [\underline{y}_i, \overline{y}_i], \qquad i = 1, \ldots, N.$$

Then, we have that

$$\begin{aligned}
\overline{\mathsf{C}}_3(\ell) &\leq \max_{x,y}\left\{x' D_P y - \ell \cdot x' D_P x\right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i],\ y_i \in [\underline{y}_i, \overline{y}_i], \quad i = 1, \ldots, N \\
&\quad + \max_{x,y}\left\{x'(P - D_P) y - \ell \cdot x'(P - D_P) x\right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i],\ y_i \in [\underline{y}_i, \overline{y}_i], \quad i = 1, \ldots, N \\
&\leq \max_{x,y}\left\{x' D_P y - \ell \cdot x' D_P x\right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i]\ y_i \in [\underline{y}_i, \overline{y}_i], \quad i = 1, \ldots, N \\
&\quad + \max_{z_1,z_2}\left\{z_1' L z_2 - \ell \cdot z_1' L z_1\right\}, \quad s.t. \quad z_{1i} \in [\underline{z}_{1i}, \overline{z}_{1i}]\ z_{2i} \in [\underline{z}_{2i}, \overline{z}_{2i}], \quad i = 1, \ldots, N \\
&= \overline{\mathsf{C}}_3^{\star}(\ell),
\end{aligned}$$

The second inequality follows because we transform $z_1 = H'x$ and $z_2 = H'y$ and transform the bounds accordingly to obtain $\underline{z}_{1i}$, $\overline{z}_{1i}$, $\underline{z}_{2i}$, and $\overline{z}_{2i}$. Note that if x^* and y^* solves

$$\max_{x,y} \left\{ x' \left(P - D_P \right) y - \lambda \cdot x' \left(P - D_P \right) x \right\}, \quad s.t. \quad x_i \in [\underline{x}_i, \overline{x}_i], \; y_i \in [\underline{y}_i, \overline{y}_i], \qquad i = 1, \ldots, N,$$

then it is also an eligible solution to

$$\max_{z_1,z_2} \left\{ z_1' \Lambda z_2 - \lambda \cdot z_1' \Lambda z_1 \right\}, \quad s.t. \quad z_{1i} \in [\underline{z}_{1i}, \overline{z}_{1i}], \; z_{2i} \in [\underline{z}_{2i}, \overline{z}_{2i}], \qquad i = 1, \ldots, N,$$

as $H'x^*$ and $H'y^*$. Thus, the latter optimization problem will produce a criterion function at least as large as the former optimization problem. Furthermore, Lemma 2 below shows how to easily solve this program with a diagonal matrix and box constraints. Now choose $\overline{B}_{3\max}$ so that $\overline{\mathsf{C}}_3^{\star}\left(\overline{B}_{3\max}\right) = 0$. Then, since $\overline{\mathsf{C}}_3\left(\overline{B}_{3\max}\right) \leq \overline{\mathsf{C}}_3^{\star}\left(\overline{B}_{3\max}\right) = 0$ then since $\overline{\mathsf{C}}_3\left(\cdot\right)$ is continuous and strictly decreasing (Dinkelbach, 1967) we have that $\overline{B}_{3\max} \geq B_{\max}$ so that $\overline{B}_{3\max}$ is a valid upper bound. We can follow similar steps to obtain a lower bound.

□

References

Armantier, O., T. Giorgio, W. van der Klauw, and B. Zafar (2016): "An overview of the Survey of Consumer Expectations." Federal Reserve Bank of New York Staff Report 800.

Bachmann, R., G. Topa, and W. van der Klaauw (2022): *Handbook of Economic Expectations*, Academic Press, 1st ed.

Bassetti, F., R. Casarin, and M. Del Negro (2024): "A Bayesian Approach for Inference on Probabilistic Surveys." Federal Reserve Bank of New York Staff Report 1025.

Becker, C., P. Duersch, T. Eife, and A. Glas (2021): "Extending the procedure of Engelberg et al. (2009) to surveys with varying interval-widths." Working Paper AWI Discussion Paper Series No. 707.

Bellemare, C., L. Bissonnette, and S. Kröger (2012): "Flexible Approximation of Subjective Expectations Using Probability Questions." *Journal of Business and Economic Statistics*, 30, 125–131.

Beresteanu, A., I. Molchanov, and F. Molinari (2011): "Sharp Identification Regions in Models with Convex Moment Predictions." *Econometrica*, 79, 1785–1821.

Bertola, G., L. Guiso, and L. Pistaferri (2005): "Uncertainty and consumer durables adjustment," *The Review of Economic Studies*, 72, 973–1007.

Bissonnette, L., J. De Bresser, A. van Soest, and H. Vonkova (2017): "Eliciting Subjective Survival Curves: Lessons from Partial Identification." *Journal of Business and Economic Statistics*, 35, 547–560.

Bull, A. D. (2012): "Honest adaptive confidence bands and self-similar functions," *Electronic Journal of Statistics*, 6, 1490–1516.

Cai, T. and M. G. Low (2004): "An adaptation theory for nonparametric confidence intervals." *Annals of Statistics*, 32, 1805–1840.

Christelis, D., D. Georgarakos, T. Jappelli, and M. van Rooij (2020): "Consumption Uncertainty and Precautionary Saving," *Review of Economics and Statistics*, 102, 148–161.

Crump, R. K., S. Eusepi, A. Tambalotti, and G. Topa (2022): "Subjective intertemporal substitution." *Journal of Monetary Economics*, 126, 118–133.

Dinkelbach, W. (1967): "On Nonlinear Fractional Programming," *Management Science*, 13, 492–498.

Dynan, K. E. (1993): "How Prudent are Consumers?" *Journal of Political Economy*, 101, 1104–1113.

Engelberg, J., C. F. Manski, and J. Williams (2009): "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters." *Journal of Business and Economic Statistics*, 27, 30–41.

García, J. A. and A. Manzanares (2007): "What Can Probability Forecasts Tell Us About Inflation Risks?" Working paper, ECB.

Giné, E. and R. Nickl (2010): "Confidence bands in density estimation," *Annals of Statistics*, 38, 1122–1170.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994): "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," in *Carnegie-Rochester Conference Series on Public Policy*, vol. 40, 59–125.

Low, M. G. (1997): "On nonparametric confidence intervals." *Annals of Statistics*, 25, 2547–2554.

Manski, C. F. and E. Tamer (2002): "Inference on Regressions with Interval Data on a Regressor or Outcome." *Econometrica*, 70, 519–546.

Merrigan, P. and M. Normandin (1996): "Precautionary Saving Motives: An Assessment from UK Time Series of Cross-Sections," *The Economic Journal*, 106, 1193–1208.

Molinari, F. (2020): "Microeconometrics with Partial Identification." in *Handbook of Econometrics*, ed. by J. J. Heckman and E. E. Leamer, Elsevier, vol. 7A, 355–486.

Stone, C. J. (1980): "Optimal rates of convergence for nonparametric estimators," *Annals of Statistics*, 8, 1348–1360.