Statistics and Computing (2024) 34:11
https://doi.org/10.1007/s11222-023-10325-0

ORIGINAL PAPER

Hybrid elicitation and quantile-parametrized likelihood

Dmytro Perepolkin[1] · Benjamin Goodrich[2] · Ullrika Sahlin[1]

Received: 4 October 2023 / Accepted: 8 October 2023 / Published online: 27 October 2023

Abstract

This paper extends the application of quantile-based Bayesian inference to probability distributions defined in terms of quantiles of observable quantities. Quantile-parameterized distributions are characterized by high shape flexibility and parameter interpretability, making them useful for eliciting information about observables. To encode uncertainty in the quantiles elicited from experts, we propose a Bayesian model based on the metalog distribution and a variant of the Dirichlet prior. We discuss the resulting hybrid expert elicitation protocol, which aims to characterize uncertainty in parameters by asking questions about observable quantities. We also compare and contrast this approach with parametric and predictive elicitation methods.

Keywords Bayesian analysis · Quantile-parameterized distributions · Quantile-based distributions · Expert knowledge elicitation · Indirect inference

Mathematics Subject Classification 62C10 · 62F15 · 62G99

1 Introduction

1.1 Parametric and predictive approach to elicitation

Bayesian parametric inference is about updating prior beliefs about the model parameters in light of new observations. The underlying assumption is that an expert's prior knowledge (or lack thereof) can be translated into a subjective probability distribution of model parameters through the process of elicitation (Winkler 1967). The direct *elicitation of parameters* represents a *structural approach* to extracting an expert's knowledge (Kadane 1980). This approach requires that the expert comprehends the model and the role a specific parameter plays within it. Unfortunately, some parameters may be abstract, challenging to interpret (such as α and β parameters in the Gamma distribution), and at times not independent, as is the case with parameters in a hierarchical model.

✉ Dmytro Perepolkin
dmytro.perepolkin@cec.lu.se

Ullrika Sahlin
ullrika.sahlin@cec.lu.se

1 Centre for Environmental and Climate Science, Lund University, Lund, Sweden

2 Applied Statistics Center, Columbia University, New York, USA

An alternative approach involves eliciting information about observable quantities, possibly conditioned on observable covariates (Kadane and Wolfson 1998), which may be more intuitive and relatable for experts. The *elicitation of predictions* aims to assess the expert's uncertainty regarding future observations (Gelman et al. 2013). Kadane and Wolfson (1998) advise against eliciting moments, with the exception of possibly the first moment (the arithmetic average). Instead, assessment should be carried out using quantiles or probabilities from the predictive distribution. The challenge with eliciting the predictive distribution is that it makes no distinction between the randomness explained by the model and the uncertainty about the parameters within it. Without this distinction, updating the expert predictions with the data coming from the new observations may be challenging.

While the non-Bayesian elicitation often stops at the quantiles or probabilities related to the expert's *predictive judgment* (Spetzler and Staël Von Holstein 1975; Morgan 2014; Keeney and von Winterfeldt 1991; Hanea et al. 2021), the Bayesian school of thought attempts to devise a method to infer the prior distribution, which could have led to the particular predictions expressed by the expert (Akbarov 2009; Hartmann et al. 2020; Winkler 1980; Kadane and Wolfson 1998; Bockting et al. 2023; Manderson and Goudie 2023). Kadane (1980) refer to this process of eliciting the predictive distribution followed by inferring the prior as the *predictive approach*, as it leverages predictions to derive the distribution of parameters.

1.2 Aims of the paper

In this paper, we propose a **hybrid elicitation** approach, which combines the elicitation of observable quantities with the elicitation of the associated uncertainty. This method combines elements of both the predictive and structural approaches to elicitation and can be employed to establish the prior distribution for a model defined by a quantile-parametrized distribution.

Quantile-parameterized distributions (QPDs) (Keelin and Powley 2011; Hadlock 2017) are parameterized by a set of quantile-probability pairs describing a random variable. As a result, the parameters in a QPD are measured on the same scale as the random variable they represent. These distributions can be utilized to model either uncertainty about future observations (predictive distribution) or the distribution of an unobservable parameter (prior distribution). For a comprehensive review of quantile-parameterized distributions, we refer to Perepolkin et al. (2023b).

Until now there has been, to the best of our knowledge, no published research on how to update the quantile parameters of a QPD in light of the new observations. This paper extends the principles of *quantile-based Bayesian inference* (Perepolkin et al. 2023a) to models parameterized by quantiles and proposes a prior distribution capable of capturing uncertainty in the quantile parameters. The proposed approach enables the elicitation and Bayesian updating of a variable quantity with minimal assumptions about the underlying model structure.

1.3 Paper structure

Section 2 introduces the method of *quantile-based inference* as proposed by Rayner and MacGillivray (2002) and Nair et al. (2020), and summarized in Perepolkin et al. (2023a). This method of inference is related to using one of the quantile-based distributions (Perepolkin et al. 2023b), which lack an explicit distribution function (CDF) and probability density function (PDF), as either prior or likelihood components within a Bayesian model. Quantile-based priors and likelihoods rely on substitutions derived from the inverse distribution function, known as the quantile function (QF).

In Sect. 3, we delve into a subclass of quantile-based distributions parameterized by sets of quantile-probability pairs (Fig. 1). We provide a brief overview of the literature concerning different methods for constructing quantile-parameterized distributions (QPDs). Our particular focus is on the quantile-based quantile-parameterized metalog distribution (Keelin 2016), chosen for its parameter flexibility.

In Sect. 4 we introduce the Bayesian model in which the likelihood is expressed using the metalog distribution, parameterized by a set of quantile-probability pairs. We demonstrate how uncertainty in these parameters can be specified through a variant of the Dirichlet distribution. Section 5 describes the elicitation of the QDirichlet prior and introduces a novel *hybrid* elicitation process for obtaining quantile-probability pairs along with the uncertainty associated with them. We illustrate our approach with excerpts from a hypothetical interview. The accompanying `qpd` R package (Perepolkin 2019), implements several quantile-parameterized distributions and includes functionality for supporting the elicitation of the Dirichlet and Connor–Mosimann distributions (Elfadaly and Garthwaite 2013).

Section 6 discusses the MCMC-based algorithm used for updating parameters in a quantile-parameterized distribution. In this paper, we employ Hamiltonian Monte Carlo algorithm in Stan, interfaced by the `cmdstanr` package in R (Gabry and Češnovar 2022). An alternative implementation using the Robust Adaptive Metropolis algorithm by Vihola (2012), interfaced by the `fmcmc` package (Vega Yon and Marjoram 2019), is available in the Supplemental Materials. The models proposed in this paper have been validated using Simulation-Based Calibration (Cook et al. 2006; Modrák et al. 2022; Talts et al. 2020). The results of the simulation studies (provided in Appendix C in Supplemental Materials) demonstrate the successful recovery of the parameter values for all widths of the posterior credible intervals.

We conclude the paper by discussion and summary of the results in Sect. 7.

2 Quantile-based Bayesian inference

The use of non-invertible quantile-based distributions as either a likelihood (Rayner and MacGillivray 2002; King 1999) or a prior (Nair et al. 2020) is not a novel concept in scientific literature. Rayner and MacGillivray (2002) described a three-step process for computing the log-likelihood of a quantile-based distribution. They applied this method to estimate the parameters of the *g-and-k* and generalized *g-and-h* distributions using maximum likelihood estimation. Similarly, Nair et al. (2020) employed quantile function substitutions to express both prior and likelihood in a quantile form. They calculated the posterior Bayes estimator of the parameters in the Govindarajulu model with uniform and generalized exponential priors. Perepolkin et al. (2023a) summarized the approaches to quantile-based inference and provided several examples of applying the principles of inference with quantile functions in both univariate and regression settings.

For a random sample $\underline{x} = \{x_1, x_2, \dots x_n\}$, the posterior distribution of θ over the parameter space Θ can be summarized as:

Fig. 1 Probability distributions, quantile-based distributions and parameterization by quantiles

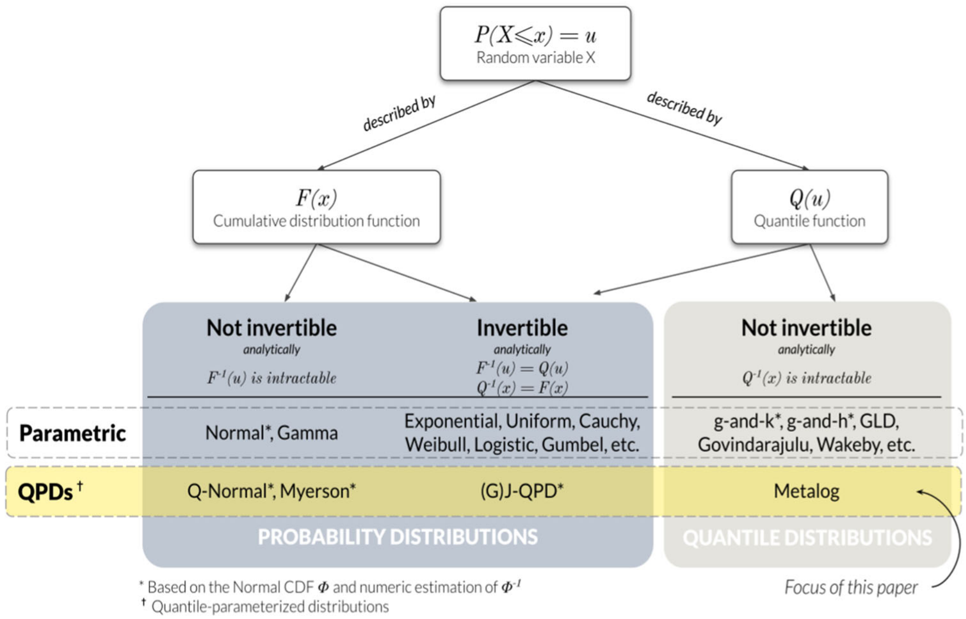

$$f(\theta|\underline{x}) \propto \mathcal{L}(\theta; \underline{x}) f(\theta) \tag{1}$$

where $f(\theta|\underline{x})$ is the posterior distribution of θ after having observed the sample $\underline{x}$, $f(\theta)$ is the prior distribution of θ, and $\mathcal{L}(\theta; \underline{x}) = \prod_{i=1}^{n} f(x_i|\theta)$ is the *density-based* form of the likelihood.

Consider a set of conditional probabilities $\underline{u} = \{u_i|\theta\} = \{F(x_i|\theta)\}$, $i = \{1, 2, \dots n\}$, corresponding to the sample $\underline{x}$ of the observable x with distribution function F given the parameter θ, called *depths*. The conditional probabilities $\underline{u}$ are degenerate random variables that are entirely determined given the observations $\underline{x}$ and the value of the parameter θ. They are called *depths*, because they indicate how "deep" a particular observation is within the distribution. Using the depths $\underline{u}$, we can calculate $\underline{Q} = \{Q_1(u_1), Q_2(u_2), \dots Q_n(u_n)|\theta\}$, where $Q(u|\theta) = F^{-1}(u|\theta)$ represents the quantile function or inverse cumulative distribution function (CDF). Since $Q(u_i|\theta) = x_i$, we can substitute $\underline{Q}$ for $\underline{x}$, and the Bayesian inference formula (1) becomes:

$$f(\theta|\underline{Q}) \propto \mathcal{L}(\theta; \underline{Q}) f(\theta) \tag{2}$$

We refer to this form of the likelihood $\mathcal{L}(\theta; \underline{Q}) = \prod_{i=1}^{n} f(Q(u_i|\theta)) = \prod_{i=1}^{n}[q(u_i|\theta)]^{-1}$ as *quantile-based* because it relies on the calculation of intermediate depths $u_i = F(x_i|\theta), i = \{1, 2, \dots n\}$. Here $[q(u_i|\theta)]^{-1}$ is reciprocal to the derivative of the quantile function $Q(u_i|\theta)$ called the *density quantile function* (Perepolkin et al. 2023a).

Both forms of the likelihood, $\mathcal{L}(\theta; \underline{Q})$ and $\mathcal{L}(\theta; \underline{x})$, are equivalent and yield the same posterior beliefs about the parameter θ (Perepolkin et al. 2023a).

3 Quantile parameterization of distributions

In this section, we consider a special class of distributions where the parameters are specified by quantile-probability pairs (Fig. 1), and see how the concept of *quantile-based inference* (Perepolkin et al. 2023a) can be applied to these distributions, as well.

A set of n quantile-probability pairs, denoted as $\mathsf{S} = \{(p_i, q_i)\}, i = \{1, 2, \dots n\}$, can be thought of as comprising a pair of ordered vectors: a vector of probabilities p and a vector of quantiles q, with $p = \{p_1, \dots p_n\}$, $p_i \in [0; 1]$, and $q = \{q_1, ..q_n\}$, $i = \{1, 2, \dots n\}$. As CDF $F(x) = p$ is a non-decreasing function, the vectors p and q are considered properly ordered iif $q_i \leq q_{i+1}, \forall q_i \in q$, and $p_i \leq p_{i+1}, \forall p_i \in p$. Additionally, the quantile-probability pairs within the set S are considered distinct iff $\forall\{(p_i, q_i)\} \in \mathsf{S}, \exists!\{(p_j, q_j)\} = \{(p_i, q_i)\}$, $j \neq i$, $i = \{1, 2, \dots n\}$, $j = \{1, 2, \dots n\}$. In this paper, we refer to the set of n distinct, properly ordered quantile-probability pairs as a size-n *quantile-probability tuple* (QPT) denoted by $\{p, q\}_n$.

3.1 SPT-parameterization

A recent review (Perepolkin et al. 2023b) describes two methods for constructing distributions parameterized by quantile-probability pairs:

- By reparameterization of existing distributions, or
- Through an optimization step, where the distribution parameters are mapped to quantiles using least squares or similar algorithms.

Distributions falling under in the first category are typically parameterized by the *symmetric percentile triplet* (SPT), a QPT of size 3. In the SPT, the middle cumulative probability $p_2 = 0.5$ represents the median, while $p_1 = 1 - p_3 = \alpha,\ \alpha \in (0, 0.5)$ (e.g. $\{0.25, 0.50, 0.75\}$ or $\{0.10, 0.50, 0.90\}$). Examples of SPT-parameterized QPDs include the Myerson distribution (Myerson 2005), the Johnson Quantile-Parameterized distribution (J-QPD) (Hadlock and Bickel 2017), and their generalizations (Perepolkin et al. 2023b; Hadlock and Bickel 2019). A special case of SPT-parameterization also exists for the metalog distribution (Keelin 2016). For the purposes of this paper, we do not consider SPT-parameterized QPDs, including the SPT-metalog, Myerson, or J-QPD, due to their strict requirement for symmetric probability parameterization.

3.2 Parameterization using implicit functions

Keelin and Powley (2011) and Powley (2013) introduce an alternative method of parametrizing a distribution by a set of quantile-probability pairs. This method relies on the finite Taylor series expansion of parameters within a known quantile function as linear functions of the cumulative probability p.

The authors created the Simple Q-Normal (SQN) distribution by taking the quantile function of a normal distribution $x \equiv \mu + \sigma\Phi^{-1}(p)$, and making the parameters μ and σ functions of p; specifically, $\mu(p) = a_1 + a_4 p$ and $\sigma(p) = a_2 + a_3 p$. In a similar vein, Keelin (2016) proposed the metalogistic (*metalog*) distribution by making the parameters μ and s in the logistic quantile function $x \equiv \mu + s\,\text{logit}(p)$ be the functions of p, i.e. $\mu = a_1 + a_4(p - 0.5) + a_5(p - 0.5)^2 + \ldots$ and $s = a_2 + a_3(p - 0.5) + a_6(p - 0.5)^2 + \ldots$. Here, μ represents the mean, s is proportional to the standard deviation such that $\sigma = s\pi/\sqrt{3}$, $\text{logit}(p) = \ln(p/(1 - p))$ is the log-odds of probability $p \in [0, 1]$ and $a_i,\ i = \{1, 2, \ldots n\}$ are real constants.

In both cases, the quantile function $Q(p)$ whose parameters also depend on p is an *implicit function*. This means that such a quantile function cannot be simply computed for arbitrary values of p. Nevertheless, with a set of n quantile-probability pairs, it is possible to determine the constants $a_i,\ i = 1, 2 \ldots n$, by solving a system of n linear equations (Keelin and Powley 2011; Powley 2013). This system can be represented as the matrix Equation (3).

$$a = \mathbb{P}^{-1}q \tag{3}$$

Keelin and Powley (2011) show the conditions under which a size-n QPT $\{p, q\}_n$ can uniquely determine the constants $a = \{a_1, \ldots a_n\}$. Additional details of the metalog distribution, including the composition of the matrix $\mathbb{P}$, can be found in Appendix A in Supplemental Materials.

The shape flexibility of the QPD increases with the number of terms added to the finite Taylor expansion of parameters within the parent distribution. To estimate the coefficients for the n-term quantile-parameterized distribution $a = \{a_1, \ldots a_n\}$, a minimum of n quantile-probability pairs is required. The order of the terms, denoted as n, is constrained by the size of the parameterizing QPT m, ($n \le m$), and concerns for overfitting. The QPT used for parameterizing a distribution can be obtained through expert elicitation or from the empirical CDF (ECDF), which is constructed from a sample of observations. The ECDF begins at zero and increments by $1/m$ at each of the m data points in the sample, representing the fraction of observations that are less than or equal to the specified value (Wasserman 2006).

Depending on the relationship between the size of the parameterizing QPT m and the number of terms n in the QPD QF we use the following terminology:

- When the size of the parameterizing QPT m equals the number of terms n in the QPD QF, i.e. $m = n$, we refer to the process of estimating the vector of coefficients $a = \{a_1 \ldots a_n\}$ as "fitting", and we call the resulting QPD "properly parameterized". In properly parameterized QPDs, the QF curve is guaranteed to pass through every QPT point. We label the n-term metalog parameterized by the n-size QPT $\{p, q\}_n$ as the **proper n-metalog**.
- When the size of the parameterizing QPT m exceeds the number of terms n in the QPD's QF (for example, when the QPT is derived from the sample ECDF and $m > n$), we refer to the process of estimating the vector of coefficients $a = \{a_1 \ldots a_n\}$ as "approximating", and we call the resulting QPD "approximated". Such approximation is typically achieved through optimization or regression, and the resulting QF curve is no longer guaranteed to pass through every QPT point. We designate the n-term metalog parameterized by the m-size QPT $\{p, q\}_m,\ m > n$ as the **approximate n-metalog**.

Given the matrix Equation (3), we have two alternative parameterizations for the proper n-metalog: it can either be directly parameterized by the coefficients $a = \{a_1, \ldots a_n\}$ (referred to as the A-parameterization) or indirectly parameterized by a QPT $\{p, q\}_n$ (referred to as the QPT-parameterization). Therefore, in this paper, when we mention the proper n-metalog, the notations $Q_{M_n}(u|a)$ and $Q_{M_n}(u|p, q)$ (where $u|a$ or $u|p, q$ represents the *depths* corresponding to the observation x) are used interchangeably. In the case where the metalog is approximated, only the A-parameterization is suitable because the number of metalog terms n is not determined by the m data points from the ECDF used to estimate the parameter vector a via Equation (3).

4 QDirichlet-metalog model

In this section, we introduce a Bayesian model with the likelihood defined by the proper n-term metalog. Since the metalog is a quantile-based distribution (Fig. 1), we employ the *quantile-based likelihood* following Equation (2). The quantile-based likelihood relies on the intermediate depths $\underline{u}|\theta$, which correspond to the sample of observations $\underline{x}$. Since a closed-form CDF for the metalog distribution is not available, we resort to numerical approximation, denoted as $u = \widehat{Q}_X^{-1}(x)$ (Perepolkin et al. 2023a).

4.1 Parameter uncertainty

To account for the uncertainty in the QPT parameters of a quantile-parameterized likelihood, such as the metalog, we must introduce uncertainty either in cumulative probabilities or the corresponding quantile values (or both). Coles and Tawn (1996) specified the prior for an extreme value model in terms of the quantile values for certain fixed cumulative probability values. Crowder (1992) suggested that the prior can be constructed based on the space of probabilities, with fixed quantiles. In a recent paper on predictive elicitation, Hartmann et al. (2020) divided the observable space into several exhaustive and mutually exclusive categories and asked experts to assign probabilities that the next observation falls into each of the categories, treating these probabilities as uncertain. They assigned a Dirichlet prior to these probability judgments.

We follow a similar approach by using the quantile values provided by the expert to partition the outcome space. We then characterize the uncertainty in the corresponding cumulative probabilities using the Dirichlet distribution (together referred to as the *QDirichlet* prior). Our method of constructing a prior distribution for the simplex Δ shares similarities with the approach adopted by Bürkner and Charpentier (2020) for modelling monotonic effects in ordinal regression. The parameter vector of the Dirichlet distribution, in conjunction with the vector of elicited quantiles, serves as hyper-parameters for the proposed QDirichlet prior, which captures the uncertainty in the parameters of the quantile-parameterized model.

4.2 The QDirichlet prior

Consider a size-n QPT $\{p, q\}_n$, consisting of a vector of probabilities p and a vector of quantile values q. Now, consider an extended vector of probabilities $b = \{0, p, 1\}$ of size $n + 2$, containing the vector p. Additionally, consider a forward difference $\Delta = \{\Delta_1 \dots \Delta_{n+1}\}$, where $\Delta_i = b_{i+1} - b_i$, $i = 1, 2, \dots (n + 1)$, which is a simplex of size $n + 1$. The simplex Δ is properly ordered iif it is based on the properly ordered vector b, and consequently, also p.

To transform the simplex Δ back into the vector of probabilities p, the cumulative sum $\Xi_1^n()$ can be used, so that $p = \Xi_1^n(\Delta) : \; p_j = \sum_{i=1}^{j} \Delta_i, \; j \in (1 \dots n)$, assuming the simplex Δ is properly ordered. If, for any reason, the simplex Δ can no longer be considered *properly ordered*, we can use an index vector of distinct values, denoted as $I = \{I_1 \dots . I_n\} : I_j = \{1, 2, \dots n\}, \; j = \{1, 2, \dots n\}$, $\exists ! I_j = I_i, \; j \neq i$. This index vector can be used to restore the proper order before accumulating the simplex Δ into the probability vector p.

To express prior uncertainty in the simplex Δ, we can use the Dirichlet distribution (Johnson et al. 1997) with a hyperparameter vector α of size $n + 1$, conditional upon the specified quantile values q. We refer to this particular variant of the Dirichlet prior as the *QDirichlet* prior, as its parameter vector α is specified in relation to the fixed quantile values q.

4.3 The metalog likelihood

We adopt the notation for quantile-based likelihoods introduced in Perepolkin et al. (2023a), where $u \overset{x}{\backsim} \dots$ should be read as "the depths u corresponding to the random variable x inversely distributed as …". Consequently, QDirichlet-Metalog model can be expressed as follows:

$$
\begin{aligned}
&u \overset{x}{\backsim} \text{Metalog}(p, q) \\
&\Delta \sim \text{Dirichlet}(\alpha | q); \\
&p = \Xi_1^n(\Delta);
\end{aligned}
\tag{4}
$$

where Δ is a simplex of size $n + 1$, $\Xi_1^n()$ is the cumulative sum operator, p is a size-n vector of cumulative probabilities and q is the corresponding size-n vector of quantiles. Furthermore, u is the depth corresponding to the observable x given the parameterizing QPT $\{p, q\}$. The depths u can be computed (typically numerically) by inverting the quantile function $\widehat{Q^{-1}}(x|p, q)$. The metalog quantile function is indirectly parameterized by the QPT $\{p, q\}_n$ through the vector of metalog coefficients a, determined by the matrix Equation (3).

In Model (4), the prior is represented by the Dirichlet distribution with hyperparameter α specifying the uncertainty in the cumulative probabilities and a vector q representing the quantile values corresponding to the sampled cumulative probabilities (QDirichlet prior). The metalog (quantile-based) likelihood parameterized by the QPT $\{p, q\}_n$ relies on depths u which can be estimated using the numerical inverse of the metalog quantile function (Perepolkin et al. 2023a).

4.4 Eliciting Dirichlet distribution

Elfadaly and Garthwaite (2013) describe a method for inferring the parameter vectors of the Dirichlet and Generalized Dirichlet (Connor–Mosimann) distributions from the conditional univariate beta distributions. In this method, the expert assesses the quartiles of the probability for each category using the elicitation of the symmetric percentile triplet (SPT) as follows:

1. The expert assesses the probability quartiles for the first category p_1.
2. The expert is then asked to assume that the median value they provided in the assessment of p_1, is in fact the correct probability (true value) for the first category.
3. Next, the expert proceeds to assess the SPT for the next category, conditional upon the previous assessment, denoted as $p_2|p_1$.
4. The three quartiles of $p_2|p_1$ are divided by $1 - p_1$ to normalize them, i.e. $p_2^*|p_1$.
5. The hyperparameters of the beta distribution representing $p_2^*|p_1$ are the determined.
6. Steps 3–5 are repeated for all categories, except the last one.

Elfadaly and Garthwaite (2013) also propose an improvement to Pratt et al. (1995)'s method of fitting the beta distribution to the elicited conditional SPTs based on the normal approximation described in Patel and Read (1996). The α and β parameters of the conditional beta-distributions are then normalized and the hyperparameter vector $\alpha = \{\alpha_1 \dots \alpha_{n+1}\}$ is estimated.

The elicitation method proposed by Elfadaly and Garthwaite (2013) can be applied to assess the uncertainty in the cumulative probabilities that parametrize the metalog likelihood. If elicitation starts with the left tail of the distribution, the first category will coincide with the first cumulative probability p_1 in the parameter vector p. Subsequent (higher) cumulative probabilities will always be conditional upon and include the median value of the lower probability. If the elicitation is performed out of order, which might be expedient to avoid anchoring effects (Spetzler and Staël Von Holstein 1975; Abbas et al. 2008), the integer index vector I of the same size can be provided along with the results of the assessment to restore the proper ordering of the simplex Δ after sampling and before its accumulation into the vector of probabilities p.

Note that the parameter vector p in the model (4) is not independent: it is paired with the vector of fixed quantile values q. There are several approaches to specifying the hyperparameter vector q.

- **Predictive distribution**. The vector q could be coming from the characterization of the prior predictive distribution. In this scenario, we could ask the expert to specify their uncertainty regarding the next observation using standard predictive elicitation techniques described in Morgan et al. (1990) or Spetzler and Staël Von Holstein (1975). Predictive elicitation results in the QPT $\{p^*, q^*\}_n$, of which vector q^* can be adopted as the true values of q, while the vector of cumulative probabilities p^* can serve as the initial values for the MCMC/HMC algorithm.
- **Hypothetical sample**. Alternatively, the vector q could be viewed as a *representative sample* from the predictive distribution. Randomly sampling the predictive distribution has the advantage that the values closer to the distribution's mode are more likely. However, if the sampled values of q are too closely spaced, fitting the Metalog to the QPT $\{p, q\}_n$ within the MCMC loop may become challenging.

The primary goal of eliciting the vector q is to *position* the prior on the data (x) scale and provide a reasonable baseline for the follow up elicitation. In fact, the hyperparameter vector q specifies the *location* of the QDirichlet prior, while the hyperparameter vector α is responsible for defining its *shape*.

5 Applications

In this section, we provide an example of *hybrid* elicitation for parameterizing the QDirichlet prior to describe the uncertainty in the $\{p, q\}_n$ parameters of the proper metalog.

5.1 Steelhead trout weights

We take a sample of 100 observations from the records of steelhead trout weights captured and released in Babine River, Canada, spanning the years of 2006-2014 (Fig. 2). The dataset has been published by Keelin (2016) and is also included in the `rmetalog` (Faber and Jung 2021) package, which is accessible on CRAN.

Our goal is to elicit prior beliefs regarding the distribution of fish weights from a hypothetical expert and subsequently update those beliefs in light of the sampled data. Sections 5.2 and 5.3 outline the elicitation process and provide details on the required diagnostics for prior specification and posterior inference.

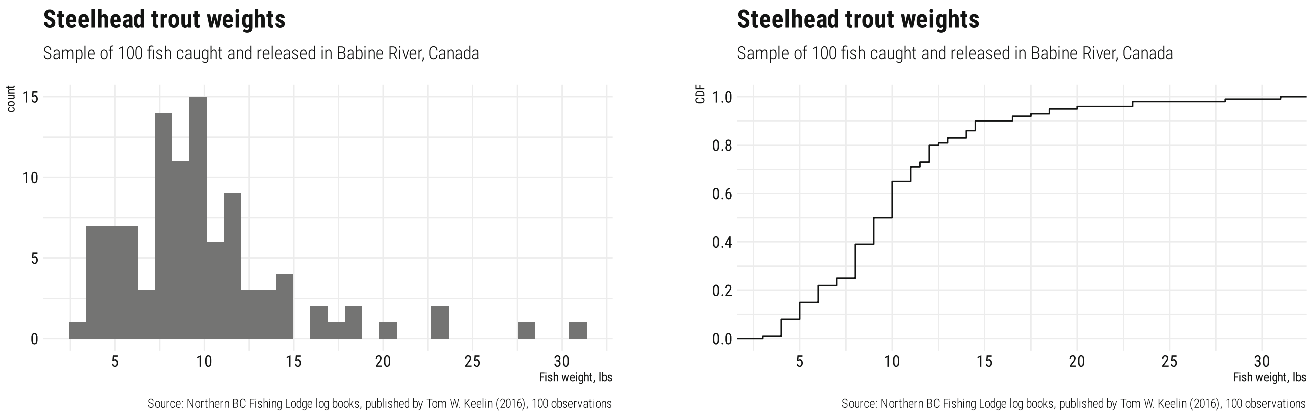

Fig. 2 Summary of the sample from the steelhead trout dataset

Table 1 Quantile-probability pairs for the predictive distribution of fish weights

Cumulative probability, p*	Weight in lbs, q*
0.1	4
0.5	9
0.9	17

5.2 Specifying location of the prior

The *hybrid* elicitation involves of two phases: the elicitation of quantile values q and the elicitation of uncertainty in the associated cumulative probabilities (i.e. potential vectors of p that could correspond to the specified q).

Imagine conducting an interview with an experienced fly-fisherman to gather information about steelhead trout weights in Canadian rivers. We begin with the elicitation of the predictive QPT using the probability-value (PV) method (Spetzler and Staël Von Holstein 1975; Abbas et al. 2008). After guiding the expert through essential preparatory steps (motivating, structuring, conditioning, encoding and verification), we elicit the following predictive QPT $\{p^*, q^*\}_3$ (Table 1).

Physical weight can be represented by non-negative values, suggesting the use of a distribution bounded on the left. To model this, we employ a semi-bounded log-metalog for the predictive QPT values (Fig. 3). Notably, the three cumulative probabilities p^* provided by the expert divide the y-axis of the CDF into four distinct bands (highlighted by different colors in Fig. 3). These bands can be seen as categories into which a weight of a randomly drawn fish could fall on the empirical CDF curve. Similar to the approach taken in Hartmann et al. (2020), we use the elicited quantile values q^* to partition the outcome space into the exhaustive and mutually exclusive categories. The widths of the bands correspond to the increments in cumulative probabilities provided by the expert, as represented by the simplex Δ in the model (Equation (4)).

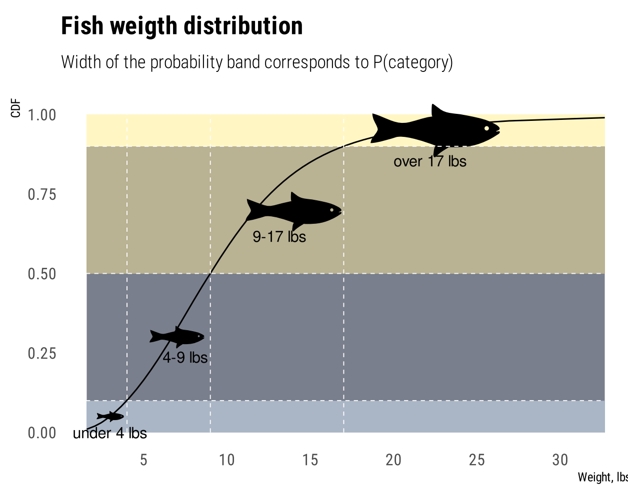

Fig. 3 Probability bands corresponding to the four fish size categories

Table 2 SPT for the count of the small fish (less than 4 lbs)

	Cumulative probability	Fish count
Small fish	0.25	70
	0.50	90
	0.75	120

All fish counts are out of the sample of 1000

In the second step of the *hybrid* elicitation we elicit uncertainty regarding these probability band widths Δ, using the values q as reference points to delineate the categories to which the random variate $p_i \in p,\ i = \{1, 2, \dots n\}$ would be associated.

5.3 Hybrid elicitation of the QDirichlet prior

The predictive QPT $\{p^*, q^*\}_n$ we elicited earlier does not inherently incorporate uncertainty (except for potential

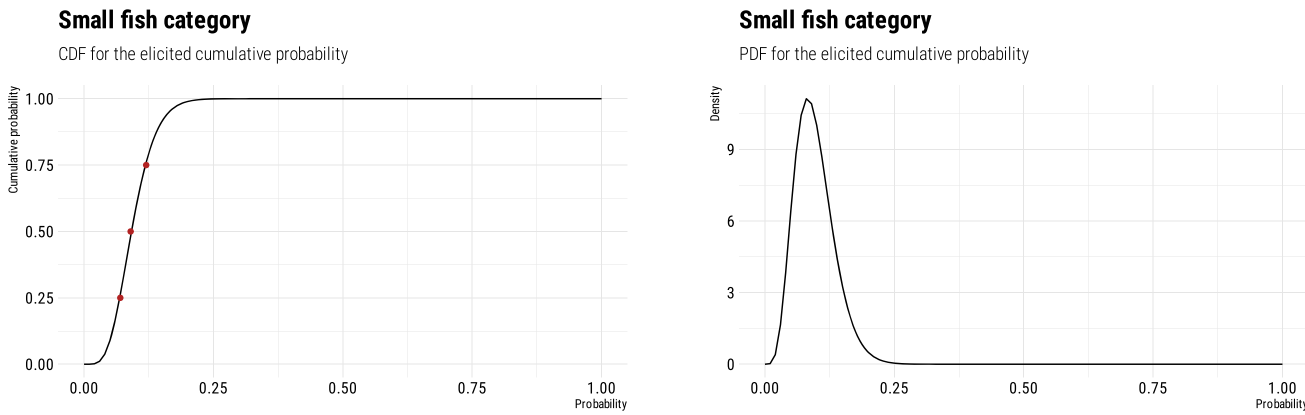

Fig. 4 Conditional beta distribution fitted to the quartiles provided by the expert

imprecision or inconsistencies in the expert's expression of belief). During this phase, we ask the expert to consider uncertainties surrounding their quantile assessments, aiming to distinguish aleatory from epistemic uncertainties (Knight 1921). A possible drawback of commencing with predictive elicitation, as we did above, is that the specified vector p^* might anchor the expert's belief to these values, potentially affecting the range of probabilities that the expert would associate with the vector q. The expert might be inclined to allow only a small and likely symmetrical variation around the initially specified p^* values, hesitating to revise their judgments. On the other hand, the specified p^* values can serve as a starting point for discussion with the expert, validating their beliefs, and deliberately challenging the expert to reassess them.

The term "aleatory" signals the use of the sampling frame, prompting us to shift from assessing the properties of population to the evaluating the properties of an imaginary sample. In this elicitation phase, we transform the cumulative probabilities into natural frequencies (Gigerenzer 2011), treating the values of $p_i \in p,\ i = \{1, 2, \dots n\}$ as proportions within a hypothetical large sample.

Recall that the expert has supplied us with a predictive QPT $\{p^*, q^*\}_3$, where $p^* = \{0.1, 0.5, 0.9\}$ and $q^* = \{4, 9, 17\}$ (Table 1).

Interviewer:
Consider a large sample of steelhead trout caught in British Columbia over the past few years, let's say 1000 fish. Based on your assessment, it's expected that around 100 fish would weigh less than 4 lbs.

The elicited predictive QPT is interpolated with a log-metalog and presented in Fig. 3. Considering the sampling frame, the curve we've drawn through the three points provided by the expert is just one of numerous empirical CDF curves that could be constructed given the inherent sampling uncertainty.

Table 3 Conditional SPTs for the counts of the fish in the 1000 fish sample

	Category*	P25	P50	P75
1	Small fish	0.07	0.090	0.12
2	Medium fish	0.34	0.410	0.50
4	Huge fish	0.05	0.075	0.15

*Small fish is under 4 lbs, Medium fish is 4–9 lbs, Huge fish is over 17 lbs

Table 4 Dirichlet parameter vector

	Category	a
1	Small fish	3.77
2	Medium fish	12.86
4	Huge fish	2.70
3	Large fish	10.72

Table 5 Connor–Mosimann parameter vectors

	Category	a	b
1	Small fish	5.87	55.24
2	Medium fish	6.77	8.01
4	Huge fish	1.32	5.25

We proceed with the elicitation by asking the expert to contemplate the fish weight cutoff of 4 lbs.

Interviewer:
Let's delve into this hypothetical sample of 1000 fish. According to your assessment, there should be approximately 100 fish weighing less than 4 lbs. We will interpret this assessment as you believing that there's about equal chance that the actual number of "small" fish (weighing less than 4 lbs) in this sample is either above or below 100. In essence, we interpret it as the median assessment. Would you like to reconsider this value?

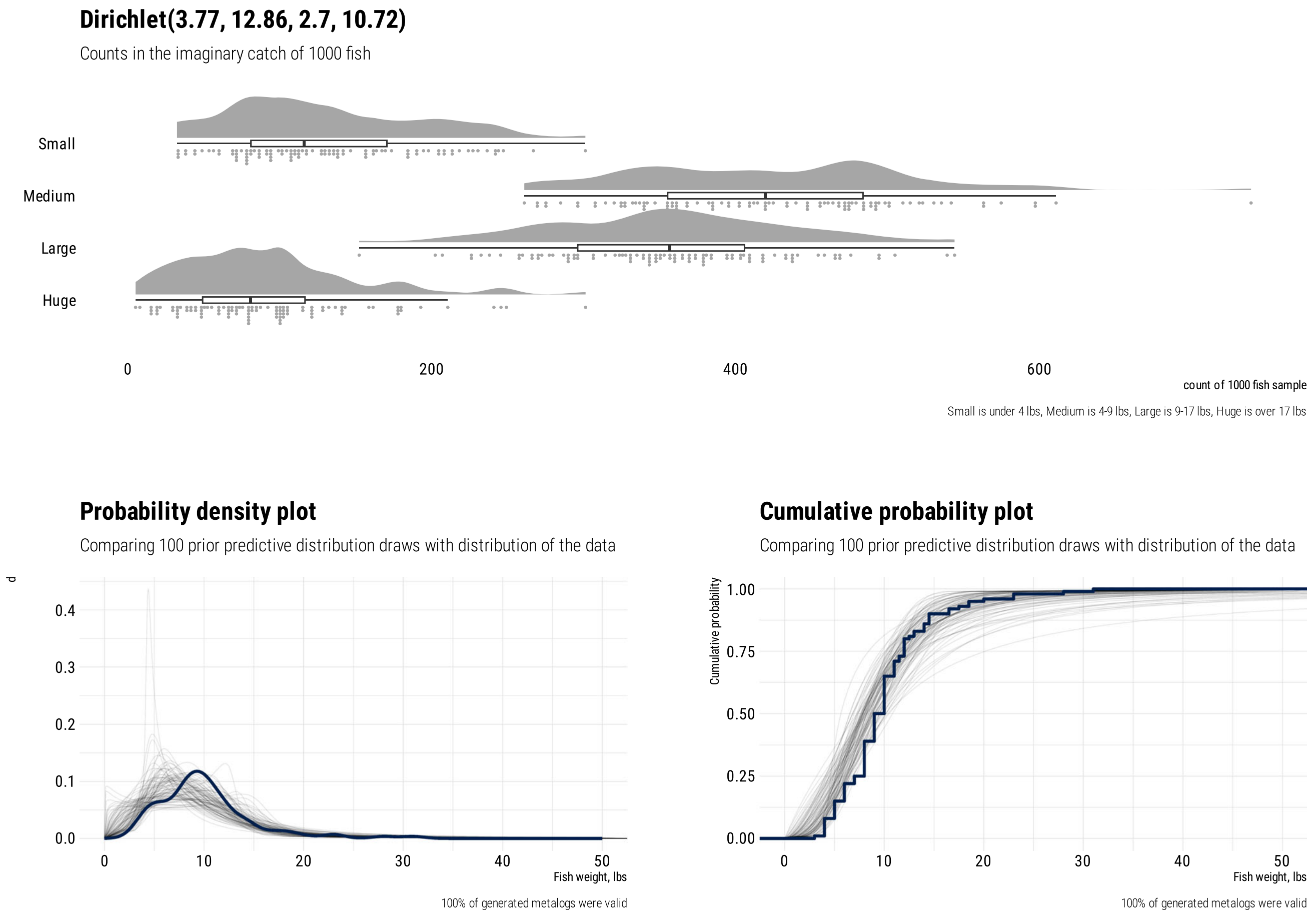

Fig. 5 Prior predictive check for Dirichlet distribution

At this stage, the expert may choose to adjust their assessment of the median. Once the median value of the first category is confirmed, we can proceed with the elicitation of the range around it. We follow the conventional *fixed probability* encoding method (Abbas et al. 2008; Spetzler and Staël Von Holstein 1975) asking the expert about the range of fish counts in the sample (which actually represent cumulative probabilities) corresponding to the quartiles or the 10th/90th percentile.

Suppose, the expert has furnished us with the revised median and the 50% Interquartile Range (IQR) around the initially assessed probability $p_1^* = 0.1$ for the count of "small" fish in the hypothetical sample of 1000, as summarized in Table 2.

From this information, we can promptly deduce the uncertainty in the "width" of our first bin. It is now characterized by a symmetric percentile triplet {0.07, 0.09, and 0.12} with $\alpha = p_1 = 1 - p_3 = 0.25$. Employing this SPT, we can fit the beta distribution (Fig. 4) using the method proposed in Elfadaly and Garthwaite (2013).

We then proceed to with the conditional elicitation of probabilities for the remaining fish weight categories and uncertainties associated with them, conditional on the median values of the previously elicited categories. During this phase, we ask the expert:

Interviewer:
Let's assume that in the sample of 1000 fish, precisely 90 were found to be small (weighing less than 4 lbs). What would be your estimate of the number of fish that would fall within the weight range of 4 to 9 lbs in such a sample?

In this question, we are soliciting the expert's input for a conditional probability distribution. Therefore, we do not hold the expert accountable for their previous assessment, where they implied that approximately half of the population would weigh 9 lbs or less (as suggested by the cumulative probability of 0.5). We anticipate that the median count would be close to 500 fish, but not necessarily an exact match. Our aim is to elicit the count of fish weighing between 4 and 9 lbs. However, if the expert prefers to provide us with the count corresponding to the "exceedance probability" (i.e., for 2 categories combined), we should subtract the median count of the first category, which is 90.

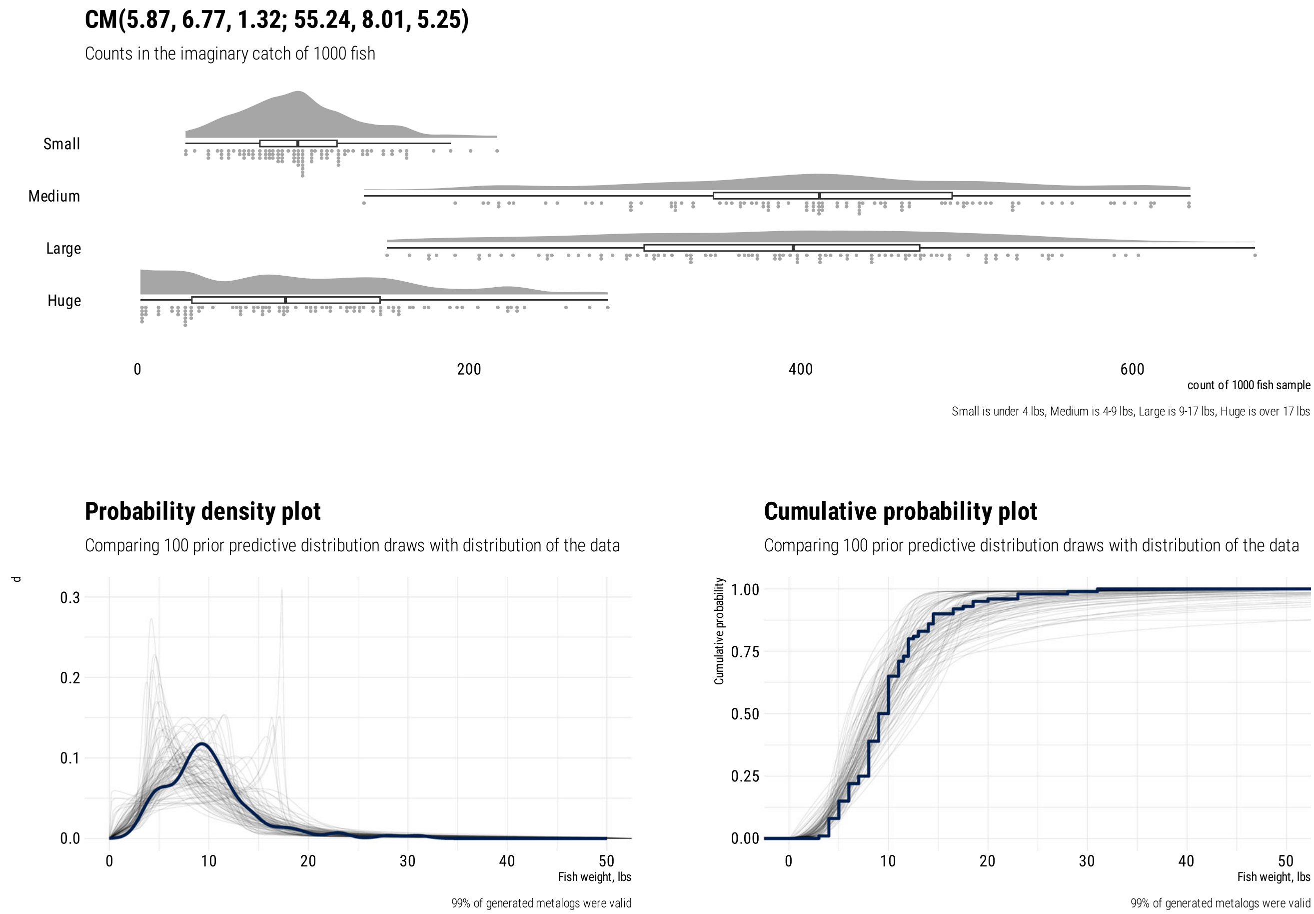

Fig. 6 Prior predictive check for Connor–Mosimann distribution

For a total of N groups, we should only need to conduct $N-1$ elicitations (a total of 3 elicitation of triplets in our case). It might be convenient to elicit the top quantile (representing the upper tail of the CDF) as $1-\sum_{i=1}^{3} p_i$) and leave the quantiles for the "Large fish" category (fish weighing between 9 and 17 lbs) to be calculated from the remaining information.

Let's assume that, after eliciting the three conditional SPTs and converting the hypothetical sample counts to probabilities, we obtain the following assessments (Table 3). Note, that the assessments for the category 3 Large fish are missing from the table. They are implied (and will be inferred from) the rest of the data.

5.4 Fitting Dirichlet and Connor–Mosimann distributions

We can utilize the elicited conditional SPTs to derive parameter vectors for the Dirichlet (Table 4) or Connor–Mosimann (Table 5) distributions following the process in Elfadaly and Garthwaite (2013). The algorithm for transforming the conditional SPTs into parameter vector(s) is implemented in the `qpd` R package (Perepolkin 2019).

The Dirichlet distribution defines a strong negative dependence between the elements of the simplex Δ, meaning that an increase in the probability of one element necessarily decreases the probability of every other element (Balakrishnan 2014). The Connor–Mosimann distribution relaxes this assumption of a strong negative correlation between the categories, allowing for a more flexible encoding of dependence between the quantiles (Wilson 2017).

5.5 Prior predictive check

Prior predictive checks are crucial for providing the expert with feedback on the elicited values and diagnosing potential issues (Gabry et al. 2019). Since uncertainties in the quantile probabilities were elicited as conditional probabilities, it is important to show to the expert the impact of the provided probability ranges on the overall multivariate distribution. This can be accomplished, for example, using marginal plots. We can draw samples from the Dirichlet distribution (Fig. 5) or the Connor–Mosimann distribution (Fig. 6) and present

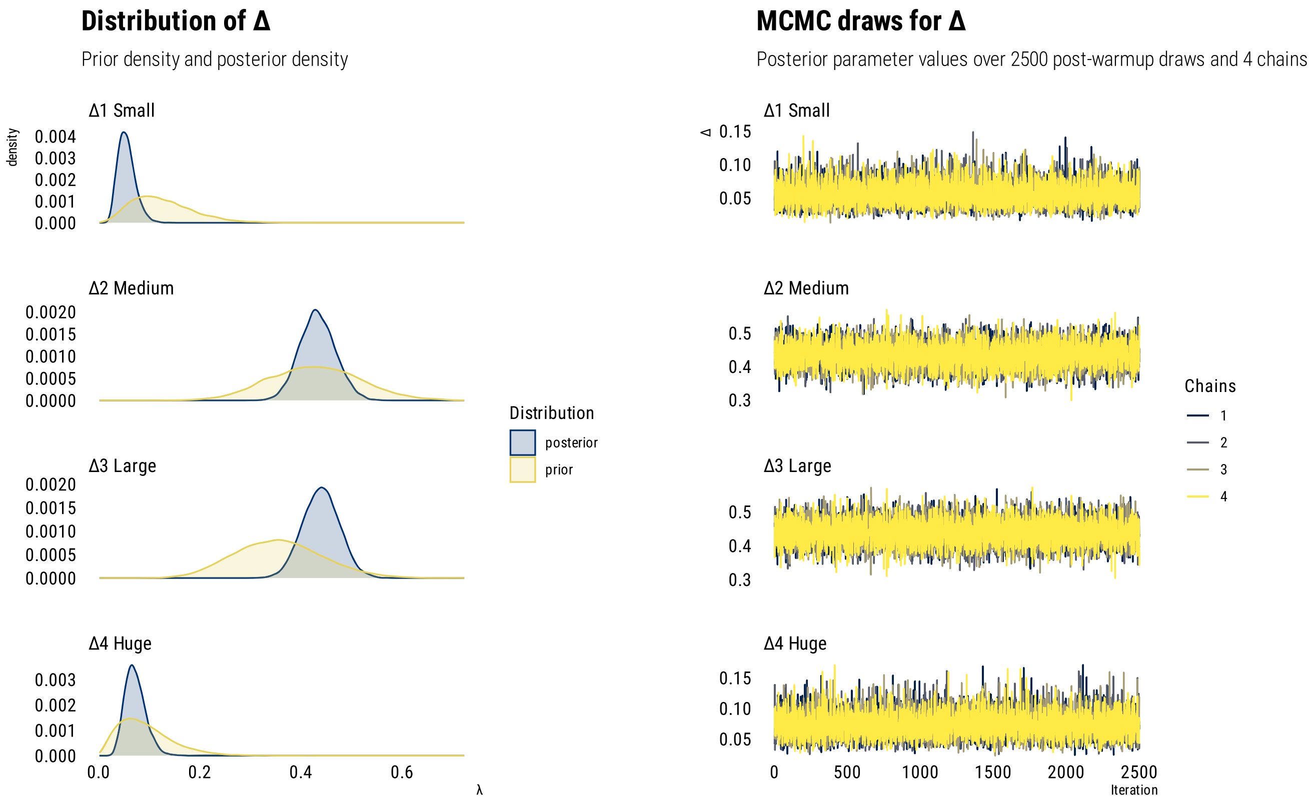

Fig. 7 Summary of the posterior draws for Δ simplex

the expert with an overview of the parameter distribution in the same format that will be used for the posterior predictive check (Fig. 8 in Sect. 6).

6 MCMC implementation

To sample from the posterior distribution of Δ,we employed the Hamiltonian Monte Carlo (HMC) algorithm in Stan, interfaced via the `cmdstanr` package (Gabry and Češnovar 2022) in R. An alternative implementation using the Robust Adaptive Metropolis algorithm by Vihola (2012), implemented in the `fmcmc` package (Vega Yon and Marjoram 2019), is provided in the Supplemental Materials

We validated the QDirichlet-Metalog model using the Simulation-Based Calibration algorithm (Cook et al. 2006; Modrák et al. 2022; Talts et al. 2020). As evident from the diagnostic plots in the Supplemental Materials Appendix C, the parameter Δ is successfully recovered for all widths of the posterior credible interval.

We have also performed Simulation-Based Calibration for the model with QCM (Generalized Dirichlet) prior. The diagnostic plots also indicate the successful recovery of the parameter vector for all widths of the posterior credible interval.

To fit the QDirichlet-Metalog model, we used 2500 post-warmup iterations and 4 chains. The posterior distribution of the parameter Δ is presented in Table 6 and Fig. 7. The results reveal a significant reduction in the uncertainty regarding the cumulative probabilities corresponding to the quantile values of 4, 9, and 17 lbs. Specifically, the lowest value 4 lbs corresponds to a cumulative probability range of 0.03−0.09, while the upper value 17 lbs corresponds to a range of 0.89−0.96, both representing 90% credible intervals.

Additionally, the posterior predictive check demonstrates a reduction in uncertainty regarding the parameter Δ. Compare the posterior predictive check in Fig. 8) with the prior predictive checks shown in Figs. 5 and 6.

7 Discussion

Over the last two decades, several probability distributions with interpretable parameters defined on the same scale as observable quantities were proposed (Myerson 2005; Keelin and Powley 2011; Hadlock and Bickel 2017). The primary goal of research into quantile-parameterized distributions is to simplify the elicitation process and make it more accessible for experts. In our proposed *hybrid* elicitation framework, tailored specifically for models with quantile-parameterized likelihoods, experts are encouraged to adopt a sampling

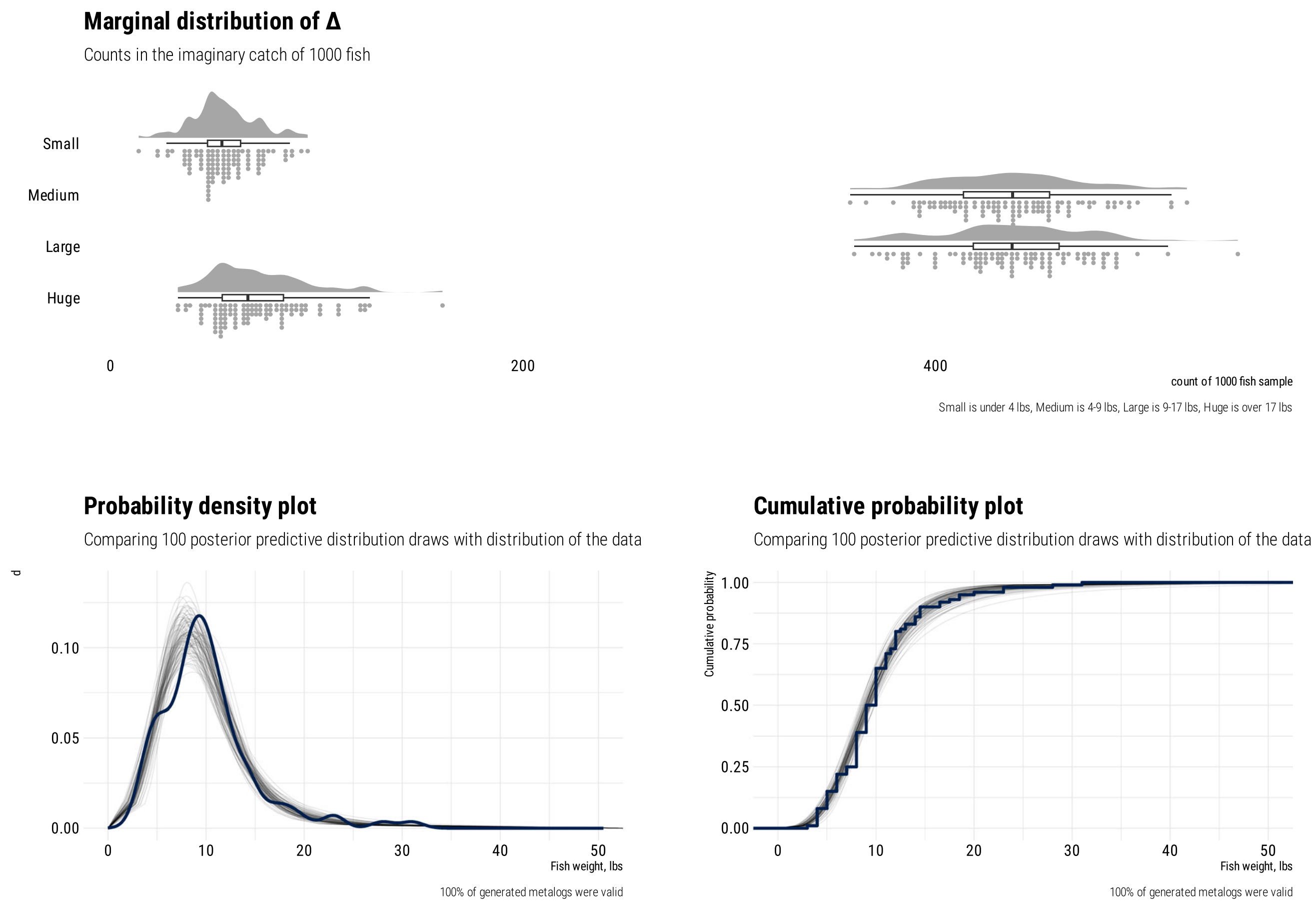

Fig. 8 Posterior predictive check for the QDirichlet-Metalog model

frame. This replaces the challenging task of expressing uncertainty about cumulative probabilities with a simpler task of expressing uncertainty about natural frequencies in a hypothetical sample (Gigerenzer 2011; Hoffrage et al. 2002, 2015).

The elicitation method for Dirichlet distribution proposed by Elfadaly and Garthwaite (2013) asks the expert to assume that the median of the previously assessed category p_{i-1} is, in fact, the true value of the probability for the category $i-1$. It then proceeds to elicit the value p_i for the next category i, conditional on this assessment. We believe that such conditioning can be made simpler if one adopts the natural frequency framework. Inputs elicited from the expert in the natural frequency frame can be easily validated through simulation, providing the expert with immediate feedback on the implications of their judgments for the model.

In Bayesian analysis, we see quantile-parameterized and parametric likelihoods as complementary. Initiating a model with a likelihood expressed by a quantile-parameterized distribution can be advantageous when only QPT judgments from experts are available, no specific choice for a parametric distribution is evident, and data is limited. As data becomes more abundant and our understanding of the data-generating process improves, a transition to a parametric likelihood can be justified.

Table 6 Posterior sample summary for the simplex Δ

Category	Mean	Median	q5	q95	rhat
Small	0.0552	0.0536	0.0306	0.0862	1.000
Medium	0.4328	0.4318	0.3763	0.4927	1.000
Huge	0.0724	0.0703	0.0427	0.1093	1.001
Large	0.4396	0.4398	0.3803	0.4986	1.000

The QDirichlet-Metalog model described in this paper can be applied in conjunction with a predictive approach to elicitation (Kadane 1980). Assuming that the only information elicited from the expert is the predictive QPT $\{p, q\}_n$, the quantiles q vector can be combined with a *uniform* Dirichlet prior, allowing the data alone to define the posterior for the simplex Δ. Given that the Dirichlet distribution is a generalization of the Beta distribution to higher dimensions, a weakly informative prior can be specified with a unit vector, i.e. Dirichlet$(1, 1, \ldots 1)$. We discuss inference using weakly informative priors in Appendix B in Supplemental Materials.

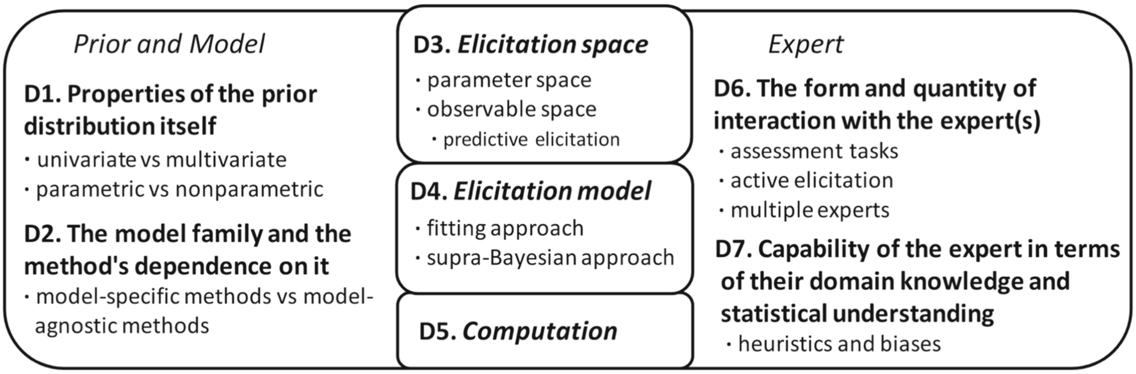

Fig. 9 Prior elicitation hypercube

Parametric elicitation aims to describe the epistemic uncertainty contained in the *parameters* of the model with the help of experts. On the other hand, predictive elicitation aims to describe the uncertainty in the *next observation* without distinguishing between the randomness in the model and the lack of knowledge about the model parameters.

Mikkola et al. (2023) proposed the prior elicitation hypercube with 7 dimensions related to the elicitation of prior distributions (Fig. 9). Following this classification, the proposed *hybrid elicitation* falls under the category of a univariate, parametric, prior-specific (D1), model-specific (D2) elicitation method, conducted in the observable space (D3). *Hybrid elicitation* leverages the approach proposed by Elfadaly and Garthwaite (2013) to derive the parameter vector(s) of the (Generalized) Dirichlet distribution (D4). This process relies on the simple arithmetic computations (D5) to transform the parameters of the conditional marginal beta distributions into the (Generalized) Dirichlet parameter vector(s). Furthermore, *hybrid elicitation* adopts an active, iterative elicitation approach (D6), requiring minal assumptions about the expert's familiarity with statistical concepts, such a detailed understanding of the underlying generative model (D7).

Hybrid elicitation begins by describing the next observation, but subsequently shifts to characterizing the uncertainty inherent in the predictive assessment itself. This is achieved by describing a hypothetical sample from the target population corresponding to the cumulative probabilities. These probabilities, in conjunction with a set of quantile values, serve as parameters within the quantile-parameterized model. Hybrid elicitation, similar to predictive elicitation, deals with observable quantities. However, like parametric (structural) elicitation, it ultimately results in characterizing the uncertainty in the model parameters. Thus, hybrid elicitation can be seen as an observation-level parametric elicitation specifically designed for for quantile-parameterized models.

Supplemental materials

Supplemental materials contain the R and Stan code for all examples used in the article. Appendix A includes the details of metalog distribution. Appendix B provides the details of the QDirichlet-metalog model with weakly informative prior. Appendix C includes the results of Simulation-Based Calibration for both models discussed in the paper.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11222-023-10325-0.

Acknowledgements The authors have no conflict of interest to declare. D Perepolkin and U Sahlin were funded by the strategic research environment Biodiversity and Ecosystem Services in Changing Climate (BECC). U Sahlin was also funded by the Crafoord Foundation (ref 20200626). B Goodrich is supported by U.S. National Science Foundations grants 2051246 and 2153019.

Author Contributions DP wrote the main manuscript text, BG provided substantial assistance with Stan implementation of quantile-based inference. US supervised the project and provided key input in model specification phase. All authors reviewed the manuscript.

Funding Open access funding provided by Lund University.

Data availability The qpd R package used in this paper is available on Github at https://github.com/dmi3kno/qpd. Contact corresponding author Dmytro Perepolkin (dmytro.perepolkin@cec.lu.se) for requests for data.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

Abbas, A.E., Budescu, D.V., Yu, H.T., et al.: A comparison of two probability encoding methods: fixed probability vs. fixed variable values. Decis. Anal. **5**(4), 190–202 (2008). https://doi.org/10.1287/deca.1080.0126

Akbarov, A.: Probability elicitation: Predictive approach. PhD thesis, University of Salford (2009)

Balakrishnan, N.: Continuous Multivariate Distributions. In: Wiley StatsRef: Statistics Reference Online. American Cancer Society, (2014) https://doi.org/10.1002/9781118445112.stat01249

Bockting, F., Radev, S.T., Bürkner, P.C.: Simulation-Based Prior Knowledge Elicitation for Parametric Bayesian Models. (2023) https://doi.org/10.48550/arXiv.2308.11672

Bürkner, P.C., Charpentier, E.: Modelling monotonic effects of ordinal predictors in Bayesian regression models. Br. J. Math. Stat. Psychol. **73**(3), 420–451 (2020). https://doi.org/10.1111/bmsp.12195

Coles, S.G., Tawn, J.A.: A Bayesian analysis of extreme rainfall data. J. Roy. Stat. Soc.: Ser. C (Appl. Stat.) **45**(4), 463–478 (1996). https://doi.org/10.2307/2986068

Cook, S.R., Gelman, A., Rubin, D.B.: Validation of software for Bayesian models using posterior quantiles. J. Comput. Graph. Stat. **15**(3), 675–692 (2006). https://doi.org/10.1198/106186006X136976

Crowder, M.: Bayesian priors based on a parameter transformation using the distribution function. Ann. Inst. Stat. Math. **44**(3), 405–416 (1992)

Elfadaly, F.G., Garthwaite, P.H.: Eliciting Dirichlet and Connor–Mosimann prior distributions for multinomial models. TEST **22**(4), 628–646 (2013). https://doi.org/10.1007/s11749-013-0336-4

Faber, I., Jung, J.: Rmetalog: the metalog distribution (2021)

Gabry, J., Češnovar, R.: Cmdstanr: R interface to 'CmdStan' (2022)

Gabry, J., Simpson, D., Vehtari, A., et al.: Visualization in Bayesian workflow. J. R. Stat. Soc. A. Stat. Soc. **182**(2), 389–402 (2019)

Gelman, A., Carlin, J.B., Stern, H.S., et al.: Bayesian data analysis. CRC Press, Cambridge (2013)

Gigerenzer, G.: What are natural frequencies? BMJ **343**, d6386 (2011). https://doi.org/10.1136/bmj.d6386

Hadlock, C.C.: Quantile-parameterized methods for quantifying uncertainty in decision analysis. PhD thesis, University of Texas, Austin, TX, (2017) https://doi.org/10.15781/T2F18SX41

Hadlock, C.C., Bickel, J.E.: Johnson Quantile-Parameterized Distributions. Decis. Anal. **14**(1), 35–64 (2017)

Hadlock, C.C., Bickel, J.E.: The generalized Johnson quantile-parameterized distribution system. Decis. Anal. **16**(1), 67–85 (2019). https://doi.org/10.1287/deca.2018.0376

Hanea, A.M., Hemming, V., Nane, G.F.: Uncertainty quantification with experts present status and research needs. Risk Anal. (2021). https://doi.org/10.1111/risa.13718

Hartmann, M., Agiashvili, G., Bürkner, P., et al.: Flexible prior elicitation via the prior predictive distribution. (2020) arXiv:2002.09868 [stat]

Hoffrage, U., Gigerenzer, G., Krauss, S., et al.: Representation facilitates reasoning: what natural frequencies are and what they are not. Cognition **84**(3), 343–352 (2002). https://doi.org/10.1016/S0010-0277(02)00050-1

Hoffrage, U., Krauss, S., Martignon, L., et al.: Natural frequencies improve Bayesian reasoning in simple and complex inference tasks. Front. Psychol. **6**, 1473 (2015)

Johnson, N.L., Kotz, S., Balakrishnan, N.: Discrete multivariate distributions. Wiley series in probability and statistics. Wiley, New York (1997)

Kadane, J., Wolfson, L.J.: Experiences in elicitation. J. Royal Stat. Soci.: Series D (The Stat.) **47**(1), 3–19 (1998). https://doi.org/10.1111/1467-9884.00113

Kadane, J.B.: Predictive and structural methods for eliciting prior distributions. Bayesian Anal. Econ. Stat. 18 (1980)

Keelin, T.W.: The metalog distributions. Decis. Anal. **13**(4), 243–277 (2016). https://doi.org/10.1287/deca.2016.0338

Keelin, T.W., Powley, B.W.: Quantile-parameterized distributions. Decis. Anal. **8**(3), 206–219 (2011). https://doi.org/10.1287/deca.1110.0213

Keeney, R., von Winterfeldt, D.: Eliciting probabilities from experts in complex technical problems. IEEE Trans. Eng. Manage. **38**(3), 191–201 (1991). https://doi.org/10.1109/17.83752

King, R.A.R.: New distributional fitting methods applied to the generalised [lambda] distribution. PhD thesis, Queensland University of Technology, Australia (1999)

Knight, F.H.: Risk, Uncertainty and Profit, vol 31. Houghton Mifflin (1921)

Manderson, A.A., Goudie, R.J.B.: Translating predictive distributions into informative priors. (2023). https://doi.org/10.48550/arXiv.2303.08528

Mikkola, P., Martin, O.A., Chandramouli, S., et al.: Prior knowledge elicitation: the past, present, and future. Bayesian Anal. **1**(1), 1–33 (2023). https://doi.org/10.1214/23-BA1381

Modrák, M., Moon, A.H., Kim, S., et al.: Simulation-based calibration checking for bayesian computation: the choice of test quantities shapes sensitivity, (2022). https://doi.org/10.48550/arXiv.2211.02383

Morgan, M.G.: Use (and abuse) of expert elicitation in support of decision making for public policy. Proc. Natl. Acad. Sci. **111**(20), 7176–7184 (2014)

Morgan, M.G., Henrion, M., Small, M.: Uncertainty: a guide to dealing with uncertainty in quantitative risk and policy analysis. Cambridge University Press, Cambridge (1990)

Myerson, R.B.: Probability models for economic decisions. Duxbury applied series. Thomson/Brooke/Cole, Belmont, CA (2005)

Nair, N.U., Sankaran, P.G., Dileepkumar, M.: Bayesian inference in quantile functions. Commun. Stat. - Theory Methods (2020). https://doi.org/10.1080/03610926.2020.1827430

Patel, J.K., Read, C.B.: Handbook of the normal distribution, vol. 150. CRC Press, Cambridge (1996)

Perepolkin, D.: Qpd: tools for quantile-parameterized distributions (2019)

Perepolkin, D., Goodrich, B., Sahlin, U.: The tenets of quantile-based inference in Bayesian models. Comput. Stat. Data Anal. **187**(107), 795 (2023). https://doi.org/10.1016/j.csda.2023.107795

Perepolkin, D., Lindström, E., Sahlin, U.: Quantile-parameterized distributions for expert knowledge elicitation. (2023b) https://doi.org/10.31219/osf.io/tq3an

Powley, B.W.: Quantile function methods for decision analysis. PhD thesis, Stanford University, Paolo Alto, CA (2013)

Pratt, J.W., Raiffa, H., Schlaifer, R.: Introduction to statistical decision theory. MIT press, Cambridge (1995)

Rayner, G.D., MacGillivray, H.L.: Numerical maximum likelihood estimation for the g-and-k and generalized g-and-h distributions. Stat. Comput. **12**(1), 57–75 (2002)

Spetzler, C.S., Staël Von Holstein, C.A.S.: Probability encoding in decision analysis. Manage. Sci. **22**(3), 340–358 (1975)

Talts, S., Betancourt, M., Simpson, D., et al.: Validating Bayesian inference algorithms with simulation-based calibration. (2020) https://doi.org/10.48550/arXiv.1804.06788

Vega-Yon, G., Marjoram, P.: Fmcmc: a friendly mcmc framework. J. Open Source Softw. (2019). https://doi.org/10.21105/joss.01427

Vihola, M.: Robust adaptive Metropolis algorithm with coerced acceptance rate. Stat. Comput. **22**(5), 997–1008 (2012)

Wasserman, L.: All of nonparametric statistics. Springer texts in statistics. Springer, New York (2006)

Wilson, E.: Fitting a modified Connor–Mosimann distribution to elicited quantiles of multinomial probabilities (2017)

Winkler, R.L.: The assessment of prior distributions in Bayesian analysis. J. Am. Stat. Assoc. **62**(319), 776–800 (1967). https://doi.org/10.1080/01621459.1967.10500894

Winkler, R.L.: Prior information, predictive distributions, and Bayesian model-building. Bayesian analysis in econometrics and statistics, pp. 95–109. North-Holland Publishing Company, Amsterdam (1980)